

A Core Calculus for Provenance

Umut A. Acar¹, Amal Ahmed², James Cheney³, and Roly Perera¹

¹ Max Planck Institute for Software Systems {umut, rolyp}@mpi-sws.org

² Indiana University amal@cs.indiana.edu

³ University of Edinburgh jcheney@inf.ed.ac.uk

Abstract. Provenance is an increasing concern due to the revolution in sharing and processing scientific data on the Web and in other computer systems. It is proposed that many computer systems will need to become provenance-aware in order to provide satisfactory accountability, reproducibility, and trust for scientific or other high-value data. To date, there is not a consensus concerning appropriate formal models or security properties for provenance. In previous work, we introduced a formal framework for provenance security and proposed formal definitions of properties called disclosure and obfuscation.

This paper develops a core calculus for provenance in programming languages. Whereas previous models of provenance have focused on special-purpose languages such as workflows and database queries, we consider a higher-order, functional language with sums, products, and recursive types and functions. We explore the ramifications of using traces based on operational derivations for the purpose of comparing other forms of provenance. We design a rich class of provenance views over traces. Finally, we prove relationships among provenance views and develop some solutions to the disclosure and obfuscation problems.

1 Introduction

Provenance, or meta-information about the origin, history, or derivation of an object, is now recognized as a central challenge in establishing trust and providing security in computer systems, particularly on the Web. Essentially, provenance management involves instrumenting a system with detailed monitoring or logging of auditable records that help explain how results depend on inputs or other (sometimes untrustworthy) sources. The security and privacy ramifications of provenance must be understood in order to safely meet the needs of users that desire provenance without introducing new security vulnerabilities or compromising the confidentiality of other users.

The lack of adequate provenance information can cause (and has caused) major problems, which we call *provenance failures* [11]. Essentially, a provenance failure can arise either from failure to *disclose* some key provenance information to users (for example, if a years-out-of-date story causes investors to panic about a company's financial stability [7]), or from failure to *obfuscate* some sensitive provenance information (for example, if a Word document is published with supposedly secret contributors' identities logged in its change history [27]). To address these problems, a number of forms of provenance have been proposed for different computational models, including *why* and *where* provenance [6], *dependency* provenance [9], and a variety of other ad hoc techniques [3, 24].

Prior work on provenance and security. Our previous work [9] appears to have been the first to explicitly relate information-flow security to a form of provenance, called *dependency provenance*. Provenance has been studied in language-based security by Cirillo et al. [13] and by Swamy et al. [26, 25]. Both focus on specifying and enforcing security policies involving provenance tracking alongside many other concerns, and not on defining provenance semantics or extraction techniques. Work on secure auditing [21, 19] and expressive programming languages for security is also related, but this work focuses on explicitly manipulating proofs of authorization or evidence about protocol or program runs rather than automatically deriving or securing provenance information.

There is also some work directly addressing security for provenance [12, 8, 15]. Chong [12] gave candidate definitions of *data security* and *provenance security* using a trace semantics, based in part on an earlier version of our trace model. Davidson et al. [15] studied privacy for provenance in scientific workflows, focusing on complexity lower bounds. Cheney [8] gave an abstract framework for provenance, proposed definitions of properties called *obfuscation* and *disclosure*, and discussed algorithms and complexity results for instances of this framework including finite automata, workflows, and the semiring model of database provenance [18].

In this paper, we build on prior work on provenance security by studying the disclosure and obfuscation properties of different forms of provenance in the context of a higher-order, pure, functional language. To illustrate what we mean by provenance, we present examples of programming with three different forms of provenance in Transparent ML (TML), a prototype implementation of the ideas of this paper.

1.1 Examples

Where-provenance. Where-provenance [6, 5] identifies at most one source location from which a part of the output was copied. For example, consider the following TML session:

```
- f [(1,2), (4,3), (5,6)];
val it = [(5,6), (3,4), (1,2)]
```

Without access to the source code, one can guess that f is doing something like

$$\text{reverse} \circ (\text{map } (\lambda(x,y).\text{if } x < y \text{ then } (x,y) \text{ else } (y,x)))$$

However, by providing where-provenance information, the system can explain whether the numbers in the result were copied from the input or constructed in some other way:

```
- trace (f [(1@L1,2@L2), (4@L3,3@L4), (5@L5,6@L6)]);
it = <trace> : ({L1:int,...}, (int*int) list) trace
- where it;
val it = [(5@L5,6), (3@L4,4), (1@L1,2)]
```

This shows that f contrives to copy the first elements of the returned pairs but construct the second components.

Dependency provenance. Dependency provenance [9] is an approach that tracks a set of all source locations on which a result depends. For example, if we have:

```
- g [(1,2,3), (4,5,6)];  
val it = [6,6] : int list
```

we again cannot tell much about what g does. By tracing and asking for dependency provenance, we can see:

```
- trace (g [(1@L1,2@L2,3@L3), (4@L4,5@L5,6@L6)]);  
val it = <trace> : ({L1:int,...}, int list) trace  
- dependency it;  
val it = [6@{L1,L2,L3}, 6@{L1,L2,L3}]
```

This suggests that g is simply summing the first triple and returning the result twice, without examining the rest of the list. We can confirm this as follows:

```
- trace (g ((1@L1,2@L2,3@L3)::[]@L));  
val it = <trace> : ({L1:int,...}, int list) trace  
- dependency it;  
val it = [6@{L1,L2,L3}]
```

The fact that L does not appear in the output confirms that g does not look further into the list.

Expression provenance. A third common form of provenance is an expression graph or tree that shows how a value was computed by primitive operations. For example, consider:

```
- (h 3, h 4, h 5)  
val it = (6,24,120);
```

We might conjecture that h is actually the factorial function. By tracing h and extracting expression provenance, we can confirm this guess (at least for the given inputs):

```
- trace (h (4@L));  
val it = <trace> : ({L:int}, int) trace  
- expression it;  
val it = 24@{L * (L-1) * (L-2) * (L-3) * 1}
```

In this case where-provenance and dependency provenance would be uninformative since the result is not copied from, and obviously depends on, the input.

This kind of provenance is used extensively in *workflow* systems often used in e-science [20], where the main program is a high-level process coordinating a number of external (and often concurrent) program or RPC calls, for example, image-processing steps or bulk data transformations, which we could model by adding primitive image-processing operations and types to our language. Thus, even though the above examples use fine-grained primitive operations, this model is also useful for coarse-grained provenance-tracking.

Provenance security. The three models of provenance above represent useful forms of provenance that might increase users’ trust or confidence that they understand the results of a program. However, if the underlying data, or the structure of the computation, is sensitive, then making this information available may lead to inadvertent vulnerabilities, by making it possible for users to infer information they cannot observe directly. This is a particular problem if we wish to disclose part of the result of a program, and provenance that justifies part of the result, while keeping other parts of the program’s execution, input, or output confidential. As a simple example, consider a program `if $x \neq 1$ then (y, y) else (z, w)` . If x is sensitive, but z and w happen to both equal 42, then it is safe to reveal the result $(42, 42)$. However, any of the above forms of provenance make it possible to distinguish which branch was taken because the two different copies of 42 in z and w will have different provenance. Thus, if the provenance information is released then a principal can infer that the second branch was taken, and hence, $x = 1$. In technical terms, we cannot *disclose* any of the above forms of provenance for the result while *obfuscating* the fact that $x = 1$.

1.2 Summary

Contributions. In this paper, we build on, and refine, the provenance security framework previously introduced by Cheney [8]. We introduce a core language with replayable execution traces for a call-by-value, higher-order functional language, and make the following technical contributions:

- Refined definitions of obfuscation and disclosure (Sec. 2).
- A core calculus defining traced execution for a pure functional programming language (Sec. 3).
- A generic provenance extraction framework that includes several previously-studied forms of provenance as instances (Sec. 4).
- An analysis of disclosure and obfuscation guarantees provided by different forms of provenance, including techniques based on slicing execution traces (Sec. 5).

Outline. Section 2 briefly recapitulates the framework introduced by Cheney [8] and refines some definitions. We present the (standard) syntax and tracing semantics of TML in Section 3. In Section 4 we introduce a framework for querying and extracting provenance views from traces, including the three models discussed above. Section 5 presents our main results about disclosure, obfuscation, and trace slicing. Section 6 presents related work and Section 7 concludes.

2 Background

We recapitulate the main components of the provenance security framework of Cheney [8]. The framework assumes a given set of *traces* \mathcal{T} , together with a collection \mathcal{Q} of possible *trace queries* $Q : \mathcal{T} \rightarrow \mathbb{B}$. These represent properties of traces that the system designer may want to protect or that legitimate users or attackers of the system may want to learn. In the previous paper, we considered refinements to take into account

the knowledge of the principals about the possible system behaviors. In this paper, we assume that all traces \mathcal{T} are considered possible by all principals, for simplicity.

For each principal A , fix a set Ω_A of the possible *provenance views* offered to A , and a function $P_A : \mathcal{T} \rightarrow \Omega_A$ mapping each trace to A 's provenance view of the trace. For the purposes of this paper, we do not consider interactions among multiple principals, so we typically consider only one principal and omit the A subscripts. We may write $(\Omega, P : \mathcal{T} \rightarrow \Omega)$ or just (Ω, P) for a provenance view. Also, we sometimes write $Q : \Omega \rightarrow \mathbb{B}$ for a provenance query, that is, a query on a provenance view.

Given this framework, we proposed the following definitions:

Definition 1 (Disclosure). *A query Q is disclosed by a provenance view (Ω, P) if for every $t, t' \in \mathcal{T}$, if $P(t) = P(t')$ then $Q(t) = Q(t')$.*

In other words, disclosure means that there can be no traces t, t' that have the same provenance view but where one satisfies the query and the other does not.

Definition 2 (Obfuscation). *A query Q is obfuscated by a provenance view (Ω, P) if for every t in \mathcal{T} , there exists $t' \in \mathcal{T}$ such that $P(t) = P(t')$ and $Q(t) \neq Q(t')$.*

Thus, obfuscation is not exactly the opposite of disclosure; instead, it means that for every trace there is another trace with the same provenance view but different Q -value. This means that a principal that has access to the provenance view but not the trace cannot be certain that Q is satisfied or not satisfied by the underlying trace.

The definitions above turn out to be too strong; in this paper we will also consider some weaker versions of disclosure and obfuscation.

Definition 3. *We say that $P : \mathcal{T} \rightarrow \Omega$ positively discloses $Q : \mathcal{T} \rightarrow \mathbb{B}$ via query $Q' : \Omega \rightarrow \mathbb{B}$ if for every t , if $Q'(P(t))$ then $Q(t)$.*

In other words, positive disclosure means that there is a query Q' on the provenance that safely overapproximates Q on the underlying trace. If $Q'(P(t))$ holds then we know $Q(t)$ holds but otherwise we may not learn anything about t .

Definition 4. *We say that P positively obfuscates $Q : \mathcal{T} \rightarrow \mathbb{B}$ if for every t satisfying Q there exists a trace t' falsifying Q such that $P(t) = P(t')$.*

In other words, positive obfuscation means that the provenance never reveals that Q holds of the trace, but it may reveal that Q fails. This weaker notion is useful for asserting that sensitive data is protected: if the sensitive data is not present in the trace then it is harmless to reveal this, but if the sensitive data is present then the provenance should hide enough information to make its presence uncertain.

Dual notions of negative disclosure and obfuscation can be defined as positive disclosure or obfuscation of $\neg Q$ respectively.

Proposition 1. *If P both positively discloses and negatively discloses Q via Q' , then P discloses Q . Similarly, if P both positively and negatively obfuscates Q then P obfuscates Q .*

In the previous paper, we gave several examples of instances of this framework. Here, for illustration, we just review one such instance, given by finite automata.

Types	$\tau ::= b \mid \tau_1 \times \tau_2 \mid \tau_1 + \tau_2 \mid \tau_1 \rightarrow \tau_2 \mid \mu\alpha.\tau \mid \alpha$
Type contexts	$\Gamma ::= [x_1 : \tau_1, \dots, x_n : \tau_n]$
Code pointers	$\kappa ::= f(x).e$
Matches	$m ::= \{\text{inl}(x_1).e_1; \text{inr}(x_2).e_2\}$
Values	$v ::= c \mid (v_1, v_2) \mid \text{inl}(v) \mid \text{inr}(v) \mid \text{roll}(v) \mid \langle \kappa, \gamma \rangle$
Expressions	$e ::= c \mid x \mid \oplus(\bar{e}) \mid \text{let } x = e_1 \text{ in } e_2$ $\mid (e_1, e_2) \mid \text{fst}(e) \mid \text{snd}(e) \mid \text{inl}(e) \mid \text{inr}(e)$ $\mid \text{roll}(e) \mid \text{unroll}(e) \mid \text{fun } \kappa \mid \text{case } e \text{ of } m \mid (e e')$
Traces	$T ::= c \mid x \mid \oplus(\bar{T}) \mid \text{let } x = T_1 \text{ in } T_2$ $\mid (T_1, T_2) \mid \text{fst}(T) \mid \text{snd}(T) \mid \text{inl}(T) \mid \text{inr}(T)$ $\mid \text{roll}(T) \mid \text{unroll}(T) \mid \text{fun } \kappa$ $\mid \text{case } T \triangleright_{\text{inl}} x.T_1 \mid \text{case } T \triangleright_{\text{inr}} x.T_2 \mid (T_1 T_2) \triangleright_{\kappa} f(x).T$
Environments	$\gamma ::= [x_1 \mapsto v_1, \dots, x_n \mapsto v_n]$

Fig. 1. Abstract syntax of Core TML.

Example 1 (Automata provenance framework). The set of traces \mathcal{T}_M of an automaton $M = (\Sigma, Q, q_0, \delta, F)$ is the set $Q(\Sigma Q)^*$ of alternating sequences of states and alphabet letters. The queries are simply regular subsets of \mathcal{T}_M . The provenance views are given by finite-state transducers. We showed that disclosure is decidable for all queries and views and that obfuscation is decidable for all queries and views whose range is finite. It is unknown whether obfuscation is decidable in the general case.

We now proceed to instantiate the framework with traces generated by a much richer language, with corresponding notions of trace query and provenance view.

3 Core Language

We will develop a core language for provenance based on a standard, typed, call-by-value, pure language, called Transparent ML, or TML. The syntax of TML types, expressions, traces, and other syntactic classes is shown in Figure 1. We include standard constructs for dealing with binary pairs, binary sums, recursive types, and recursive functions; more general constructs such as records, datatypes, or simultaneous recursive functions can of course be handled without difficulty. In $f(x).e$, both f and x are variable names; f is the name of the recursively defined function while x is the name of the argument.

We abbreviate functional terms of the form $f(x).e$ using the letter κ , when convenient; similarly, we often abbreviate the expression $\{\text{inl}(x_1).e_1; \text{inr}(x_2).e_2\}$ as m . We sometimes refer to κ or m as a *code pointer* or *match pointer* respectively; in a fixed program, there are a fixed finite number of such terms and so we can share them instead of explicitly copying them when used in traces.

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

$$\frac{}{\gamma, x \Downarrow \gamma(x), x} \quad \frac{}{\gamma, c \Downarrow c, c} \quad \frac{}{\gamma, \mathbf{fun} \kappa \Downarrow \langle \kappa, \gamma \rangle, \mathbf{fun} \kappa} \quad \frac{\gamma, \bar{e} \Downarrow \bar{v}, \bar{T}}{\gamma, \oplus(\bar{e}) \Downarrow \hat{\oplus}(\bar{v}), \oplus(\bar{T})}$$

$$\frac{\frac{\gamma, e_1 \Downarrow v_1, T_1 \quad \gamma, e_2 \Downarrow v_2, T_2}{\gamma, (e_1, e_2) \Downarrow (v_1, v_2), (T_1, T_2)} \quad \frac{\gamma, e \Downarrow (v_1, v_2), T}{\gamma, \mathbf{fst}(e) \Downarrow v_1, \mathbf{fst}(T)}}{\gamma, \mathbf{inl}(e) \Downarrow \mathbf{inl}(v), \mathbf{inl}(T)} \quad \frac{\frac{\gamma, e_1 \Downarrow v_1, T_1 \quad \gamma[x \mapsto v_1], e_2 \Downarrow v_2, T_2}{\gamma, \mathbf{let} x = e_1 \mathbf{in} e_2 \Downarrow v_2, \mathbf{let} x = T_1 \mathbf{in} T_2}}{\gamma, \mathbf{roll}(e) \Downarrow \mathbf{roll}(v), \mathbf{roll}(T)} \quad \frac{\gamma, e \Downarrow \mathbf{roll}(v), T}{\gamma, \mathbf{unroll}(e) \Downarrow v, \mathbf{unroll}(T)}$$

$$\frac{\frac{(\mathbf{inl}(x_1).e_1 \in m) \quad \gamma, e \Downarrow \mathbf{inl}(v), T \quad \gamma[x_1 \mapsto v], e_1 \Downarrow v_1, T_1}{\gamma, \mathbf{case} e \mathbf{of} m \Downarrow v_1, \mathbf{case} T \triangleright_{\mathbf{inl}} x_1.T_1}}{\gamma, \mathbf{case} e \mathbf{of} m \Downarrow v_2, \mathbf{case} T \triangleright_{\mathbf{inr}} x_2.T_2}} \quad \frac{(\mathbf{inr}(x_2).e_2 \in m) \quad \gamma, e \Downarrow \mathbf{inr}(v), T \quad \gamma[x_2 \mapsto v], e_2 \Downarrow v_2, T_2}{\gamma, \mathbf{case} e \mathbf{of} m \Downarrow v_2, \mathbf{case} T \triangleright_{\mathbf{inr}} x_2.T_2}}$$

$$\frac{(\kappa = f(x).e) \quad \frac{\gamma, e_1 \Downarrow \langle \kappa, \gamma' \rangle, T_1}{\gamma, e_2 \Downarrow v_2, T_2} \quad \gamma'[f \mapsto \langle \kappa, \gamma' \rangle, x \mapsto v_2], e \Downarrow v, T}{\gamma, (e_1 e_2) \Downarrow v, (T_1 T_2) \triangleright_{\kappa} f(x).T}}$$

Fig. 2. Dynamic semantics of Core TML: selected rules for expression evaluation.

3.1 Dynamic Semantics

We augment a standard large-step operational semantics for TML by adding a parameter T , which records a trace of the evaluation of the expression. The judgment $\gamma, e \Downarrow v, T$ says that in environment γ , expression e evaluates to value v with trace T . (Traces were defined in Figure 1.) Many of the trace forms are isomorphic to the corresponding expression forms. The exceptions are the case and application evaluation rules. In either case, the first argument is evaluated to determine what expression to evaluate to obtain the final result. For case traces, we record the trace of the case scrutinee and the taken branch. Traces can contain free variables and so we re-bind the variable in the trace of the taken branch. Similarly for an application we record the traces of the function subexpression and the argument subexpression, and also record the trace of the evaluation of the body of the function. Again, since the body trace can mention the function and argument names as free variables, we re-bind these variables.

We want to emphasize at this point that we do not necessarily expect that implementations routinely construct fully detailed traces along the above lines. Rather, the trace semantics is proposed here as a candidate for the most detailed form of provenance we will consider. Recording and compressing or filtering relevant information from traces in an efficient way is beyond the scope of this paper.

Trace Replay. We equip traces with a semantics that relates them to expressions. We write $\gamma, T \rightsquigarrow v$ for the *replay* relation that reruns a trace on an environment (possibly different from the one originally used to construct T). The rules for most trace forms are the same as the standard rules for evaluating the corresponding expression forms.

$$\boxed{\gamma, T \rightsquigarrow v}$$

$$\frac{\gamma, T \rightsquigarrow \mathbf{inl}(v) \quad \gamma[x_1 \mapsto v], T_1 \rightsquigarrow v_1}{\gamma, \mathbf{case } T \triangleright_{\mathbf{inl}} x_1.T_1 \rightsquigarrow v_1} \quad \frac{\gamma, T \rightsquigarrow \mathbf{inr}(v) \quad \gamma[x_2 \mapsto v], T_2 \rightsquigarrow v_2}{\gamma, \mathbf{case } T \triangleright_{\mathbf{inr}} x_2.T_2 \rightsquigarrow v_2}$$

$$\frac{\gamma, T_1 \rightsquigarrow \langle \kappa, \gamma' \rangle \quad \gamma, T_2 \rightsquigarrow v_2, T_2' \quad \gamma'[f \mapsto \langle \kappa, \gamma' \rangle, x \mapsto v_2], T \rightsquigarrow v}{\gamma, (T_1 T_2) \triangleright_{\kappa} f(x).T \rightsquigarrow v}$$

Fig. 3. Dynamic semantics of Core TML: selected rules for trace replay.

Figure 3 shows the rules for replaying case and application traces. Essentially, these rules require that the same control flow branches are taken as in the original run. If the input environment is different enough that the same branches cannot be taken, then replay fails.

3.2 Basic Properties of Traces

In this section, we identify key properties of traces, including type safety, and the consistency and fidelity properties that characterize how traces record the evaluation of an expression.

Determinacy and Type Safety. We employ a standard type system for expressions, with straightforward extensions to handle traces. Determinacy of typechecking and type-safety can be established for expression evaluation and trace replay. Types do not play a significant role in the main technical results, however, so we elide the details.

Consistency and Fidelity. We say that a trace T is *consistent* with an environment γ if there exists v such that $\gamma, T \rightsquigarrow v$. Evaluation produces consistent traces, and replaying a trace on the same input yields the same value:

Theorem 1 (Consistency). *If $\gamma, e \Downarrow v, T$ then $\gamma, T \rightsquigarrow v$.*

Furthermore, the trace produced by evaluation is faithful to the original expression, in the sense that whenever the trace can be successfully replayed on a different input, the result (and its trace) is the same as what we would obtain by rerunning e from scratch, and the resulting trace is the same as well. We call this property *fidelity*.

Theorem 2 (Fidelity). *If $\gamma, e \Downarrow v, T$ and $\gamma', T \rightsquigarrow v'$ then $\gamma', e \Downarrow v', T$.*

4 Provenance Views and Trace Queries

4.1 Provenance Extraction

Many previous approaches to provenance can be viewed as performing a form of *annotation propagation*. The idea is to decorate the input with annotations (often, initially, unique identifiers) and propagate the annotations through the evaluation. For example, in where-provenance, annotations are optional tags that can be thought of as pointers showing where output data was copied from in the source [6, 5]. Other techniques, such

$$\begin{aligned}
F(x, \hat{\gamma}) &= \hat{\gamma}(x) \\
F(\mathbf{let} \ x = T_1 \ \mathbf{in} \ T_2, \hat{\gamma}) &= F(T_2, \hat{\gamma}[x \mapsto F(T_1, \hat{\gamma})]) \\
F(c, \hat{\gamma}) &= c^{F_c} \\
F(\oplus(T_1, \dots, T_n), \hat{\gamma}) &= (\hat{\oplus}(v_1, \dots, v_n))^{F_{\oplus}(a_1, \dots, a_n)} \\
&\quad \text{where } v_i^{a_i} = F(T_i, \hat{\gamma}) \\
F((T_1, T_2), \hat{\gamma}) &= (F(T_1, \hat{\gamma}), F(T_2, \hat{\gamma}))^{\perp} \\
F(\mathbf{fst}(T), \hat{\gamma}) &= v_1^{F_1(a,b)} \\
&\quad \text{where } (v_1^b, \hat{v}_2)^a = F(T, \hat{\gamma}) \\
F(\mathbf{inl}(T), \hat{\gamma}) &= \mathbf{inl}(F(T, \hat{\gamma}))^{\perp} \\
F((\mathbf{case} \ T) \triangleright_{\mathbf{inl}} \ x.T_1, \hat{\gamma}) &= v^{F_L(a,b)} \\
&\quad \text{where } \mathbf{inl}(\hat{v})^a = F(T, \hat{\gamma}) \\
&\quad \text{and } v^b = F(T_1, \hat{\gamma}[y \mapsto \hat{v}]) \\
F(\mathbf{fun} \ \kappa, \hat{\gamma}) &= \langle \kappa, \hat{\gamma} \rangle^{F_{\kappa}} \\
F((T_1 \ T_2) \triangleright_{\kappa} \ f(x).T, \hat{\gamma}) &= v^{F_{\text{app}}(a,b)} \\
&\quad \text{where } \langle \kappa, \hat{\gamma}' \rangle^a = F(T_1, \hat{\gamma}) \\
&\quad \text{and } \hat{v}_2 = F(T_2, \hat{\gamma}) \\
&\quad \text{and } v^b = F(T, \hat{\gamma}'[f \mapsto \langle \kappa, \hat{\gamma}' \rangle^a, x \mapsto \hat{v}_2])
\end{aligned}$$

Fig. 4. Generic extraction (selected rules).

as why-, how-, dependency, and workflow provenance, can also be defined in terms of annotation propagation [18, 17, 4, 10].

Based on this observation, we define a provenance extraction framework in which values are decorated with annotations and extraction functions take traces and return annotated values that can be interpreted as useful provenance information. We apply this framework to specify several concrete annotation schemes and extraction functions.

Extraction framework. Let A be an arbitrary set of *annotations* a , which we usually assume includes a blank annotation \perp and a countably infinite set of identifiers $\ell \in \text{Loc}$, called *locations*. We define A -*annotated values* \hat{v} (or just *annotated values*, when A is clear) using the following grammar:

$$\begin{aligned}
\hat{v} &::= v^a & \hat{\gamma} &::= [x_1 \mapsto \hat{v}_1, \dots, x_n \mapsto \hat{v}_n] \\
v &::= c \mid (\hat{v}_1, \hat{v}_2) \mid \mathbf{inl}(\hat{v}) \mid \mathbf{inr}(\hat{v}) \mid \langle \kappa, \hat{\gamma} \rangle \mid \mathbf{roll}(\hat{v})
\end{aligned}$$

We write $\hat{\gamma}$ for *annotated environments* mapping variables to annotated values. We define an erasure function $|\hat{v}|$ that maps each annotated value to an ordinary value by erasing the annotations. Similarly, $|\hat{\gamma}|$ is the ordinary environment obtained by erasing the annotations from the values of $\hat{\gamma}$.

We will define a family of provenance extraction functions $F(T, \hat{\gamma})$ that take a trace T and an environment $\hat{\gamma}$ and return an annotated value. Each such F can be specified by giving the following annotation-propagation functions:

$$\begin{aligned}
F_c, F_{\kappa} &: A \\
F_1, F_2, F_L, F_R, F_{\text{app}}, F_{\text{unroll}} &: A \times A \rightarrow A \\
F_{\oplus} &: A \times \dots \times A \rightarrow A
\end{aligned}$$

Each function shows how the annotations involved in the corresponding computational step propagate to the result. For example, $F_1(a, b)$ gives the annotation on the result of a `fst`-projection, where a is the annotation on the pair and b is the annotation of the first element. Figure 4 shows how to propagate annotations through a trace given basic annotation-propagation functions.

Remark 1. The extraction framework hard-wires the behavior of certain operations such as `let`, `inl()`, `inr()`, and application. It would also be possible to extend the framework to provide to customize their behavior; however, this functionality is not needed by any of the forms of provenance in this paper, and it is not clear whether there are natural provenance models that require them.

Theorem 3. *Every generic provenance extraction function is compatible with replay: that is, for any $\hat{\gamma}, T, v$, if $|\hat{\gamma}|, T \rightsquigarrow v$ then $|F(T, \hat{\gamma})| = v$.*

Where-provenance. Where-provenance can be defined via an annotation-propagating semantics where annotations are either labels ℓ or the blank annotation \perp . We define the where-provenance semantics $W(T, \hat{\gamma})$ using the following annotation-propagation functions:

$$\begin{aligned} W_c, W_\kappa &= \perp \\ W_1, W_2, W_L, W_R, W_{\text{app}}, W_{\text{unroll}} &= \lambda(x, y).y \\ W_\oplus &= \lambda(a_1, \dots, a_n).\perp \end{aligned}$$

Essentially, these functions preserve the annotations of data that are copied, and annotate computed or constructed data with \perp . This semantics is similar to that in Buneman et al. [5] and previous treatments of where-provenance in databases, adapted to TML.

Theorem 4. *Suppose $|\hat{\gamma}|, e \Downarrow v', T$. If an annotated value v^a appears in $W(T, \hat{\gamma})$ with annotation $a \neq \perp$, then v^a is an exact copy (including any nested annotations) of a part of $\hat{\gamma}$.*

Expression provenance. To model expression provenance, we consider *expression annotations* t consisting of labels ℓ , blanks \perp , constants c , or primitive function applications $\oplus(t_1, \dots, t_n)$. We define expression-provenance extraction $E(T, \hat{\gamma})$ in much the same way as W , with the following differences:

$$E_c = c \quad E_\oplus(t_1, \dots, t_n) = \oplus(t_1, \dots, t_n)$$

It would also be straightforward to define a translation from traces to provenance graphs (for example, Open Provenance Model graphs [23]).

The correctness property for expression provenance states that the expression annotation correctly recomputes the value it annotates.

Theorem 5. *Suppose $|\hat{\gamma}|, e \Downarrow v', T$, where each subvalue in $\hat{\gamma}$ is annotated with a copy of itself. If an annotated value v^e appears in $E(T, \hat{\gamma})$ with $e \neq \perp$, then e is a closed expression evaluating to $|v|$.*

Dependency provenance. To extract dependency provenance, we will use annotations ϕ that are sets of source locations $\{\ell_1, \dots, \ell_n\}$. Initial annotations consist of distinct singleton sets $\{\ell\}$. We define $D(T, \hat{\gamma})$ using the following propagation functions:

$$\begin{aligned} D_c, D_\kappa &= \emptyset \\ D_1, D_2, D_L, D_R, D_{\text{app}}, D_{\text{unroll}} &= \lambda(x, y). x \cup y \\ D_\oplus &= \lambda(a_1, \dots, a_n). a_1 \cup \dots \cup a_n \end{aligned}$$

This semantics is based on the dynamic provenance tracking semantics given by Cheney et al. [9], generalized to TML.

This definition satisfies the *dependency-correctness* property introduced in [9]. This property requires an auxiliary relation \approx_ℓ that says that two annotated values are equal except possibly at parts labeled by ℓ , whose straightforward definition we omit. Then we can show:

Theorem 6. *Suppose $|\hat{\gamma}|, e \Downarrow v, T$ and $\hat{\gamma}' \approx_\ell \hat{\gamma}$ and $|\hat{\gamma}'|, e \Downarrow v', T'$. Then we have $D(T, \hat{\gamma}) \approx_\ell D(T', \hat{\gamma}')$.*

This says that the label of a value in the input propagates to all parts of the output where changing the value can have an impact on the result.

Path annotations. For annotations to be useful when the full input is unavailable, we consider annotations where the locations ℓ are *paths* that uniquely address parts of the input environment. We write $\text{path}(\gamma)$ for the environment γ with each component annotated with the path to that component. More generally, we define $\text{path}_p(\gamma)$ as $[x_1 := \text{path}_{p.x}(\gamma(x_1)), \dots, x_n := \text{path}_{p.x_n}(\gamma(x_n))]$ where $\text{path}_p(v)$ is defined as follows:

$$\begin{aligned} \text{path}_p(c) &= c^p \\ \text{path}_p((v_1, v_2)) &= (\text{path}_{p.1}(v_1), \text{path}_{p.2}(v_2))^p \\ \text{path}_p(\text{inl}(v)) &= \text{inl}(\text{path}_{p.1}(v))^p \\ \text{path}_p(\text{inr}(v)) &= \text{inr}(\text{path}_{p.1}(v))^p \\ \text{path}_p(\langle \kappa, \gamma \rangle) &= \langle \kappa, \text{path}_p(\gamma) \rangle^p \end{aligned}$$

For example, $\text{path}([x = (1, 2), y = \text{inl}(4)]) = [x = (1^{x.1}, 2^{x.2})^x, y = \text{inl}(4^{y.\text{inl}})^y]$.

4.2 Patterns, partial traces, and trace queries

We introduce patterns for values, environments and traces. The syntax of patterns (pattern environments) is similar to that of values (respectively environments), extended with special *holes*:

$$\begin{aligned} p &::= c \mid (p_1, p_2) \mid \text{inl}(p) \mid \text{inr}(p) \mid \text{roll}(p) \mid \langle \kappa, \rho \rangle \mid \diamond \mid \square \\ \rho &::= [x_1 \mapsto p_1, \dots, x_n \mapsto p_n] \end{aligned}$$

Patterns actually denote binary relations on values. The hole symbol \square denotes the total relation, while the exact-match symbol \diamond denotes the identity relation. (The \diamond pattern is used in backward disclosure slicing.)

$$\boxed{v \approx_p v}$$

$$\frac{}{v \approx_{\square} v'} \quad \frac{}{v \approx_{\diamond} v} \quad \frac{}{c \approx_c c} \quad \frac{v_1 \approx_{p_1} v'_1 \quad v_2 \approx_{p_2} v'_2}{(v_1, v_2) \approx_{(p_1, p_2)} (v'_1, v'_2)}$$

$$\frac{v \approx_p v' \quad C \in \{\text{inl}, \text{inr}, \text{roll}\}}{C(v) \approx_{C(p)} C(v')} \quad \frac{\gamma \approx_{\rho} \gamma'}{\langle \kappa, \gamma \rangle \approx_{\langle \kappa, \rho \rangle} \langle \kappa, \gamma' \rangle}$$

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

Fig. 5. Equality modulo patterns.

$$\begin{aligned}
\square \sqcup p &= p \sqcup \square = p & \diamond \sqcup p &= p \sqcup \diamond = p[\diamond/\square] \\
(p_1, p_2) \sqcup (p'_1, p'_2) &= (p_1 \sqcup p'_1, p_2 \sqcup p'_2) & c \sqcup c &= c \\
C(p) \sqcup C(p') &= C(p \sqcup p') & C &\in \{\text{inl}, \text{inr}, \text{roll}\} \\
\langle \kappa, \rho \rangle \sqcup \langle \kappa, \rho' \rangle &= \langle \kappa, \rho \sqcup \rho' \rangle \\
(\rho \sqcup \rho')(x) &= \begin{cases} \rho(x) \sqcup \rho'(x) & x \in \text{dom}(\rho) \cup \text{dom}(\rho') \\ \rho(x) & x \in \text{dom}(\rho) \setminus \text{dom}(\rho') \\ \rho'(x) & x \in \text{dom}(\rho') \setminus \text{dom}(\rho) \end{cases}
\end{aligned}$$

Fig. 6. Least upper bounds of patterns and environments.

We say that v matches v' modulo p (written $v \approx_p v'$) if v and v' match the structure of p , and are equal at corresponding positions denoted by \diamond . Moreover, we say $p \sqsubseteq v$ if $v \approx_p v$, and we write $p \sqcup p'$ for the least upper bound (join) of two patterns. Rules defining \approx_p and \sqcup are given in Figures 5 and 6.

When $p \sqsubseteq v$, we write $v|_p$ for the pattern obtained by replacing all of the \diamond -holes in p with the corresponding values in v . For example, $(1, 2)|_{(\diamond, \square)} = (1, \square)$. Similarly, we write $\gamma|_{\rho}$ for $[x_1 = \gamma(x_1)|_{\rho(x_1)}, \dots, x_n = \gamma(x_n)|_{\rho(x_n)}]$.

We also consider partial traces, usually written S , which are trace expressions where some subexpressions have been replaced with \square :

$$S ::= \dots \mid \square$$

As with patterns, we write $S \sqsubseteq T$ to indicate that T matches S , that is, S can be made equal to T by filling in some holes.

For the purpose of disclosure and obfuscation analysis, we will consider the “traces” to be triples (γ, T, v) where T is consistent with γ and v , that is, $\gamma, T \curvearrowright v$. We refer to such a triple as a consistent triple. Furthermore, to analyze forms of provenance based on annotation we will consider consistent annotated triples $(\hat{\gamma}, T, \hat{v})$ where $\hat{v} = F(T, \hat{\gamma})$. We consider trace or provenance queries built out of partial values and partial traces. We will later also consider queries derived from different forms of provenance, based on annotated triples.

Definition 5. 1. Let $\phi(\gamma)$ be a predicate on input environments. An input query $\text{IN}_{\gamma}.\phi(\gamma)$ is defined as $\{(\gamma, T, v) \mid \gamma, T \curvearrowright v \text{ and } \phi(\gamma)\}$. As a special case, we write IN_{ρ} for $\text{IN}_{\gamma}.\{\rho \sqsubseteq \gamma\}$.

2. Let $\phi(v)$ be a predicate on output values. An output query $\text{OUT}v.\phi(v)$ is defined as $\{(\gamma, T, v) \mid \gamma, T \curvearrowright v \text{ and } \phi(v)\}$. As a special case, we write $\text{OUT}_p = \text{OUT}v.(p \sqsubseteq v)$.

5 Disclosure and obfuscation analysis

5.1 Disclosure

We first consider properties disclosed by various forms of provenance considered above. Both where-provenance and expression provenance disclose useful information about the input. Dependency provenance does not disclose input information in an easy-to-analyze way, but is useful for obfuscation, as discussed later.

For where-provenance, we consider input queries $Q_{v_0, \ell} = \text{IN}\gamma. (\gamma.\ell = v_0)$ and output queries $Q'_{v_0, \ell} = \text{OUT}v. (v.\ell = v_0)$, where ℓ is a path and v_0 is a value. Such a query tests whether the value at a given path in γ or v matches the provided value.

Theorem 7. *The provenance view $(\gamma, T, v) \mapsto \text{W}(T, \text{path}(\gamma))$ positively discloses $Q_{v_0, \ell}$ via $Q'_{v_0, \ell}$.*

For expression-provenance, let $\gamma(t)$ be the result of evaluating t in γ with all paths ℓ replaced by their values $\gamma.\ell$ in γ . We consider queries $Q_{t, v_0} = \text{IN}\gamma. (\gamma(t) = v_0)$, where t is an expression provenance annotation and v_0 is a value. Such a query tests whether evaluating an expression e over γ yields the specified value. For example, $\text{IN}\gamma. x.1 + y.2 = 4$ holds for $\gamma = [x = (1, 2), y = (2, 3)]$, because $\gamma(x.1) + \gamma(y.2) = 1 + 3 = 4$. We also consider output queries $Q'_{t, v_0} = \text{OUT}v. (\hat{v}_1^t \text{ appears in } v \text{ and } |\hat{v}_1| = v_0)$, that simply test whether an annotated copy of v_0 appears in the output with annotation t .

Theorem 8. *The provenance view $(\gamma, T, v) \mapsto \text{E}(T, \text{path}(\gamma))$ positively discloses Q_{t, v_0} via Q'_{t, v_0} .*

Expression provenance and where-provenance are also related in the following sense:

Theorem 9. *Where-provenance is computable from expression-provenance.*

Proof. Where-provenance annotations can be extracted from expression-provenance annotations by mapping locations ℓ to themselves and all other expressions to \perp .

Note that this implies that for a function like “factorial”, the where-provenance of the output is always \perp . Hence, any query disclosed by where-provenance is disclosed by expression-provenance, and any query obfuscated by expression-provenance is also obfuscated by where-provenance.

We now consider a form of *trace slicing* that takes a partial output value and removes information from the input and trace that is not needed to disclose the output. We show that such *disclosure slices* also disclose generic provenance views (Theorem 11). Thus, disclosure slices form a quite general form of provenance in their own right.

$$\boxed{p, T \xrightarrow{\text{disc}} S, \rho}$$

$$\begin{array}{c}
\frac{}{\square, T \xrightarrow{\text{disc}} \square, []} \quad \frac{}{p, x \xrightarrow{\text{disc}} x, [x \mapsto p]} \quad \frac{}{p, c \xrightarrow{\text{disc}} c, []} \quad \frac{}{\langle \kappa, \rho \rangle, \text{fun } \kappa \xrightarrow{\text{disc}} \text{fun } \kappa, \rho} \\
\frac{}{\langle \kappa', \rho \rangle, \text{fun } \kappa \xrightarrow{\text{disc}} \text{fun } \kappa', \square} \quad \frac{p_2, T_2 \xrightarrow{\text{disc}} S_2, \rho_2[x \mapsto p_1] \quad p_1, T_1 \xrightarrow{\text{disc}} S_1, \rho_1}{p_2, \text{let } x = T_1 \text{ in } T_2 \xrightarrow{\text{disc}} \text{let } x = S_1 \text{ in } S_2, \rho_1 \sqcup \rho_2} \\
\frac{\diamond, T_1 \xrightarrow{\text{disc}} S_1, \rho_1 \quad \dots \quad \diamond, T_n \xrightarrow{\text{disc}} S_n, \rho_n}{p, \oplus(T_1, \dots, T_n) \xrightarrow{\text{disc}} \oplus(S_1, \dots, S_n), \rho_1 \sqcup \dots \sqcup \rho_n} \\
\frac{p_1, T_1 \xrightarrow{\text{disc}} S_1, \rho_1 \quad p_2, T_2 \xrightarrow{\text{disc}} S_2, \rho_2}{(p_1, p_2), (T_1, T_2) \xrightarrow{\text{disc}} (S_1, S_2), \rho_1 \sqcup \rho_2} \quad \frac{(p, \square), T \xrightarrow{\text{disc}} S, \rho}{p, \text{fst}(T) \xrightarrow{\text{disc}} \text{fst}(S), \rho} \\
\frac{p, T \xrightarrow{\text{disc}} S, \rho}{\text{inl}(p), \text{inl}(T) \xrightarrow{\text{disc}} \text{inl}(S), \rho} \quad \frac{p, T \xrightarrow{\text{disc}} S, \rho}{\text{inl}(p), \text{inr}(T) \xrightarrow{\text{disc}} \text{inr}(\square), []} \\
\frac{p, T \xrightarrow{\text{disc}} S, \rho}{\text{roll}(p), \text{roll}(T) \xrightarrow{\text{disc}} \text{roll}(S), \rho} \quad \frac{\text{roll}(p), T \xrightarrow{\text{disc}} S, \rho}{p, \text{unroll}(T) \xrightarrow{\text{disc}} \text{unroll}(S), \rho} \\
\frac{p_1, T_1 \xrightarrow{\text{disc}} S_1, \rho_1[x_1 \mapsto p] \quad \text{inl}(p), T \xrightarrow{\text{disc}} S, \rho}{p_1, \text{case } T \triangleright_{\text{inl}} x_1.T_1 \xrightarrow{\text{disc}} \text{case } S \triangleright_{\text{inl}} x_1.S_1, \rho \sqcup \rho_1} \\
\frac{p, T \xrightarrow{\text{disc}} S, \rho[f \mapsto p_1, x \mapsto p_2] \quad p_1 \sqcup \langle \kappa, \rho \rangle, T_1 \xrightarrow{\text{disc}} S_1, \rho_1 \quad p_2, T_2 \xrightarrow{\text{disc}} S_2, \rho_2}{p, (T_1 T_2) \triangleright_{\kappa} f(x).T \xrightarrow{\text{disc}} (S_1 S_2) \triangleright_{\kappa} f(x).S, \rho_1 \sqcup \rho_2} \\
\frac{\text{fvs}(\kappa) = \{x_1, \dots, x_n\}}{\diamond, \text{fun } \kappa \xrightarrow{\text{disc}} \text{fun } \kappa, [x_1 \mapsto \diamond, \dots, x_n \mapsto \diamond]} \quad \frac{\diamond, T_1 \xrightarrow{\text{disc}} S_1, \rho_1 \quad \diamond, T_2 \xrightarrow{\text{disc}} S_2, \rho_2}{\diamond, (T_1, T_2) \xrightarrow{\text{disc}} (S_1, S_2), \rho_1 \sqcup \rho_2} \\
\frac{\diamond, T \xrightarrow{\text{disc}} S, \rho}{\diamond, \text{inl}(T) \xrightarrow{\text{disc}} \text{inl}(S), \rho} \quad \frac{\diamond, T \xrightarrow{\text{disc}} S, \rho}{\diamond, \text{inr}(T) \xrightarrow{\text{disc}} \text{inr}(S), \rho} \quad \frac{\diamond, T \xrightarrow{\text{disc}} S, \rho}{\diamond, \text{roll}(T) \xrightarrow{\text{disc}} \text{roll}(S), \rho}
\end{array}$$

Fig. 7. Disclosure slicing (selected rules).

Definition 6. Let $\gamma, T \Downarrow v$, and suppose $S \sqsubseteq T$ and $\rho \sqsubseteq \gamma$. We say (ρ, S) is a disclosure slice with respect to partial value p if for all $\gamma' \sqsupseteq \rho$ and $T' \sqsupseteq S$ such that if $\gamma', T' \rightsquigarrow v'$, we have $p \sqsubseteq v$ iff $p \sqsubseteq v'$.

Note that by this definition, minimal disclosure slices exist (since there are finitely many slices) but need not be unique. For example, both $\square \vee \text{true}$ and $\text{true} \vee \square$ are disclosure slices showing that $\text{true} \vee \text{true}$ evaluates to true , but $\square \vee \square$ is not a disclosure slice.

Figure 7 shows selected rules defining a disclosure slicing judgment $p, T \xrightarrow{\text{disc}} \rho, S$. Basically, the idea is to push a partial value backwards through a trace to obtain a partial input environment and trace slice. The partial input environment is needed to handle local variables in traces; for example, in the rule for `let`, we first slice through the body of the `let`, then identify the partial value showing the needed parts of the `let`-bound value, and use that to slice backwards through the first subtrace. Slicing for application

traces is similar, but more complex due to the closure environments. Note also that the special \diamond patterns are used to slice backwards through primitive operations even when we do not know the values of the inputs or results. (Another possibility is to annotate the traces of primitive operations with these values.)

Lemma 1. *If $\gamma, T \curvearrowright v$ then for any $p \sqsubseteq v$ there exists $S \sqsubseteq T$ and $\rho \sqsubseteq \gamma$ such that $p, T \xrightarrow{\text{disc}} S, \rho$. Moreover, there is a unique least such S and ρ .*

We choose a function $\text{Disc}_p(\gamma, T, v)$ on consistent triples (γ, T, v) whose value is $(\gamma|_\rho, S)$ where $p, T \xrightarrow{\text{disc}} S, \rho$ and S, ρ are the least slices (obtained by determinizing the slicing algorithm by applying the first, hole-propagating rule greedily before any other rule). The idea is that we slice using the rules in Figure 7 and then transform ρ by filling in all \diamond -holes with the corresponding values in γ . Note that the v parameter is irrelevant and is included only so that Disc_p is a uniform function from consistent triples to slices.

To prove the correctness of this disclosure slicing algorithm, we need a stronger notion of equivalence. Recall the definition of $v \approx_p v'$ as shown in Figure 5. Using this relation, we can prove the correctness of the slicing relation as follows:

Lemma 2. *Assume $\gamma, T \curvearrowright v$ and $p, T \xrightarrow{\text{disc}} S, \rho$.*

1. *If $p \sqsubseteq v$ then for all $\gamma' \approx_\rho \gamma$ and $T' \sqsupseteq S$, if $\gamma', T' \Downarrow v'$ then $v' \approx_p v$.*
2. *If $p \not\sqsubseteq v$ then for all $\gamma' \approx_\rho \gamma$ and $T' \sqsupseteq S$, if $\gamma', T' \Downarrow v'$ then $p \not\sqsubseteq v'$.*

Proof. Both parts follow by induction on the structure of slicing derivations.

Correctness follows as a consequence of the above two properties.

Theorem 10. *Disc_p discloses OUT_p .*

Proof. Suppose $\text{Disc}_p(\gamma, T, v) = \text{Disc}_p(\gamma', T', v')$. Then $(\gamma|_\rho, S) = (\gamma'|_{\rho'}, S')$ where $p, T \xrightarrow{\text{disc}} S, \rho$ and $p, T' \xrightarrow{\text{disc}} S', \rho'$. Hence, $S = S'$ and $\gamma|_\rho = \gamma'|_{\rho'}$, which in turn implies $\gamma \approx_\rho \gamma'$. Suppose that $\text{OUT}_p(\gamma, T, v)$ holds; that is, $p \sqsubseteq v$. Then by Lemma 2(1), $v' \approx_p v$ so $p \sqsubseteq v'$. Conversely, suppose $p \not\sqsubseteq v$. Then by Lemma 2(2), we have $p \not\sqsubseteq v'$.

Disclosure from slices. Finally, we link disclosure for value patterns to disclosure for generic provenance views. Essentially, we show that for any F , the disclosure slice for p positively discloses the F -provenance annotations of values matching p . Informally, this means that disclosure slices provide a highly general form of provenance specialized to a part of the output: one can compute and reveal the disclosure slice and others can then compute any generic provenance view from the slice, without rerunning the original computation or consulting input data or subtraces that are dropped in the slice.

Theorem 11. *Assume $|\hat{\gamma}|, T \curvearrowright v$ and $p \sqsubseteq v$. Suppose $p, T \xrightarrow{\text{disc}} S, \rho$. Suppose that F is a generic extraction function. Then the annotations associated with p in $F(T, \hat{\gamma})$ can be correctly extracted from S using only input parts needed by ρ . That is, suppose we have $\hat{\gamma} \approx_\rho \hat{\gamma}'$ (lifting \approx_ρ to annotated values in the obvious way) and $T' \sqsupseteq S$, where $|\hat{\gamma}'|, T' \curvearrowright v'$. Then we have $F(T, \hat{\gamma}) \approx_p F(T', \hat{\gamma}')$.*

5.2 Obfuscation

We now consider obfuscating properties of traces. We first consider what can be obfuscated by the standard provenance views. Where-provenance, essentially, obfuscates anything that can never be copied to the output or affect the control flow of something that is copied to the output. Similarly, expression provenance obfuscates any part of the input that never participates in expression annotations. In both cases, we can potentially learn about parts of the input that affected control flow, however. For example, `if $x = 1$ then 1 else y` does not obfuscate the value of x in either model, provided y comes from the input, since we can inspect the annotation of the result to determine that $x = 1$ or $x \neq 1$.

Given that we want to ensure obfuscation, we consider conservative techniques that accept (or construct) only provenance views that successfully obfuscate, but may reject some views or construct views that are unnecessarily opaque.

There are several ways to erase information from traces (or other provenance views) to ensure obfuscation of input properties. One way is to re-use the static analysis of dependency provenance (in [9], for example) to identify parts of the output that suffice to make it impossible to guess sensitive parts of the input. Alternatively we can use dynamic dependency provenance to increase precision, by propagating dependency tracking information from the input to the output.

This is similar to using static analysis or dynamic labels for information flow security; the difference is one of emphasis. In information flow security, we usually identify high- or low-security *locations* and try to certify that high-security data does not affect the computation of low-security data; here, instead, we identify a high-security *property* of the trace (e.g. that the input satisfies a certain formula) and try to determine what parts of the output do not depend on sensitive inputs, and hence can be safely included in the provenance view. However, these techniques do not provide guidance about what parts of the *trace* can be safely included in the provenance view.

Here, we develop an alternative approach based on traces. Consider a pattern $\rho \sqsubseteq \gamma$ that erases all information that is considered sensitive. We construct an *obfuscation slice* by re-evaluating T on ρ as much as possible, to compute a sliced trace S and partial value p . We erase parts of T and of the original output value that depend on the erased parts of ρ . Thus, any part of the trace or output value that remains in the obfuscation slice is irrelevant to the sensitive part of the input, and cannot be used to guess it.

Figure 8 shows a syntactic algorithm for computing obfuscation slices as described above. Many rules are essentially generalizations of the rules for evaluation to allow for partial inputs, outputs and traces. The rules of interest, near the bottom of the figure, show how to handle attempts to compute that encounter holes in places where a value constructor is expected. When this happens, we essentially propagate the hole result and return a hole trace. This may be unnecessarily aggressive for some cases, but is necessary for the case and application traces where the trace form gives clues about the control flow.

We define $\text{Obf}_\rho(\gamma, T, v)$ as (p, S) where $\rho, T \xrightarrow{\text{obf}} p, S$. We can show that this is total for well-formed, \diamond -free traces and input environments.

Lemma 3. *If $\gamma, T \rightsquigarrow v$ and $\rho \sqsubseteq \gamma$ and $\rho, T \xrightarrow{\text{obf}} p, S$ then $p \sqsubseteq v$ and $S \sqsubseteq T$.*

$$\boxed{\rho, T \xrightarrow{\text{obf}} p, S}$$

$$\begin{array}{c}
\frac{}{\rho, x \xrightarrow{\text{obf}} \rho(x), x} \quad \frac{}{\rho, c \xrightarrow{\text{obf}} \rho(x), c} \quad \frac{}{\rho, \text{fun } \kappa \xrightarrow{\text{obf}} \langle \kappa, \rho \rangle, \text{fun } \kappa} \\
\frac{\rho, T_1 \xrightarrow{\text{obf}} p_1, S_1 \quad \rho[x \mapsto p_1], T_2 \xrightarrow{\text{obf}} p_2, S_2}{\rho, \text{let } x = T_1 \text{ in } T_2 \xrightarrow{\text{obf}} p_2, \text{let } x = S_1 \text{ in } S_2} \\
\frac{\rho, T_1 \xrightarrow{\text{obf}} v_1, S_1 \quad \dots \quad \rho, T_n \xrightarrow{\text{obf}} v_n, S_n}{\rho, \oplus(T_1, \dots, T_n) \xrightarrow{\text{obf}} \oplus(v_1, \dots, v_n), \oplus(S_1, \dots, S_n)} \\
\frac{\rho, T_1 \xrightarrow{\text{obf}} p_1, S_1 \quad \rho, T_2 \xrightarrow{\text{obf}} p_2, S_2}{\rho, (T_1, T_2) \xrightarrow{\text{obf}} (p_1, p_2), (S_1, S_2)} \quad \frac{\rho, T \xrightarrow{\text{obf}} (p_1, p_2), S}{\rho, \text{fst}(T) \xrightarrow{\text{obf}} p_1, \text{fst}(S)} \\
\frac{\rho, T \xrightarrow{\text{obf}} p, S}{\rho, \text{inl}(T) \xrightarrow{\text{obf}} \text{inl}(p), \text{inl}(S)} \quad \frac{\rho, T \xrightarrow{\text{obf}} \text{inl}(p), S \quad \rho[x_1 \mapsto p], T_1 \xrightarrow{\text{obf}} p_1, S_1}{\rho, \text{case } T \triangleright_{\text{inl}} x_1.T_1 \xrightarrow{\text{obf}} p_1, \text{case } S \triangleright_{\text{inl}} x_1.S_1} \\
\frac{\rho, T \xrightarrow{\text{obf}} p, S}{\rho, \text{roll}(T) \xrightarrow{\text{obf}} \text{roll}(p), \text{roll}(S)} \quad \frac{\rho, T \xrightarrow{\text{obf}} \text{roll}(p), S}{\rho, \text{unroll}(T) \xrightarrow{\text{obf}} p, \text{unroll}(S)} \\
\frac{\rho, T_1 \xrightarrow{\text{obf}} \langle \kappa, \rho_0 \rangle, S_1 \quad \rho, T_2 \xrightarrow{\text{obf}} p_2, S_2 \quad \rho[f \mapsto \langle \kappa, \rho_0 \rangle, x \mapsto p_2], T \xrightarrow{\text{obf}} p, S}{\rho, (T_1 T_2) \triangleright_{\kappa} f(x).T \xrightarrow{\text{obf}} p, (S_1 S_2) \triangleright_{\kappa} f(x).S} \\
\frac{\rho, T_i \xrightarrow{\text{obf}} \square, S_i}{\rho, \oplus(T_1, \dots, T_n) \xrightarrow{\text{obf}} \square, \square} \quad \frac{\rho, T \xrightarrow{\text{obf}} \square, S}{\rho, \text{fst}(T) \xrightarrow{\text{obf}} \square, \square} \quad \frac{\rho, T \xrightarrow{\text{obf}} \square, S}{\rho, \text{unroll}(T) \xrightarrow{\text{obf}} \square, \square} \\
\frac{\rho, T \xrightarrow{\text{obf}} \square, S}{\rho, \text{case } T \triangleright_{\text{inl}} x_1.T_1 \xrightarrow{\text{obf}} \square, \square} \quad \frac{\rho, T_1 \xrightarrow{\text{obf}} \square, S_1}{\rho, (T_1 T_2) \triangleright_{\kappa} f(x).T \xrightarrow{\text{obf}} \square, \square}
\end{array}$$

Fig. 8. Obfuscation slicing (selected rules).

Lemma 4. *If $\gamma, e \Downarrow v, T$ and $\rho \sqsubseteq \gamma$ and $\rho, T \xrightarrow{\text{obf}} p, S$ then for all $\gamma' \sqsupseteq \rho$, if $\gamma', e \Downarrow v', T'$ then $\rho, T' \xrightarrow{\text{obf}} p, S$.*

Theorem 12. *For traces generated by terminating expressions, and ρ with holes of nonsingular types, and $\rho' \sqsupset \rho$, we have Obf_ρ positively obfuscates $\text{IN}_{\rho'}$.*

Proof. Suppose $\text{IN}_{\rho'}$ holds of (γ, T, v) where $\rho' \sqsupset \rho$. Then $\rho \sqsubset \rho' \sqsubseteq \gamma$. Moreover, since the inclusion is strict, ρ must contain holes that can be replaced with different values, so there exists another $\gamma' \sqsupseteq \rho$ that differs from ρ' . Since T was generated by a terminating expression, we know that $\gamma', e \Downarrow v', T'$ can be derived for some v', T' . By Lemma 4 (and the easily-verified determinacy of $\xrightarrow{\text{obf}}$) we know that $\rho, T' \xrightarrow{\text{obf}} p, S$, hence $\text{Obf}_\rho(\gamma', T', v') = (p, S) = \text{Obf}_\rho(\gamma, T, v)$, as required.

5.3 Discussion

The analysis in section 5.1 gives novel characterizations of what information is disclosed by where-provenance and expression provenance. Essentially, where-provenance

discloses information about what parts of the input are copied to the output, while expression provenance additionally discloses information about how parts of the input can be combined to compute parts of the output. The analysis in section 5.1 also shows (in a formal sense) that where-provenance and expression provenance are closely related: one can obtain where-provenance from expressions simply by erasure. Moreover, we can obtain a number of intermediate provenance models based on transductions over expression-provenance annotations.

The disclosure slicing algorithm is based on an interesting insight (which we are exploring in concurrent work on slicing): at a technical level, the information we need for program comprehension via slicing (to understand how a program has evaluated its inputs to produce outputs) is quite similar to what we need for provenance. Although we have identified connections between provenance and slicing before [9], our disclosure and obfuscation slicing algorithms provide further evidence of this close connection.

6 Related Work

There is a huge, and growing, literature on provenance [3, 10, 24, 22], but there is little work on formal models of provenance and none on provenance in a general-purpose higher-order language. Due to space limits, we confine our comparison to closely related work on formal techniques for provenance, and on related ideas in programming languages and language-based security. We refer the interested reader to the aforementioned surveys for more information on provenance in workflows and databases, and to [11, 8] for further discussion of prior work on provenance security.

Provenance. This work differs from previous work on provenance in databases in several important ways. First, we consider a general purpose, higher-order language, whereas previous work considers database query languages of limited expressiveness (e.g., monotone query languages), which include unordered collection types with monadic iteration operations but not sum types, recursive types or first-class functions. Second, we aim to record traces adequate to answer a wide range of provenance queries in this general setting, whereas previous work has focused on particular kinds of queries (e.g., where-provenance [6, 5], why-provenance [6], how-provenance [18, 17]).

Provenance has also been studied extensively for scientific workflow systems [3, 24, 14]. Most work in this area describes the provenance tracking behavior of a system through examples and does not give a formal semantics that could be used to prove correctness properties. An exception is Hidders et al. [20], which is the closest workflow provenance work to ours. They model workflows using a core database query language extended with nondeterministic, external function calls, and partially formalize a semantics of *runs*, or sets of triples (γ, e, v) labeling an operational derivation tree.

Other related topics. Our trace model is partly inspired by previous work on self-adjusting computation [1], where execution traces are used to efficiently recompute functional programs under arbitrary modifications to their inputs. Provenance-like ideas have also appeared in the context of bidirectional computation [2]. Dimoulas et al. study correctness properties for *blame* in contracts based on semantic properties that, they

suggest, may be related to provenance [16]. However, to our knowledge no formal relationships between provenance and self-adjusting computation, bidirectional computation, or blame have been developed.

7 Conclusions

While the importance of understanding provenance and its security characteristics has been widely documented, to date there has been little work on formal modeling of either provenance or provenance security. In this paper, we elaborate upon the ideas introduced in previous work [8], by instantiating the formal framework proposed there with a general-purpose functional programming language and a natural notion of execution traces. We showed how more conventional forms of provenance can be extracted from such traces via a generic provenance extraction mechanism. Furthermore, we studied the key notions of disclosure and obfuscation in this context. In the process we identified weaker *positive* and *negative* variants of disclosure and obfuscation, based on the observation that the original definitions seem too strong to be satisfied often in practice. Our main results include algorithms for *disclosure slicing*, which traverses a trace backwards to retain information needed to certify how an output was produced, and *obfuscation slicing*, which reruns a trace on partial input (excluding sensitive parts of the input), yielding a partial trace and partial output that excludes all information that could help a principal learn sensitive data.

To summarize, our main contribution is the development of a general model of provenance in the form of a core calculus that instruments runs of programs with detailed execution traces. We validated the design of this calculus by showing that traces generalize other known forms of provenance and by studying their disclosure and obfuscation properties. There are many possible avenues for future work, including:

- identifying richer languages for defining trace queries or provenance views
- developing and implementing effective algorithms for trace slicing, and relating these to program slicing
- extending trace and provenance models to handle references, exceptions, input/output, concurrency, nondeterminism, communication, etc.

References

1. U. A. Acar, G. E. Blelloch, and R. Harper. Adaptive functional programming. *ACM Trans. Program. Lang. Syst.*, 28(6):990–1034, 2006.
2. A. Bohannon, J. N. Foster, B. C. Pierce, A. Pilkiewicz, and A. Schmitt. Boomerang: resourceful lenses for string data. In *POPL*, pages 407–419, New York, NY, USA, 2008. ACM.
3. R. Bose and J. Frew. Lineage retrieval for scientific data processing: a survey. *ACM Comput. Surv.*, 37(1):1–28, 2005.
4. P. Buneman, J. Cheney, W.-C. Tan, and S. Vansummeren. Curated databases. In *PODS*, pages 1–12, 2008.
5. 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.

6. P. Buneman, S. Khanna, and W. Tan. Why and where: A characterization of data provenance. In *ICDT*, number 1973 in LNCS, pages 316–330, 2001.
7. S. Carey and G. Rogow. UAL shares fall as old story surfaces online. *Wall Street Journal*, September 2008. <http://online.wsj.com/article/SB122088673-738010213.html>.
8. J. Cheney. A formal framework for provenance security. In *CSF*, pages 281–293. IEEE, 2011.
9. J. Cheney, A. Ahmed, and U. A. Acar. Provenance as dependency analysis. *Mathematical Structures in Computer Science*, 21(6):1301–1337, 2011.
10. 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.
11. J. Cheney, S. Chong, N. Foster, M. Seltzer, and S. Vansummeren. Provenance: A future history. In *OOPSLA Companion (Onward! 2009)*, pages 957–964, 2009.
12. S. Chong. Towards semantics for provenance security. In *Workshop on the Theory and Practice of Provenance*, 2009. Informal online proceedings: <http://www.usenix.org/events/tapp09/>.
13. A. Cirillo, R. Jagadeesan, C. Pitcher, and J. Riely. Tapido: Trust and authorization via provenance and integrity in distributed objects. In *ESOP*, volume 4960 of LNCS, pages 208–223, 2008.
14. S. B. Davidson and J. Freire. Provenance and scientific workflows: challenges and opportunities. In *SIGMOD*, pages 1345–1350, New York, NY, USA, 2008.
15. S. B. Davidson, S. Khanna, T. Milo, D. Panigrahi, and S. Roy. Provenance views for module privacy. In *PODS*, pages 175–186, 2011.
16. C. Dimoulas, R. B. Findler, C. Flanagan, and M. Felleisen. Correct blame for contracts: no more scapegoating. In *POPL*, pages 215–226, New York, NY, USA, 2011. ACM.
17. J. N. Foster, T. J. Green, and V. Tannen. Annotated XML: queries and provenance. In *PODS*, pages 271–280, 2008.
18. T. J. Green, G. Karvounarakis, and V. Tannen. Provenance semirings. In *PODS*, pages 31–40, 2007.
19. N. Guts, C. Fournet, and F. Z. Nardelli. Reliable evidence: auditability by typing. In *ESORICS*, pages 168–183, 2009.
20. J. Hidders, N. Kwasnikowska, J. Sroka, J. Tyszkiewicz, and J. Van den Bussche. A formal model of dataflow repositories. In *DILS*, volume 4544 of LNCS, pages 105–121, 2007.
21. L. Jia, J. A. Vaughan, K. Mazurak, J. Zhao, L. Zarko, J. Schorr, and S. Zdancewic. Aura: a programming language for authorization and audit. In *ICFP*, pages 27–38, New York, NY, USA, 2008.
22. L. Moreau. The foundations for provenance on the web. *Foundations and Trends in Web Science*, 2(2–3), 2010.
23. L. Moreau et al. The open provenance model core specification (v1.1). *Future Generation Computer Systems*, 27(6):743–756, June 2010.
24. Y. Simmhan, B. Plale, and D. Gannon. A survey of data provenance in e-science. *SIGMOD Record*, 34(3):31–36, 2005.
25. N. Swamy, J. Chen, C. Fournet, P.-Y. Strub, K. Bhargavan, and J. Yang. Secure distributed programming with value-dependent types. In *ICFP*, pages 266–278, 2011.
26. N. Swamy, B. J. Corcoran, and M. Hicks. Fable: A language for enforcing user-defined security policies. In *IEEE Symposium on Security and Privacy*, pages 369–383, 2008.
27. S. Varghese. UK government gets bitten by Microsoft Word. *Sydney Morning Herald*, July 2003. <http://www.smh.com.au/articles/2003/07/02/-1056825430340.html>.