Edinburgh Research Explorer

# Database Queries that Explain their Work

**Link:**
Link to publication record in Edinburgh Research Explorer

**Document Version:**
Preprint (usually an early version)

OPEN ACCESS

# Database Queries that Explain their Work

James Cheney

University of Edinburgh
jcheney@inf.ed.ac.uk

Amal Ahmed

Northeastern University
amal@ccs.neu.edu

Umut A. Acar

Carnegie Mellon University &
INRIA-Rocquencourt
umut@cs.cmu.edu

## Abstract

Provenance for database queries or scientific workflows is often motivated as providing *explanation*, increasing understanding of the underlying data sources and processes used to compute the query, and *reproducibility*, the capability to recompute the results on different inputs, possibly specialized to a part of the output. Many provenance systems claim to provide such capabilities; however, most lack formal definitions or guarantees of these properties, while others provide formal guarantees only for relatively limited classes of changes. Building on recent work on provenance traces and slicing for functional programming languages, we introduce a detailed tracing model of provenance for multiset-valued Nested Relational Calculus, define trace slicing algorithms that extract subtraces needed to explain or recompute specific parts of the output, and define query slicing and differencing techniques that support explanation. We state and prove correctness properties for these techniques and present a proof-of-concept implementation in Haskell.

*Keywords*    provenance, database queries, slicing

## 1. Introduction

Over the past decade, the use of complex computer systems in science has increased dramatically: databases, scientific workflow systems, clusters, and cloud computing based on frameworks such as MapReduce [16] or PigLatin [26] are now routinely used for scientific data analysis. With this shift to computational science based on (often) unreliable components and noisy data comes decreased transparency, and an increased need to understand the results of complex computations by auditing the underlying processes.

This need has motivated work on *provenance* in databases, scientific workflow systems, and many other settings [7, 25]. There is now a great deal of research on extending such systems with rich provenance-tracking features. Generally, these systems aim to provide high-level explanations intended to aid the user in understanding how a computation was performed, by recording and presenting additional "trace" information.

Over time, two distinct approaches to provenance have emerged: (1) the use of *annotations* propagated through database queries to illustrate *where-provenance* linking results to source data [7], *lineage* or *why-provenance* linking result records to sets of witnessing input records [15], or *how-provenance* describing how results were produced via algebraic expressions [20], and (2) the use of graphical *provenance traces* to illustrate how workflow computations construct final results from inputs and configuration parameters [5, 21, 29]. However, to date few systems formally specify the semantics of provenance or give formal guarantees characterizing how provenance "explains" results.

For example, scientists often conduct parameter sweeps to search for interesting results. The provenance trace of such a computation may be large and difficult to navigate. Once the most promising results have been identified, a scientist may want to extract just that information that is needed to explain the result, without showing all of the intermediate search steps or uninteresting results. Conversely, if the results are counterintuitive, the scientist may want to identify the underlying data that contributed to the anomalous result. Missier et al. [24] introduced the idea of a "golden trail", or a subset of the provenance trace that explains, or allows reproduction of, a high-value part of the output. They proposed techniques for extracting "golden trails" using recursive Datalog queries over provenance graphs; however, they did not propose definitions of correctness or reproducibility.

It is a natural question to ask how we know when a proposed solution, such as Missier et al.'s "golden trail" queries, correctly explains or can be used to correctly reproduce the behavior of the original computation. It seems to have been taken for granted that simple graph traversals suffice to at least overapproximate the desired subset of the graph. As far as we know, it is still an open question how to define and prove such correctness properties for most provenance techniques. In fact, these properties might be defined and formalized in a number of ways, reflecting different modeling choices or requirements. In any case, in the absence of clear statements and proofs of correctness, claims that different forms of provenance "explain" or allow "reproducibility" are difficult to evaluate objectively.

The main contribution of this paper is to formalize and prove the correctness of an approach to fine-grained provenance for database queries. We build on our approach developed in prior work, which we briefly recapitulate. Our approach is based on analogies between the goals of provenance tracking for databases and workflows, and those of classical techniques for program comprehension and analysis, particularly *program slicing* [31] and *information flow* [28]. Both program slicing and information flow rely critically on notions of *dependence*, such as the familiar control-flow and data-flow dependences in programming languages.

We previously introduced a provenance model for NRC (including difference and aggregation operations) called *dependency provenance* [12], and showed how it can be used to compute *data slices*, that is, subsets of the input to the query that include all of the information relevant to a selected part of the output. Some other forms of provenance for database query languages, such as how-provenance [20], satisfy similar formal guarantees that can be used to predict how the output would change under certain classes of input changes, specifically those expressible by semiring homomorphisms. For example, Amsterdamer et al.'s system [3] is based on the semiring provenance model, so the effects of deletions on parts of the output can be predicted by inspecting their provenance, but other kinds of changes are not supported.

More recently, we proposed an approach to provenance called *self-explaining computation* [11] and explored it in the context of a general-purpose functional programming language [1, 27]. In this approach, detailed execution traces are used as a form of

provenance. Traces explain results in the sense that they can be *replayed* to recompute the results, and they can be *sliced* to obtain smaller traces that provide more concise explanations of parts of the output. Trace slicing also produces a slice of the input showing what was needed by the trace to compute the output. Moreover, other forms of provenance can be extracted from traces (or slices), and we also showed that traces can be used to compute program slices efficiently through lazy evaluation. Finally, we showed how traces support *differential slicing* techniques that can highlight the differences between program runs in order to explain and precisely localize bugs in the program or errors in the input data.

Our long-term vision is to develop self-explaining computation techniques covering all components used in day-to-day scientific practice. Databases are probably the single most important such component. Since our previous work already applies to a general-purpose programming language, one way to proceed would be to simply implement an interpreter for NRC in this language, and inherit the slicing behavior from that. However, without some further inlining or optimization, this naive strategy would yield traces that record both the behavior of the NRC query and its interpreter, along with internal data structures and representation choices whose details are (intuitively) irrelevant to understanding the high-level behavior of the query.

### 1.1 Technical overview

In this paper, we develop a tracing semantics and trace slicing techniques tailored to NRC (over a multiset semantics). This semantics evaluates a query $Q$ over an input database (i.e. environment $\gamma$ mapping relation names to table values), yielding the usual result value $v$ as well as a *trace* $T$. Traces are typically large and difficult to decipher, so we consider a scenario where a user has run $Q$, inspected the results $v$, and requests an explanation for a part of the result, such as a field value of a single record. As in our previous work for functional programs, we use partial values with "holes" $\square$ to describe parts of the output that are to be explained. For example, if the result of a program is just a pair $(1, 2)$ then the pattern $(1, \square)$ can be used to request an explanation for just the first component. Given a partial value $p$ matching the output, our approach computes a "slice" consisting of a partial trace and a partial input environment, where components not necessary for recomputing the explained output part $p$ have been deleted.

The main technical contribution of this paper over our previous work [1, 27] is its treatment of tracing and slicing for collections. There are two underlying technical challenges; we illustrate both (and our solutions) via a simple example query $Q = \sigma_{A<B}(R) \cup \rho_{A\mapsto B, B\mapsto A}(\sigma_{A\geq B}(R))$ over a table $R$ with attributes $A, B$. Here, $\sigma_\phi$ is relational selection of all tuples satisfying a predicate $\phi$ and $\rho_{A\mapsto B, B\mapsto A}$ is renaming. Thus, $Q$ simply swaps the fields of records where $A \geq B$, and leaves other records alone.

The first challenge is how to address elements of multisets reliably across different executions and support propagation of addresses in the output backwards towards the input. Our solution is to use a mildly enriched semantics in which multiset elements carry explicit labels; that is, we view multisets of elements from $X$ as functions $I \to X$ from some index set to $X$. For example, if $R$ is labeled as follows and we use this enriched semantics to evaluate the above query $Q$ on $R$, we get a result:

$$R = \begin{array}{c|ccc} id & A & B & C \\ \hline [r_1] & 1 & 2 & 7 \\ [r_2] & 2 & 3 & 8 \\ [r_3] & 4 & 3 & 9 \end{array} \quad Q(R) = \begin{array}{c|ccc} id & A & B & C \\ \hline [1,r_1] & 1 & 2 & 7 \\ [1,r_2] & 2 & 3 & 8 \\ [2,r_3] & 3 & 4 & 9 \end{array}$$

where in each case the $id$ column contains a distinct index $r_i$. In $Q(R)$, the first '1' or '2' in each index indicates whether the row was generated by the left or right subexpression in the union ('$\cup$').

In general, we use sequences of natural numbers $[i_1, \ldots, i_n] \in \mathbb{N}^*$ as indices, and we maintain a stronger invariant: the set of indexes used in a multiset must form a *prefix code*. We define a semantics for NRC expressions over such collections that is fully deterministic and does not resort to generation of fresh intermediate labels; in particular, the union operation adjusts the labels to maintain distinctness.

The second technical challenge involves extending *patterns* for partial collections. In our previous work, any subexpression of a value can be replaced by a hole. This works well in a conventional functional language, where typical values (such as lists and trees) are essentially initial algebras built up by structural induction. However, when we consider unordered collections such as bags, our previous approach becomes awkward.

For example, if we want to use a pattern to focus on only the $B$ field of the second record in query result $Q(R)$, we can only do this by deleting the other record values, and the smallest such pattern is $\{[1, r_1].\square, [2, r_2].\langle A{:}\square, B{:}3, C{:}\square\rangle, [2, r_3].\square\}$. This is tolerable if there are only a few elements, but if there are hundreds or millions of elements and we are only interested in one, this is a significant overhead. Therefore, we introduce *enriched patterns* that allow us to replace entire subsets with holes, for example, $p' = \{[2, r_2].\langle B{:}3; \square\rangle\} \dot\cup \square$. Enriched patterns are not just a convenience; we show experimentally that they allow traces and slices to be smaller (and computed faster) by an order of magnitude or more.

### 1.2 Outline

The rest of this paper is structured as follows. Section 2 reviews the (multiset-valued) Nested Relational Calculus and presents our tracing semantics and deterministic labeling scheme. Section 3 presents trace slicing, including a simple form of patterns. Section 4 shows how to enrich our pattern language to allow for partial record and partial set patterns, which considerably increase the expressiveness of the pattern language, leading to smaller slices. Section 5 presents the query slicing algorithm and shows how to compute differential slices. Section 6 presents additional examples and discussion. Section 7 presents our implementation demonstrating the benefits of laziness and enriched patterns for trace slicing. Section 8 discusses related and future work and Section 9 concludes.

Due to space limitations, and in order to make room for examples and high-level discussion, some (mostly routine) formal details and proofs are relegated to the appendix.

## 2. Traced Evaluation for NRC

The nested relational calculus [9] is a simply-typed core language with collection types that can express queries on nested data similar to those of SQL on flat relations, but has simpler syntax and cleaner semantics. In this section, we show how to extend the ideas and machinery developed in our previous work on traces and slicing for functional languages [27] to NRC. Developing formal foundations for tracking provenance in the presence of unordered collections presents a number of challenges not encountered in the functional programming setting, as we will explain.

### 2.1 Syntax and Dynamic Semantics

Figure 1 presents the abstract syntax of NRC expressions, values, and traces. The expression $\emptyset$ denotes the empty collection, $\{e\}$ constructs a singleton collection, and $e_1 \cup e_2$ takes the (multiset) union of two collections. The operation $\texttt{sum } e$ computes the sum of a collection of integers, while the predicate $\texttt{empty } e$ tests whether the collection denoted by $e$ is empty. Additional aggregation operations such as count, maximum and average can easily be accommodated. Finally, the comprehension operation $\bigcup\{e' \mid x \in e\}$ iterates over

$$\boxed{\gamma, e \Downarrow v, T}$$

$$\frac{}{\gamma, \mathtt{c} \Downarrow \mathtt{c}, \mathtt{c}} \qquad \frac{\gamma, e_1 \Downarrow \mathtt{c}_1, T_1 \quad \cdots \quad \gamma, e_n \Downarrow \mathtt{c}_n, T_n}{\gamma, \mathtt{f}(e_1,\ldots,e_n) \Downarrow \hat{\mathtt{f}}(\mathtt{c}_1,\ldots,\mathtt{c}_n), \mathtt{f}(T_1,\ldots,T_n)} \qquad \frac{}{\gamma, x \Downarrow \gamma(x), x} \qquad \frac{\gamma, e_1 \Downarrow v_1, T_1 \quad \gamma[x \mapsto v_1], e_2 \Downarrow v_2, T_2}{\gamma, \mathtt{let}\ x = e_1\ \mathtt{in}\ e_2 \Downarrow v_2, \mathtt{let}\ x = T_1\ \mathtt{in}\ T_2}$$

$$\frac{\gamma, e_1 \Downarrow v_1, T_1 \quad \cdots \quad \gamma, e_n \Downarrow v_n, T_n}{\gamma, \langle A_1{:}e_1,\ldots,A_n{:}e_n\rangle \Downarrow \langle A_1{:}v_1,\ldots,A_n{:}v_n\rangle, \langle A_1{:}T_1,\ldots,A_n{:}T_n\rangle} \qquad \frac{\gamma, e \Downarrow \langle A_1{:}v_1,\ldots,A_n{:}v_n\rangle, T}{\gamma, e.A_i \Downarrow v_i, T.A_i}$$

$$\frac{\gamma, e \Downarrow \mathtt{true}, T \quad \gamma, e_1 \Downarrow v_1, T_1}{\gamma, \mathtt{if}(e,e_1,e_2) \Downarrow v_1, \mathtt{if}(T,e_1,e_2) \triangleright_{\mathtt{true}} T_1} \qquad \frac{\gamma, e \Downarrow \mathtt{false}, T \quad \gamma, e_2 \Downarrow v_2, T_2}{\gamma, \mathtt{if}(e,e_1,e_2) \Downarrow v_2, \mathtt{if}(T,e_1,e_2) \triangleright_{\mathtt{false}} T_2} \qquad \frac{}{\gamma, \emptyset \Downarrow \emptyset, \emptyset} \qquad \frac{\gamma, e \Downarrow v, T}{\gamma, \{e\} \Downarrow \{\epsilon.v\}, \{T\}}$$

$$\frac{\gamma, e_1 \Downarrow v_1, T_1 \quad \gamma, e_2 \Downarrow v_2, T_2}{\gamma, e_1 \cup e_2 \Downarrow 1 \cdot v_1 \uplus 2 \cdot v_2, T_1 \cup T_2} \qquad \frac{\gamma, e \Downarrow v, T \quad \gamma, x \in v, e' \Downarrow^* v', \Theta}{\gamma, \bigcup\{e' \mid x \in e\} \Downarrow v', \bigcup\{e' \mid x \in T\} \triangleright \Theta} \qquad \frac{\gamma, e \Downarrow \{\ell_1.v_1,\ldots,\ell_n.v_n\}, T}{\gamma, \mathtt{sum}\ e \Downarrow v_1 \hat{+} \ldots \hat{+} v_n, \mathtt{sum}\ T}$$

$$\frac{\gamma, e \Downarrow v, T \quad v = \emptyset}{\gamma, \mathtt{empty}\ e \Downarrow \mathtt{true}, \mathtt{empty}\ T} \qquad \frac{\gamma, e \Downarrow v, T \quad v \neq \emptyset}{\gamma, \mathtt{empty}\ e \Downarrow \mathtt{false}, \mathtt{empty}\ T}$$

$$\boxed{\gamma, x \in v, e \Downarrow^* v, \Theta}$$

$$\frac{}{\gamma, x \in \emptyset, e \Downarrow^* \emptyset, \emptyset} \qquad \frac{\gamma, x \in v_1, e \Downarrow^* v_1', \Theta_1 \quad \gamma, x \in v_2, e \Downarrow^* v_2', \Theta_2}{\gamma, x \in v_1 \uplus v_2, e \Downarrow^* v_1' \uplus v_2', \Theta_1 \uplus \Theta_2} \qquad \frac{\gamma[x \mapsto v], e \Downarrow v', T}{\gamma, x \in \{\ell.v\}, e \Downarrow^* \ell \cdot v', \{\ell.T\}}$$

**Figure 2.** Traced evaluation.

| Operations | $f ::= + \mid - \mid * \mid / \mid = \mid < \mid \leq \mid \cdots$ |
|---|---|
| Expressions | $e ::= c \mid f(e_1,\ldots,e_n) \mid x \mid \mathtt{let}\ e = x\ \mathtt{in}\ e'$ |
| | $\mid\ \langle A_1 : e_1,\ldots,A_n : e_n\rangle \mid e.A \mid \mathtt{if}(e,e',e'')$ |
| | $\mid\ \emptyset \mid \{e\} \mid e_1 \cup e_2 \mid \bigcup\{e' \mid x \in e\}$ |
| | $\mid\ \mathtt{empty}\ e \mid \mathtt{sum}\ e \mid \cdots$ |
| Labels | $\ell ::= \ell.\ell' \mid \epsilon \mid n$ |
| Values | $v ::= c \mid \langle A_1 : v_1,\ldots A_n : v_n\rangle$ |
| | $\mid\ \emptyset \mid \{\ell_1.v_1,\ldots,\ell_n.v_n\}$ |
| Environments | $\gamma ::= [x_1 \mapsto v_1,\ldots,x_n \mapsto v_n]$ |
| Traces | $T ::= \cdots \mid \mathtt{if}(T,e',e'') \triangleright_{\mathtt{true}} T \mid \mathtt{if}(T,e',e'') \triangleright_{\mathtt{false}} T$ |
| | $\mid\ \bigcup\{e \mid x \in T\} \triangleright \Theta$ |
| Trace Sets | $\Theta ::= \{\ell_1.T_1,\ldots,\ell_n.T_n\}$ |
| Types | $\tau ::= \mathtt{int} \mid \mathtt{bool} \mid \langle A_1 : \tau_1,\ldots,A_n : \tau_n\rangle \mid \{\tau\}$ |
| Type Contexts | $\Gamma ::= x_1 : \tau_1,\ldots,x_n : \tau_n$ |

**Figure 1.** NRC expressions, values, traces, and types.

the collection obtained by evaluating $e$, evaluating $e'(x)$ with $x$ bound to each element of the collection in turn, and returning a collection containing the union of all of the results. We sometimes consider pairs $(e_1, e_2)$, a special case of records $\langle \#_1 : e_1, \#_2 : e_2 \rangle$ using two designated field names $\#_1$ and $\#_2$. Many trace forms are similar to those for expressions; only the differences are shown.

Labels are sequences $\ell = [i_1,\ldots,i_n] \in \mathbb{N}^*$, possibly empty. The empty sequence is written $\epsilon$, and labels can be concatenated $\ell \cdot \ell'$; concatenation is associative. Record field names are written $A, B, A_1, A_2, \ldots$.

Values in NRC include constants $c$, which we assume include at least booleans and integers. Record values are essentially partial functions from field names to values, written $\langle A_1 : v_1,\ldots,A_n : v_n \rangle$. Collection values are essentially partial, finite-domain functions from labels in $\mathbb{N}^*$ to values, which we write $\{\ell_1.v_1,\ldots,\ell_n.v_n\}$. Since they denote functions, collections and records are identified up to reordering of their elements, and their field names or labels are always distinct. We write $\ell \cdot v$ for the operation that prepends $\ell$ to each of the labels in a set $v$, that is,

$$\ell \cdot \{\ell_1.v_1,\ldots,\ell_n.v_n\} = \{\ell \cdot \ell_1.v_1,\ldots,\ell \cdot \ell_n.v_n\}.$$

Other operations on labels and labeled collections will be introduced in due course.

The labels on the elements of a collection provide us with a persistent address for a particular element of the collection. This capability is essential when asking and answering provenance queries about parts of the source or output data, and when tracking fine-grained dependencies.

Both expressions and traces are subject to a type system. NRC types include collection types $\{\tau\}$ which are often taken to be sets, bags (multisets), or lists, though in this paper, we consider multiset collections only. However, types do not play a significant role in this paper so the typing rules are omitted. For expressions, the typing judgment $\Gamma \vdash e : \tau$ is standard and the typing rules for trace well-formedness $\Gamma \vdash T : \tau$ are presented in Appendix A.

***Traced evaluation*** NRC traces include a trace form corresponding to each of the expressions described above. The structure of the traces is best understood by inspecting the typing rules (Figure 12) and the traced evaluation rules (Figure 2), which define a judgment $\gamma, e \Downarrow v, T$ indicating that evaluating an expression $e$ in environment $\gamma$ yields a value $v$ and a trace $T$. We assume an environment $\Sigma$ associating constants and function symbols with their types, and write $\hat{f}$ or $\hat{+}$ for the semantic operations corresponding to $f$ or $+$, and so on. In most cases, the trace form is similar to the expression form; for example the trace of a constant or variable is a constant trace $c$, the trace of a primitive operation $f(e_1,\ldots,e_n)$ is a primitive operation trace $f(T_1,\ldots,T_n)$ applied to the traces $T_i$ of the arguments $e_i$, the trace of a record expression is a trace record constructor $\langle A_1 : T_1,\ldots,A_n : T_n \rangle$, and the trace of a field projection $e.A$ is a trace $T.A$. Also, the trace of a let-binding is a let-binding trace $\mathtt{let}\ x = T_1\ \mathtt{in}\ T_2$, where $x$ is bound in $T_2$. In these cases, the traces mimic the expression structure.

The traced evaluation rules for conditionals illustrate that traces differ from expressions in recording control flow decisions. The trace of a conditional is a conditional trace $\mathtt{if}(T,e_1,e_2) \triangleright_b T'$ where $T$ is the trace of the conditional test, $b$ is the Boolean value of the test $e_1$, and $T'$ is the trace of the taken branch. The expressions $e_1$ and $e_2$ are not strictly necessary but retained to preserve structural similarity to the original expression.

The trace of $\emptyset$ is a constant trace $\emptyset$. To evaluate a singleton-collection constructor $\{e\}$, we evaluate $e$ to obtain a value $v$ and return the singleton $\{\epsilon.v\}$ with empty label $\epsilon$. We return the sin-

gleton trace $\{T\}$ recording the trace for the evaluation of the element. To evaluate the union of two expressions, we evaluate each one and take the semantic union (written $\uplus$) of the resulting collections, with a '1' or '2' concatenated onto the beginning of each label to reflect whether each element came from the first or second part of the union; the union trace $T_1 \cup T_2$ records the traces for the evaluation of the two subexpressions. For $\mathtt{sum}\ e$, evaluating $e$ yields a collection of numbers whose sum we return, together with a sum trace $\mathtt{sum}\ T$ recording the trace for evaluation of $e$. Evaluation of emptiness tests $\mathtt{empty}\ e$ is analogous, yielding a trace $\mathtt{empty}\ T$.

To evaluate a comprehension $\bigcup\{e' \mid x \in e\}$, we first evaluate $e$, which yields a collection $v$ and trace $T$, and then (using auxiliary judgment $\gamma, x \in v, e' \Downarrow^* v', \Theta$) evaluate $e'$ repeatedly with $x$ bound to each element $v_i$ of the collection $v$ to get resulting values $v_i'$ and corresponding traces $T_i'$. We return a new collection $v' = \{\ell_1 \cdot v_1', \ldots, \ell_n \cdot v_n'\}$; similarly we return a labeled set of traces $\Theta = \{\ell_1.T_1, \ldots, \ell_n.T_n\}$. (Analogously to values, trace sets are essentially finite partial functions from labels to traces). For each of these collections, we prepend the appropriate label $\ell_i$ of the corresponding input element.

A technical point of note is that the resulting trace $T_i$ may contain free occurrences of $x$. As in our trace semantics for functional programs, these variables serve as markers in $T_i$ that will be critical for the trace replay semantics. The comprehension trace records, using the notation $\bigcup\{e' \mid x \in T\} \triangleright \{\ell_1.T_1, \ldots, \ell_n.T_n\}$, that the trace $T$ was used to compute a multiset $v$, and $x$ was bound to each element $\ell_i.v_i$ in $v$ in turn, with trace $T_i$ showing how the corresponding subset of the result was computed. The comprehension trace also records the expression $e'$ and bound variable $x$, which are again not strictly necessary but preserve the structural similarity to the original expression.

At this point it is useful to provide some informal motivation for the labeling semantics, compared for example to other semantics that use annotations or labels as a form of provenance. We do not view the labels themselves as provenance; instead, they provide a useful infrastructure for traces, which do capture a form of provenance. Moreover, by calculating the label of each part of an intermediate or final result deterministically (given the labels on the input), we provide a way to reliably refer to parts of the output, which otherwise may be unaddressable in a multiset-valued semantics. This is essential for supporting compositional slicing for operations such as let-binding or comprehension, where the output of one subexpression becomes a part of the input for another.

A central point of our semantics is that evaluation preserves the property that labels uniquely identify the elements of each multiset. This naturally assumes that the labels on the input collections are distinct. In fact, a stronger property is required: evaluation preserves the property that set labels form a *prefix code*. In the following, we write $x \le y$ to indicate that sequence $x$ is a prefix of sequence $y$.

**Definition 2.1.** *A prefix code over $\Sigma$ is a set of sequences $L \subseteq \Sigma^*$ such that for every $x, y \in L$, if $x \le y$ then $x = y$. A sub-prefix code of a prefix code $L$ is a prefix code $L'$ such that for all $x \in L$ there exists $y \in L'$ such that $y \le x$. We write $L' \le L$ to indicate that $L'$ is a sub-prefix code of $L$. We say that $L$ and $L'$ are prefix-disjoint when no element of $L$ is a prefix of an element of $L'$ and vice versa.*

Let $v$ be a collection $v = \{\ell_1.v_1, \ldots, \ell_n.v_n\}$. We define the domain of $v$ to be $\mathrm{dom}(v) = \{\ell_1, \ldots, \ell_n\}$. . We say that a value or value environment is *prefix-labeled* if for every collection $v$ occurring in it, the labels $\ell_1, \ldots, \ell_n$ are distinct and $\mathrm{dom}(v)$ is a prefix code. Similarly, we say that a trace is prefix-labeled if every labeled trace set $\Theta = \{\ell_1.T_1, \ldots, \ell_n.T_n\}$ is prefix-labeled.



**Figure 3.** Trace replay.

**Theorem 2.2.** *If $\gamma$ is prefix-labeled and $\gamma, e \Downarrow v, T$ then $v$ and $T$ are both prefix-labeled. Moreover, if $\gamma$ and $v$ are prefix-labeled and $\gamma, x \in v, e \Downarrow^* v', \Theta$ then $v'$ and $\Theta$ are prefix-labeled, and in addition $\mathrm{dom}(\Theta) = \mathrm{dom}(v) \le \mathrm{dom}(v')$.*

*Proof.* By induction on derivations. The key cases are those for union, which is straightforward, and for comprehensions. For the latter case we need the second part, to show that whenever $\gamma$ and $v$ are prefix-labeled, if $\gamma, x \in v, e \Downarrow^* v', \Theta$ then $v'$ and $\Theta$ are prefix-labeled, and in addition $\mathrm{dom}(\Theta) = \mathrm{dom}(v)$ and $\mathrm{dom}(v)$ is a sub-prefix code of $\mathrm{dom}(v')$. $\square$

The prefix code property is needed later in the slicing algorithms, when we will need it to match elements of collections produced by comprehensions with corresponding elements of the trace set $\Theta$. From now on, we assume that all values, environments, and traces are prefix-labeled, so any labeled set is assumed to have the prefix code property.

We will use the following query as a running example.

$$Q = \bigcup\{\mathtt{if}(x.B = 3, \{\langle A{:}x.A, B{:}x.C\rangle\}, \{\}) \mid x \in R\}$$

This is a simple selection query; it identifies records in $R$ that have $B$-value of 3, and returns record $\langle A{:}x.A, B{:}x.C\rangle$ containing $x$'s $A$ value and its $C$ value renamed to $B$. The result of $Q$ on the input $R$ in the introduction is $Q(R) = \{[r_2].\langle A{:}2, B{:}8\rangle, [r_3].\langle A{:}4, B{:}9\rangle\}$, and the trace is:

$$T = \bigcup\{\_ \mid x \in R\} \triangleright \{ \quad [r_1].\mathtt{if}(x.B = 3, \_, \_) \triangleright_{\mathtt{false}} \{\},$$
$$[r_2].\mathtt{if}(x.B = 3, \_, \_) \triangleright_{\mathtt{true}} \{\_\},$$
$$[r_3].\mathtt{if}(x.B = 3, \_, \_) \triangleright_{\mathtt{true}} \{\_\} \quad \}$$

where _ indicates omitted (easily inferrable) subexpressions.

***Replay*** We introduce a judgment $\gamma, T \curvearrowright v$ for *replaying* a trace on a (possibly different) environment $\gamma$. The rules for replaying NRC traces are presented in Figure 3. Many of the rules are straightforward or analogous to the corresponding evaluation rules. Here we only discuss the replay rules for conditional and comprehension traces.

For the conditional rules, the basic idea is as follows. If replaying the trace $T$ of the test yields the same boolean value $b$ as recorded in the trace, we replay the trace of the taken branch. If the test yields a different value, then replay fails.

To replay a comprehension trace $\bigcup \{e \mid x \in T\} \rhd \Theta$, rule RCOMP first replays trace $T$ to update the set of elements $v$ over which we will iterate. We define a separate judgment $\gamma, x \in v, \Theta \curvearrowright^* v'$ to iterate over set $v$ and replay traces on the corresponding elements. For elements $\ell_i \in \mathrm{dom}(\Theta)$, we replay the corresponding trace $\Theta(\ell_i)$. Replay fails if $v$ contains any labels not present in $\Theta$.

Replaying a trace can fail if either the branch taken in a conditional test differs from that recorded in the trace, or the intermediate set obtained from rerunning a comprehension trace includes values whose labels are not present in the trace. This means, in particular, that changes to the input can be replayed if they only change base values or delete set elements, but changes leading to additions of new labels to sets involved in comprehensions typically cannot be replayed.

Returning to the running example, suppose we change field $B$ of row $r_1$ of $R$ from 2 to 5. This change has no effect on the control flow choice taken in $Q$, and replaying the trace $T$ succeeds. Likewise, changing the $A$ or $C$ field of any column of $R$ has no effect, since these values do not affect the control flow choices in $T$. However, changing the $B$ field of a row to 3 (or changing it from 3 to something else) means that replay will fail.

## 2.2 Key properties

Before moving on to consider trace slicing, we identify some properties that formalize the intuition that traces are consistent with and faithfully record execution.

Evaluation and traced evaluation are deterministic:

**Proposition 2.3** (Determinacy)**.** *If* $\gamma, e \Downarrow v, T$ *and* $\gamma, e \Downarrow v', T'$ *then* $v = v'$ *and* $T = T'$. *If* $\gamma, T \curvearrowright v$ *and* $\gamma, T \curvearrowright v'$ *then* $v = v'$.

When the trace is irrelevant, we write $\gamma, e \Downarrow v$ to indicate that $\gamma, e \Downarrow v, T$ for some $T$.

Traced evaluation is type-safe and produces well-typed traces, and trace replay is also type-safe.

**Theorem 2.4.** *If* $\gamma : \Gamma$ *and* $\Gamma \vdash e : \tau$ *and* $\gamma, e \Downarrow v, T$ *then* $v : \tau$ *and* $\Gamma \vdash T : \tau$. *If* $\gamma : \Gamma$ *and* $\Gamma \vdash T : \tau$ *and* $\gamma, T \curvearrowright v$ *then* $v : \tau$.

Traces can be represented using pointers to share common subexpressions; using this DAG representation, traces can be stored in space polynomial in the input. (This sharing happens automatically in our implementation in Haskell.)

**Proposition 2.5.** *For a fixed* $e$*, if* $\gamma, e \Downarrow v, T$ *then the sizes of* $v$ *and of the DAG representation of* $T$ *are at most polynomial in* $|\gamma|$.

*Proof.* Most cases are straightforward. The only non-trivial case is for comprehensions, where we need a stronger induction hypothesis: if $\gamma, x \in v, e \Downarrow^* v', \Theta$ then the sizes of $v'$ and of the DAG representation of $\Theta$ are at most polynomial in $|\gamma|$. □

Furthermore, traced evaluation produces a trace that replays to the same value as the original expression run on the original environment. We call this property *consistency*.

**Proposition 2.6** (Consistency)**.** *If* $\gamma, e \Downarrow v, T$ *then* $\gamma, T \curvearrowright v$.

Finally, trace replay is faithful to ordinary evaluation in the following sense: if $T$ is generated by running $e$ in $\gamma$ and we successfully replay $T$ on $\gamma'$ then we obtain the same value (and same trace) as if we had rerun $e$ from scratch in $\gamma'$, and vice versa:

**Proposition 2.7** (Fidelity)**.** *If* $\gamma, e \Downarrow v, T$*, then for any* $\gamma', v'$ *we have* $\gamma', e \Downarrow v', T$ *if and only if* $\gamma', T \curvearrowright v'$.

Observe that consistency is a special case of fidelity (with $\gamma' = \gamma, v' = v$). Moreover, the "if" direction of fidelity holds even though replay can fail, because we require that $\gamma', e \Downarrow v', T$ holds for the *same trace* $T$. If $\gamma', e \Downarrow v', T'$ is derivable but only for a different trace $T'$, then replay fails.

## 3. Trace Slicing

The goal of the trace slicing algorithm we consider is to remove information from a trace and input that is not needed to recompute a part of the output. To accommodate these requirements, we introduce traces and values with *holes* and more generally, we consider *patterns* that represent relations on values capturing possible changes. In this section, we limit attention to pairs, and consider slicing for a class of *simple* patterns. We consider records and more expressive *enriched* patterns in the next section.

We extend traces with holes $\square$

$$T \quad ::= \quad \cdots \mid \square$$

and define a subtrace relation $\sqsubseteq$ that is essentially a syntactic precongruence on traces and trace sets such that $\square \sqsubseteq T$ holds and $\Theta \subseteq \Theta'$ implies $\Theta \sqsubseteq \Theta'$. (The definition of $\sqsubseteq$ is shown in full in the companion technical report.) Intuitively, holes denote parts of traces we do not care about, and we can think of a trace $T$ with holes as standing for a set of possible complete traces $\{T' \mid T \sqsubseteq T'\}$ representing different ways of filling in the holes.

The syntax of *simple patterns* $p$, *set patterns* $\mathbf{sp}$, and *pattern environments* $\rho$ is:

$$\begin{aligned}
p &::= \quad \square \mid \diamondsuit \mid \mathsf{c} \mid (p_1, p_2) \mid \mathbf{sp} \\
\mathbf{sp} &::= \quad \emptyset \mid \{\ell_1.p_1, \ldots, \ell_n.p_n\} \\
\rho &::= \quad [x_1 \mapsto p_1, \ldots, x_n \mapsto p_n]
\end{aligned}$$

Essentially, a pattern is a value with holes in some positions. The meaning of each pattern is defined through a relation $\approx_p$ that says when two values are equivalent with respect to a pattern. This relation is defined in Figure 5. A hole $\square$ indicates that the part of the value is unimportant, that is, for slicing purposes we don't care about that part of the result. Its associated relation $\approx_\square$ relates any two values. An *identity pattern* $\diamondsuit$ is similar to a hole: it says that the value is important but its exact value is not specified, and its associated relation $\approx_\diamondsuit$ is the identity relation on values. Complete set patterns $\{\ell_1.p_1, \ldots, \ell_n.p_n\}$ specify the labels and patterns for all of the elements of a set; that is, such a pattern relates only sets that have exactly the labeled elements specified and whose corresponding values match according to the corresponding patterns.

We define the union of two set patterns as $\{\overline{\ell_i.p_i}\} \uplus \{\overline{\ell'_i.p'_i}\} = \{\overline{\ell_i.p_i}, \overline{\ell'_i.p'_i}\}$ provided their domains $\vec{\ell_i}$ and $\vec{\ell'_i}$ are prefix-disjoint. We define a (partial) least upper bound operation on patterns $p \sqcup p'$ such that for any $v, v'$ we have $v \approx_{p \sqcup p'} v'$ if and only if $v \approx_p v'$ and $v \approx_{p'} v'$; the full definition is shown in Figure 4. We define the partial ordering $p \sqsubseteq p'$ as $p \sqcup p' = p'$. We say that a value $v$ *matches* pattern $p$ if $p \sqsubseteq v$. Observe that this implies $v \approx_p v$. We extend the $\sqcup$ and $\sqsubseteq$ operations to pattern environments $\rho$ pointwise, that is, $(\rho \sqcup \rho')(x) = \rho(x) \sqcup \rho(x')$.

We define several additional operations on patterns that are needed for the slicing algorithm. Consider the following *singleton*

$$\begin{aligned}
\square \sqcup p = p \sqcup \square &= p \\
\Diamond \sqcup p = p \sqcup \Diamond &= p[\Diamond/\square] \\
\mathsf{c} \sqcup \mathsf{c} &= \mathsf{c} \\
(p_1, p_2) \sqcup (p_1', p_2') &= (p_1 \sqcup p_1', p_2 \sqcup p_2') \\
\{\overline{\ell_i.p_i}\} \sqcup \{\overline{\ell_i.p_i'}\} &= \{\overline{\ell_i.p_i \sqcup p_i'}\}
\end{aligned}$$

where:

$$\begin{aligned}
\Diamond[\Diamond/\square] = \square[\Diamond/\square] &= \Diamond \\
\mathsf{c}[\Diamond/\square] &= \mathsf{c} \\
(p_1, p_2)[\Diamond/\square] &= (p_1[\Diamond/\square], p_2[\Diamond/\square]) \\
\{\overline{\ell_i.p_i}\}[\Diamond/\square] &= \{\overline{\ell_i.p_i[\Diamond/\square]}\}
\end{aligned}$$

**Figure 4.** Least upper bound for simple patterns

---

$\boxed{v \approx_p v'}$

$$\dfrac{}{v \approx_\square v'} \qquad \dfrac{}{v \approx_\Diamond v} \qquad \dfrac{}{\mathsf{c} \approx_\mathsf{c} \mathsf{c}} \qquad \dfrac{v_1 \approx_{p_1} v_1' \quad v_1 \approx_{p_2} v_2'}{(v_1, v_2) \approx_{(p_1, p_2)} (v_1', v_2')}$$

$$\dfrac{v_i \approx_{p_i} v_i' \quad (i \in \{1, \dots, n\})}{\{\ell_1.v_1, \dots, \ell_n.v_n\} \approx_{\{\ell_1.p_1, \dots, \ell_n.p_n\}} \{\ell_1.v_1', \dots, \ell_n.v_n'\}}$$

$$\gamma \approx_\rho \gamma' \iff \forall x \in dom(\rho).\ \gamma(x) \approx_{\rho(x)} \gamma'(x)$$

**Figure 5.** Simple pattern equivalence.

---

*extraction* operation $p.\epsilon$ and *label projection* operation $p[\ell]$:

$$\begin{aligned}
(\{\epsilon.p\}).\epsilon &= p & \square.\epsilon = \square \quad \Diamond.\epsilon = \Diamond \\
\mathbf{sp}[\ell] &= \{\ell'.v \mid \ell.\ell'.v \in \mathbf{sp}\} & \square[\ell] = \square \quad \Diamond[\ell] = \Diamond
\end{aligned}$$

These operations are only used for the above cases; they have no effect on constant or pair patterns. For sets, $p[\ell]$ extracts the subset of $p$ whose labels start with $\ell$, truncating the initial prefix $\ell$, while if $p$ is $\square$ or $\Diamond$ then again $p[\ell]$ returns the same kind of hole. Moreover, $dom(p[\ell])$ is a prefix code if $dom(p)$ is.

Suppose $L$ is a prefix code. We define *restriction* of a set pattern $p$ to $L$ as follows:

$$\mathbf{sp}|_L = \{\ell.\ell'.p \in \mathbf{sp} \mid \ell \in L\} \qquad \square|_L = \square \qquad \Diamond|_L = \Diamond$$

It is easy to see that $dom(p|_L) \subseteq dom(p)$ so $dom(p|_L)$ is a prefix code if $dom(p)$ is, so this operation is well-defined on collections:

**Lemma 3.1.** *If* $\mathbf{sp}$ *is prefix-labeled and* $L$ *is a prefix code then* $\mathbf{sp}[\ell]$ *and* $\mathbf{sp}|_L$ *are prefix-labeled.*

We show other properties of patterns in Appendix B.

***Backward Slicing*** The rules for backward slicing are given in Figure 6. The judgment $p, T \searrow \rho, T'$ slices trace $T$ with respect to a pattern $p$ to yield the slice $T'$ and sliced input environment $\rho$. The sliced input environment records what parts of the input are needed to produce $p$; this is needed for slicing operations such as let-binding or comprehensions. The main new ideas are in the rules for collection operations, particularly comprehensions. The slicing rules SConst, SPrim, SVar, SLet, SPair, SProj$_i$, and SIf follow essentially the same idea as in our previous work [1, 27]. We focus discussion on the new cases, but we review the key ideas for these operations here in order to make the presentation self-contained.

The rules for collections use labels and set pattern operations to effectively undo the evaluation of the set pattern. The rule SEmpty is essentially the same as the constant rule. The rule SSng uses the singleton extraction operation $p.\epsilon$ to obtain a pattern describing the single element of a singleton set value matching $p$. The rule SUnion uses the two projections $p[1]$ and $p[2]$ to obtain the

patterns describing the subsets obtained from the first and second subtraces in the union pattern, respectively.

The rule SComp uses a similar idea to let-binding. We slice the trace set $\Theta$ using an auxiliary judgment $p, x.\Theta \searrow^* \rho, \Theta_0, p_0$, obtaining a sliced input environment, sliced trace set $\Theta_0$, and pattern $p_0$ describing the set of values to which $x$ was bound. We then use $p_0$ to slice backwards through the subtrace $T$ that constructed the set. The auxiliary slicing judgment for trace sets has three rules: a trivial rule SEmpty$^*$ when the set is empty, rule SSng$^*$ that uses label projection $p[\ell]$ to handle a singleton trace set, and a rule SUnion$^*$ handling larger trace sets by decomposing them into subsets. This is essentially a structural recursion over the trace set, and is deterministic even though the rules can be used to decompose the trace sets in many different ways, because $\uplus$ is associative. Rule SUnion$^*$ also requires that we restrict the set pattern to match the domains of the corresponding trace patterns.

The slicing rules SSum and SEmptyP follow the same idea as for primitive operations at base type: we require that the whole set value be preserved exactly. As discussed by Perera et al. [27] with primitive operations, this is a potential source of overapproximation, since (for example) for an emptiness test, all we really need is to preserve the number of elements in the set, not their values. The last rule shows how to slice pair patterns when the pattern is $\Diamond$: we slice both of the subtraces by $\Diamond$ and combine the results.

It is straightforward to show that the slicing algorithm is well-defined for consistent traces. That is, if $\gamma, T \curvearrowright v$ and $p \sqsubseteq v$ then there exists $\rho, S$ such that $p, T \searrow \rho, S$ holds, where $\rho \sqsubseteq \gamma$ and $S \sqsubseteq T$. We defer the correctness theorem for slicing using simple patterns to the end of the next section, since it is a special case of correctness for slicing using enriched patterns.

Continuing our running example, consider the pattern $p = \{[r_2].\langle A{:}\square, B{:}8\rangle, [r_3].\square\}$. The slice of $T$ with respect to this pattern is of the form:

$$T' = \bigcup\{\_\mid x \in R\} \triangleright \{\quad \begin{aligned} &[r_1].\square, \\ &[r_2].\mathtt{if}(x.B = 3, \_, \_) \triangleright_{\mathtt{true}} \{\_\}, \\ &[r_3].\square \quad \} \end{aligned}$$

and the slice of $R$ is $R' = \{[r_1].\square, [r_2].\langle A{:}\square, B{:}3, C{:}8\rangle, [r_3].\square\}$. (That is, $R'$ is the value of $\rho(R)$, where $\rho$ is the pattern environment produced by slicing $T$ with respect to $p$.) Observe that the value of $A$ is not needed but the value of $B$ must remain 3 in order to preserve the control flow behavior of the trace on $r_2$. The holes indicate that changes to $r_1$ and $r_3$ in the input cannot affect $r_2$. However, a trace matching $T'$ cannot be replayed if any of $r_1, r_2, r_3$ are deleted from the input or if the $A$ field is removed from an input record, because the replay rules require all of the labels mentioned in collections or records in $T'$ to be present. We now turn our attention to enriched patterns, which mitigate these drawbacks.

## 4. Enriched Patterns

So far we have considered only complete set patterns of the form $\{\ell_1.p_1, \dots, \ell_n.p_n\}$. These patterns relate pairs of values that have exactly $n$ elements labeled $\ell_1, \dots, \ell_n$, each of which matches $p_1, \dots, p_n$ respectively.

This is awkward, as we already can observe in our running example above: to obtain a slice describing how one record was computed, we need to use a set pattern that lists all of the indexes in the output. Moreover, the slice with respect to such a pattern may also include labeled subtraces explaining why the other elements exist in the output (and no others). This information seems intuitively irrelevant to the value at $\ell_1$, and can be a major overhead if the collection is large. We also have considered only binary pairs; it would be more convenient to support record patterns directly.

$$\boxed{p, T \searrow \rho, S}$$

$$\frac{}{\Box, T \searrow [], \Box} \text{ SHole} \qquad \frac{}{p, \mathtt{c} \searrow [], \mathtt{c}} \text{ SConst} \qquad \frac{\Diamond, T_1 \searrow \rho_1, S_1 \quad \cdots \quad \Diamond, T_n \searrow \rho_n, S_n}{p, \mathtt{f}(T_1, \ldots, T_n) \searrow \rho_1 \sqcup \cdots \sqcup \rho_n, \mathtt{f}(S_1, \ldots, S_n)} \text{ SPrim} \qquad \frac{}{p, x \searrow [x \mapsto p], x} \text{ SVar}$$

$$\frac{p_2, T_2 \searrow \rho_2[x \mapsto p_1], S_2 \quad p_1, T_1 \searrow \rho_1, S_1}{p_2, \mathtt{let}\ x = T_1\ \mathtt{in}\ T_2 \searrow \rho_1 \sqcup \rho_2, \mathtt{let}\ x = S_1\ \mathtt{in}\ S_2} \text{ SLet} \qquad \frac{p_1, T_1 \searrow \rho_1, S_1 \quad p_2, T_2 \searrow \rho_2, S_2}{(p_1, p_2), (T_1, T_2) \searrow \rho_1 \sqcup \rho_2, (S_1, S_2)} \text{ SPair} \qquad \frac{(p, \Box), T \searrow \rho, S}{p, T.\#_1 \searrow \rho, S.\#_1} \text{ SProj}_1$$

$$\frac{(\Box, p), T \searrow \rho, S}{p, T.\#_2 \searrow \rho, S.\#_2} \text{ SProj}_2 \qquad \frac{p, T' \searrow \rho', S' \quad b, T \searrow \rho, S}{p, \mathtt{if}(T, e_1, e_2) \rhd_b T' \searrow \rho' \sqcup \rho, \mathtt{if}(S, e_1, e_2) \rhd_b S'} \text{ SIf} \qquad \frac{}{p, \emptyset \searrow [], \emptyset} \text{ SEmpty} \qquad \frac{p.\epsilon, T \searrow \rho, S}{p, \{T\} \searrow \rho, \{S\}} \text{ SSng}$$

$$\frac{p[1], T_1 \searrow \rho_1, S_1 \quad p[2], T_2 \searrow \rho_2, S_2}{p, T_1 \ \cup\ T_2 \searrow \rho_1 \sqcup \rho_2, S_1 \ \cup\ S_2} \text{ SUnion} \qquad \frac{p, x.\Theta \searrow^* \rho', \Theta', p' \quad p', T \searrow \rho, S}{p, \bigcup\{e \mid x \in T\} \rhd \Theta \searrow \rho \sqcup \rho', \bigcup\{e \mid x \in S\} \rhd \Theta'} \text{ SComp}$$

$$\frac{\Diamond, T \searrow \rho, S}{p, \mathtt{sum}\ T \searrow \rho, \mathtt{sum}\ S} \text{ SEmptyP} \qquad \frac{\Diamond, T \searrow \rho, S}{p, \mathtt{empty}\ T \searrow \rho, \mathtt{empty}\ S} \text{ SSum} \qquad \frac{\Diamond, T_1 \searrow \rho_1, S_1 \quad \Diamond, T_2 \searrow \rho_2, S_2}{\Diamond, (T_1, T_2) \searrow \rho_1 \sqcup \rho_2, (S_1, S_2)} \text{ SDiamond}$$

$$\boxed{p, x.\Theta \searrow^* \rho, \Theta_0, p_0}$$

$$\frac{}{\emptyset, x.\emptyset \searrow^* [], \emptyset, \emptyset} \text{ SEmpty}^* \qquad \frac{p[\ell], T \searrow \rho[x \mapsto p_0], S}{p, x.\{\ell.T\} \searrow^* \rho, \{\ell.S\}, \{\ell.p_0\}} \text{ SSng}^*$$

$$\frac{p|_{\mathrm{dom}(\Theta_1)}, x.\Theta_1 \searrow^* \rho_1, \Theta_1', p_1 \quad p|_{\mathrm{dom}(\Theta_2)}, x.\Theta_2 \searrow^* \rho_2, \Theta_2', p_2}{p, x.\Theta_1 \uplus \Theta_2 \searrow^* \rho_1 \sqcup \rho_2, \Theta_1' \uplus \Theta_2', p_1 \uplus p_2} \text{ SUnion}^*$$

**Figure 6.** Backward trace slicing.

In this section we sketch how to enrich the language of patterns to allow for *partial set* and *partial record patterns*, as follows:

$$
\begin{aligned}
p \quad &::= \quad \Box \mid \Diamond \mid \mathtt{c} \mid \mathbf{sp} \mid \mathbf{rp} \\
\mathbf{rp} \quad &::= \quad \langle\rangle \mid \langle \overline{A_i : p_i} \rangle \mid \langle \overline{A_i : p_i}; \Box \rangle \mid \langle \overline{A_i : p_i}; \Diamond \rangle \\
\mathbf{sp} \quad &::= \quad \emptyset \mid \{\overline{\ell_i.p_i}\} \mid \{\overline{\ell_i.p_i}\} \dot{\cup} \Box \mid \{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond
\end{aligned}
$$

Record patterns are of the form $\langle \overline{A_i : p_i} \rangle$, listing the fields and the patterns they must match, possibly followed by $\Box$ or $\Diamond$, which the remainder of the record must match. The pattern $\{\overline{\ell_i.p_i}\} \dot{\cup} \Box$ stands for a set with labeled elements matching patterns $p_1, \ldots, p_n$, plus some additional elements whose values we don't care about. For example, we can use the pattern $\{\ell_1.p_1\} \dot{\cup} \Box$ to express interest in why element $\ell_1$ matches $p_1$, when we don't care about the rest of the set. The second partial pattern, $\{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond$, has similar behavior, but it says that the rest of the sets being considered must be equal. For example, $\{\ell.\Box\} \dot{\cup} \Diamond$ says that the two sets are equal except possibly at label $\ell$. This pattern is needed mainly in order to ensure that we can define a least upper bound on enriched patterns, since we cannot express $(\{\ell_1.v_1, \ldots, \ell_n.p_n\} \dot{\cup} \Box) \sqcup \Diamond$ otherwise.

We define $\mathrm{dom}(\{\overline{\ell_i.p_i}\} \dot{\cup} \Box) = \mathrm{dom}(\{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond) = \{\ell_1, \ldots, \ell_n\}$. Disjoint union of enriched patterns $\mathbf{sp} \uplus \mathbf{sp}'$ is defined only if the domains are prefix-disjoint, so that the labels of the result still form a prefix code; this operation is defined in Figure 7.

We now extend the definitions of $p[\Diamond/\Box]$ and $\sqcup$ to account for extended patterns. We extend the $[\Diamond/\Box]$ substitution operation as follows:

$$
\begin{aligned}
\langle \overline{A_i : p_i} \rangle[\Diamond/\Box] \quad &= \quad \langle \overline{A_i : p_i[\Diamond/\Box]} \rangle \\
\langle \overline{A_i : p_i}; \Box \rangle[\Diamond/\Box] \quad &= \quad \langle \overline{A_i : p_i[\Diamond/\Box]}; \Diamond \rangle \\
\langle \overline{A_i : p_i}; \Diamond \rangle[\Diamond/\Box] \quad &= \quad \langle \overline{A_i : p_i[\Diamond/\Box]}; \Diamond \rangle \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Box)[\Diamond/\Box] \quad &= \quad \{\overline{\ell_i.p_i[\Diamond/\Box]}\} \dot{\cup} \Diamond \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond)[\Diamond/\Box] \quad &= \quad \{\overline{\ell_i.p_i[\Diamond/\Box]}\} \dot{\cup} \Diamond
\end{aligned}
$$

We handle the additional cases of the $\sqcup$ operation in Figure 8, and extend the $\eqsim_p$ relation as shown in Figure 9. Note that taking the least upper bound of a partial pattern with a complete pattern yields a complete pattern, while taking the least upper bound of partial patterns involving $\Diamond$ again relies on the $[\Diamond/\Box]$ substitution operation.

$$
\begin{aligned}
\Diamond \uplus \Box = \Box \uplus \Diamond = \Box \uplus \Box \quad &= \quad \Box \\
\Diamond \uplus \Diamond \quad &= \quad \Diamond \\
\mathbf{sp} \uplus \emptyset = \emptyset \uplus \mathbf{sp} \quad &= \quad \mathbf{sp} \\
\{\overline{\ell_i.p_i}\} \uplus \Box = \Box \uplus \{\overline{\ell_i.p_i}\} \quad &= \quad \{\overline{\ell_i.p_i}\} \dot{\cup} \Box \\
\{\overline{\ell_i.p_i}\} \uplus \Diamond = \Diamond \uplus \{\overline{\ell_i.p_i}\} \quad &= \quad \{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond \\
\{\overline{\ell_i.p_i}\} \uplus (\{\overline{\ell_j'.p_j'}\} \dot{\cup} \Box) \quad &= \quad (\{\overline{\ell_i.p_i}\} \uplus \{\overline{\ell_j'.p_j'}\}) \dot{\cup} \Box \\
\{\overline{\ell_i.p_i}\} \uplus (\{\overline{\ell_j'.p_j'}\} \dot{\cup} \Diamond) \quad &= \quad (\{\overline{\ell_i.p_i}\} \uplus \{\overline{\ell_j'.p_j'}\}) \dot{\cup} \Diamond \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Box) \uplus (\{\overline{\ell_j'.p_j'}\} \dot{\cup} \Box) \quad &= \quad (\{\overline{\ell_i.p_i}\} \uplus \{\overline{\ell_j'.p_j'}\}) \dot{\cup} \Box \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond) \uplus (\{\overline{\ell_j'.p_j'}\} \dot{\cup} \Box) \quad &= \quad (\{\overline{\ell_i.p_i}\} \uplus \{\overline{\ell_j'.p_j'}\}) \dot{\cup} \Box \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Box) \uplus (\{\overline{\ell_j'.p_j'}\} \dot{\cup} \Diamond) \quad &= \quad (\{\overline{\ell_i.p_i}\} \uplus \{\overline{\ell_j'.p_j'}\}) \dot{\cup} \Box \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond) \uplus (\{\overline{\ell_j'.p_j'}\} \dot{\cup} \Diamond) \quad &= \quad (\{\overline{\ell_i.p_i}\} \uplus \{\overline{\ell_j'.p_j'}\}) \dot{\cup} \Diamond
\end{aligned}
$$

**Figure 7.** Enriched pattern union

We extend the singleton extraction $p.\epsilon$, label projection $p[\ell]$, and restriction $p|_L$ operations on set patterns as follows:

$$
\begin{aligned}
(\{\epsilon.p\} \dot{\cup} \Diamond).\epsilon = (\{\epsilon.p\} \dot{\cup} \Box).\epsilon \quad &= \quad p \\
\ell \cdot (\{\overline{\ell_i.p_i}\} \dot{\cup} \Box) \quad &= \quad (\ell \cdot \{\overline{\ell_i.p_i}\}) \uplus \Box \\
\ell \cdot (\{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond) \quad &= \quad (\ell \cdot \{\overline{\ell_i.p_i}\}) \uplus \Diamond \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Box)[\ell] \quad &= \quad (\{\overline{\ell_i.p_i}\}[\ell]) \uplus \Box \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond)[\ell] \quad &= \quad (\{\overline{\ell_i.p_i}\}[\ell]) \uplus \Diamond \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Box)|_L \quad &= \quad (\{\overline{\ell_i.p_i}\}|_L) \uplus \Box \\
(\{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond)|_L \quad &= \quad (\{\overline{\ell_i.p_i}\}|_L) \uplus \Diamond
\end{aligned}
$$

Note that in many cases, we use the disjoint union operation $\uplus$ on the right-hand side; this ensures, for example, that we never produce results of the form $\emptyset \dot{\cup} \Box$ or $\emptyset \dot{\cup} \Diamond$; these are normalized to $\Box$ and $\Diamond$ respectively, and this normalization reduces the number of corner cases in the slicing algorithm.

We define a record pattern projection operation $p.A$ as follows:

$$
\begin{aligned}
\langle A_1{:}p_1, \ldots, A_n{:}p_n \rangle.A_i \quad &= \quad p_i \qquad \Box.A = \Box \\
\langle A_1{:}p_1, \ldots, A_n{:}p_n; \_ \rangle.A_i \quad &= \quad p_i \qquad \Diamond.A = \Diamond \\
\langle A_1{:}p_1, \ldots, A_n{:}p_n; \Box \rangle.B \quad &= \quad \Box \qquad (B \notin \{A_1, \ldots, A_n\}) \\
\langle A_1{:}p_1, \ldots, A_n{:}p_n; \Diamond \rangle.B \quad &= \quad \Diamond \qquad (B \notin \{A_1, \ldots, A_n\})
\end{aligned}
$$

We extend the slicing judgment to accommodate these new patterns in Figure 10. The rules SRec and SProj$_A$ are similar to those for pairs, except that we use the field projection operation in the case for a record trace, and we use partial record patterns

$$\{\overline{\ell_i.p_i}, \overline{\ell'_j.q_j}\} \sqcup (\{\overline{\ell_i.p'_i}\} \dot{\cup} \Box) \;=\; \{\overline{\ell_i.p_i \sqcup p'_i}, \overline{\ell'_j.q_j}\}$$

$$\{\overline{\ell_i.p_i}, \overline{\ell'_j.q_j}\} \sqcup (\{\overline{\ell_i.p'_i}\} \dot{\cup} \Diamond) \;=\; \{\overline{\ell_i.p_i \sqcup p'_i}, \overline{\ell'_j.q_j[\Diamond/\Box]}\}$$

$$(\{\overline{\ell_i.p_i}, \overline{\ell'_j.q_j}\} \dot{\cup} \Box) \sqcup (\{\overline{\ell_i.p'_i}, \overline{\ell''_k.r_k}\} \dot{\cup} \Box) \;=\; \{\overline{\ell_i.p_i \sqcup p'_i}, \overline{\ell'_j.q_j}, \overline{\ell''_k.r_k}\} \dot{\cup} \Box$$

$$(\{\overline{\ell_i.p_i}, \overline{\ell'_j.q_j}\} \dot{\cup} \Box) \sqcup (\{\overline{\ell_i.p'_i}, \overline{\ell''_k.r_k}\} \dot{\cup} \Diamond) \;=\; \{\overline{\ell_i.p_i \sqcup p'_i}, \overline{\ell'_j.q_j[\Diamond/\Box]}, \overline{\ell''_k.r_k}\} \dot{\cup} \Diamond$$

$$(\{\overline{\ell_i.p_i}, \overline{\ell'_j.q_j}\} \dot{\cup} \Diamond) \sqcup (\{\overline{\ell_i.p'_i}, \overline{\ell''_k.r_k}\} \dot{\cup} \Diamond) \;=\; \{\overline{\ell_i.p_i \sqcup p'_i}, \overline{\ell'_j.q_j[\Diamond/\Box]}, \overline{\ell''_k.r_k[\Diamond/\Box]}\} \dot{\cup} \Diamond$$

$$\langle \overline{A_i : p_i}, \overline{B_j : q_j} \rangle \sqcup (\langle \overline{A_i : p'_i}; \Box \rangle) \;=\; \langle \overline{A_i : p_i \sqcup p'_i}, \overline{B_j : q_j} \rangle$$

$$\langle \overline{A_i : p_i}, \overline{B_j : q_j} \rangle \sqcup (\langle \overline{A_i : p'_i}; \Diamond \rangle) \;=\; \langle \overline{A_i : p_i \sqcup p'_i}, \overline{B_j : q_j[\Diamond/\Box]} \rangle$$

$$(\langle \overline{A_i : p_i}, \overline{B_j : q_j}; \Box \rangle) \sqcup (\langle \overline{A_i : p'_i}, \overline{C_k : r_k}; \Box \rangle) \;=\; \langle \overline{A_i : p_i \sqcup p'_i}, \overline{B_j : q_j}, \overline{C_k : r_k}; \Box \rangle$$

$$(\langle \overline{A_i : p_i}, \overline{B_j : q_j}; \Box \rangle) \sqcup (\langle \overline{A_i : p'_i}, \overline{C_k : r_k}; \Diamond \rangle) \;=\; \langle \overline{A_i : p_i \sqcup p'_i}, \overline{B_j : q_j[\Diamond/\Box]}, \overline{C_k : r_k}; \Diamond \rangle$$

$$(\langle \overline{A_i : p_i}, \overline{B_j : q_j}; \Diamond \rangle) \sqcup (\langle \overline{A_i : p'_i}, \overline{C_k : r_k}; \Diamond \rangle) \;=\; \langle \overline{A_i : p_i \sqcup p'_i}, \overline{B_j : q_j[\Diamond/\Box]}, \overline{C_k : r_k[\Diamond/\Box]}; \Diamond \rangle$$

**Figure 8.** Least upper bound for enriched patterns (excluding some symmetric cases).

$$\frac{v_1 \approx_{p_1} v'_1 \quad \cdots \quad v_n \approx_{p_n} v'_n}{\langle \overline{A_i : v_i} \rangle \approx_{\overline{\langle A_i : p_i \rangle}} \langle \overline{A_i : v'_i} \rangle}$$

$$\frac{v_1 \approx_{p_1} v'_1 \quad \cdots \quad v_n \approx_{p_n} v'_n}{\langle \overline{A_i : v_i}, \overline{B_j : w_j} \rangle \approx_{\overline{\langle A_i : p_i \rangle; \Box}} \langle \overline{A_i : v'_i}, \overline{C_k : w'_k} \rangle}$$

$$\frac{v_1 \approx_{p_1} v'_1 \quad \cdots \quad v_n \approx_{p_n} v'_n}{\langle \overline{A_i : v_i}, \overline{B_j : w_j} \rangle \approx_{\overline{\langle A_i : p_i \rangle; \Diamond}} \langle \overline{A_i : v'_i}, \overline{B_j : w_j} \rangle}$$

$$\frac{v_1 \approx_{p_1} v'_1 \quad \cdots \quad v_n \approx_{p_n} v'_n}{\{\overline{\ell_i.v_i}\} \uplus v \approx_{\{\overline{\ell_i.p_i}\} \dot{\cup} \Box} \{\overline{\ell_i.v'_i}\} \uplus v'}$$

$$\frac{v_1 \approx_{p_1} v'_1 \quad \cdots \quad v_n \approx_{p_n} v'_n}{\{\overline{\ell_i.v_i}\} \uplus v \approx_{\{\overline{\ell_i.p_i}\} \dot{\cup} \Diamond} \{\overline{\ell_i.v'_i}\} \uplus v}$$

**Figure 9.** Enriched pattern equivalence

$$\boxed{p, T \searrow \rho, S}$$

$$\frac{\langle A:p; \Box \rangle, T \searrow \rho, S}{p, T.A \searrow \rho, S.A} \; \text{SRec}$$

$$\frac{p.A_1, T_1 \searrow \rho_1, S_1 \quad \cdots \quad p.A_n, T_n \searrow \rho_n, S_n}{p, \langle A_1{:}T_1, \ldots, A_n{:}T_n \rangle \searrow \rho_1 \sqcup \cdots \sqcup \rho_n, \langle A_1{:}S_1, \ldots, A_n{:}S_n \rangle} \; \text{SProj}_A$$

$$\boxed{p, x.\Theta \searrow^* \rho, \Theta', p_0}$$

$$\frac{}{\Box, x.\Theta \searrow^* [], \emptyset, \Box} \; \text{SHole}^* \qquad \frac{}{\Diamond, x.\emptyset \searrow^* [], \emptyset, \emptyset} \; \text{SDiamond}^*$$

**Figure 10.** Backward trace slicing over enriched patterns.

$\langle A : p; \Box \rangle$ in the case for a field projection trace. The added rules SHole$^*$ and SDiamond$^*$ handle the possibility that a partial pattern reduces to $\Box$ or $\Diamond$ through projection; we did not need to handle this case earlier because a simple set pattern is either a hole (which could be handled by the rule $\Box, T \searrow [], \Box$) or a complete set pattern showing all of the labels of the result.

Both simple and extended patterns satisfy a number of lemmas that are required to prove the correctness of trace slicing.

**Lemma 4.1** (Properties of union and restriction)**.**

1. If $p_1 \sqsubseteq v_1$ and $p_2 \sqsubseteq v_2$ and $v_1 \approx_{p_1} v'_1$ and $v_2 \approx_{p_2} v'_2$ then $v_1 \uplus v_2 \approx_{p_1 \uplus p_2} v'_1 \uplus v'_2$, provided all of these disjoint unions are defined.
2. If $p \sqsubseteq v_1 \uplus v_2$ and $L_1 \le \operatorname{dom}(v_1)$ and $L_2 \le \operatorname{dom}(v_2)$ and $L_1, L_2$ are prefix-disjoint, then $p|_{L_1} \sqsubseteq v_1$ and $p|_{L_2} \sqsubseteq v_2$.

**Lemma 4.2** (Projection and $\sqsubseteq$)**.**

1. If $p \sqsubseteq \{\epsilon.v\}$ then $p.\epsilon \sqsubseteq v$.
2. If $p \sqsubseteq 1 \cdot v_1 \uplus 2 \cdot v_2$ then $p[1] \sqsubseteq v_1$ and $p[2] \sqsubseteq v_2$.
3. If $p \sqsubseteq \ell \cdot v$ then $p[\ell] \sqsubseteq v$.
4. If $p \sqsubseteq \langle \overline{A_i : v_i} \rangle$ then $p.A_i \sqsubseteq v_i$.

**Lemma 4.3** (Projection and $\approx_p$)**.**

1. If $p \sqsubseteq \{\epsilon.v\}$ and $v \approx_{p.\epsilon} v'$ then $\{\epsilon.v\} \approx_p \{\epsilon.v'\}$.
2. If $p \sqsubseteq 1 \cdot v_1 \uplus 2 \cdot v_2$ and $v_1 \approx_{p[1]} v'_1$ and $v_2 \approx_{p[2]} v'_2$ then $1 \cdot v_1 \uplus 2 \cdot v_2 \approx_p 1 \cdot v'_1 \uplus 2 \cdot v'_2$.
3. If $p \sqsubseteq \ell \cdot v$ and $v \approx_{p[\ell]} v'$ then $\ell \cdot v \approx_p \ell \cdot v'$.
4. If $p \sqsubseteq \langle \overline{A_i : v_i} \rangle$ and $v_1 \approx_{p.A_1} v'_1, \ldots, v_n \approx_{p.A_n} v'_n$ then $\langle \overline{A_i : v_i} \rangle \approx_p \langle \overline{A_i : v'_i} \rangle$.

Proofs are collected in Appendix B.

We now state the key correctness property for slicing. Intuitively, it says that if we slice $T$ with respect to output pattern $p$, obtaining a slice $\rho$ and $S$, then $p$ will be reproduced on recomputation under any change to the input and trace that is consistent with the slice — formally, that means that the changed trace $T'$ must match the sliced trace $S$, and the changed input $\gamma'$ must match $\gamma$ modulo $\rho$.

**Theorem 4.4** (Correctness of Slicing)**.**

1. Suppose $\gamma, T \curvearrowright v$ and $p \sqsubseteq v$ and $p, T \searrow \rho, S$. Then for all $\gamma' \approx_\rho \gamma$ and $T' \sqsupseteq S$ such that $\gamma', T' \curvearrowright v'$ we have $v' \approx_p v$.
2. Suppose $\gamma, x \in v_0, \Theta \curvearrowright^* v$ and $p \sqsubseteq v$ and $p, x.\Theta_0 \searrow \rho, \Theta'_0, p_0$, where $\Theta_0 \subseteq \Theta$. Then for all $\gamma' \approx_\rho \gamma$ and $v'_0 \approx_{p_0} v_0$ and $\Theta' \sqsupseteq \Theta'_0$ such that $\gamma', x \in v_0, \Theta' \curvearrowright^* v'$ we have $v' \approx_p v$.

Returning to our running example, we can use the enriched pattern $p' = \{[r_2].\langle B{:}8; \Box \rangle\} \dot{\cup} \Box$ to indicate interest in the $B$ field of $r_2$, without naming the other fields of the row or the other row indexes. Slicing with respect to this pattern yields the following slice:

$$T'' = \bigcup \{\_ \mid x \in R\} \triangleright \{[r_2].\mathtt{if}(x.B = 3, \_, \_) \triangleright_{\mathtt{true}} \{\_\}\}$$

and the slice of $R$ is $R'' = \{[r_2].\langle B{:}3, C{:}8; \Box \rangle\} \dot{\cup} \Box$. This indicates that (as before) the values of $r_1$ and $r_3$ and of the $A$ field of $r_2$

are irrelevant to $r_2$ in the result; unlike $T'$, however, we can also potentially replay if $r_1$ and $r_3$ have been deleted from $R$. Likewise, we can replay if the $A$ field has been removed from a record, or if some other field such as $D$ is added. This illustrates that enriched patterns allow for smaller slices than simple patterns, with greater flexibility concerning possible updates to the input.

A natural question is whether slicing computes the (or a) smallest possible answer. Because our definition of correct slices is based on recomputation, minimal slices are not computable, by a straightforward reduction from the undecidability of minimizing dependency provenance [1, 12].

## 5. Query and Differential Slicing

We now adapt slicing techniques to provide explanations in terms of query expressions, and show how to use differences between slices to provide precise explanations for parts of the output.

### 5.1 Query slicing

Our previous work [27] gave an algorithm for extracting a *program slice* from a trace. We now adapt this idea to queries. A trace slice shows the parts of the trace that need to be replayed in order to compute the desired part of the output; similarly, a *query slice* shows the part of the query expression that is needed in order to ensure that the desired part of the output is recomputed. As with traces, we allow holes $\square$ in programs to allow deleting subexpressions, and define $\sqsubseteq$ as a syntactic precongruence such that $\square \sqsubseteq e$. We also define a least upper bound operation $e \sqcup e'$ on partial query expressions in the obvious way, so that $\square \sqcup e = e$.

We define a judgment $p, T \searrow \rho, e$ that traverses $T$ and "unevaluates" $p$, yielding a partial input environment $\rho$ and partial program $e$. The rules are illustrated in Figure 11. Many of the rules are similar to those for trace slicing; the main differences arise in the cases for conditionals and comprehensions, where we collapse the sliced expressions back into expressions, possibly inserting holes or merging sliced expressions obtained by running the same code in different ways (as in a comprehension that contains a conditional).

Again, it is straightforward to show that if $\gamma, e \Downarrow v, T$ and $p \sqsubseteq v$ then there exist $\rho, e'$ such that $p, T \searrow \rho, e'$, where $\rho \sqsubseteq \gamma$ and $e' \sqsubseteq e$. The essential correctness property for query slices is similar to that for trace slices: again, we require that rerunning any sufficiently similar query on a sufficiently similar input produces a result that matches $p$. The proof of this result is in Appendix D.

**Theorem 5.1** (Correctness of Query Slicing).

1. *Suppose $\gamma, T \curvearrowright v$ and $p \sqsubseteq v$ and $p, T \searrow \rho, e$. Then for all $\gamma' \eqsim_\rho \gamma$ and $e' \sqsupseteq e$ such that $\gamma', e' \Downarrow v'$ we have $v' \eqsim_p v$.*
2. *Suppose $\gamma, x \in v_0, \Theta \curvearrowright^* v$ and $p \sqsubseteq v$ and $p, \Theta \searrow \rho, e_0, p_0$. Then for all $\gamma' \eqsim_\rho \gamma$ and $v'_0 \eqsim_{p_0} v_0$ and $e'_0 \sqsupseteq e_0$ such that $\gamma', x \in v'_0, e'_0 \Downarrow^* v'$ we have $v' \eqsim_p v$.*

Combining with the consistency property (Proposition 2.6), we have:

**Corollary 5.2.** *Suppose $\gamma, e \Downarrow v, T$ and $p \sqsubseteq v$ and $p, T \searrow \rho, e'$. Then for all $\gamma' \eqsim_\rho \gamma$ and $e'' \sqsupseteq e'$ such that $\gamma', e'' \Downarrow v'$, we have $v \eqsim_p v'$.*

Continuing our running example, the query slice for the pattern $p'$ considered above is

$$Q' = \bigcup \{ \texttt{if}(x.B = 3, \{\langle A:\square, B:x.C\rangle\}, \{\}) \mid x \in R \}$$

since only the computation of the $A$ field in the output is irrelevant to $p'$.

$\boxed{p, T \searrow \rho, e}$

$$\overline{\square, T \searrow [], \square} \qquad \overline{p, \texttt{c} \searrow [], \texttt{c}}$$

$$\frac{\diamond, T_1 \searrow \rho_1, e_1 \quad \cdots \quad \diamond, T_n \searrow \rho_n, e_n}{p, \texttt{f}(T_1, \ldots, T_n) \searrow \rho_1 \sqcup \cdots \sqcup \rho_n, \texttt{f}(e_1, \ldots, e_n)}$$

$$\overline{p, x \searrow [x \mapsto p], x}$$

$$\frac{p_2, T_2 \searrow \rho_2[x \mapsto p_1], e_2 \quad p_1, T_1 \searrow \rho_1, e_1}{p_2, \texttt{let } x = T_1 \texttt{ in } T_2 \searrow \rho_1 \sqcup \rho_2, \texttt{let } x = e_1 \texttt{ in } e_2}$$

$$\frac{p_1, T_1 \searrow \rho_1, e_1 \quad \cdots \quad p_n, T_n \searrow \rho_n, e_n}{\langle A_1{:}p_1, \ldots, A_n{:}p_n\rangle, \langle A_1{:}T_1, \ldots, A_n{:}T_n\rangle \searrow \rho_1 \sqcup \cdots \sqcup \rho_n, \langle A_1{:}e_1, \ldots, A_n{:}e_n\rangle}$$

$$\frac{p.A, T \searrow \rho, S}{p, T.A \searrow \rho, S.A}$$

$$\frac{p_1, T_1 \searrow \rho_1, e'_1 \quad \texttt{true}, T \searrow \rho, e'}{p_1, \texttt{if}(T, e_1, e_2) \triangleright_{\texttt{true}} T_1 \searrow \rho_1 \sqcup \rho, \texttt{if}(e', e'_1, \square)}$$

$$\frac{p_2, T_2 \searrow \rho_2, e'_2 \quad \texttt{false}, T \searrow \rho, e'}{p_2, \texttt{if}(T, e_1, e_2) \triangleright_{\texttt{false}} T_2 \searrow \rho_2 \sqcup \rho, \texttt{if}(e', \square, e'_2)} \qquad \overline{\emptyset, \emptyset \searrow [], \emptyset}$$

$$\frac{p[\epsilon], T \searrow \rho, e}{p, \{T\} \searrow \rho, \{e\}} \qquad \frac{p[1], T_1 \searrow \rho_1, e_1 \quad p[2], T_2 \searrow \rho_2, e_2}{p, T_1 \cup T_2 \searrow \rho_1 \sqcup \rho_2, e_1 \cup e_2}$$

$$\frac{\diamond, T \searrow \rho, e}{p, \texttt{sum } T \searrow \rho, \texttt{sum } e} \qquad \frac{\diamond, T \searrow \rho, e}{p, \texttt{empty } T \searrow \rho, \texttt{empty } e}$$

$$\frac{p, x.\Theta \searrow^* \rho', e', p_0 \quad p_0, T \searrow \rho, e'_0}{p, \bigcup\{e \mid x \in T\} \triangleright \Theta \searrow \rho \sqcup \rho', \bigcup\{e' \mid x \in e'_0\}}$$

$$\frac{\diamond, T_1 \searrow \rho_1, e_1 \quad \cdots \quad \diamond, T_n \searrow \rho_n, e_n}{\diamond, \langle A_1 : T_1, \ldots, A_n : T_n\rangle \searrow \rho_1 \sqcup \cdots \sqcup \rho_n, \langle A_1 : e_1, \ldots, A_n : e_n\rangle}$$

$$\overline{\diamond, \emptyset \searrow [], \emptyset}$$

$\boxed{p, x.\Theta \searrow \rho, e, p'}$

$$\overline{\square, x.\Theta \searrow^* [], \square, \square} \qquad \overline{\emptyset, x.\emptyset \searrow^* [], \square, \emptyset} \qquad \overline{\diamond, x.\emptyset \searrow^* [], \square, \emptyset}$$

$$\frac{p[\ell], T \searrow \rho[x \mapsto p_0], e}{p, x.\{\ell.T\} \searrow^* \rho, e, \{\ell.p_0\}}$$

$$\frac{p|_{\text{dom}(\Theta_1)}, x.\Theta_1 \searrow^* \rho_1, e_1, p_1 \quad p|_{\text{dom}(\Theta_2)}, x.\Theta_2 \searrow^* \rho_2, e_2, p_2}{p, x.\Theta_1 \uplus \Theta_2 \searrow^* \rho_1 \sqcup \rho_2, e_1 \sqcup e_2, p_1 \uplus p_2}$$
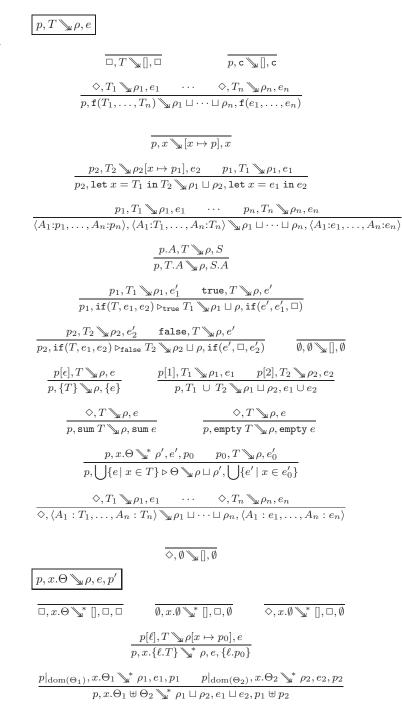
**Figure 11.** Unevaluation (selected rules).

### 5.2 Differential slicing

We consider a *pattern difference* to be a pair of patterns $(p, p')$ where $p \sqsubseteq p'$. Intuitively, a pattern difference selects the part of a value present in the outer component $p'$ and not in the inner component $p$. For example, the pattern difference $(\langle B{:}\square; \square\rangle, \langle B{:}8; \square\rangle)$ selects the value 8 located in the $B$ component of a record. We can also write this difference as $\langle B{:}\boxed{8}; \square\rangle$, using $\boxed{?}$ to highlight the boundary between the inner and outer pattern. Trace and query pattern differences are defined analogously.

It is straightforward to show by induction that slicing is monotonic in both arguments:

**Lemma 5.3** (Monotonicity). *If $p \sqsubseteq p'$ and $T \sqsubseteq T'$ and $p', T' \searrow \rho', S'$ then there exist $\rho, S$ such that $p, T \searrow \rho, S$ and $\rho \sqsubseteq \rho'$ and $S \sqsubseteq S'$. In addition, $p, S' \searrow \rho, S$.*

This implies that given a pattern difference and a trace, we can compute a trace difference using the following rule:

$$\frac{p_2, T \searrow \rho_2, S_2 \qquad p_1, S_2 \searrow \rho_1, S_1}{(p_1, p_2), T \searrow (\rho_1, \rho_2), (S_1, S_2)}$$

It follows from monotonicity that $\rho_1 \sqsubseteq \rho_2$ and $S_1 \sqsubseteq S_2$, thus, pattern differences yield trace differences. Furthermore, the second part of monotonicity implies that we can compute the smaller slice $S_1$ from the larger slice $S_2$, rather than re-traverse $T$. It is also possible to define a simultaneous differential slicing judgment, as an optimization to ensure we only traverse the trace once.

Query slicing is also monotone, so differential query slices can be obtained in exactly the same way. Revisiting our running example one last time, consider the differential pattern $\{[r_2].\langle B: \boxed{8}; \Box \rangle\} \dot\cup \Box$. The differential query slice for the pattern $p'$ considered above is

$$Q'' = \bigcup \{ \text{if}(x.B = 3, \{\langle A:\Box, B:\boxed{x.C}\rangle\}, \{\}) \mid x \in R \}$$

## 6. Examples and Discussion

In this section we present some more complex examples to illustrate key points.

***Renaming*** Recall the swapping query from the introduction, written in NRC as

$$Q_1 = \bigcup \{\{\text{if}(x.A > x.B, \langle A:x.B, B:x.A\rangle, \langle x\rangle)\} \mid x \in R\}$$

This query illustrates a key difference between our approach and the how-provenance model of Green et al. [20]. As discussed in [13], renaming operations are ignored by how-provenance, so the how-provenance annotations of the results of $Q_1$ are the same as for a query that simply returns $R$. In other words, the choice to swap the fields when $A > B$ is not reflected in the how-provenance, which shows that it is impossible to extract where-provenance (or traces) from how-provenance. Extracting where-provenance from traces appears straightforward, extending our previous work [1].

This example also illustrates how traces and slices can be used for partial recomputation. The slice for output pattern $\{[1, r_1].\langle B:2\rangle\} \dot\cup \Box$, for example, will show that this record was produced because the $A$ component of $\langle A:1, B:2, C:7\rangle$ at index $[r_1]$ in the input was less than or equal to the $B$ component. Thus, we can replay after any change that preserves this ordering information.

***Union*** Consider query

$$Q_2 = \bigcup \{\{\langle B:x.B\rangle\} \mid x \in R\} \cup \{\langle B:3\rangle\}$$

that projects the $B$ fields of elements of $R$ and adds another copy of $\langle B:3\rangle$ to the result. This yields

$$Q_2(R) = \{[1, r_1].\langle B:2\rangle, [1, r_2].\langle B:3\rangle, [1, r_3].\langle B:3\rangle, [2].\langle B:3\rangle\}$$

This illustrates that the indexes may not all have the same length, but still form a prefix code. If we slice with respect to $\{[1, r_2].\langle B:3\rangle\} \dot\cup \Box$ then the query slice is:

$$Q_2' = \bigcup \{\{\langle B:x.B\rangle\} \mid x \in R\} \cup \Box$$

and $R' = \{[r_2].\langle B:3; \Box\rangle\} \dot\cup \Box$ whereas if we slice with respect to $\{[2].\langle B:3\rangle\} \dot\cup \Box$ then the query slice is $Q_2'' = \Box \cup \{\langle B:3\rangle\}$ and $R'' = \Box$, indicating that this part of the result has no dependence on the input.

A related point: one may wonder whether it makes sense to select a particular copy of $\langle B:3\rangle$ in the output, since in a conventional multiset, multiple copies of the same value are indistinguishable. We believe it is important to be able to distinguish different copies of a value, which may have different explanations. Treating $n$ copies of a value as a single value with multiplicity $n$ would obscure this distinction and force us to compute the slices of all of the copies even if only a single explanation is required. This is why we have chosen to work with indexed sets, rather than pure multisets.

***Joins*** So far all examples have involved a single table $R$. Consider a simple join query

$$Q_3 = \{\langle A:x.A, B:y.C\rangle \mid x \in R, y \in S, x.B = y.B\}$$

and consider the following table $S$, and the result $Q_3(R, S)$.

| | $id$ | $B$ | $C$ | | | $id$ | $A$ | $B$ |
|---|---|---|---|---|---|---|---|---|
| $S =$ | $[s_1]$ | 2 | 4 | $Q_3(R, S) =$ | | $[r_1, s_1]$ | 1 | 4 |
| | $[s_2]$ | 3 | 4 | | | $[r_2, s_2]$ | 2 | 4 |
| | $[s_3]$ | 4 | 5 | | | $[r_3, s_2]$ | 4 | 5 |

The full trace of this query execution is as follows:

$$
\begin{aligned}
T_3 = \bigcup\{\llcorner\lrcorner \mid x \in R\} \rhd \{ \\
[r_1].\bigcup\{\llcorner\lrcorner \mid y \in S\} \rhd \{ \quad & [s_1].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{true}} \{\llcorner\}, \\
& [s_2].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{false}} \{\}, \\
& [s_3].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{false}} \{\}\}, \\
[r_2].\bigcup\{\llcorner\lrcorner \mid y \in S\} \rhd \{ \quad & [s_1].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{false}} \{\}, \\
& [s_2].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{true}} \{\llcorner\}, \\
& [s_3].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{false}} \{\}\}, \\
[r_3].\bigcup\{\llcorner\lrcorner \mid y \in S\} \rhd \{ \quad & [s_1].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{false}} \{\}, \\
& [s_2].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{true}} \{\llcorner\}, \\
& [s_3].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{false}} \{\}\}\}
\end{aligned}
$$

Slicing with respect to $\{[r_1, s_1].\langle A:1; \Box\rangle, [r_2, s_2].\langle B:4; \Box\rangle\} \dot\cup \Box$ yields trace slice

$$
\begin{aligned}
T_3' = \bigcup\{\llcorner\lrcorner \mid x \in R\} \rhd \{ \\
[r_1].\bigcup\{\llcorner\lrcorner \mid y \in S\} \rhd \{ \quad & [s_1].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{true}} \{\llcorner\}, \\
[r_2].\bigcup\{\llcorner\lrcorner \mid y \in S\} \rhd \{ \quad & [s_2].\text{if}(x.B = y.B, \llcorner, \lrcorner) \rhd_{\text{true}} \{\llcorner\}\}
\end{aligned}
$$

and input slice $R_3' = \{[r_1].\langle A:1, B:2; \Box\rangle, [r_2].\langle B:3; \Box\rangle\}$ and $S_3' = \{[s_1].\langle B:2; \Box\rangle, [s_2].\langle B:3; C:4\rangle\}$.

***Workflows*** NRC expressions can be used to represent workflows, if primitive operations are added representing the workflow steps [2, 21]. To illustrate query slicing and differential slicing for a workflow-style query, consider the following more complex query:

$$Q_4 = \{f(x, y) \mid x \in T, y \in T, z \in U, p(x, y), q(x, y, z)\}$$

where $f$ computes some function of $x$ and $y$ and $p$ and $q$ are selection criteria. Here, we assume $T$ and $U$ are collections of data files and $f, p, q$ are additional primitive operations on them. This query exercises most of the distinctive features of our approach; we can of course translate it to the NRC core calculus used in the rest of the paper. If $dom(T) = \{t_1, \ldots, t_{10}\}$ and $dom(U) = \{u_1, \ldots, u_{10}\}$ then we might obtain result $\{[t_3, t_4, u_5].v_1, [t_6, t_8, u_{10}].v_2\}$. If we focus on the value $v_1$ using the pattern $\{[t_3, t_4, u_5].v_1\} \dot\cup \Box$, then the program slice we obtain is $Q_4$ itself, while the data slice might be $T' = \{[t_3].\diamond, [t_4].\diamond\} \dot\cup \Box, U' = \{[u_5].\diamond\} \dot\cup \Box$, indicating that if the values at $t_3, t_4, u_5$ are held fixed then the end result will still be $v_1$. The trace slice is similar, and shows that $v_1$ was computed by applying $f$ with $x$ bound to the value at $t_3$ in $T$, $y$ bound to $t_4$, and $z$ bound to $u_5$, and that $p(x, y)$ and $q(x, y, z)$ succeeded for these values.

If we consider a differential slice using pattern difference $\{[t_3, t_4, u_5].\boxed{v_1}\} \dot\cup \Box$ then we obtain the following program difference:

$$Q_4^\delta = \bigcup \{\boxed{f(x, y)} \mid x \in T, y \in T, z \in U, p(x, y), q(x, y, z)\}$$

This shows that most of the query is needed to ensure that the result at $[t_3, t_4, u_5]$ is produced, but the subterm $f(x, y)$ is only needed to compute the value $v_1$. This can be viewed as a query-based explanation for this part of the result.

## 7. Implementation

To validate our design and experiment with larger examples, we extended our Haskell implementation Slicer of program slicing for functional programs [27] with the traces and slicing techniques presented in this paper. We call the resulting system NRCSlicer; it supports a free combination of NRC and general-purpose functional programming features. NRCSlicer interprets expressions in-memory without optimization. As reported previously for Slicer, we have experimented with several alternative tracing and slicing strategies, which use Haskell's lazy evaluation strategy in different ways. The alternatives we consider here are:

- *eager*: the trace is fully computed during evaluation.
- *lazy*: the value is computed eagerly, but the trace is computed lazily using Haskell's default lazy evaluation strategy.

To evaluate the effectiveness of enriched patterns, we measured the time needed for the eager and lazy techniques to trace and slice the workflow example $Q_4$ in the previous section. We considered a instantiation of the workflow where the data values are simply integers and with input tables $T, U = \{1, \ldots, 50\}$, and defined the operations $f(x, y)$ as $x * y$, $p(x, y)$ as $x < y$, and $q(x, y, z)$ as $x^2 + y^2 = z^2$. This is not a realistic workflow, and we expect that the time to evaluate the basic operations of a realistic workflow following this pattern would be much larger. However, the overheads of tracing and slicing do not depend on the execution time of primitive operations, so we can still draw some conclusions from this simplistic example.

The comprehension iterates over $50^3 = 125,000$ triples, producing 20 results. We considered simple and enriched patterns selecting a single element of the result. We measured evaluation time, and the overhead of tracing, trace slicing, and query slicing. The experiments were conducted on a MacBook Pro with 2GB RAM and a 2.8GHz Intel Core Duo, using GHC version 7.4.

|  | eval | trace | slice | qslice |
|---|---|---|---|---|
| eager-simple | 0.5 | 1.5 | 2.5 | 1.6 |
| eager-enriched | 0.5 | 1.5 | <0.1 | <0.1 |
| lazy-simple | 0.5 | 0.7 | 1.3 | 1.7 |
| lazy-enriched | 0.5 | 0.7 | <0.1 | <0.1 |

The times are in seconds. The "eval" column shows the time needed to compute the result without tracing. The "trace", "slice", and "qslice" columns show the added time needed to trace and compute slices. The full traces in each of these runs have over 2.1 million nodes; the simple pattern slices are almost as large, while the enriched pattern slices are only 95 nodes. For this example, slicing is over an order of magnitude faster using enriched patterns. The lazy tracing approach required less total time both for tracing and slicing (particularly for simple patterns). Thus, Haskell's built-in lazy evaluation strategy offers advantages by avoiding explicitly constructing the full trace in memory when it is not needed; however, there is still room for improvement. Again, however, for an actual workflow involving images or large data files, the evaluation time would be much larger, dwarfing the time for tracing or slicing.

Our implementation is a proof-of-concept that evaluates queries in-memory via interpretation, rather than compilation; further work would be needed to adapt our approach to support fine-grained provenance for conventional database systems. Nevertheless, our experimental results do suggest that the lazy tracing strategy and use of enriched patterns can effectively decrease the overhead of tracing, making it feasible for in-memory execution of workflows represented in NRC.

## 8. Related and future work

Program slicing has been studied extensively [17, 30, 31], as has the use of execution traces, for example in dynamic slicing. Our work contrasts with much of this work in that we regard the trace and underlying data as being of interest, not just the program. Some of our previous work [12] identified analogies between program slicing and provenance, but to our knowledge, there is no other prior work on slicing in databases.

Lineage and why-provenance were motivated semantically in terms of identifying *witnesses*, or parts of the input needed to ensure that a given part of the output is produced by a query. Early work on lineage in relational algebra [15] associates each output record with a witness. Buneman et al. studied a more general notion called why-provenance that maps an output part to a collection of witnesses [7, 8]. This idea was generalized further to the *how-provenance* or *semiring* model [18, 20], based on using algebraic expressions as annotations; this approach has been extended to handle some forms of negation and aggregation [4, 19]. Semiring homomorphisms commute with query evaluation; thus, homomorphic changes to the input can be performed directly on the output without re-running the query. However, this approach only applies to changes describable as semiring homomorphisms, such as deletion.

Where-provenance was also introduced by Buneman et al. [7, 8]. Although the idea of tracking where input data was copied from is natural, it is nontrivial to characterize semantically, because where-provenance does not always respect semantic equivalence. In later work, Buneman et al. [6] studied where-provenance for the pure NRC and characterized its expressiveness for queries and updates. It would be interesting to see whether their notion of *expressive completeness* for where-provenance could be extended to richer provenance models, such as traces, possibly leading to an implementation strategy via translation to plain NRC.

Provenance has been studied extensively for scientific workflow systems [5, 29], but there has been little formal work on the semantics of workflow provenance. The closest work to ours is that of Hidders et al. [21], who model workflows by extending the NRC with nondeterministic, external function calls. They sketch an operational semantics that records *runs* that contain essentially all of the information in a derivation tree, represented as a set of triples. They also suggest ways of extracting *subruns* from runs, but their treatment is partial and lacks strong formal guarantees analogous to our results.

There have been some attempts to reconcile the database and workflow views of provenance; Hidders et al. [21] argued for the use of Nested Relational Calculus (NRC) as a unifying formalism for both workflow and database operations, and subsequently Kwasnikowska and Van den Bussche [22] showed how to map this model to the Open Provenance Model. Acar et al. [2] later formalized a graph model of provenance for NRC. The most advanced work in this direction appears to be that of Amsterdamer et al. [3], who combined workflow and database styles of provenance in the context of the PigLatin system (a MapReduce variant based on nested relational queries). Lipstick allows analyzing the impact of restricted hypothetical changes (such as deletion) on parts of the output, but to our knowledge no previous work provides a formal guarantee about the impact of changes other than deletion.

In our previous work [12], we introduced *dependency provenance*, which conservatively over-approximates the changes that can take place in the output if the input is changed. We developed definitions and techniques for dependency provenance in full NRC including nonmonotone operations (empty, sum) and primi-

tive functions. Dependency provenance cannot predict exactly how the output will be affected by a general modification to the source, but it can guarantee that some parts of the output will not change if certain parts of the input are fixed. Our notion of equivalence modulo a pattern is a generalization of the *equal-except-at* relation used in that work. Motivated by dependency provenance, an earlier technical report [10] presented a model of traced evaluation for NRC and proved elementary properties such as fidelity. However, it did not investigate slicing techniques, and used nondeterministic label generation instead of our deterministic scheme; our deterministic approach greatly simplifies several aspects of the system, particularly for slicing.

There are several intriguing directions for future work, including developing more efficient techniques for traced evaluation and slicing that build upon existing database query optimization capabilities. It appears possible to translate multiset queries so as to make the labels explicit, since a fixed given query increases the label depth by at most a constant. Thus, it may be possible to evaluate queries with label information but without tracing first, then gradually build the trace by slicing backwards through the query, re-evaluating subexpressions as necessary. Other interesting directions include the use of slicing techniques for security, to hide confidential input information while disclosing enough about the trace to permit recomputation, and the possibility of extracting other forms of provenance from traces, as explored in the context of functional programs in prior work [1].

## 9. Conclusion

The importance of provenance for transparency and reproducibility is widely recognized, yet there has been little explicit discussion of correctness properties formalizing intuitions about how provenance is to provide reproducibility. In self-explaining computation, traces are considered to be explanations of a computation in the sense that the trace can be used to recompute (parts of) the output under hypothetical changes to the input. This paper develops the foundations of self-explaining computation for database queries, by defining a tracing semantics for NRC, proposing a formal definition of correctness for tracing (fidelity) and slicing, and defining a correct (though potentially overapproximate) algorithm for trace slicing. Trace slicing can be used to obtain smaller "golden trail" traces that explain only a part of the input or output, and explore the impact of changes in hypothetical scenarios similar to the original run. At a technical level, the main contributions are the careful use of prefix codes to label multiset elements, and the development of enriched patterns that allow more precise slices. Our design is validated by a proof-of-concept implementation that shows that laziness and enriched patterns can significantly improve performance for small (in-memory) examples.

In the near term, we plan to combine our work on self-explaining functional programs [27] and database queries (this paper) to obtain slicing and provenance models for programming languages with query primitives, such as F# [14] or Links [23]. Ultimately, our aim is to extend self-explaining computation to programs that combine several execution models, including workflows, databases, conventional programming languages, Web interaction, or cloud computing.

## References

[1] U. A. Acar, A. Ahmed, J. Cheney, and R. Perera. A core calculus for provenance. *Journal of Computer Security*, 21:919–969, 2013. Full version of a POST 2012 paper.

[2] U. A. Acar, P. Buneman, J. Cheney, N. Kwasnikowska, J. Van den Bussche, and S. Vansummeren. A graph model of data and workflow provenance. In *TAPP*, 2010.

[3] Y. Amsterdamer, S. B. Davidson, D. Deutch, T. Milo, J. Stoyanovich, and V. Tannen. Putting lipstick on pig: Enabling database-style workflow provenance. *PVLDB*, 5(4):346–357, 2011.

[4] Y. Amsterdamer, D. Deutch, and V. Tannen. Provenance for aggregate queries. In *PODS*, pages 153–164. ACM, 2011.

[5] R. Bose and J. Frew. Lineage retrieval for scientific data processing: a survey. *ACM Comput. Surv.*, 37(1):1–28, 2005.

[6] P. Buneman, J. Cheney, and S. Vansummeren. On the expressiveness of implicit provenance in query and update languages. *ACM Transactions on Database Systems*, 33(4):28, November 2008.

[7] P. Buneman, S. Khanna, and W. Tan. Why and where: A characterization of data provenance. In *ICDT*, number 1973 in LNCS, pages 316–330. Springer, 2001.

[8] P. Buneman, S. Khanna, and W. Tan. On propagation of deletions and annotations through views. In *PODS*, pages 150–158, 2002.

[9] P. Buneman, S. A. Naqvi, V. Tannen, and L. Wong. Principles of programming with complex objects and collection types. *Theor. Comp. Sci.*, 149(1):3–48, 1995.

[10] J. Cheney, U. A. Acar, and A. Ahmed. Provenance traces. *CoRR*, arXiv.org/abs/0812.0564, 2008.

[11] J. Cheney, U. A. Acar, and R. Perera. Toward a theory of self-explaining computation. In *In search of elegance in the theory and practice of computation: a Festschrift in honour of Peter Buneman*, number 8000 in LNCS, pages 193–216. Springer, 2013.

[12] J. Cheney, A. Ahmed, and U. A. Acar. Provenance as dependency analysis. *Mathematical Structures in Computer Science*, 21(6):1301–1337, 2011.

[13] J. Cheney, L. Chiticariu, and W. C. Tan. Provenance in databases: Why, how, and where. *Foundations and Trends in Databases*, 1(4):379–474, 2009.

[14] J. Cheney, S. Lindley, and P. Wadler. A practical theory of language-integrated query. In *ICFP*, pages 403–416, New York, NY, USA, 2013. ACM.

[15] Y. Cui, J. Widom, and J. L. Wiener. Tracing the lineage of view data in a warehousing environment. *ACM Trans. Database Syst.*, 25(2):179–227, 2000.

[16] J. Dean and S. Ghemawat. MapReduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.

[17] J. Field and F. Tip. Dynamic dependence in term rewriting systems and its application to program slicing. *Information and Software Technology*, 40(11–12):609–636, 1998.

[18] J. N. Foster, T. J. Green, and V. Tannen. Annotated XML: queries and provenance. In *PODS*, pages 271–280, 2008.

[19] F. Geerts and A. Poggi. On database query languages for $K$-relations. *J. Applied Logic*, 8(2):173–185, 2010.

[20] T. J. Green, G. Karvounarakis, and V. Tannen. Provenance semirings. In *PODS*, pages 31–40, 2007.

[21] J. Hidders, N. Kwasnikowska, J. Sroka, J. Tyszkiewicz, and J. Van den Bussche. A formal model of dataflow repositories. In *DILS*, 2007.

[22] N. Kwasnikowska and J. Van den Bussche. Mapping the NRC dataflow model to the open provenance model. In *IPAW*, pages 3–16, 2008.

[23] S. Lindley and J. Cheney. Row-based effect types for database integration. In *TLDI*, pages 91–102. ACM Press, 2012.

[24] P. Missier, B. Ludäscher, S. Dey, M. Wang, T. McPhillips, S. Bowers, M. Agun, and I. Altintas. Golden trail: Retrieving the data history that matters from a comprehensive provenance repository. *International Journal of Digital Curation*, 7(1):139–150, 2011.

[25] L. Moreau. The foundations for provenance on the web. *Foundations and Trends in Web Science*, 2(2–3), 2010.

[26] C. Olston, B. Reed, U. Srivastava, R. Kumar, and A. Tomkins. Pig latin: a not-so-foreign language for data processing. In *SIGMOD*, pages 1099–1110. ACM, 2008.

[27] R. Perera, U. A. Acar, J. Cheney, and P. B. Levy. Functional programs that explain their work. In *ICFP*, pages 365–376. ACM, 2012.

[28] A. Sabelfeld and A. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, 2003.

[29] Y. Simmhan, B. Plale, and D. Gannon. A survey of data provenance in e-science. *SIGMOD Record*, 34(3):31–36, 2005.

[30] F. Tip. A survey of program slicing techniques. *J. Prog. Lang.*, 3(3), 1995.

[31] M. Weiser. Program slicing. In *ICSE*, pages 439–449. IEEE Press, 1981.

$$\boxed{\Gamma \vdash T : \tau}$$

$$\frac{n \in \mathbb{N}}{\Gamma \vdash n : \texttt{int}} \qquad \frac{b \in \{\texttt{true},\texttt{false}\}}{\Gamma \vdash b : \texttt{bool}} \qquad \frac{\texttt{f} : (b_1,\ldots,b_n) \to b \in \Sigma \quad \Gamma \vdash T_1 : b_1 \quad \cdots \quad \Gamma \vdash T_n : b_n}{\Gamma \vdash \texttt{f}(T_1,\ldots,T_n) : b} \qquad \frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau}$$

$$\frac{\Gamma \vdash T_1 : \tau_1 \quad \Gamma, x : \tau_1 \vdash T_2 : \tau_2}{\Gamma \vdash \texttt{let } x = T_1 \texttt{ in } T_2 : \tau_2} \qquad \frac{\Gamma \vdash T_1 : \tau_1 \quad \cdots \quad \Gamma \vdash T_n : \tau_n}{\Gamma \vdash \langle A_1 : T_1, \ldots, A_n : T_n \rangle : \langle A_1 : \tau_1, \ldots, A_n : \tau_n \rangle} \qquad \frac{\Gamma \vdash T : \langle A_1 : \tau_1, \ldots, A_n : \tau_n \rangle}{\Gamma \vdash T.A_i : \tau_i}$$

$$\frac{b \in \{\texttt{true},\texttt{false}\} \quad \Gamma \vdash T : \texttt{bool} \quad \Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \tau \quad \Gamma \vdash T' : \tau}{\Gamma \vdash \texttt{if}(T, e_1, e_2) \triangleright_b T' : \tau} \qquad \frac{}{\Gamma \vdash \emptyset : \{\tau\}} \qquad \frac{\Gamma \vdash T : \tau}{\Gamma \vdash \{T\} : \{\tau\}}$$

$$\frac{\Gamma \vdash T : \{\tau\} \quad \Gamma \vdash T' : \{\tau\}}{\Gamma \vdash T \cup T' : \{\tau\}} \qquad \frac{\Gamma \vdash T : \{\texttt{int}\}}{\Gamma \vdash \texttt{sum } T : \texttt{int}} \qquad \frac{\Gamma \vdash T : \{\tau\}}{\Gamma \vdash \texttt{empty } T : \texttt{bool}} \qquad \frac{\Gamma, x : \tau \vdash e : \{\tau'\} \quad \Gamma \vdash T : \{\tau\} \quad \Gamma, x \in \{\tau\} \vdash \Theta : \{\tau'\}}{\Gamma \vdash \bigcup\{e \mid x \in T\} \triangleright \Theta : \{\tau'\}}$$

$$\boxed{\Gamma, x \in \{\tau\} \vdash \Theta : \{\tau'\}}$$

$$\frac{}{\Gamma, x \in \{\tau\} \vdash \emptyset : \{\tau'\}} \qquad \frac{\Gamma, x : \tau \vdash T : \{\tau'\}}{\Gamma, x \in \{\tau\} \vdash \{\ell.T\} : \{\tau'\}} \qquad \frac{\Gamma, x \in \{\tau\} \vdash \Theta_1 : \{\tau'\} \quad \Gamma, x \in \{\tau\} \vdash \Theta_2 : \{\tau'\}}{\Gamma, x \in \{\tau\} \vdash \Theta_1 \cup \Theta_2 : \{\tau'\}}$$

**Figure 12.** Well-typed traces.

$$\boxed{T' \sqsubseteq T}$$

$$\frac{}{\square \sqsubseteq T} \qquad \frac{}{T \sqsubseteq T} \qquad \frac{T'' \sqsubseteq T' \quad T' \sqsubseteq T}{T'' \sqsubseteq T} \qquad \frac{T'_1 \sqsubseteq T_1 \quad \cdots \quad T'_n \sqsubseteq T_n}{\texttt{f}(T'_1,\ldots,T'_n) \sqsubseteq \texttt{f}(T_1,\ldots,T_n)} \qquad \frac{T'_1 \sqsubseteq T_1 \quad T'_2 \sqsubseteq T_2}{\texttt{let } x = T'_1 \texttt{ in } T'_2 \sqsubseteq \texttt{let } x = T_1 \texttt{ in } T_2}$$

$$\frac{T'_1 \sqsubseteq T_1 \quad \cdots \quad T'_n \sqsubseteq T_n}{\langle A_1 : T'_1, \ldots, A_n : T'_n \rangle \sqsubseteq \langle A_1 : T_1, \ldots, A_n : T_n \rangle} \qquad \frac{T' \sqsubseteq T}{T'.A \sqsubseteq T.A} \qquad \frac{T'_0 \sqsubseteq T_0 \quad T' \sqsubseteq T}{\texttt{if}(T'_0, e_1, e_2) \triangleright_b T' \sqsubseteq \texttt{if}(T_0, e_1, e_2) \triangleright_b T} \qquad \frac{}{\emptyset \sqsubseteq \emptyset}$$

$$\frac{T' \sqsubseteq T}{\{T'\} \sqsubseteq \{T\}} \qquad \frac{T'_1 \sqsubseteq T_1 \quad T'_2 \sqsubseteq T_2}{T'_1 \cup T'_2 \sqsubseteq T_1 \cup T_2} \qquad \frac{T' \sqsubseteq T}{\texttt{sum } T' \sqsubseteq \texttt{sum } T} \qquad \frac{T' \sqsubseteq T}{\texttt{empty } T' \sqsubseteq \texttt{empty } T} \qquad \frac{T' \sqsubseteq T \quad \Theta' \sqsubseteq \Theta}{\bigcup\{e \mid x \in T'\} \triangleright \Theta' \sqsubseteq \bigcup\{e \mid x \in T\} \triangleright \Theta}$$

$$\Theta' \sqsubseteq \Theta \iff \forall \ell \in \mathrm{dom}(\Theta').\Theta'(\ell) \sqsubseteq \Theta(\ell)$$

**Figure 13.** Subtrace relation.

## A. Auxiliary definitions

Figure 12 summarizes the typing rules for traces. Figure 13 defines the subtrace relation.

## B. Proofs of pattern properties

We prove the required properties for enriched patterns. The corresponding properties for the sublanguage of simple patterns follow immediately since simple patterns are closed under the relevant operations.

**Lemma B.1.** *If $p, p' \sqsubseteq v$ then $p \sqcup p'$ exists and is the least upper bound of $p$ and $p'$.*

*Proof.* If both $p$ and $p'$ match some pattern $q$, then it is straightforward to show by induction on $q$ that $p \sqcup p'$ is defined and $p \sqcup p' \sqsubseteq q$. Specifically, if $p$ or $p'$ is $\square$ or $\diamond$ then we are done; if $p$ and $p'$ are both constants then we are done; otherwise, in each case, the toplevel structure of $p$ and $p'$ must match $q$, so that we can apply one of the rules for $\sqcup$ on smaller terms that match part of $q$. When $q = v$, the desired result follows. The second part (that $\sqcup$ is a least upper bound) also follows directly since clearly, $p, p' \sqsubseteq p \sqcup p'$ and if $p, p' \sqsubseteq q$ then a similar argument shows that $p \sqcup p' \sqsubseteq q$. $\square$

**Lemma B.2.** *For any $v, v', p$, $v \approx_{p[\diamond/\square]} v'$ holds if and only if $v \approx_p v'$ holds and $v = v'$.*

*Proof.* Straightforward induction on the derivation of $v \approx_{p[\diamond/\square]} v'$. $\square$

**Lemma B.3.** *For any $v, v', p, p'$, we have $v \approx_{p \sqcup p'} v'$ if and only if $v \approx_p v'$ and $v \approx_{p'} v'$. Moreover, $p \sqsubseteq p'$ if and only if for all $v, v'$, we have $v \approx_{p'} v'$ implies $v \approx_p v'$.*

*Proof.* For the first part, we proceed by induction on the total size of $p, p'$. The cases where one of $p, p'$ is $\square$ or $\diamond$ are straightforward; for $\diamond$ we also need Lemma B.2. The cases involving constants, pairs, or complete set and record patterns are also straightforward.

There are several similar cases involving partial set or record patterns. We illustrate two representative cases:

1. If $p = \{\overline{\ell_i.p_i}, \overline{\ell'_i.q_i}\}$ and $p' = \{\overline{\ell_i.p'_i}\} \,\dot{\cup}\, \square$ then $p \sqcup p' = \{\overline{\ell_i.p_i \sqcup p'_i}, \overline{\ell'_i.q_i}\}$. First, suppose $v \approx_{p \sqcup p'} v'$. This means $v = \{\overline{\ell_i.v_i}, \overline{\ell'_i.w_i}\}$ and $v' = \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w'_i}\}$, where $v_i \approx_{p_i \sqcup p'_i} v'_i$ and $w_i \approx_{q_i} w'_i$. Therefore, by induction, $v_i \approx_{p_i} v'_i$ and $v_i \approx_{p'_i} v'_i$, so we can conclude that $\{\overline{\ell_i.v_i}, \overline{\ell'_i.w_i}\} \approx_p \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w'_i}\}$ and $\{\overline{\ell_i.v_i}, \overline{\ell'_i.w_i}\} \approx_{p'} \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w'_i}\}$, as required.

14

Conversely, if we assume $v \approx_p v'$ and $v \approx_{p'} v'$, then we must have $v = \{\overline{\ell_i.v_i}, \overline{\ell'_i.w_i}\}$ and $v' = \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w'_i}\}$, where $v_i \approx_{p_i} v'_i$ and $v_i \approx_{p'_i} v'_i$ and $w_i \approx_{q_i} w'_i$. Thus, by induction we have $v_i \approx_{p_i \sqcup p'_i} v'_i$ so we can conclude $\{\overline{\ell_i.v_i}, \overline{\ell'_i.w_i}\} \approx_{p \sqcup p'} \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w'_i}\}$.

2. If $p = \{\overline{\ell_i.p_i}, \overline{\ell'_i.q_i}\} \mathbin{\dot\cup} \square$ and $p' = \{\overline{\ell_i.p'_i}, \overline{\ell''_i.r_i}\} \mathbin{\dot\cup} \diamond$ then $p \sqcup p' = \{\overline{\ell_i.p_i \sqcup p'_i}, \overline{\ell'_i.q_i[\diamond/\square]}, \overline{\ell''_i.r_i}\} \mathbin{\dot\cup} \diamond$. First, suppose $v \approx_{p \sqcup p'} v'$. This means $v = \{\overline{\ell_i.v_i}, \overline{\ell'_i.w_i}, \overline{\ell''_i.u_i}\} \uplus w_0$ and $v' = \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w'_i}, \overline{\ell''_i.u'_i}\} \uplus w_0$, where $v_i \approx_{p_i \sqcup p'_i} v'_i$ and $w_i \approx_{q_i[\diamond/\square]} w'_i$ and $u_i \approx_{r_i} u'_i$. Therefore, by induction, $v_i \approx_{p_i} v'_i$ and $v_i \approx_{p'_i} v'_i$, and we also have $w_i \approx_{q_i} w'_i$ and $w_i = w'_i$, so we can conclude that $\{\overline{\ell_i.v_i}, \overline{\ell'_i.w_i}, \overline{\ell''_i.u_i}\} \uplus w_0 \approx_p \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w'_i}, \overline{\ell''_i.u'_i}\} \uplus w_0$ and $\{\overline{\ell_i.v'_i}, \overline{\ell'_i.w_i}, \overline{\ell''_i.u_i}\} \uplus w_0 \approx_{p'} \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w'_i}, \overline{\ell''_i.u'_i}\} \uplus w_0$, as required. Conversely, if we assume $v \approx_p v'$ and $v \approx_{p'} v'$, then we must have $v = \{\overline{\ell_i.v_i}, \overline{\ell'_i.w_i}, \overline{\ell''_i.u_i}\} \uplus v_0$ and $v' = \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w'_i}, \overline{\ell''_i.u'_i}\} \uplus v'_0$ where $v_i \approx_{p_i} v'_i$ and $v_i \approx_{p'_i} v'_i$ and $w_i \approx_{q_i} w'_i$ and $u_i \approx_{r_i} u'_i$. In addition, we must have that $w_i = w'_i$ and $v_0 = v'_0$ since $v$ and $v'$ must be equal at all labels not in $\overline{\ell_i}, \overline{\ell''_i}$. Thus, by induction we have $v_i \approx_{p_i \sqcup p'_i} v'_i$ and (using Lemma B.2) we can also easily show that $w_i \approx_{q_i[\diamond/\square]} w_i$, so we can conclude that $\{\overline{\ell_i.v_i}, \overline{\ell'_i.w_i}, \overline{\ell''_i.u'_i}\} \uplus v_0 \approx_{p \sqcup p'} \{\overline{\ell_i.v'_i}, \overline{\ell'_i.w_i}, \overline{\ell''_i.u_i}\} \uplus v_0$.

The second part follows immediately from the definition of $p \sqsubseteq p'$ as $p \sqcup p' = p'$. $\qquad\square$

**Lemma B.4** (Properties of union and restriction).

1. If $p_1 \sqsubseteq v_1$ and $p_2 \sqsubseteq v_2$ and $v_1 \approx_{p_1} v'_1$ and $v_2 \approx_{p_2} v'_2$ then $v_1 \uplus v_2 \approx_{p_1 \uplus p_2} v'_1 \uplus v'_2$, provided all of these disjoint unions are defined.
2. If $p \sqsubseteq v_1 \uplus v_2$ and $L_1 \le \mathrm{dom}(v_1)$ and $L_2 \le \mathrm{dom}(v_2)$ and $L_1, L_2$ are prefix-disjoint, then $p|_{L_1} \sqsubseteq v_1$ and $p|_{L_2} \sqsubseteq v_2$.

*Proof.* For part 1, assume $p_1 \sqsubseteq v_1$, $p_2 \sqsubseteq v_2$, $v_1 \approx_{p_1} v'_1$ and $v_2 \approx_{p_2} v'_2$, and assume that the domains of $p_1$ and $p_2$, $v_1$ and $v_2$, and $v'_1$ and $v'_2$ are prefix-disjoint respectively, so that the unions exist. There are several cases. If $p_1$ or $p_2$ is $\square$ then the conclusion is immediate. If both are $\diamond$ then $v_1 = v'_1$ and $v_2 = v'_2$ so $v_1 \uplus v_2 = v'_1 \uplus v'_2$.

Most of the remaining cases are straightforward; we illustrate with the case $p_1 = \{\overline{\ell_i.p_i}\} \mathbin{\dot\cup} \square$ and $p_2 = \{\overline{\ell'_i.q_i}\} \mathbin{\dot\cup} \diamond$. In this case,

$$
\begin{aligned}
v_1 &= \{\overline{\ell_i.v_i^1}\} \uplus v_0^1 \\
v_2 &= \{\overline{\ell'_i.v_i^2}\} \uplus v_0^2 \\
v'_1 &= \{\overline{\ell_i.w_i^1}\} \uplus w_0^1 \\
v'_2 &= \{\overline{\ell'_i.w_i^2}\} \uplus w_0^2
\end{aligned}
$$

and we also know that $v_i^1 \approx_{p_i} w_i^1$ and $v_i^2 \approx_{q_i} w_i^2$ for each $i$. Therefore,

$$
\begin{aligned}
\{\overline{\ell_i.v_i^1}\} \uplus v_0^1 \uplus \{\overline{\ell'_i.v_i^2}\} \uplus v_0^2 \quad &= \quad \{\overline{\ell_i.v_i^1}, \overline{\ell'_i.v_i^2}\} \uplus v_0^1 \uplus v_0^2 \\
&\approx_{\{\overline{\ell_i.p_i}, \overline{\ell'_i.q_i}\} \mathbin{\dot\cup} \square} \quad \{\overline{\ell_i.w_i^1}, \overline{\ell'_i.w_i^2}\} \uplus w_0^1 \uplus w_0^2 \\
&= \quad \{\overline{\ell_i.w_i^1}\} \uplus w_0^1 \uplus \{\overline{\ell'_i.w_i^1}\} \uplus w_0^2
\end{aligned}
$$

For part 2, assume $p \sqsubseteq v_1 \uplus v_2$ and $L_i \le \mathrm{dom}(v_i)$ for $i \in \{1, 2\}$. We proceed by case analysis on $p$. The cases $p = \square$ and $p = \diamond$ are immediate since $\square|_L = \square$ and $\diamond|_L = \diamond$. If $p = \{\overline{\ell_i.p_i}\}$ is a complete set pattern, then $p|_{L_1}$ selects just those elements of $p$ that have a prefix in $L_i$, and since every element of $v_q$ has its label's prefix in $L_1$ we must have that $v_1 = \{\overline{\ell'_i.v'_i}\}$ where $p_i \sqsubseteq v'_i$ for each $i$ and $p|_{L_1} = \{\overline{\ell'_i.p'_i}\}$, which is what we need to show. The cases for $p = \{\overline{\ell_i.p_i}\} \mathbin{\dot\cup} \square$ and $p = \{\overline{\ell_i.p_i}\} \mathbin{\dot\cup} \diamond$ are similar.

A symmetric argument suffices to show $p|_{L_2} \sqsubseteq v_2$. $\qquad\square$

**Lemma B.5** (Projection and $\sqsubseteq$).

1. If $p \sqsubseteq \{\epsilon.v\}$ then $p.\epsilon \sqsubseteq v$.
2. If $p \sqsubseteq 1 \cdot v_1 \uplus 2 \cdot v_2$ then $p[1] \sqsubseteq v_1$ and $p[2] \sqsubseteq v_2$.
3. If $p \sqsubseteq \ell \cdot v$ then $p[\ell] \sqsubseteq v$.
4. If $p \sqsubseteq \langle \overline{A_i : v_i} \rangle$ then $p.A_i \sqsubseteq v_i$.

*Proof.* For part (1), suppose $p \sqsubseteq \{\epsilon.v\}$. We proceed by case analysis on $p$. The cases for $\square$ and $\diamond$ are trivial. If $p = \{\epsilon.p'\}$ then $p.\epsilon = p'$ so the conclusion follows. The cases for $p = \{\epsilon.p'\} \mathbin{\dot\cup} \square$ or $p = \{\epsilon.p'\} \mathbin{\dot\cup} \diamond$ are similar.

For part (2), suppose $p \sqsubseteq 1 \cdot v_1 \uplus 2 \cdot v_2$. Suppose $v_1 = \{\overline{\ell_i.v_i}\}$ and $v_2 = \{\overline{\ell'_i.v'_i}\}$. If $p$ is a complete set pattern, then it must be of the form $\{\overline{1.\ell_i.p_i}, \overline{2.\ell'_i.q_i}\}$, where $p_i \sqsubseteq v_i$ and $q_i \sqsubseteq v'_i$. The desired conclusion follows since $p[1] = \{\overline{\ell_i.p_i}\}$, and a symmetric argument shows that $p[2] \sqsubseteq v_2$. The cases for partial set patterns are similar, since the $\square$ or $\diamond$ is preserved by the projection operation.

For part (3), suppose $p \sqsubseteq \ell \cdot v$. The proof is analogous to the previous case.

For part (4), suppose $p \sqsubseteq \langle \overline{A_i : v_i} \rangle$. If $p = \diamond$ or $\square$, the result is immediate; otherwise, $p$ is a record pattern. If it is a total record pattern $\langle \overline{A_i : p_i} \rangle$ then clearly $p.A_i = p_i \sqsubseteq v_i$. Otherwise, it is a partial pattern, in which case either $p.A_i$ is a pattern $p_i$ mentioned in $p$, in which case we are done, or $p.A_i = \diamond$ or $p.A_i = \square$, and the conclusion follows immediately. $\qquad\square$

**Lemma B.6** (Projection and $\approx_p$).

1. If $p \sqsubseteq \{\epsilon.v\}$ and $v \approx_{p.\epsilon} v'$ then $\{\epsilon.v\} \approx_p \{\epsilon.v'\}$.
2. If $p \sqsubseteq 1 \cdot v_1 \uplus 2 \cdot v_2$ and $v_1 \approx_{p[1]} v'_1$ and $v_2 \approx_{p[2]} v'_2$ then $1 \cdot v_1 \uplus 2 \cdot v_2 \approx_p 1 \cdot v'_1 \uplus 2 \cdot v'_2$.
3. If $p \sqsubseteq \ell \cdot v$ and $v \approx_{p[\ell]} v'$ then $\ell \cdot v \approx_p \ell \cdot v'$.
4. If $p \sqsubseteq \langle \overline{A_i : v_i} \rangle$ and $v_1 \approx_{p.A_1} v'_1, \dots, v_n \approx_{p.A_n} v'_n$ then $\langle \overline{A_i : v_i} \rangle \approx_p \langle \overline{A_i : v'_i} \rangle$.

*Proof.* For part (1), suppose $p \sqsubseteq \{\epsilon.v\}$ and $v \eqsim_{p.\epsilon} v'$. If $p = \square$ or $\diamond$ then the result is immediate (this is the case for all three parts of the lemma). If $p = \{\epsilon.p'\}$, $p = \{\epsilon.p'\} \dot\cup \square$, or $p = \{\epsilon.p'\} \dot\cup \diamond$ then $p.\epsilon = p'$ so $v \eqsim_{p'} v'$. We can conclude $\{\epsilon.v\} \eqsim_p \{\epsilon.v'\}$.

For part (2), suppose $p \sqsubseteq 1 \cdot v_1 \uplus 2 \cdot v_2$ and $v_1 \eqsim_{p[1]} v_1'$ and $v_2 \eqsim_{p[2]} v_2'$. As usual, the cases $p = \square$ and $p = \diamond$ are trivial. Suppose $v_1 = \{\overline{\ell_i.v_i'}\}$ and $v_2 = \{\overline{\ell_i'.v_i''}\}$. If $p$ is a complete set pattern it must be of the form $\{\overline{1.\ell_i.p_i}, \overline{2.\ell_i'.q_i}\}$, and $v_1 \eqsim_{\{\overline{\ell_i.p_i}\}} v_1'$ and $v_2 \eqsim_{\{\overline{\ell_i'.q_i}\}} v_2'$. It is straightforward to show that $1 \cdot v_1 \eqsim_{\{\overline{1.\ell_i.p_i}\}} 1 \cdot v_1'$ and $2 \cdot v_2 \eqsim_{\{\overline{2.\ell_i'.q_i}\}} 2 \cdot v_2'$, so by previous results we have $1 \cdot v_1 \uplus 2 \cdot v_2 \eqsim_p 1 \cdot v_1' \uplus 2 \cdot v_2'$.

The cases for partial patterns follow the same reasoning, making use of the fact that projection preserves the partial pattern.

For part (3), suppose $p \sqsubseteq \ell \cdot v$. The argument is similar to the previous case.

For part (4) suppose $p \sqsubseteq \langle \overline{A_i : v_i} \rangle$ and $v_i \eqsim_{p.A_i} v_i'$ for each $i$. If $p$ is $\square$, $\diamond$, or a complete record pattern then the conclusion is immediate. Otherwise, if $p = \langle \overline{B_i : q_i} \rangle \dot\cup \square$, the conclusion is immediate since each component of the records $\langle \overline{A_i : v_i} \rangle$ and $\langle \overline{A_i : v_i'} \rangle$ either match the appropriate $q_i$ or need not match because of the hole. Finally, if $p = \langle \overline{B_i : q_i} \rangle \dot\cup \diamond$, the conclusion follows since each pair of corresponding components of the records $\langle \overline{A_i : v_i} \rangle$ and $\langle \overline{A_i : v_i'} \rangle$ either match the appropriate $q_i$ or are equal because $p.A_i = \diamond$ if $A_i$ is not among the $B_j$. $\square$

## C. Proof of correctness of trace slicing

**Lemma C.1.** *If $\gamma, x \in v_1, \Theta \curvearrowright^* v_2$ and $v_1' \subseteq v_1$ then there exists $v_2' \subseteq v_2$ such that $\gamma, x \in v_1', \Theta \curvearrowright^* v_2'$.*

*Proof.* The proof is straightforward by induction on derivations. The cases for $v_1 = \emptyset$ and $v_1 = \{\ell.v\}$ are immediate; if $v_1 = w_1 \uplus w_2$ then we proceed by induction using the subsets $w_1' = w_1 \cap v_1'$ and $w_2' = w_2 \cap v_1'$. $\square$

We prove correctness of the full trace slicing algorithm, with enriched patterns, since the correctness for simple patterns follows as a special case.

**Theorem C.2** (Correctness of Slicing)**.**

1. *Suppose $\gamma, T \curvearrowright v$ and $p \sqsubseteq v$ and $p, T \searrow \rho, S$. Then for all $\gamma' \eqsim_\rho \gamma$ and $T' \sqsupseteq S$ such that $\gamma', T' \curvearrowright v'$ we have $v' \eqsim_p v$.*
2. *Suppose $\gamma, x \in v_0, \Theta \curvearrowright^* v$ and $p \sqsubseteq v$ and $p, x.\Theta_0 \searrow^* \rho, \Theta_0', p_0$, where $\Theta_0 \subseteq \Theta$. Then for all $\gamma' \eqsim_\rho \gamma$ and $v_0' \eqsim_{p_0} v_0$ and $\Theta' \sqsupseteq \Theta_0'$ such that $\gamma', x \in v_0', \Theta' \curvearrowright^* v'$ we have $v' \eqsim_p v$.*

*Proof.* For part (1), the proof is by induction on the structure of slicing derivations, using inversion to extract information from other derivations. The cases for variables, constants, primitive operations, and let-binding are exactly as in previous work [1]. The cases for conditionals are similar to the those for variant types and case constructs in previous work.

We show the cases for records and set operations, which are new to this paper (records are handled similarly to pairs in our previous work, so the cases for pairs are omitted).

- Field projection. If the last step in the slicing derivation is

$$\frac{\langle A_i{:}p; \square \rangle, T \searrow \rho, S}{p, T.A_i \searrow \rho, S.A_i}$$

then the evaluation derivation must be of the form

$$\frac{\gamma, T \curvearrowright \langle A_1{:}v_1, \ldots, A_n{:}v_n \rangle}{\gamma, T.A_i \curvearrowright v_i}$$

Let $\gamma' \eqsim_\rho \gamma$ and $T' \sqsupseteq S.A_i$ be given, where $\gamma', T' \curvearrowright v'$. Then $T'$ must have the form $T''.A_i$ for some $T'' \sqsupseteq S$ so the replay derivation is of the form

$$\frac{\gamma', T'' \curvearrowright \langle A_1{:}v_1', \ldots, A_n{:}v_n' \rangle}{\gamma', T''.A_i \curvearrowright v_i'}$$

The induction hypothesis applies since it is easy to show that $\langle A_i : p; \square \rangle \sqsubseteq \langle A_1{:}v_1, \ldots, A_n{:}v_n \rangle$. Therefore

$$\langle A_1{:}v_1', \ldots, A_n{:}v_n' \rangle \eqsim_{\langle A_i{:}p;\square \rangle} \langle A_1{:}v_1, \ldots, A_n{:}v_n \rangle .$$

From this it is obvious that $v_i' \eqsim_p v_i$.
- Record. If the last step in the slicing derivation is

$$\frac{p.A_1, T_1 \searrow \rho_1, S_1 \quad \cdots \quad p.A_n, T_n \searrow \rho_n, S_n}{p, \langle A_1{:}T_1, \ldots, A_n{:}T_n \rangle \searrow \rho_1 \sqcup \cdots \sqcup \rho_n, \langle A_1{:}S_1, \ldots, A_n{:}S_n \rangle}$$

then the evaluation derivation must be of the form

$$\frac{\gamma, T_1 \curvearrowright v_1 \quad \cdots \quad \gamma, T_n \curvearrowright v_n}{\gamma, \langle A_1{:}T_1, \ldots, A_n{:}T_n \rangle \curvearrowright \langle A_1{:}v_1, \ldots, A_n{:}v_n \rangle}$$

Let $\gamma' \eqsim_{\rho_1 \sqcup \cdots \sqcup \rho_n} \gamma$ and $T' \sqsupseteq \langle A_1{:}S_1, \ldots, A_n{:}S_n \rangle$ be given, where $\gamma', T' \curvearrowright v'$. Then $T'$ must have the form $\langle A_1{:}T_1', \ldots, A_n{:}T_n' \rangle$, where $T_i' \sqsupseteq S_i$ for each $i$, so the replay derivation is of the form

$$\frac{\gamma', T_1' \curvearrowright v_1' \quad \cdots \quad \gamma', T_n' \curvearrowright v_n'}{\gamma', \langle A_1{:}T_1', \ldots, A_n{:}T_n' \rangle \curvearrowright \langle A_1{:}v_1', \ldots, A_n{:}v_n' \rangle}$$

By Lemma B.5 we know $p.A_i \sqsubseteq v_i$ for each $i$, and $\gamma' \approx_{\rho_i} \gamma$ for each $i$, so by induction $v_i' \approx_{p.A_i} v_i$ for each $i$. Using Lemma B.6 we can conclude that $\langle A_1{:}v_1', \ldots, A_n{:}v_n' \rangle \approx_p \langle A_1{:}v_1, \ldots, A_n{:}v_n \rangle$.

- Empty set. This case is trivial, similar to the usual case for constants.
- Singleton. This case follows immediately from the relevant properties of $p.\epsilon$, using similar reasoning to the record projection case.
- Union. If the last step in the slicing derivation is

$$\frac{p[1], T_1 \searrow \rho_1, S_1 \qquad p[2], T_2 \searrow \rho_2, S_2}{p, T_1 \cup T_2 \searrow \rho_1 \sqcup \rho_2, S_1 \cup S_2}$$

then the evaluation derivation must be of the form

$$\frac{\gamma, T_1 \curvearrowright v_1 \qquad \gamma, T_2 \curvearrowright v_2}{\gamma, T_1 \cup T_2 \curvearrowright 1 \cdot v_1 \uplus 2 \cdot v_2}$$

Let $\gamma' \approx_{\rho_1 \sqcup \rho_2} \gamma$ and $T' \sqsupseteq T_1 \cup T_2$ be given, where $\gamma', T' \curvearrowright v'$. Then $T'$ must have the form $T_1' \cup T_2'$ where $T_i' \sqsupseteq S_i$ so the replay derivation is of the form

$$\frac{\gamma', T_1' \curvearrowright v_1' \qquad \gamma', T_2' \curvearrowright v_2'}{\gamma', T_1' \cup T_2' \curvearrowright 1 \cdot v_1' \uplus 2 \cdot v_2'}$$

By Lemma B.5 we know $p[1] \sqsubseteq v_1$ and $p[2] \sqsubseteq v_2$, and $\gamma' \approx_{\rho_i} \gamma$, so by induction we have $v_1' \approx_{p[1]} v_1$ and $v_2' \approx_{p[2]} v_2$, and therefore by Lemma B.6 we can conclude $1 \cdot v_1' \uplus 2 \cdot v_2' \approx_p 1 \cdot v_1 \uplus 2 \cdot v_2$.

- Sum and emptiness. In both cases, since the slice ensures that the whole argument to the sum or emptiness test is preserved, the argument is straightforward. For example, for emptiness suppose the derivation is of the form:

$$\frac{\diamond, T \searrow \rho, S}{p, \mathtt{empty}\ T \searrow \rho, \mathtt{empty}\ S}$$

Then there are two cases. If replay derivation is of the form

$$\frac{\gamma, T \curvearrowright \emptyset}{\gamma, \mathtt{empty}\ T \curvearrowright \mathtt{true}}$$

Suppose $\gamma' \approx_\rho \gamma$ and $T' \sqsupseteq \mathtt{empty}\ S$ with $\gamma', T' \curvearrowright v'$. Then $T'$ is of the form $\mathtt{empty}\ T''$ with $T'' \sqsupseteq S$, so the replay derivation must be of the form:

$$\frac{\gamma', T'' \curvearrowright v''}{\gamma', \mathtt{empty}\ T'' \curvearrowright v'}$$

By induction, $v'' \approx_\diamond \emptyset$, which implies $v'' = \emptyset$ so $v' = \mathtt{true} \approx_p \mathtt{true}$ (since $p \sqsubseteq \mathtt{true}$). The cases where the argument to $\mathtt{empty}\ T$ evaluates to a nonempty set, and for $\mathtt{sum}\ T$, are similar.

- Comprehension. If the derivation is of the form

$$\frac{p, x.\Theta \searrow^* \rho', \Theta_0, p_0 \qquad p_0, T \searrow \rho, S}{p, \bigcup\{e \mid x \in T\} \triangleright \Theta \searrow \rho \sqcup \rho', \bigcup\{e \mid x \in S\} \triangleright \Theta_0}$$

then the replay derivation must be of the form

$$\frac{\gamma, T \curvearrowright v_0 \qquad \gamma, x \in v_0 \Theta \curvearrowright^* v}{\gamma, \bigcup\{e \mid x \in T\} \triangleright \Theta \curvearrowright v}$$

Suppose $\gamma' \approx_{\rho \sqcup \rho'} \gamma$ and $T' \sqsupseteq \bigcup\{e \mid x \in S\} \triangleright \Theta_0$ with $\gamma', T' \curvearrowright v'$. Then $T'$ must be of the form $\bigcup\{e \mid x \in T''\} \triangleright \Theta'$ for some $T'' \sqsupseteq S$ and $\Theta' \sqsupseteq \Theta_0$ so the replay derivation must be of the form

$$\frac{\gamma', T'' \curvearrowright v_0' \qquad \gamma', x \in v_0', \Theta' \curvearrowright^* v'}{\gamma', \bigcup\{e \mid x \in T''\} \triangleright \Theta' \curvearrowright v'}$$

Since $p_0 \sqsubseteq v_0$ and $\gamma' \approx_\rho \gamma$ and $\gamma' \approx_{\rho'} \gamma$ we have $v_0' \approx_{p_0} v_0$ by induction. Thus, by the second induction hypothesis, since $p \sqsubseteq v$ and $\gamma' \approx_{\rho'} \gamma$ and $v_0' \approx_{p_0} v_0$ we have $v' \approx_p v$.

For part (2), the proof is again by induction on the structure of derivations.

- If the slicing derivation is of the form

$$\overline{\emptyset, x.\emptyset \searrow^* [], \emptyset, \emptyset}$$

then the conclusion is immediate, since rerunning an empty trace set always yields the empty set.
- If the slicing derivation is of the form

$$\overline{\diamond, x.\emptyset \searrow^* [], \emptyset, \emptyset}$$

then the conclusion is immediate as before, since rerunning an empty trace set always yields the empty set.

- If the slicing derivation is of the form

$$\overline{\Box, x.\Theta \searrow^* [], \emptyset, \Box}$$

then the conclusion is immediate, since any two values match according to $\Box$.
- If the slicing derivation is of the form

$$\frac{p[\ell], T \searrow \rho[x \mapsto p_0], S}{p, x.\{\ell.T\} \searrow^* \rho, \{\ell.S\}, \{\ell.p_0\}}$$

then observe that $\Theta \supseteq \{\ell.T\}$ by assumption, so $\Theta(\ell) = T$. So, the replay derivation must have the form

$$\frac{\ell \in \mathrm{dom}(\Theta) \qquad \gamma[x \mapsto v_0], T \curvearrowright v}{\gamma, x \in \{\ell.v_0\}, \Theta \curvearrowright^* \ell \cdot v}$$

Now suppose that $\gamma' \approx_\rho \gamma$ and $v_0' \approx_{\{\ell.p_0\}} \{\ell.v_0\}$ and $\Theta' \sqsupseteq \{\ell.S\}$ are given where $\gamma', x \in v_0', \Theta' \curvearrowright^* v'$. It follows that $v_0' = \{\ell.v_0''\}$ and $v_0'' \approx_{p_0} v_0$, so $\gamma'[x \mapsto v_0''] \approx_{\rho[x \mapsto p_0]} \gamma[x \mapsto v_0]$. Moreover, by inversion the derivation must have the form:

$$\frac{\ell \in \mathrm{dom}(\Theta') \qquad \gamma'[x \mapsto v_0''], \Theta'(\ell) \curvearrowright v'}{\gamma', x \in \{\ell.v_0''\}, \Theta' \curvearrowright^* \ell \cdot v'}$$

Since by Lemma B.5 $p[\ell] \sqsubseteq v$ and $\Theta'(\ell) \sqsupseteq S$ (which holds because $\Theta' \sqsupseteq \{\ell.S\}$), we have by induction that $v' \approx_{p[\ell]} v$, and using Lemma B.6 we can conclude that $\ell \cdot v' \approx_p \ell \cdot v$, as desired.
- Suppose the slicing derivation is of the form:

$$\frac{p|_{\mathrm{dom}(\Theta_1)}, x.\Theta_1 \searrow^* \rho_1, \Theta_1', p_1 \qquad p|_{\mathrm{dom}(\Theta_2)}, x.\Theta_2 \searrow^* \rho_2, \Theta_2', p_2}{p, x.\Theta_1 \uplus \Theta_2 \searrow^* \rho_1 \sqcup \rho_2, \Theta_1' \uplus \Theta_2', p_1 \uplus p_2}$$

and suppose that $\gamma, x \in v_0, \Theta \curvearrowright^* v$ where $\Theta \supseteq \Theta_1 \uplus \Theta_2$. Suppose that $\gamma' \approx_{\rho_1 \sqcup \rho_2} \gamma$ and $v_0' \approx_{p_1 \uplus p_2} v_0$ and $\Theta' \sqsupseteq \Theta_1' \uplus \Theta_2'$ are given, where $\gamma', x \in v_0', \Theta' \curvearrowright^* v'$. We need to show that $v' \approx_p v$.
Since $v_0' \approx_{p_1 \uplus p_2} v_0$, it is straightforward to show that there must exist $v_1, v_2, v_1', v_2'$ such that $v_1 \uplus v_2 = v_0, v_1' \uplus v_2' = v_0', v_1 \approx_{p_1} v_1$ and $v_2' \approx_{p_2} v_2$. Furthermore, by Lemma C.1 we know that $\gamma, x \in v_i, \Theta \curvearrowright^* w_i$ and $\gamma', x \in v_i', \Theta' \curvearrowright^* w_i'$ for some $w_1, w_2, w_1', w_2'$. Therefore, we can conclude that:

$$\frac{\gamma, x \in v_1, \Theta \curvearrowright^* w_1 \qquad \gamma, x \in v_2, \Theta \curvearrowright^* w_2}{\gamma, x \in v_1 \uplus v_2, \Theta \curvearrowright^* w_1 \uplus w_2} \qquad \frac{\gamma', x \in v_1', \Theta' \curvearrowright^* w_1' \qquad \gamma', x \in v_2', \Theta' \curvearrowright^* w_2'}{\gamma', x \in v_1' \uplus v_2', \Theta' \curvearrowright^* w_1' \uplus w_2'}$$

Furthermore, since $v_1 \uplus v_2 = v_0$, by determinacy we know that $w_1 \uplus w_2 = v$ and similarly $w_1' \uplus w_2 = v'$. By induction since $p_i \sqsubseteq v_i$ and $\Theta_i \subseteq \Theta$, we know that $w_i' \approx_{p|_{\mathrm{dom}(\Theta_i)}} w_i$, so we know that

$$v' = w_1' \uplus w_2' \approx_{p|_{\mathrm{dom}(\Theta_i)} \uplus p|_{\mathrm{dom}(\Theta_2)}} w_1 \uplus w_2 = v .$$

To conclude, since $p \sqsubseteq w_1 \uplus w_2$ and $\mathrm{dom}(\Theta_1) \leq w_1$ and $\mathrm{dom}(\Theta_2) \leq w_2$, it follows that $p = p|_{\mathrm{dom}(\Theta_1)} \uplus p|_{\mathrm{dom}(\Theta_2)}$, so we can conclude $w_1 \uplus w_2 \approx_p w_1' \uplus w_2'$ as desired.

This exhausts all cases and completes the proof. $\qquad\Box$

## D. Proof of correctness of query slicing

**Lemma D.1.** *If $\gamma, x \in v_1, e \Downarrow^* v_2$ and $v_1' \subseteq v_1$ then there exists $v_2' \subseteq v_2$ such that $\gamma, x \in v_1', e \Downarrow^* v_2'$.*

*Proof.* The proof is straightforward by induction on derivations. The cases for $v_1 = \emptyset$ and $v_1 = \{\ell.v\}$ are immediate; if $v_1 = w_1 \uplus w_2$ then we proceed by induction using the subsets $w_1' = w_1 \cap v_1'$ and $w_2' = w_2 \cap v_1'$. $\qquad\Box$

We prove Theorem 5.1 by strengthening the induction hypothesis as follows:

**Theorem D.2** (Correctness of Query Slicing)**.**

1. *Suppose $\gamma, T \curvearrowright v$ and $p \sqsubseteq v$ and $p, T \searrow_{\!\!\mathsf{q}} \rho, e$. Then for all $\gamma' \approx_\rho \gamma$ and $e' \sqsupseteq e$ such that $\gamma', e' \Downarrow v'$ we have $v' \approx_p v$.*
2. *Suppose $\gamma, x \in v_0, \Theta \curvearrowright^* v$ and $p \sqsubseteq v$ and $p, \Theta_0 \searrow_{\!\!\mathsf{q}} \rho, e_0, p_0$, where $\Theta_0 \subseteq \Theta$. Then for all $\gamma' \approx_\rho \gamma$ and $v_0' \approx_{p_0} v_0$ and $e_0' \sqsupseteq e_0$ such that $\gamma', x \in v_0', e_0' \Downarrow^* v'$ we have $v' \approx_p v$.*

*Proof.* The proof is by induction on the structure of query slicing derivations. Many of the cases are essentially the same as for trace slicing. The cases for conditionals are straightforward, since in either case the sliced trace and environment retain enough information to force the same branch to be taken on recomputation. We show the details of the cases involving conditionals and comprehensions.
For part (1), we consider a conditional and comprehension rule:

- If the slicing derivation is of the form:

$$\frac{p_1, T_1 \searrow_{\!\!\mathsf{q}} \rho_1, e_1' \qquad \mathtt{true}, T \searrow_{\!\!\mathsf{q}} \rho, e'}{p_1, \mathtt{if}(T, e_1, e_2) \triangleright_{\mathtt{true}} T_1 \searrow_{\!\!\mathsf{q}} \rho_1 \sqcup \rho, \mathtt{if}(e', e_1', \Box)}$$

then the replay derivation must be of the form:

$$\frac{\gamma, T \curvearrowright \texttt{true} \qquad \gamma, T_1 \curvearrowright v_1}{\gamma, \texttt{if}(T, e_1, e_2) \triangleright_{\texttt{true}} T_1 \curvearrowright v_1}$$

Suppose $\gamma' \approx_{\rho_1 \sqcup \rho} \gamma$ and $e'' \sqsupseteq \texttt{if}(e', e_1', \square)$ are given where $\gamma', e' \Downarrow v'$. Then $e'' = \texttt{if}(e_0'', e_1'', e_2'')$, where $e_0'' \sqsupseteq e'$ and $e_1'' \sqsupseteq e_1'$, so there are two cases for the evaluation derivation. If it has the form

$$\frac{\gamma', e_0' \Downarrow \texttt{false} \qquad \gamma', e_2' \Downarrow v_2'}{\gamma', \texttt{if}(e_0', e_1', e_2') \Downarrow v_2'}$$

then since $\gamma' \approx_\rho \gamma$ and $\gamma' \approx_{\rho_1} \gamma$, by induction we would have that $\texttt{false} \approx_{\texttt{true}} \texttt{true}$, which is absurd. So this case cannot arise. Otherwise, the derivation must have the form:

$$\frac{\gamma', e_0' \Downarrow \texttt{true} \qquad \gamma', e_1' \Downarrow v_1'}{\gamma', \texttt{if}(e_0', e_1', e_2') \Downarrow v_1'}$$

Since $\gamma' \approx_\rho \gamma$ and $\gamma' \approx_{\rho_1} \gamma$, by induction we have that $v_1' \approx_{p_1} v_1$ as desired.

- Comprehension. If the derivation is of the form

$$\frac{p, x.\Theta \searrow^* \rho', e_1', p_0 \qquad p_0, T \searrow \rho, e_0'}{p, \bigcup\{e \mid x \in T\} \triangleright \Theta \searrow \rho \sqcup \rho', \bigcup\{e_1' \mid x \in e_0'\}}$$

then the replay derivation must be of the form

$$\frac{\gamma, T \curvearrowright v_0 \qquad \gamma, x \in v_0, \Theta \curvearrowright^* v}{\gamma, \bigcup\{e \mid x \in T\} \triangleright \Theta \curvearrowright v}$$

Suppose $\gamma' \approx_{\rho \sqcup \rho'} \gamma$ and $e'' \sqsupseteq \bigcup\{e_1' \mid x \in e_0'\}$ with $\gamma', e'' \Downarrow v'$. Then $e''$ must be of the form $\bigcup\{e_1'' \mid x \in e_0''\}$ for some $e_1'' \sqsupseteq e_1'$ and $e_0'' \sqsupseteq e_0'$, so the evaluation derivation must be of the form

$$\frac{\gamma', e_0'' \Downarrow v_0' \qquad \gamma', x \in v_0', e_1'' \Downarrow^* v'}{\gamma', \bigcup\{e_1'' \mid x \in e_0''\} \Downarrow v'}$$

Since $p_0 \sqsubseteq v_0$ and $\gamma' \approx_\rho \gamma$ and $\gamma' \approx_{\rho'} \gamma$ we have $v_0' \approx_{p_0} v_0$ by induction. Thus, by the second induction hypothesis, since $p \sqsubseteq v$ and $\gamma' \approx_{\rho'} \gamma$ and $v_0' \approx_{p_0} v_0$ we have $v \approx_p v'$.

For part (2), we consider the singleton and union rules:

- If the slicing derivation is of the form

$$\frac{p[\ell], T \searrow \rho[x \mapsto p_0], e'}{p, x.\{\ell.T\} \searrow^* \rho, e', \{\ell.p_0\}}$$

then recall that by assumption, $\Theta(\ell) \sqsupseteq \{\ell.T\}$, so $\Theta(\ell) = T$, so the replay derivation must have the form

$$\frac{\ell \in \text{dom}(\Theta) \qquad \gamma[x \mapsto v_0], T \curvearrowright v}{\gamma, x \in \{\ell.v_0\}, \Theta \curvearrowright^* \ell \cdot v}$$

Now suppose that $\gamma' \approx_\rho \gamma$ and $v_0' \approx_{\{\ell.p_0\}} \{\ell.v_0\}$ and $e'' \sqsupseteq e'$ are given where $\gamma', x \in v_0', e'' \Downarrow^* v'$. It follows that $v_0' = \{\ell.v_0''\}$ and $v_0'' \approx_{p_0} v_0$, so $\gamma'[x \mapsto v_0''] \approx_{\rho[x \mapsto p_0]} \gamma[x \mapsto v_0]$. Moreover, by inversion the derivation must have the form:

$$\frac{\gamma'[x \mapsto v_0''], e'' \Downarrow v'}{\gamma', x \in \{\ell.v_0''\}, e'' \Downarrow^* \ell \cdot v'}$$

Since by Lemma B.5 $p[\ell] \sqsubseteq v$ we have by induction that $v' \approx_{p[\ell]} v$, and using Lemma B.6 we can conclude that $\ell \cdot v' \approx_p \ell \cdot v$, as desired.

- Suppose the slicing derivation is of the form:

$$\frac{p|_{\text{dom}(\Theta_1)}, x.\Theta_1 \searrow^* \rho_1, e_1', p_1 \qquad p|_{\text{dom}(\Theta_2)}, x.\Theta_2 \searrow^* \rho_2, e_2', p_2}{p, x.\Theta_1 \uplus \Theta_2 \searrow^* \rho_1 \sqcup \rho_2, e_0', p_1 \uplus p_2}$$

and suppose that $\gamma, x \in v_0, \Theta \curvearrowright^* v$ where $\Theta \sqsupseteq \Theta_1 \uplus \Theta_2$. Suppose that $\gamma' \approx_{\rho_1 \sqcup \rho_2} \gamma$ and $v_0' \approx_{p_1 \uplus p_2} v_0$ and $e_0' \sqsupseteq e_0$ are given, where $\gamma', x \in v_0', e_0' \Downarrow^* v'$. We need to show that $v' \approx_p v$.

Since $v_0' \approx_{p_1 \uplus p_2} v_0$, it is straightforward to show that there must exist $v_1, v_2, v_1', v_2'$ such that $v_1 \uplus v_2 = v_0$, $v_1' \uplus v_2' = v_0'$, $v_1 \approx_{p_1} v_1$ and $v_2' \approx_{p_2} v_2$. Furthermore, by Lemmas C.1 and D.1 we know that $\gamma, x \in v_i, \Theta \curvearrowright^* w_i$ and $\gamma', x \in v_i', e_0' \Downarrow^* w_i'$ for some $w_1, w_2, w_1', w_2'$. Therefore, we can conclude that:

$$\frac{\gamma, x \in v_1, \Theta \curvearrowright^* w_1 \qquad \gamma, x \in v_2, \Theta \curvearrowright^* w_2}{\gamma, x \in v_1 \uplus v_2, \Theta \curvearrowright^* w_1 \uplus w_2} \qquad \frac{\gamma', x \in v_1', e_0' \Downarrow^* w_1' \qquad \gamma', x \in v_2', e_0' \Downarrow^* w_2'}{\gamma', x \in v_1' \uplus v_2', e_0' \Downarrow^* w_1' \uplus w_2'}$$

Furthermore, since $v_1 \uplus v_2 = v_0$, by determinacy we know that $w_1 \uplus w_2 = v$ and similarly $w_1' \uplus w_2' = v'$. By induction since $p_i \sqsubseteq v_i$ and $\Theta_i \subseteq \Theta$, we know that $w_i' \approx_{p|_{\text{dom}(\Theta_i)}} w_i$, so we know that

$$v' = w_1' \uplus w_2' \approx_{p|_{\text{dom}(\Theta_i)} \uplus p|_{\text{dom}(\Theta_2)}} w_1 \uplus w_2 = v \ .$$

To conclude, since $p \sqsubseteq w_1 \uplus w_2$ and $\mathrm{dom}(\Theta_1) \leq w_1$ and $\mathrm{dom}(\Theta_2) \leq w_2$, it follows that $p = p|_{\mathrm{dom}(\Theta_1)} \uplus p|_{\mathrm{dom}(\Theta_2)}$, so we can conclude $w_1 \uplus w_2 \eqsim_p w_1' \uplus w_2'$ as desired.

This exhausts all cases and completes the proof. $\qquad\square$