

# A Core Calculus for Provenance Inspection

Wilmer Ricciotti

Laboratory for Foundations of Computer Science

University of Edinburgh

[research@wilmer-ricciotti.net](mailto:research@wilmer-ricciotti.net)

August 17, 2017

## Abstract

Recent research has been devoting increasing attention to provenance, or information describing the origin, derivation, and history of data, due to its relevance to critical issues including transparency, privacy, and security. Engineering a software system to make it provenance-aware by means of ad-hoc instrumentation requires a substantial effort: the development of general-purpose infrastructure is thus very important to achieve the goal of making provenance widely available. In this article we describe a core functional language equipped with a provenance-aware semantics that is sufficiently generic to accommodate many notions of provenance proposed in the literature. While existing proposals typically treat provenance views and provenance extraction as second-class, extralinguistic mechanisms, in our work provenance views are expressed as standard programs and provenance data can be reflected into the language, allowing for programs that inspect their own provenance.

## 1 Introduction

Not all information can be trusted: we know from experience that some sources are unreliable; but even trustworthy sources can make mistakes, which is why, to decide whether some information can be believed, we need to consider how it was processed. This shows the importance of an accurate description of the origin, history, and derivation of some information, which has been known for decades as *provenance* [20, 2], but has clearly been studied by philosophers of knowledge for a much longer time. Provenance is essential to the development of scientific knowledge: raw experimental data is useless without a thorough explanation of how it was obtained, allowing the public to reproduce the experiment (at least in principle) and confirm it or contradict it. Data collection and processing was originally a human activity (and scientists had to keep an accurate record of it); by means of computer systems, it is now possible to process much greater quantities of data in an efficient way.

Computer systems do not, however, provide the provenance of the data they process for free: propagating it from the input to the output, while logging all intermediate operations, requires additional code; and more code means not only more effort, but also more occasions to introduce bugs. For these reasons, a line of research has devoted its attention to automated techniques to record provenance. Many of these techniques are designed to be employed in interesting but restricted settings (especially scientific workflow computation [18, 4, 11, 21] and databases [10, 13, 14, 3, 6]).

In a 2013 paper [1], Acar et al. have laid the basis for tracking provenance in a general-purpose programming language, by proposing a core functional calculus (*Transparent ML*, or *TML*) equipped with a provenance extraction framework. They provide a *tracing* big-step semantics, which records the history (or *trace*) of computation; by analysing this trace, the extraction framework is able to “replay” the computation, while at the same time propagating provenance annotations from the source expression to the final value. They show that *TML* and its provenance framework are expressive enough to extract many forms of provenance.

Despite this, in *TML* provenance remains a second class citizen: the extraction framework is an external layer combining several user-provided functions (*provenance views*) to process annotations and an execution trace and return the provenance; neither the views, nor the annotations, nor the execution trace are expressed as *TML* terms. Secondly, traces appear as an unneeded intermediate step, as the extraction framework is required to replay them, essentially executing the program a second time. Furthermore, *TML* traces closely reflect the big-step semantics of the language, limiting implementations to a call-by-value strategy and essentially requiring provenance inspection to take place after the execution of the program, and making *self-inspecting programs* impossible.

## 1.1 Summary

We take inspiration from TML and develop a more tightly integrated *provenance inspection calculus* (PIC), in which provenance is expressed as, and manipulated by first-class terms of the language. In this paper, we provide the following contributions:

- a core functional calculus with special constructs to manipulate syntactic locations (represented as lists of integers)
- language types for provenance labellings and provenance *transducers* together with a rich set of combinators
- a small-step, traceless semantics associating to each computation step the corresponding provenance transducer
- a simplified framework to define provenance views as language expressions
- an *inspection* construct allowing reflective reasoning on provenance

**Outline.** Section 2 introduces some basic notions about provenance. The syntax and semantics of PIC is introduced in Section 3, along with the provenance combinators. In Section 4 we provide a framework for the simplified definition of provenance views. Section 5 integrates PIC with an inspection construct and discusses the full language with an extensive example. Finally in Section 6 we draw conclusions and discuss future work.

## 2 Background

Provenance is most easily defined informally, as a form of metadata describing the origin or history of an artifact, or the process by which it has been constructed. As soon as we try to make this concept more formal, we realize that it depends on choices which are, to a certain degree, arbitrary: the origin of data can include an identifier of the person or organization that provided it; a description of the technical equipment that produced it, along with its accuracy; a timestamp; localization information from a GPS sensor; etc. Likewise, the computational process manipulating data can be described with various levels of precision. Since the choice of provenance information is largely a human task, we do not expect a single, universally acceptable notion of provenance. It is however clear that any notion of provenance, to be mechanized, requires facilities to annotate input and output data with metadata, along with a way to analyze the execution history of a software system.

The core calculus TML [1] provides a big-step tracing semantics, with evaluation judgments in the following form:<sup>1</sup>

$$M \Downarrow V, t$$

The judgment evaluates  $M$  and produces a result  $V$  and a trace  $t$ , which is a data structure describing the computation that produced  $V$  (with a slight approximation,  $t$  can be considered a linearized version of the evaluation tree). In this setting, the evaluation of an arithmetic expression  $N_1 + N_2$  takes the following form:

$$\frac{N_1 \Downarrow 2, t_1 \quad N_2 \Downarrow 3, t_2 \quad \hat{\dagger}(2, 3) = 5}{N_1 + N_2 \Downarrow 5, t_1 + t_2}$$

where  $t_1 + t_2$  is the trace expressing the evaluation of a sum whose two arguments have histories described, respectively, by  $t_1$  and  $t_2$ , and  $\hat{\dagger}$  is the meta-language sum. The evaluation of a larger term yields a more complex trace:

$$\frac{\frac{N_1 \Downarrow 2, t_1 \quad N_2 \Downarrow 3, t_2 \quad \hat{\dagger}(2, 3) = 5}{N_1 + N_2 \Downarrow 5, t_1 + t_2} \quad N_3 \Downarrow 1, t_3}{\langle N_1 + N_2, N_3 \rangle \Downarrow \langle 5, 1 \rangle, \langle t_1 + t_2, t_3 \rangle} \\ \mathbf{fst}(\langle N_1 + N_2, N_3 \rangle) \Downarrow 5, \mathbf{fst}(\langle t_1 + t_2, t_3 \rangle)$$

TML provides a framework to extract various forms of provenance from traces: the user annotates the source term with primitive provenance information and passes it, along with the execution trace, to a *provenance view* which, basically, replays the computation while at the same time propagating the annotations. Views are TML's way of allowing a user to define several notions of provenance: if we annotate the original term as  $\mathbf{fst}(\langle N_1^{a_1} + N_2^{a_2}, N_3^{a_3} \rangle^{a_4})$ , and we want the provenance of the final value to reflect that of all the portions of the input on which it depends (*dependency provenance* [9]), it is possible to define, by structural recursion on the execution trace, a view that will return  $5^{a_1 \cup a_2 \cup a_4}$ . Notice that the framework is an external facility: annotations do not have to be TML terms, and views are defined as set-theoretic functions (although it is implied that they should be computable). Other forms of provenance that can be expressed in TML include:

<sup>1</sup>For simplicity, we omit evaluation environments.

Prim. types	$B$	$::= \mathbf{1} \mid \mathbf{B} \mid \mathbf{N} \mid \dots$
Types	$\sigma, \tau$	$::= B \mid \sigma \times \tau \mid \sigma \supset \tau \mid [\sigma] \mid \sim\sigma$
Contexts	$\Gamma$	$::= [x_1 : \tau_1, \dots, x_n : \tau_n]$
Expressions	$L, M, N$	$::= c \mid x \mid M N \mid f(x).M \mid \oplus(\vec{M})$ $\mid \bullet \mid \sim M \mid \epsilon \mid M \# N \mid \text{case } M \text{ of } \pi$ $\mid \text{fst}(M) \mid \text{snd}(M) \mid \langle M, N \rangle \mid \rho(\vec{x}.M)$
Patterns	$\pi$	$::= \{\mathbf{tt} \mapsto M; \mathbf{ff} \mapsto N\}$ $\mid \{\bullet \mapsto M; \sim x \mapsto N\} \mid \{\epsilon \mapsto M; x \# y \mapsto N\}$
Views	$F$	$::= \langle M_\beta, M_{\text{fst}}, M_{\text{snd}}, M_\bullet, M_\sim,$ $M_\epsilon, M_\#, M_{\text{tt}}, M_{\text{ff}}, M_\oplus \rangle$
Values	$U, V$	$::= c \mid x \mid \bullet \mid \sim U \mid \epsilon \mid U \# V \mid \langle U, V \rangle$ $\mid f(x).M \mid \rho(\vec{x}.M)$

Figure 1: Syntax of the language

- where-provenance [7], whose labels identify the input location from which each portion of an expression was copied;
- expression provenance [1], whose labels are expression graphs or trees describing how a value was computed in terms of primitive operations.

Motivated by the limitations of TML, we set out to embed the extraction framework into the language itself: as we started investigating the issue of allowing provenance extraction within programs, we realized that evaluating a sub-program to obtain its trace and then processing its trace with a provenance view, essentially evaluating it a second time, made little sense; it should be possible to compute the provenance applying a view immediately, without materializing the execution trace as an intermediate step. To accomplish this task, we decided to replace traces with *transducers*, or functions that transform provenance according to a certain execution history. A major obstacle in defining transducers is the need to undo the substitutions created by the evaluation of functions or pattern matching: to address it, we equip the language with a special *relocation* operator. One property of transducers is that they can be combined by functional composition to obtain the transducer corresponding to express a transitivity rule. This allows us to adopt a small-step semantics, which can be used to reason on the provenance of non-terminating programs.<sup>2</sup>

### 3 A provenance-aware calculus

We define the *provenance inspection calculus* PIC, a pure functional language equipped with a reduction semantics and extended with facilities to annotate expressions with provenance and propagate it during the computation. Its syntax, shown in Fig. 1, includes a standard type system, with primitive types  $B$  (e.g. the singleton type  $\mathbf{1}$ , booleans  $\mathbf{B}$ , and natural numbers  $\mathbf{N}$ ), pairs, functions, and lists and optional values on a base type  $\sigma$  (respectively  $[\sigma]$  and  $\sim\sigma$ ).

The definition of expressions is standard. We parametrize the syntax and semantics over primitive constants  $c$  and primitive operations  $\oplus$ , whose arguments must be of primitive types. We assume constants and primitive operations including booleans  $\mathbf{tt}$  and  $\mathbf{ff}$ , and natural numbers  $0, 1, 2, \dots$ , along with standard arithmetic operations and the equality test on natural numbers  $\stackrel{?}{=}$ .

We fix a multi-sorted signature  $\Sigma$ , containing type declarations for constants (e.g.:  $c : \sigma$ ) and primitive operations ( $\oplus : \vec{\tau}_m \supset \tau'$ , where  $\vec{\tau}_m$  is a sequence of  $m$  types  $\tau_1, \dots, \tau_m$ ). Each operation symbol  $\oplus$  be associated to a metalanguage function  $\hat{\oplus}$  from values to values, expressing its semantics, thus defining an algebra for  $\Sigma$ .

Expressions  $f(x).M$  represent recursive functions  $f$  with formal argument  $x$  and body  $M$ , where both  $f$  and  $x$  are bound in  $M$ ; non-recursive lambda-expressions  $\lambda x.M$  are treated as recursive functions  $\_ (x).M$ , under the condition that the irrelevant function name  $\_$  is free in  $M$ .

For optional values and lists, we use the following notation: an expression of type  $\sim\sigma$  can be obtained either by encapsulating some  $M$  of type  $\sigma$  using the syntax  $\sim M$ , or by means of the “none” constructor represented by  $\bullet$ ; list constructors include the empty list  $\epsilon$  and the *cons* operation  $M \# N$ . For explicit lists, we will use the notation  $[M_1; \dots; M_k]$ , which is syntactic sugar for  $M_1 \# \dots \# M_k \# \epsilon$ .

Two case analysis constructs are used to decompose lists and optional values; similarly, case analysis also supports reasoning on booleans. We will allow ourselves to use an extended case analysis notation, which should

<sup>2</sup>A small-step semantics of TML was originally considered for this precise reason, but left for future work. In the words of its proponents, “moving to a small-step semantics [seems] likely to complicate the trace semantics (and subsequent analysis) considerably”.



where we have abused the cons notation to apply it to sets of terms  $\mathcal{S}$  as follows:

$$n \# \mathcal{S} \triangleq \{n \# M : M \in \mathcal{S}\}$$

A lookup meta-operation  $M|_\ell$  computes the subterm of  $M$  to which the location  $\ell$  points; if, while scanning  $M$ , the operation reaches a variable before consuming the whole location, it returns the variable together with the remaining, unconsumed part of the location. The location lookup operation is defined in Fig. 2.

In all other cases, the location is inconsistent with the term being scanned and lookup is undefined.

Given a location  $\ell_N$  valid for  $N$ , we can obtain the corresponding location  $\ell$  valid for  $\langle M, N \rangle$  as  $\ell = 2 \# \ell_N$ . More generally, given a composite  $n$ -ary expression  $k(M_1, \dots, M_n) = k(\overrightarrow{M_n})$ , we can map locations for  $M_i$  into locations for  $k(\overrightarrow{M_n})$  by means of a *location injection*  $\mathfrak{J}_{n,i}^{\text{Loc}} = \lambda \ell. i \# \ell$ . Dually, a location for  $k(\overrightarrow{M_n})$  can be processed by case analysis as follows: we map the root location  $\epsilon$  to a given  $N$  of type  $\tau$ ; moreover, for each  $i = 1, \dots, n$ , we provide a function  $p_i$  from  $\text{Loc}$  to a fixed type  $\tau$  which is used to map locations  $i \# \ell$  for the  $i$ -th subexpressions to  $p_i \ell$ . This case analysis, which we call *location elimination*, can be expressed as a term of the language, which we will refer to as  $\mathfrak{E}_n^{\text{Loc}}$ .

### 3.2 Labels and provenance labellings

A provenance labelling for an expression  $M$  is a function that maps all locations valid for  $M$  to expressions of an arbitrary option type  $\sim\tau$ , where  $\bullet$  is taken to represent a default, non-informative label. We will denote the type of  $\tau$ -valued labellings by  $\text{Prov}_\tau = \text{Loc} \supset \sim\tau$ ; we will also allow ourselves to write  $\text{Prov}$  when the intended  $\tau$  is obvious. We will use the metavariables  $a$  and  $\tilde{a}$ , possibly decorated with subscripts or superscripts, to denote base labels and optional labels respectively. In general, we will define provenance labellings only for the locations in  $\text{locs}(M)$ , even though our type system does not guarantee that labellings will be applied consistently: we assume that if a labelling is not explicitly defined for a location, it returns  $\bullet$ . The default labelling  $\lambda \bullet$  will be denoted by  $\perp$ .

The labelling for a composite  $n$ -ary expression  $k(\overrightarrow{M_n})$  can be obtained by putting together labellings  $p_i$  for every sub-expression  $M_i$ , and providing an additional label  $a$  for the root. The corresponding  $n$ -ary *labelling injection* operator  $\mathfrak{J}_n^{\text{Prov}}$  has the same definition of the location elimination  $\mathfrak{E}_n^{\text{Loc}}$ . Dually, a labelling for a composite expression  $k(\overrightarrow{M_n})$  induces a projection labelling for each of its sub-expressions  $M_i$ . The *labelling projections*  $\mathfrak{E}_{n,i}^{\text{Prov}}$  defined as  $\lambda p.p \circ \mathfrak{J}_{n,i}^{\text{Loc}}$  compose a labelling for  $k(\overrightarrow{M_n})$  with the  $i$ -th injection.

### 3.3 Provenance transducers

The way a labelling is transformed by evaluation allows us to express different kinds of provenance, including where-provenance, dependency provenance, and expression provenance: whenever an expression  $M$  reduces to  $N$ , there should be a way to transform any labelling for  $M$  into a labelling for  $N$ . This transformation will be expressed intra-linguistically, by means of functions of type  $\text{Prov}_\tau \supset \text{Prov}_\tau$  which we call *provenance transducers*; this type will be also written  $\text{Trans}_\tau$  for short. Rather than defining an ad-hoc transducer for each possible reduction, which would be of little use, we will give rules to synthesize transducers based on the reduction rules of the language: for example, we will have a transducer for the reduction of function application, two transducers for the two reductions of case analysis over option types, and similarly two transducers for case analysis over lists, etc.

In the special case of congruence rules, the corresponding transducer will not be provided by the user: the language provides default transducers that reflect the administrative nature of these rules. For example, given a transducer  $T$  for the reduction  $M \hookrightarrow M'$ , the transducer for the congruence reduction  $M N \hookrightarrow M' N$  can and should be obtained by lifting  $T$ .

$$\begin{array}{lll}
x|_\ell = x, \ell & M|_\epsilon = M, \epsilon & \oplus(\overrightarrow{M_n})|_{i\#\ell} = M_i|_\ell \\
(f(x).M)|_{1\#\ell} = M|_\ell & (M_1 M_2)|_{i\#\ell} = M_i|_\ell & \langle M_1, M_2 \rangle|_{i\#\ell} = M_i|_\ell \\
\text{fst}(M)|_{1\#\ell} = M|_\ell & \text{snd}(M)|_{1\#\ell} = M|_\ell & (\sim M)|_{1\#\ell} = M|_\ell \\
(M_1 \# M_2)|_{i\#\ell} = M_i|_\ell & & 
\end{array}$$

$$\text{case } M_1 \text{ of } \{- \mapsto M_2; - \mapsto M_3\}|_{i\#\ell} = M_i|_\ell$$

Figure 2: Location lookup

Location injection:	$\text{Loc} \supset \text{Loc}$ $\mathfrak{J}_{n.i}^{\text{Loc}} := \lambda \ell. i \# \ell$ ( $n$ is irrelevant)
Location elimination:	$(\text{Loc} \supset \tau) \supset \dots \supset (\text{Loc} \supset \tau) \supset \tau \supset \text{Loc} \supset \tau$ $\mathfrak{E}_n^{\text{Loc}} := \lambda p_1, \dots, p_n, a, \ell. \text{case } \ell \text{ of } \{\epsilon \mapsto a; 1 \# \ell' \mapsto p_1 \ell'; \dots; n \# \ell' \mapsto p_n \ell'\}$
Provenance injection:	$\text{Prov}_\tau \supset \dots \supset \text{Prov}_\tau \supset \sim \tau \supset \text{Prov}_\tau$ $\mathfrak{J}_n^{\text{Prov}} := \mathfrak{E}_n^{\text{Loc}}$
Provenance projection:	$\text{Prov}_\tau \supset \text{Prov}_\tau$ $\mathfrak{E}_{n.i}^{\text{Prov}} := \lambda p. p \circ \mathfrak{J}_{n.i}^{\text{Loc}}$
Transducer congruence:	$\text{Trans}_\tau \supset \text{Trans}_\tau$ $\mathfrak{J}_{n.i} := \lambda t, p. \mathfrak{J}_n^{\text{Prov}} (\mathfrak{E}_{n.1}^{\text{Prov}} p) \dots (\mathfrak{E}_{n.(i-1)}^{\text{Prov}} p) (t (\mathfrak{E}_{n.i}^{\text{Prov}} p)) (\mathfrak{E}_{n.(i+1)}^{\text{Prov}} p) \dots (\mathfrak{E}_{n.n}^{\text{Prov}} p) (p \epsilon)$

Figure 3: Definition of the provenance combinators

In general, given a reduction  $M_i \hookrightarrow M'_i$  and a corresponding transducer  $T$ , the transducer for the congruence  $k(M_1, \dots, M_i, \dots, M_n) \hookrightarrow k(M_1, \dots, M'_i, \dots, M_n)$  must perform the following actions:

1. obtain a labelling  $p$  for the full unreduced expression;
2. apply labelling projections to obtain the labellings  $p_1, \dots, p_n$  for each of the sub-expressions  $M_1, \dots, M_n$ ;
3. apply  $T$  to  $p_i$  to obtain a labelling for  $M'_i$ ;
4. use the labelling injection  $\mathfrak{J}_n^{\text{Prov}}$  to compose the new labelling for  $M'_i$  with the old labellings for the unchanged  $M_j$  ( $j \neq i$ ).

The above actions are expressed by the operator  $\mathfrak{J}_{n.i}$  ( $i$ -th *transducer congruence* of order  $n$ ), defined in terms of  $\mathfrak{J}_n^{\text{Prov}}$  and  $\mathfrak{E}_{n.j}^{\text{Prov}}$ .

Fig. 3 summarizes the definition of location injections and elimination, labelling injection and projections, and transducer congruences. We will refer to these operations, collectively, by the name of *provenance combinators*.

### 3.4 Semantics

We now use the notions introduced in the previous sections to formally define a provenance-aware small-step semantics: its rules are shown in Fig. 4. The reduction judgment  $F \vdash M \xrightarrow{T} N$  of provenance-aware semantics involves a provenance transducer  $T$ , which has type  $\text{Trans}_\tau$  for some  $\tau$ , and relates the provenance of  $M$  to that of  $N$ . As anticipated, some of the reduction rules make use of the relocation operator  $\rho(\dots)$ : to understand why, it is sufficient to see that a transducer ultimately needs to translate every location in the reduct into a location in the redex. Sometimes this translation can be expressed as an elementary program: for instance, when reducing a pair projection

$$\text{fst } \langle R, S \rangle \hookrightarrow R$$

a location  $\ell$  for  $R$  corresponds to  $1 \# 1 \# \ell$  in  $\text{fst } \langle R, S \rangle$ . However, for reduction rules involving substitutions, the correspondence is not so simple. In the case of a function application reduction

$$(f(x).M) N \hookrightarrow M[f(x).M, N/f, x]$$

the only way to map locations in the reduct to locations in the redex is by means of a structural analysis of the expressions involved in the reduction, which is not allowed by standard language constructs. We thus extend the language with relocation expressions  $\rho(\vec{x}.M)$  which allow us to map redex locations to unreduced locations by essentially “undoing” the substitution.

Relocation operations are functions with type  $\text{Loc} \supset \text{Loc}$ : when  $\rho(\vec{x}_n.M)$  is applied to a location  $\ell$  in  $\text{locs}(M[\vec{N}_n/\vec{x}_n])$ , it first computes the lookup  $M|_\ell$ : if the lookup yields  $x_i, \ell'$ , for  $i = 1, \dots, n$  – i.e. one of the variables declared by the relocation, together with an unconsumed location –  $\ell'$  represents the path to the subexpression of  $N_i$  which corresponds to location  $\ell$  in the substituted expression; in this case, the relocation maps  $\ell$  to  $i \# \ell'$ . If instead the lookup  $M|_\ell$  returns another subterm of  $M$  together with  $\epsilon$ , the relocation maps  $\ell$  to  $(n+1) \# \ell$ ; in all the other cases,  $\ell$  is an invalid location for  $M[\vec{N}_n/\vec{x}_n]$ , and we return a (meaningless) default location  $\epsilon$ .

Reduction rules

$$\begin{array}{c}
\frac{}{F \vdash (f(x).M) N \xrightarrow{F_\beta (\beta \rho(f,x.M))} M[f(x).M, N/f, x]} \quad \frac{}{F \vdash \oplus(\vec{V}_n) \xrightarrow{F_\oplus \vec{V}_n} \hat{\oplus}(\vec{V}_n)} \\
\frac{}{F \vdash \mathbf{fst}(\langle M, N \rangle) \xrightarrow{F_{\mathbf{fst}}} M} \quad \frac{}{F \vdash \mathbf{snd}(\langle M, N \rangle) \xrightarrow{F_{\mathbf{snd}}} N} \quad \frac{(\mathbf{tt} \mapsto N) \in m}{F \vdash \mathbf{case} \mathbf{ff} \text{ of } m \xrightarrow{F_{\mathbf{tt}}} N} \quad \frac{(\mathbf{ff} \mapsto N) \in m}{F \vdash \mathbf{case} \mathbf{ff} \text{ of } m \xrightarrow{F_{\mathbf{ff}}} N} \\
\frac{(\bullet \mapsto N) \in m}{F \vdash \mathbf{case} \bullet \text{ of } m \xrightarrow{F_\bullet} N} \quad \frac{(\sim x \mapsto N) \in m}{F \vdash \mathbf{case} \sim M \text{ of } m \xrightarrow{F_\sim (\beta_\sim \rho(x.N))} N[M/x]} \\
\frac{(\epsilon \mapsto N) \in m}{F \vdash \mathbf{case} \epsilon \text{ of } m \xrightarrow{F_\epsilon} N} \quad \frac{(h \# t \mapsto N) \in m}{F \vdash \mathbf{case} M_1 \# M_2 \text{ of } m \xrightarrow{F_\# (\beta_\# \rho(h,t.N))} N[M_1, M_2/h, t]} \\
\frac{M|_\ell = M', \epsilon \quad M' \notin \vec{x}_n}{F \vdash \rho(\vec{x}_n.M) \ell \xrightarrow{\lambda \perp} (n+1) \# \ell} \quad \frac{M|_\ell = x_i, \ell' \quad 1 \leq i \leq n}{F \vdash \rho(\vec{x}_n.M) \ell \xrightarrow{\lambda \perp} i \# \ell'} \\
\frac{M|_\ell = N, \ell' \quad \ell' \neq \epsilon \quad \text{for all } i \text{ s.t. } 1 \leq i \leq n : N \neq x_i}{F \vdash \rho(\vec{x}_n.M) \ell \xrightarrow{\lambda \perp} \epsilon}
\end{array}$$

Congruence rules

$$\begin{array}{c}
\frac{F \vdash M_i \xrightarrow{T} M'_i}{F \vdash \oplus(\vec{M}_n) \xrightarrow{\mathfrak{J}_{n,i} T} \oplus(M_1, \dots, M'_i, \dots, M_n)} \quad \frac{F \vdash M \xrightarrow{T} M'}{F \vdash M N \xrightarrow{\mathfrak{J}_{2,1} T} M' N} \quad \frac{F \vdash N \xrightarrow{T} N'}{F \vdash M N \xrightarrow{\mathfrak{J}_{2,2} T} M N'} \\
\frac{F \vdash M \xrightarrow{T} M'}{F \vdash \langle M, N \rangle \xrightarrow{\mathfrak{J}_{2,1} T} \langle M', N \rangle} \quad \frac{F \vdash N \xrightarrow{T} N'}{F \vdash \langle M, N \rangle \xrightarrow{\mathfrak{J}_{2,2} T} \langle M, N' \rangle} \\
\frac{F \vdash M \xrightarrow{T} M'}{F \vdash \mathbf{fst}(M) \xrightarrow{\mathfrak{J}_{1,1} T} \mathbf{fst}(M')} \quad \frac{F \vdash M \xrightarrow{T} M'}{F \vdash \mathbf{snd}(M) \xrightarrow{\mathfrak{J}_{1,1} T} \mathbf{snd}(M')} \\
\frac{F \vdash M \xrightarrow{T} M'}{F \vdash \mathbf{case} M \text{ of } m \xrightarrow{\mathfrak{J}_{3,1} T} \mathbf{case} M' \text{ of } m}
\end{array}$$

Figure 4: Provenance-carrying semantics

Thanks to relocations, we can map the locations of a beta-reduced expression back to the original function application. For instance, if  $\ell \in \mathbf{locs}(M[f(x).M, N/f, x])$ ,  $\rho(f, x.M) \ell$  will reduce to one among  $1 \# \ell_1$  (where  $\ell_1 \in \mathbf{locs}(f(x).M)$ ),  $2 \# \ell_2$  (where  $\ell_2 \in \mathbf{locs}(N)$ ), or  $3 \# \ell_3$  (where  $\ell_3 \in \mathbf{locs}(M)$ ). Mapping these locations to locations in  $\mathbf{locs}((f(x).M) N)$  is then a matter of a simple case analysis, which we provide as a relocation helper  $\beta$  (with type  $(\mathbf{Loc} \supset \mathbf{Loc}) \supset \mathbf{Loc} \supset \mathbf{Loc}$ ):

$$\beta := \lambda r, \ell. \mathbf{case} r \ell \text{ of } \{3 \# \ell_3 \mapsto 1 \# 1 \# \ell_3; \_ \mapsto \ell\}$$

Then for all  $N$ , the expression  $\beta \rho(f, x.M)$  maps locations in  $\mathbf{locs}(M[f(x).M, N/f, x])$  to locations in  $\mathbf{locs}((f(x).M) N)$ .

Non-recursive function applications  $(\lambda x.M) N \hookrightarrow M[x \mapsto N]$  are a special case of recursive functions, thus the same operator can be used to express their relocation function as well.

The view  $F$  is, in essence, a collection of basic transducers that correspond to the basic reduction rules. This is especially clear in the rules for the contraction of pair projections and the simpler forms of **case** reductions, which simply return the transducers  $F_{\mathbf{fst}}$ ,  $F_{\mathbf{snd}}$ ,  $F_\bullet$ ,  $F_\epsilon$ ,  $F_{\mathbf{tt}}$ , and  $F_{\mathbf{ff}}$ . The reduction of primitive operations is similar, save for the fact that we allow a transducer  $F_\oplus$  to receive the actual arguments of  $\oplus$  as parameters; furthermore, notice that primitive operations only reduce when applied to values.

The rule for function application must express the propagation of provenance under substitution; this requires us to produce a transducer capable of mapping locations in the substitution  $M[f(x).M, N/f, x]$  back to their label in the unreduced application  $(f(x).M) N$ . To accomplish this task,  $F_\beta$  will receive the relocation function  $\beta \rho(f, x.M)$  as a parameter. Case analysis can reduce to substituted expressions as well, which explains why  $F_\sim$  and  $F_\#$  receive relocations  $\beta_\sim \rho(x.N)$  and  $\beta_\# \rho(h, t.N)$  as parameters; the concrete definition of  $\beta_\sim$  and

Expression typing

$$\begin{array}{c}
\frac{c : \tau \in \Sigma}{\Gamma \vdash c : \tau} \quad \frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau} \quad \frac{\Gamma \vdash \vec{M} : \vec{\sigma} \quad \oplus : \vec{\sigma} \supset \tau \in \Sigma}{\Gamma \vdash \oplus(\vec{M}) : \tau} \quad \frac{\Gamma, \vec{x} : \vec{\sigma} \vdash M : \tau}{\Gamma \vdash \rho(\vec{x}.M) : \text{Loc} \supset \text{Loc}} \\
\\
\frac{\Gamma \vdash M : \sigma \quad \Gamma \vdash N : \tau}{\Gamma \vdash \langle M, N \rangle : \sigma \times \tau} \quad \frac{\Gamma \vdash M : \sigma \times \tau}{\Gamma \vdash \text{fst}(M) : \sigma} \quad \frac{\Gamma \vdash M : \sigma \times \tau}{\Gamma \vdash \text{snd}(M) : \tau} \\
\\
\frac{\Gamma, f : \sigma \supset \tau, x : \sigma \vdash M : \tau}{\Gamma \vdash f(x).M : \sigma \supset \tau} \quad \frac{\Gamma \vdash M : \sigma \supset \tau \quad \Gamma \vdash N : \sigma}{\Gamma \vdash M N : \tau} \quad \frac{\Gamma \vdash M : \mathbf{B} \quad \Gamma \vdash N : \tau \quad \Gamma \vdash R : \tau}{\Gamma \vdash \text{case } M \text{ of } \{\mathbf{tt} \mapsto N; \mathbf{ff} \mapsto R\} : \tau} \\
\\
\frac{}{\Gamma \vdash \bullet : \sim \sigma} \quad \frac{\Gamma \vdash M : \sigma}{\Gamma \vdash \sim M : \sim \sigma} \quad \frac{\Gamma \vdash M : \sim \sigma \quad \Gamma \vdash N : \tau \quad \Gamma, x : \sigma \vdash R : \tau}{\Gamma \vdash \text{case } M \text{ of } \{\bullet \mapsto N; \sim x \mapsto R\} : \tau} \\
\\
\frac{}{\Gamma \vdash \epsilon : [\sigma]} \quad \frac{\Gamma \vdash M : \sigma \quad \Gamma \vdash N : [\sigma]}{\Gamma \vdash M \# N : [\sigma]} \quad \frac{\Gamma \vdash M : [\sigma] \quad \Gamma \vdash N : \tau \quad \Gamma, h : \sigma, t : [\sigma] \vdash R : \tau}{\Gamma \vdash \text{case } M \text{ of } \{\epsilon \mapsto N; h \# t \mapsto R\} : \tau}
\end{array}$$

Provenance view typing

$$\frac{\text{for all } \kappa : \Gamma \vdash F_\kappa : \mathcal{T}_\sigma(\kappa)}{\Gamma \vdash F : \mathcal{T}_\sigma}$$

Figure 5: Typing rules

$$\begin{aligned}
\mathcal{T}_\sigma(\beta) &= (\text{Loc} \supset \text{Loc}) \supset \text{Trans}_\sigma \\
\mathcal{T}_\sigma(\sim) &= (\text{Loc} \supset \text{Loc}) \supset \text{Trans}_\sigma \\
\mathcal{T}_\sigma(\#) &= (\text{Loc} \supset \text{Loc}) \supset \text{Trans}_\sigma \\
\mathcal{T}_\sigma(\oplus) &= \vec{\tau} \supset \text{Trans}_\sigma && (\text{if } \oplus : \vec{\tau} \supset \tau' \in \Sigma) \\
\mathcal{T}_\sigma(-) &= \text{Trans}_\sigma
\end{aligned}$$

Figure 6: The type of views

$\beta_\#$  should not be too surprising:

$$\begin{aligned}
\beta_\sim &:= \lambda r, \ell. \text{case } r \ell \text{ of} \\
&\quad \{1 \# \ell_1 \mapsto 1 \# 1 \# \ell_1; 2 \# \ell_2 \mapsto 3 \# \ell_2; - \mapsto \ell\} \\
\beta_\# &:= \lambda r, \ell. \text{case } r \ell \text{ of} \\
&\quad \{1 \# \ell_1 \mapsto 1 \# 1 \# \ell_1; 2 \# \ell_2 \mapsto 1 \# 2 \# \ell_2; - \mapsto \ell\}
\end{aligned}$$

(notice that in  $\beta_\sim$  the last branch of the case analysis can only be reached when  $\ell$  is  $\epsilon$  or an invalid location).

Finally, three rules define the computational behaviour of the relocater  $\rho$ . The first rule applies to locations  $\ell$  that, within  $M[\vec{N}_n/\vec{x}_n]$ , reference a proper subterm of  $M$  that is not one of the substituted variables: this is converted to the location  $(n+1)\#\ell$ . The second one explains what to do when the lookup of location  $\ell$  within  $M$  references the  $i$ -th substituted variable, possibly with an unconsumed  $\ell'$  (which is a right-hand sublist of the original  $\ell$ ): the reduced expression is  $\ell'$ , prepended by the natural number  $i$ . Finally, the third rule provides the evaluation of a relocation applied to an invalid location: this situation is meaningless, but we reduce to  $\epsilon$ , with no special meaning other than it being a canonical location.

Since relocations are usually tools to extract provenance, and not expressions whose provenance should be extracted, for both reduction rules involving  $\rho$  we provide a forgetful transducer  $\lambda_{\cdot}\perp$ .

Congruence rules use the transducer congruences from Fig. 3. The generic form of a congruence is that of an  $n$ -ary composite expression  $k(M_1, \dots, M_i, \dots, M_n)$ , which reduces to  $k(M_1, \dots, M'_i, \dots, M_n)$  whenever  $M_i$  reduces to  $M'_i$ . Given a transducer  $T$  from  $M_i$  to  $M'_i$  we can obtain the congruence transducer by means of the combinator  $\mathcal{J}_{n,i}$ .

### 3.5 Type system

The type system of PIC (Fig. 5) is defined in a standard way. Aside from mentioning that it is parametric on the signature  $\Sigma$  used to type constants and primitive operations, the only rule that deserves to be mentioned here is the one concerning relocations. Relocations, which are functions from  $\text{Loc}$  to  $\text{Loc}$ , are well typed provided that their argument is well-typed in a properly extended context.



Even though provenance views as such are not PIC-expressions, they are used in reductions and must be well-typed to function properly. To typecheck a view, we use an auxiliary definition  $\mathcal{T}_\sigma$  (Fig. 6): each of the transducers  $F_\kappa$  within a view  $F$  must have type  $\mathcal{T}_\sigma(\kappa)$ , where  $\sigma$  is a type of provenance labels and  $\kappa$  identifies one of the basic reduction rules. Each  $\mathcal{T}_\sigma(\kappa)$  returns a provenance transducer of type  $\text{Trans}_\sigma$ , which may be parametrized by a relocation (for  $\kappa = \beta, \sim, \#$ ) or by the actual arguments of the primitive operation being reduced (for  $\kappa = \oplus$ ). Type-safety for PIC follows immediately.

**Theorem 1** (Preservation and progress). *1. If  $\Gamma \vdash M : \sigma$ ,  $\Gamma \vdash F : \mathcal{T}_\tau$ , and  $F \vdash M \xrightarrow{T} M'$ , then  $\Gamma \vdash M' : \sigma$  and  $\Gamma \vdash T : \text{Trans}_\tau$ .*  
*2. If  $\vdash M : \sigma$  and  $M$  is not a value, then for all  $F, \tau$  s.t.  $\vdash F : \mathcal{T}_\tau$ , there exist  $T, M'$  such that  $F \vdash M \xrightarrow{T} M'$ .*

## 4 Provenance views

The provenance-carrying semantics we defined allows us to choose a view  $F$  with great flexibility. Consider for example the reduction:

$$F \vdash \mathbf{fst}(\langle M, N \rangle) \xrightarrow{F_{\mathbf{fst}}} M$$

Here, the transducer  $F_{\mathbf{fst}}$  is allowed to transform the input labelling without any restriction: the most obvious choice would be to propagate the labels of the first element of the source pair, but other choices are possible.  $F$  could erase all labels, or add a static label to arbitrary subterms of  $M$ ; it could even provide a new labelling of  $M$  using labels taken from  $N$ , even though  $N$  is not part of the reduced term. The only limitation is that  $F_{\mathbf{fst}}$  be a transducer expressible as a language function.

The generality with which provenance views can be defined is one of the strengths of our approach, but it comes at a cost: views will usually be defined by pattern matching on locations, and the transducers for function applications and case analysis will often have to use a relocation function explicitly, which is rather cumbersome. For many of the notions of provenance we are interested in, such a generality is not needed and is only an element of confusion. Consider, for example, a  $\mathbf{fst}$ -projection labelled with three distinguished provenance annotations in the form  $\mathbf{fst}(\langle M^{\tilde{a}_1}, N^{\tilde{a}_2} \rangle^{\tilde{a}_3})$  (for better readability, we represent provenance annotations as superscripts, rather than providing an explicit function expression from locations to annotations). When reducing the projection, we might want the outermost label of the reduced  $M$  to depend only on those  $\tilde{a}_i$  (ignoring the annotations of  $N$ ) and the inner subexpressions of  $M$  to keep the same labels that they had before the reduction. In other words, we would like the labelling to be transformed as follows:

$$\mathbf{fst}(\langle M^{\tilde{a}_1}, N^{\tilde{a}_2} \rangle^{\tilde{a}_3}) \rightsquigarrow M^{v_{\mathbf{fst}} \tilde{a}_3 \tilde{a}_1 \tilde{a}_2}$$

where  $v_{\mathbf{fst}}$  is some expression with type  $\sim\tau \supset \sim\tau \supset \sim\tau \supset \sim\tau$ , for an annotation type  $\sim\tau$ .

Similarly, when reducing  $((f(x).M^{\tilde{a}_1})^{\tilde{a}_2} N^{\tilde{a}_3})^{\tilde{a}_4}$ , we often expect the reduced term to be, essentially,  $M^{\tilde{a}_1}$  where free occurrences of  $f$  and  $x$  have been replaced by  $(f(x).M^{\tilde{a}_1})^{\tilde{a}_2}$  and  $N^{\tilde{a}_3}$ ; the outermost label of the reduced term could additionally depend on  $\tilde{a}_2$  and  $\tilde{a}_4$  (which annotate AST nodes that are destroyed by the reduction). A transducer  $F_\beta$  satisfying these conditions could be defined as:

$$\lambda r, p, \ell. \text{ case } \ell \text{ of } \{ \epsilon \mapsto v_\beta (p \ \epsilon) (p \ [1]) (p \ (r \ \epsilon)); \_ \mapsto p \ (r \ \ell) \}$$

where  $v_\beta$  is an arbitrary term combining three annotations into one,  $(p \ \epsilon)$  and  $(p \ [1])$  return the annotations corresponding to  $\tilde{a}_4$  and  $\tilde{a}_2$  in the example. The argument  $r$  receives a relocation function in the form  $\beta \ \rho(f, x.M)$ , provided by the standard provenance-carrying semantics, which is used to retrieve the “natural” annotation arising from the substitution of a labelled term into another labelled term.

These examples show that while manipulating raw provenance labellings is tricky, we can keep things simple by providing functions like  $v_{\mathbf{fst}}$  and  $v_\beta$ , which combine a few distinguished annotations and propagate all the others.

Based on this considerations, we provide the following *lifting framework* that lifts functions  $v_\kappa$  on provenance annotations to the corresponding provenance transducers  $F_\kappa$ :

$$\begin{aligned}
F_\beta(v_\beta) &:= \lambda r, p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_\beta (p \epsilon) (p [1]) (p (r \epsilon)); - \mapsto p (r \ell)\} \\
F_{\text{fst}}(v_{\text{fst}}) &:= \lambda p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_{\text{fst}} (p \epsilon) (p [1]) (p [1; 1]); - \mapsto p (1 \# 1 \# \ell)\} \\
F_{\text{snd}}(v_{\text{snd}}) &:= \lambda p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_{\text{snd}} (p \epsilon) (p [1]) (p [1; 2]); - \mapsto p (1 \# 2 \# \ell)\} \\
F_{\text{tt}}(v_{\text{tt}}) &:= \lambda p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_{\text{tt}} (p \epsilon) (p [1]) (p [2]); - \mapsto p (2 \# \ell)\} \\
F_{\text{ff}}(v_{\text{ff}}) &:= \lambda p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_{\text{ff}} (p \epsilon) (p [1]) (p [3]); - \mapsto p (3 \# \ell)\} \\
F_\bullet(v_\bullet) &:= \lambda p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_\bullet (p \epsilon) (p [1]) (p [2]); - \mapsto p (2 \# \ell)\} \\
F_\sim(v_\sim) &:= \lambda r, p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_\sim (p \epsilon) (p [1]) (p (r \epsilon)); - \mapsto p (r \ell)\} \\
F_\epsilon(v_\epsilon) &:= \lambda p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_\epsilon (p \epsilon) (p [1]) (p [2]); - \mapsto p (2 \# \ell)\} \\
F_\#(v_\#) &:= \lambda r, p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_\# (p \epsilon) (p [1]) (p (r \epsilon)); - \mapsto p (r \ell)\} \\
F_\oplus(v_\oplus) &:= \lambda \vec{x}_n, p, \ell. \text{ case } \ell \text{ of} \\
&\quad \{\epsilon \mapsto v_\oplus \vec{x}_n (p \epsilon) (p [1]) (p [2]) \cdots (p [n]) \\
&\quad \quad ; - \mapsto \epsilon\}
\end{aligned}$$

$$F(v) := \langle F_\beta(v_\beta), \dots, F_\oplus(v_\oplus) \rangle$$

## 4.1 Where-provenance

Using the lifting framework, we define a where-provenance view  $W$  by means of label-propagating transducers. The view can use any optional type of labels  $\sim\tau$ . When a reduction copies certain data, the corresponding transducer will propagate the provenance annotations along with the data; when a reduction transforms the term beyond recognizability (for example, when reducing primitive operations), there is no label to propagate, therefore the corresponding transducer will produce the default, uninformative label  $\bullet$ .

$$\begin{aligned}
w_\oplus &:= \lambda \vec{x}_n, \tilde{a}, \tilde{a}_n. \bullet \\
w_\kappa &:= \lambda -, -, \tilde{a}. \tilde{a} & (\kappa \neq \oplus) \\
W &:= F(w)
\end{aligned}$$

*Example 1.* Let  $f := \lambda x, y. \text{case } x \stackrel{?}{=} y \text{ of } \{\text{tt} \mapsto x; \text{ff} \mapsto y + 1\}$ . We map  $f$  2 to a list  $[1; 2; 3]$ :

$$\text{map } (f \ 2)[1; 2; 3] \xrightarrow{*} [2; 2; 4]$$

Now consider the following provenance labelling  $L$  assigning annotations  $\tilde{a}, \tilde{a}_1, \tilde{a}_2, \tilde{a}_3$  as follows:

$$\text{map } (f \ 2^{\tilde{a}}) [1^{\tilde{a}_1}; 2^{\tilde{a}_2}; 3^{\tilde{a}_3}]$$

(subexpressions without an annotation are considered to be labelled by  $\bullet$ ). A where-provenance reduction acts as follows:

$$\text{map } (f \ 2^{\tilde{a}}) [1^{\tilde{a}_1}; 2^{\tilde{a}_2}; 3^{\tilde{a}_3}] \xrightarrow{*} [2; 2^{\tilde{a}}; 4]$$

In the output list, only the second element is annotated with  $a$ , meaning that it was copied from the argument of  $f$ ; the other two elements were not copied, but obtained by incrementing previous values in the same position, and thus receive the  $\bullet$  annotation.

As a general rule, after each execution step we expect subexpressions labelled with  $\sim a$  to be copied from parts of the original expression having the same label. The actual well-behavedness property of where-provenance, in our setting, is less naïve than that: consider the reduction

$$(\lambda x.(\lambda y.(x y)^a)) N \hookrightarrow \lambda y.(N y)^a$$

The subexpression  $(N y)$  is indeed derived from the similarly labelled  $(x y)$ , but the two are only equal up to the substitution  $[N/x]$ . The relation between the two expressions can be made formal:

**Definition 1.**  $M \sqsubseteq N \iff \exists s.M[s] \xrightarrow{*} N$ , where  $s$  is a substitution.

**Lemma 1.** 1.  $\sqsubseteq$  is an order relation;

2. if  $M \sqsubseteq N$  and  $N \xrightarrow{*} R$ , then  $M \sqsubseteq R$ .

We can then prove that  $W$  satisfies the following well-behavedness property of where-provenance:

**Theorem 2** (well-behavedness of  $W$ ). *Suppose  $W \vdash M \xrightarrow{T} N$ : then for all labellings  $L$  and  $\ell \in \text{locs}(N)$ , if  $T L \ell \rightsquigarrow \sim a$ , there exists  $\ell' \in \text{locs}(M)$  such that  $L \ell' \rightsquigarrow \sim a$  and  $M|_{\ell'} \sqsubseteq N|_{\ell}$ .*

## 4.2 Expression provenance

To define expression provenance, we allow provenance labels to be simplified syntax trees with constants or primitive annotations as leaves and primitive operation symbols as inner nodes. More formally, we assume a type of labels  $\tau$  containing the following elements:

$$a ::= \bar{b} \mid k_c \mid k_{\oplus}(\bar{a}_n)$$

where  $b$  belongs to a type of basic annotations  $\tau_b$ , and  $k_c$  and  $k_{\oplus}$ , defined for all constants  $c$  and all  $n$ -ary basic operations  $\oplus$ , are reified representations of language expressions. As usual,  $\tau$  will be wrapped in an optional type  $\sim\tau$ .

Accordingly, a numeric value, say 2, could have several possible labels: a label  $\sim\bar{b}$  means that 2 was copied from a part of the input with the same label; a label  $\sim k_+(a_1, a_2)$  means it was obtained by adding together two values originally labelled with  $\sim a_1$  and  $\sim a_2$ ; the label  $\bullet$ , as usual, provides no information; finally the label  $\sim k_2$ , merely indicates a literal 2, but does not specify its provenance otherwise, therefore it can be considered a variant of  $\bullet$ .

The expression provenance view can be defined by means of a simple variation on where-provenance. The only transducer that needs a different definition, unsurprisingly, is the one associated with the evaluation of primitive operations ( $e_{\oplus}$ ):

$$\begin{aligned} e_{\oplus} &:= \lambda \vec{x}_n, \tilde{a}, \vec{a}_n. k_{\oplus}(\overrightarrow{\tilde{a}_n * x_n}) \\ e_{\kappa} &:= \lambda -, -, \tilde{a}. \tilde{a} \end{aligned} \quad (\kappa \neq \oplus)$$

$$E := F(e)$$

In the definition, we use the following  $*$  operation:

$$(\tilde{a}) * x = \begin{cases} \sim k_x & \text{if } \tilde{a} = \bullet \\ \tilde{a} & \text{else} \end{cases}$$

where  $k_x$  is the label corresponding to the constant  $x$  (remember that primitive operations are reduced when applied to constants, thus we must have  $x = c$  for some  $c$ ). This allows us to conjure an informative annotation for constants that have not been given a provenance.

*Example 2.* We start with the same labelling as in the where-provenance example:

$$\text{map } (f \ 2^{\tilde{a}}) [1^{\tilde{a}_1}; 2^{\tilde{a}_2}; 3^{\tilde{a}_3}]$$

If we evaluate the term according to expression provenance, we obtain the following labelled value:

$$[2^{k_+(\tilde{a}_1, k_1)}; 2^{\tilde{a}}; 4^{k_+(\tilde{a}_3, k_1)}]$$

Just like in where-provenance,  $2^{\tilde{a}}$  indicates a term that was copied from the input; the term  $2^{k_+(\tilde{a}_1, k_1)}$ , instead, has been obtained by adding together  $1^{\tilde{a}_1}$  and an unannotated literal 1.

Expression provenance annotations should allow us to recompute the annotated expression. To make this formal, let us consider metalanguage functions  $h$  mapping basic annotation values in  $\tau_b$  to arbitrary language values. We then define the extension of such an  $h$  to full annotations as the metalanguage function  $\hat{h}$  mapping values in  $\tau$  to arbitrary language values, as follows:

$$\begin{aligned}\hat{h}(\bar{b}) &= h(b) \\ \hat{h}(k_c) &= c \\ \hat{h}(k_{\oplus}(\tilde{a}_1, \dots, \tilde{a}_n)) &= \hat{\oplus}(\hat{h}(\tilde{a}_1), \dots, \hat{h}(\tilde{a}_n))\end{aligned}$$

We then introduce *consistent mappings* (a modified version of the analogous concept used in [1]) as those functions  $h$  which agree with a certain annotated expression.

**Definition 2.** Let  $M$  be a term, and  $L$  a provenance labelling for  $M$  of type  $\text{Loc} \supset \sim\tau$ . We say that  $h$  is a *consistent mapping* for  $L \triangleright M$  (and write  $h \Vdash L \triangleright M$ ) if and only if for all  $\ell \in \text{locs}(M)$  there exist an annotation  $\tilde{a}$  such that  $L \ell \rightsquigarrow \tilde{a}$  and, if  $\tilde{a} = \sim a'$ , then  $\hat{h}(a') \sqsubseteq M|_{\ell}$ .

Finally, we prove that reduction under the view  $E$  preserves consistent mappings.

**Theorem 3** (well-behavedness of  $E$ ). *If  $h \Vdash L \triangleright M$  and  $E \vdash M \xrightarrow{T} N$ , then  $h \Vdash T L \triangleright N$ .*

### 4.3 Dependency provenance

Dependency provenance associates to each expression location an annotation containing a set of labels. A full definition of dependency provenance would therefore require us to encode sets as a type of the language, providing at the same time an implementation of standard set operations. We could for example use the list encoding of sets, which is a simple programming exercise; for the purposes of this paper, we will abstract from the actual definition of the type of dependency annotations, assuming that a sound implementation of set operations like union and membership test exists. In particular, we assume that the default annotation  $\bullet$  be interpreted as the empty set of labels  $\emptyset$  rather than the “none” optional value.

The dependency view then merely amounts to taking the union of all the dependency sets involved in a reduction:

$$\begin{aligned}d_{\oplus} &:= \lambda \vec{x}_n, \vec{a}_{n+1}. \bigcup \vec{a}_{n+1} \\ d_{\kappa} &:= \lambda \tilde{a}_1, \tilde{a}_2, \tilde{a}_3. \tilde{a}_1 \cup \tilde{a}_2 \cup \tilde{a}_3 \quad (\kappa \neq \oplus)\end{aligned}$$

$$D := F(d)$$

*Example 3.* We adapt our example to dependency provenance by extending the language of provenance labels to accept sets of annotations:

$$\text{map } \langle f \ 2^{\{a\}}, [1^{\{a_1\}}; 2^{\{a_2\}}; 3^{\{a_3\}}] \rangle$$

(where subexpressions without an annotation are considered to be labelled by  $\emptyset$ ). After evaluation, we get the following dependency provenance labelling:

$$[2^{\{a_1, a\}}; 2^{\{a_2, a\}}; 4^{\{a_3, a\}}]$$

The new labelling reflects the fact that each element of the list depends from both the corresponding element in the source list, and the argument to the function  $f$ .

To state the well-behavedness property for the view  $D$  we use a relation  $L \triangleright M \approx_a L' \triangleright M'$  that holds when the annotated terms  $L \triangleright M$  and  $L' \triangleright M'$  are *quasi-equal*, up to subterms annotated with a label  $a$  (e.g.  $\langle 1, 2^{\mathcal{S}} \rangle \approx_a \langle 1, 3^{\mathcal{S}'} \rangle$  provided that  $a \in \mathcal{S} \cap \mathcal{S}'$ ). The idea is that, if starting with  $L \triangleright M$  we perform a reduction step  $D \vdash M \xrightarrow{T} N$  whose redex is entirely guarded by the label  $a$ , we obtain a term that is still quasi-equal to  $L' \triangleright M'$ . If instead the redex is not guarded by  $a$ , then we can find a reduction  $D \vdash M' \xrightarrow{T'} N'$  such that  $T L \triangleright N \approx_a T' L' \triangleright N'$ .

To indicate reductions guarded and not guarded by an annotation, we use the notation:

$$\begin{aligned}D \vdash (L \triangleright M) &\xrightarrow{T}_a N \\ D \vdash (L \triangleright M) &\xrightarrow{T}_{\bar{a}} N\end{aligned}$$

(a formal definition is given in the appendix).

Expressions  $L, M, N ::= \dots \mid \iota(F, L \triangleright M)$   
Views  $F ::= \langle M_\beta, M_{\text{fst}}, M_{\text{snd}}, M_\bullet, M_\sim, M_\epsilon, M_\#, M_\oplus, M_\iota \rangle$   
 $\text{locs}(\iota(F, L \triangleright M)) = \{\epsilon\}$

$$\frac{G \vdash M \xrightarrow{T} N}{F \vdash \iota(G, L \triangleright M) \xrightarrow{I} \iota(G, T L \triangleright N)} \quad \frac{\text{FV}(V) = \emptyset}{F \vdash \iota(G, L \triangleright V) \xrightarrow{F_\iota} \langle V, L \rangle}$$

$$\frac{\Gamma \vdash F : \mathcal{T}_\sigma \quad \Gamma \vdash L : \text{Prov}_\sigma \quad \Gamma \vdash M : \tau}{\Gamma \vdash \iota(F, L \triangleright M) : \tau \times \text{Prov}_\sigma}$$

$$\mathcal{T}_\sigma(\iota) = \text{Trans}_\sigma$$

$$F(v_\iota) = \lambda p, \ell. \text{case } \ell \text{ of } \{\epsilon \mapsto v_\iota (p \ \epsilon); \_ \mapsto \bullet\}$$

Figure 7: Syntax, semantics, and typing of inspection

**Theorem 4** (well-behavedness of  $D$ ). *Suppose  $L \triangleright M \approx_a L' \triangleright M'$ . Then:*

1. *If  $D \vdash (L \triangleright M) \xrightarrow{T}_a N$ , we have  $T L \triangleright N \approx_a L' \triangleright M'$ .*
2. *If  $D \vdash (L \triangleright M) \xrightarrow{T}_{\bar{a}} N$ , there exist  $T', N'$  such that  $D \vdash M' \xrightarrow{T'} N'$  and  $T L \triangleright N \approx_a T' L' \triangleright N'$ .*

Admittedly, this property is weaker than the *dependency-correctness* of [9], as we do not consider redexes that are *partially guarded* by  $a$ , such as  $(f(x).M)^{\{a\}} N$ . We believe that full dependency-correctness also holds, but cannot be proved without resorting to a more complex argument involving well-typedness; we thus leave this proof as future work.

## 5 Self-inspection

The language described in Section 3, together with the view definition framework of Section 4, can be compared to TML. Although we have simplified the type system by renouncing to recursive types, the two languages are quite similar; the provenance extraction framework can also be compared to our view definition framework. Our approach, however, does provide a few enhancements:

- our small-step semantics is not limited to call-by-value evaluation, like their big-step definition, but accommodates a variety of evaluation strategies;
- provenance views and provenance data are expressed in the same language as the programs whose provenance is being considered;
- in our language, we need not take the extra step of materializing the execution trace of a program before extracting its provenance;
- as a consequence, to compute the provenance of a program we do not have to replay its execution trace.

The most important reason why these features matter is that they make it possible for our language to be extended with an introspective inspection operation, which we will denote by the following syntax:

$$\iota(F, L_M \triangleright M)$$

The intended semantics of inspections, informally, is that  $M$  will be evaluated to  $V$ ; at the same time, the provenance view  $F$  will be applied to  $L_M$  to obtain a new labelling  $L_V$  for  $V$ ; finally, evaluation will produce the pair  $\langle V, L_V \rangle$ , making both the value and its provenance available for the rest of the program.

Provenance inspection is introduced by means of the extension described in Figure 7. From the point of view of location lookup, inspections are treated like opaque boxes: the `locs` operation only returns the singleton  $\{\epsilon\}$  when applied to an inspection, so access to its syntactic subterms is not allowed.

The typing rule checks that the type of the provenance view  $F$  agree with that of the provenance labelling  $L$ ; the term under inspection can be of any type.

Two different reduction rules are applied depending on whether the term under inspection has been fully evaluated or not. When the term  $M$  under inspection is not a value, it can be reduced to  $N$ , by means of the

first reduction rule, using the view  $G$  specified by the inspection; this returns a transducer  $T$  that we apply to the initial provenance labelling  $L$  to obtain a new labelling  $T L$  suitable for  $N$ .

When, after a certain number  $n$  of reduction steps, the inspected term becomes a value  $V$ , with a certain provenance

$L' = T_n (T_{n-1} \cdots (T_0 L) \cdots)$ ; then the inspection can be concluded by a final step, using the second rule, which merely returns the value  $V$  together with its labelling  $L'$ . The premise requiring that  $V$  should be a closed term is needed because free variables might be replaced by any expression, including expressions containing redexes, and in such a case the inspection should continue; however, since we do not allow reduction inside binders, this check is only necessary if we are interested in evaluating open terms.

Unlike TML, the syntax of PIC allows free nesting of provenance inspections. For this reason we also need to explain what an external observer can see when an inspection is performed. Here we need to make a choice and decide whether inner reductions can be observed or not. Since inspections come with a local provenance labelling, they cannot be easily reconciled with a non-local provenance view; then the simplest policy, which we adopt, is to make inner reductions *not* observable: this is obtained by returning the identity  $I = \lambda x.x$  as the transducer for inner reductions. The final inspection step, instead, can be observed: it employs the transducer  $F_L$ , which is defined in the outer view  $F$ .

The presence of nested inspections has implications for privacy and confidentiality: on one hand, inner inspections decide whether to share provenance information with external observers, and to what extent; on the other hand, the external observer is aware of hidden computations thanks to the transducer  $F_L$ . This interaction is not unexpected: implications of provenance for confidentiality, privacy, availability, and other security properties have been investigated by some previous work ([8, 5, 12, 16]). We will now elaborate further by considering an example of security application.

## 5.1 Example: inspection of dynamic linking

In the previous sections we have discussed provenance as a property of data, and how it evolves during the execution of a program, but we have not provided concrete examples of how provenance can be used in a concrete setting.

We will here provide a demonstration of how provenance can be used to detect possible security issues. Our example considers dynamic linking, a technique widely available in modern operating systems, which allows the linking of object code from several sources to happen at load time (or in certain cases at run time), rather than compile time. When an executable file is built from source code containing references to functions defined in a shared (or *dynamic-link*) library, those references remain unresolved; instead, special directives are added to the executable file header, including a record of the required library functions (sometimes called *import table*). When loading the executable file, the operating system is expected to find the required shared libraries and load them into memory beside the program code, and link the two by resolving the function references declared in the import table.

Dynamic linking allows a system to avoid duplication of frequently used code, both on disk and in RAM (when two processes using the same library are running concurrently, the memory image of the library can be shared across the two addressing spaces by means of paging). While 30 years ago programs usually employed less than 5 dynamic-link libraries, today's software systems can require several even hundreds of libraries provided by different developers. Libraries from trusted sources may coexist with libraries from untrusted sources, and they may use each other's services.

To model a simple dynamic linker, we assume a type  $[\mathbf{N}] \supset \mathbf{N}$  for shared library functions – an acceptable assumption, given that executable file formats often enforce a loose or trivial typing discipline for these functions. A store of type  $[\mathbf{id} \times ([\mathbf{N}] \supset \mathbf{N})]$  associates function identifiers to the corresponding shared function.

A program comprises executable code, which we model as a function of type  $[[\mathbf{N}] \supset \mathbf{N}] \supset [\mathbf{N}] \supset \mathbf{N}$ , and an import table, which in our setting is simply a list of identifiers of shared functions that should be loaded. The code receives as its first parameter a list of concrete functions matching the imported identifiers from import table, while the second argument must be a list of input parameters for the program; at the end, a natural number is returned.

The goal of the dynamic linker is to serve the list of imported functions to the program code. If the store and the import table are both ordered by increasing values of identifiers, this can be achieved by the following procedure:

```

dynalink := λstore, code, imptab.
  let imps := filter_map (λ⟨x, f⟩. x ∈ imptab) (λ⟨x, f⟩. f) store
  in code imps

```

where `filter_map`  $p f l$  ignores the elements of the list  $l$  that do not satisfy the predicate  $p$ , but is otherwise the same as `map`  $f l$ .

We now consider a toy program whose purpose is to store passwords to permanent memory. The program employs *salt* and a hash function to make dictionary attacks unfeasible in case a malicious agent gained access to the password file:

```
savepw := λimps, pw.
    let salt := $RND [] in
    $WRITE [salt];
    $WRITE [$HASH [(salt||pw)]]
```

where the semi-colon is a binary operator on natural numbers simply returning the second argument, and `||` combines two natural numbers into one, for example by concatenating their bit representations, or by means of Cantor's encoding of pairs. `HASH`, `RND`, and `WRITE` are identifiers for three imported functions: the import table for `savepw` is thus the following:

```
[HASH; RND; WRITE]
```

Finally, we write `$ID` as syntactic sugar for an imported function call; concretely, this involves a lookup by ordinal number in the import list *imps*: e.g. `$WRITE = nth imps 2` (where `nth` returns the element of a list referred to by a zero-based index).

Suppose that we do not have access to the source code of `savepw`: we can still read its import table and see that the program requires the use of a cryptographic hash function, a random number generator, and persistent storage. This is not suspicious, but for increased trust we may want to inspect the operation of `savepw` and verify that only a properly encrypted password has been stored.

For this purpose, we use the following provenance label type:

```
~label ::= • | ~private | ~random | ~rndpriv | ~breach
         | ~standard | ~unsafe | ~oneway
```

The first four labels are used to express the provenance of data: the default `•` labels data that is public or irrelevant as far as confidentiality goes; data labelled by `~private`, on the contrary, is confidential; `~random` is used for non-private data coming from a random source; finally, `~rndpriv` labels data containing both private and random information. A label `~breach`, also used to annotate data, is only generated when a potential security breach is detected.

The last three labels are used to annotate functions:

- `~standard` is used for pure functions that manipulate the input in an unknown way; they can throw away part of the input, but may not create private data out of thin air: since the randomized part of the input may not be preserved, this label combines with labelled data as follows

```
~standard(•) = •
~standard(~random) = •
~standard(~private) = ~private
~standard(~rndpriv) = ~private
~standard(~breach) = ~breach
```

- `~unsafe` annotates functions that should not receive private data, because they contain untrusted code with side effects (e.g. writing to disk, sending data over a network, displaying information on a terminal): it behaves as follows

```
~unsafe(•) = •
~unsafe(~random) = •
~unsafe(~private) = ~breach
~unsafe(~rndpriv) = ~breach
~unsafe(~breach) = ~breach
```

- `~oneway` is similar to `standard` but is used with functions known to be *one-way*; thus when applied to randomized private data, it produces non-confidential output:

```
~oneway(•) = •
~oneway(~private) = ~private
~oneway(~breach) = ~breach
~oneway(~random) = •
~oneway(~rndpriv) = •
```

This annotation propagation policy can be implemented as a provenance view  $S$ . The labelling  $L_{\text{store}}$  for the shared library functions will usually be fixed, and we assume that the concrete functions `hash`, `rnd`, `write` are labelled by `oneway`, `random`, `unsafe`. Given the password `pw` (which, for simplicity, will be a natural number), we make it confidential by defining its labelling as

$$L_{\text{pw}} := \lambda \ell. \text{case } \ell \text{ of } \{ \epsilon \mapsto \sim\text{private}; \_ \mapsto \bullet \}$$

As previously mentioned, an outer inspection cannot see what happens within an inner one, so an additional task of  $S$ , besides propagating annotations, is to flag inner inspections as possible security breaches, by means of the transducer

$$S_\iota := \lambda \_ . \_ . \text{breach}$$

Now we can apply the dynamic linker to `savepw` and try running it by means of the syntax:

$$M = \text{dynalink store savepw [HASH;RND;WRITE] pw}$$

where the provenance labelling for  $M$  can be obtained by combining  $L_{\text{store}}$ ,  $L_{\text{pw}}$ , and trivial labellings for all the other parts of  $M$ :

$$L = \mathcal{J}_2^{\text{Prov}} (\mathcal{J}_2^{\text{Prov}} (\mathcal{J}_2^{\text{Prov}} (\mathcal{J}_2^{\text{Prov}} \perp L_{\text{store}} \bullet) \perp \bullet) \perp \bullet) L_{\text{pw}} \bullet$$

Finally, we perform an inspection:

$$\iota(S, L \triangleright M) \xrightarrow{*} \langle 0, L' \rangle$$

where 0 is the value returned by the call to `WRITE`, and  $L' \in \xrightarrow{*} \sim\text{public}$  confirms that no security breaches were detected.

## 6 Conclusions

The language PIC that we have described in this paper represents provenance by means of labelling functions. This seems to be a natural approach in a calculus allowing provenance to be manipulated as a first-class expression because in this way all the provenance annotations for the same term are gathered in the same place. Other approaches based on annotation propagation would scatter this information across various subterms, and would thus require additional effort to extract the provenance and separate it from the term it describes.

Provenance as labelling functions also seems to have a relatively elegant, albeit slightly low-level, theory of combinators, which guides the definition of the provenance-aware semantics. Although defining provenance views by combining annotations locally (as in TML) is simpler than doing so by handling whole labellings (as required by our semantics), we have provided a simplified framework to bridge this gap.

Although PIC is a pure functional language, we envision extending it with imperative constructs such as references and exceptions. Experience in the related area of program slicing ([19]) tells us that such an extension would not be straightforward, but might provide a good starting point.

We have chosen to make provenance inspection local: an outer inspection cannot observe the computation happening within an inner inspection; this ensures a certain confidentiality of metadata, which cannot be shared by accident with the public. However, this is not the only possible way to handle nested inspections: a *transparent* provenance inspection could also be defined using the following evaluation rule:

$$\frac{G \vdash M \xrightarrow{G[\mathfrak{R}]} N}{F \vdash \iota(G, L \triangleright M) \xrightarrow{\mathcal{J}_{1.1} (F[\mathfrak{R}])} \iota(G, T \ L \triangleright N)}$$

where  $\mathfrak{R}$  identifies the reduction rule used in the premise, and  $F[\mathfrak{R}]$  returns the transducer for rule  $\mathfrak{R}$  according to the view  $F$ : this allows reduction within  $\iota$  to be treated like any other congruence. Apart from a slight notational complication, this extension appears straightforward.

Another possibility, given  $\iota(G, L \triangleright M)$  is to make  $M$  invisible to an external observer, but to allow an outer inspection to provide a labelling for  $G$  and  $L$ . Indeed, since a provenance view contains standard language functions, it needs not be statically determined: it is perfectly admissible to receive provenance transducers as the arguments of a program, and use them to construct a view to be used in an inspection. In other words, views may have a non-trivial provenance, which could provide a reason to inspect the *provenance of provenance*. This concept appears not to be entirely new (see e.g. [15, 17]); further investigation of relations between nested inspections and provenance of provenance will be the subject of future work.



## 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.
- [2] David A. Bearman. The power of the principle of provenance. *Archivaria*, 21:1427, 1985.
- [3] Deepavali Bhagwat, Laura Chiticariu, Wang-Chiew Tan, and Gaurav Vijayvargiya. An annotation management system for relational databases. *VLDB Journal*, 14(4):373–396, 2005.
- [4] Rajendra Bose and James Frew. Lineage retrieval for scientific data processing: a survey. *ACM Comput. Surv.*, 37(1):1–28, 2005.
- [5] Uri Braun, Avraham Shinnar, and Margo I. Seltzer. Securing provenance. In *HotSec*, 2008.
- [6] Peter Buneman, James Cheney, and Stijn Vansummeren. On the expressiveness of implicit provenance in query and update languages. *ACM Transactions on Database Systems*, 33(4):28, November 2008.
- [7] Peter Buneman, Sanjeev Khanna, and Wang-Chiew Tan. Why and where: A characterization of data provenance. In *ICDT*, number 1973 in LNCS, pages 316–330, 2001.
- [8] J. Cheney. A formal framework for provenance security. In *Proceedings of the 24th IEEE Computer Security Foundations Symposium (CSF)*, pages 281–293. IEEE, 2011.
- [9] James Cheney, Amal Ahmed, and Umut a. Acar. Provenance as dependency analysis. *Mathematical Structures in Comp. Sci.*, 21(6):1301–1337, December 2011.
- [10] Yingwei Cui, Jennifer Widom, and Janet L. Wiener. Tracing the lineage of view data in a warehousing environment. *ACM Trans. Database Syst.*, 25(2):179–227, 2000.
- [11] Susan B. Davidson and Juliana Freire. Provenance and scientific workflows: Challenges and opportunities. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD ’08, pages 1345–1350, New York, NY, USA, 2008. ACM.
- [12] Susan B. Davidson, Sanjeev Khanna, Sudeepa Roy, and Sarah Cohen Boulakia. Privacy issues in scientific workflow provenance. In *Proceedings of the 1st International Workshop on Workflow Approaches to New Data-centric Science*, Wands ’10, pages 3:1–3:6, New York, NY, USA, 2010. ACM.
- [13] Boris Glavic and Gustavo Alonso. Provenance for nested subqueries. In *Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology*, EDBT ’09, pages 982–993, New York, NY, USA, 2009. ACM.
- [14] Boris Glavic, Renée J. Miller, and Gustavo Alonso. Using sql for efficient generation and querying of provenance information. In *In Search of Elegance in the Theory and Practice of Computation*, pages 291–320, 2013.
- [15] Kwan Hee Han, Seock Kyu Yoo, and Bohyun Kim. Qualitative and quantitative analysis of workflows based on the uml activity diagram and petri net. *WSEAS Trans. Info. Sci. and App.*, 6(7):1249–1258, July 2009.
- [16] Ragib Hasan, Radu Sion, and Marianne Winslett. Introducing secure provenance: Problems and challenges. In *Proceedings of the 2007 ACM Workshop on Storage Security and Survivability*, StorageSS ’07, pages 13–18, New York, NY, USA, 2007. ACM.
- [17] Zachary Hensley, Jibonananda Sanyal, and Joshua New. Provenance in sensor data management. *Queue*, 11(12):50:50–50:63, December 2013.
- [18] Jan Hidders, Natalia Kwasnikowska, Jacek Sroka, Jerzy Tyszkiewicz, and Jan Van den Bussche. A formal model of dataflow repositories. In *Proceedings of the 4th International Conference on Data Integration in the Life Sciences*, DILS’07, pages 105–121, Berlin, Heidelberg, 2007. Springer-Verlag.
- [19] Wilmer Ricciotti, Jan Stolarek, Roly Perera, and James Cheney. Imperative functional programs that explain their work. In *ICFP 2017*, 2017. In press.
- [20] T. R. Schellenberg. The principle of provenance and modern records in the united states. *The American Archivist*, 28(1):39–41, 1965.
- [21] Yogesh Simmhan, Beth Plale, and Dennis Gannon. A survey of data provenance in e-science. *SIGMOD Record*, 34(3):31–36, 2005.

## A Proofs

**Lemma 2.** For all  $\ell \in \text{locs}(R[\overrightarrow{M_n}/\overrightarrow{x_n}])$  (where the  $x_i$  are distinct), one of the following holds:

- $R|_\ell = x_i, \ell'$  and  $R[\overrightarrow{M_n}/\overrightarrow{x_n}]|_\ell = M_i|_{\ell'}$ , for some  $\ell' \in \text{locs}(M_i)$
- $R|_\ell = R', \epsilon$  and  $R[\overrightarrow{M_n}/\overrightarrow{x_n}]|_\ell = R'[\overrightarrow{M_n}/\overrightarrow{x_n}], \epsilon$ , for some  $R'$ .

*Proof.* By structural induction on  $R$ , followed by a case analysis on  $\ell \in \text{locs}(R[\overrightarrow{M_n}/\overrightarrow{x_n}])$ . For the recursive cases, we only consider functions, but the other cases are similar.

- $R = x_i$   
then  $\ell \in \text{locs}(M_i)$  and  $R|_\ell = x_i|_\ell = x, \ell$ . We complete the proof by proving  $x_i[\overrightarrow{M_n}/\overrightarrow{x_n}]|_\ell = M_i|_\ell$ , which is trivial by definition of substitution.
- $R = y \notin \overrightarrow{x_n}$   
then  $\ell \in \text{locs}(y)$ , which implies  $\ell = \epsilon$  and  $R|_\ell = y|_\epsilon = y, \epsilon$ . Finally,  $y[\overrightarrow{M_n}/\overrightarrow{x_n}]|_\epsilon = y[\overrightarrow{M_n}/\overrightarrow{x_n}], \epsilon$ , which completes the proof.
- $R = g(y).S$ , where  $g, y \notin \overrightarrow{x_n}, \overrightarrow{M_n}$   
then  $\ell = \epsilon$  or  $\ell = 1.\ell_1$ , where  $\ell_1 \in \text{locs}(S[\overrightarrow{M_n}/\overrightarrow{x_n}])$ . In the first case, we have  $g(y).S|_\epsilon = g(y).S, \epsilon$  and  $(g(y).S)[\overrightarrow{M_n}/\overrightarrow{x_n}]|_\epsilon = (g(y).S)[\overrightarrow{M_n}/\overrightarrow{x_n}], \epsilon$ , which proves the thesis. In the other case,  $g(y).S|_{1.\ell_1} = S|_{\ell_1}$  and  $(g(y).S)[\overrightarrow{M_n}/\overrightarrow{x_n}]|_{1.\ell_1} = S[\overrightarrow{M_n}/\overrightarrow{x_n}]|_{\ell_1}$ , and the thesis follows by induction hypothesis. □

**Lemma 3.** For all  $\ell \in \text{locs}(M[f(x).M, N/f, x])$ , there exists  $\ell' \in \text{locs}(f(x).M) \ N$  such that  $\beta \rho(f, x.M) \ell \hookrightarrow \ell'$  and  $(f(x).M) \ N|_{\ell'} \sqsubseteq M[f(x).M, N/f, x]|_\ell$ .

*Proof.* By cases on  $M|_\ell$ , using Lemma 2, we get three cases:

- $M|_\ell = f, \ell_1$ ,  $M[f(x).M, N/f, x]|_\ell = f(x).M|_{\ell_1}$ , and  $\ell_1 \in \text{locs}(f(x).M)$ : we easily prove that  $\beta \rho(f, x.M) \ell \xrightarrow{*} 1 \# \ell_1$  and  $(f(x).M) \ N|_{1 \# \ell_1} = f(x).M|_{\ell_1} = M[f(x).M, N/f, x]|_\ell$ , thus *a fortiori*  $(f(x).M) \ N|_{1 \# \ell_1} \sqsubseteq M[f(x).M, N/f, x]|_\ell$ .
- $M|_\ell = x, \ell_1$ ,  $M[f(x).M, N/f, x]|_\ell = N|_{\ell_1}$ , and  $\ell_1 \in \text{locs}(N)$ : we easily prove that  $\beta \rho(f, x.M) \ell \xrightarrow{*} 2 \# \ell_1$  and  $(f(x).M) \ N|_{2 \# \ell_1} = N|_{\ell_1} = M[f(x).M, N/f, x]|_\ell$ , thus *a fortiori*  $(f(x).M) \ N|_{2 \# \ell_1} \sqsubseteq M[f(x).M, N/f, x]|_\ell$ .
- $M[f(x).M, N/f, x]|_\ell = M'[f(x).M, N/f, x], \epsilon$  and  $M|_\ell = M', \epsilon$ : we easily prove that  $\beta \rho(f, x.M) \ell \xrightarrow{*} 1 \# 1 \# \ell$ ,  $(f(x).M) \ N|_{1.1.\ell} = M', \epsilon$ , and  $M[f(x).M, N/f, x]|_\ell = M'[f(x).M, N/f, x], \epsilon$ ; consequently  $M|_\ell \sqsubseteq M'[f(x).M, N/f, x]$ , which proves the thesis. □

**Lemma 4.** For all  $\ell \in \text{locs}(N[M/x])$ , there exists  $\ell' \in \text{locs}(\text{case } \sim M \text{ of } m)$  such that  $\beta_{\sim} \rho(x.N) \ell \xrightarrow{*} \ell'$  and, if  $\sim x \mapsto N \in m$ , then  $\text{case } \sim M \text{ of } m|_{\ell'} \sqsubseteq N[M/x]|_\ell$ .

For all  $\ell \in \text{locs}(N[M_1, M_2/x, y])$ , there exists  $\ell' \in \text{locs}(\text{case } M_1 \# M_2 \text{ of } m)$  such that  $\beta_{\#} \rho(x.N) \ell \xrightarrow{*} \ell'$  and, if  $x \# y \mapsto N \in m$ , then  $\text{case } M_1 \# M_2 \text{ of } m|_{\ell'} \sqsubseteq N[M_1, M_2/x, y]|_\ell$ .

*Proof.* Follows the same principle as Lemma 3. □

**Lemma 5.** Assume that for all  $\ell_1 \in \text{locs}(M')$  and for all  $J$  there exists  $\ell_2 \in \text{locs}(M)$  s.t.  $T \ J \ \ell_1 \xrightarrow{*} J \ \ell_2$  and  $M|_{\ell_2} \sqsubseteq M'|_{\ell_1}$ . Then for all  $\ell_1 \in \text{locs}(M' \ N)$  and for all  $J$  there exists  $\ell_2 \in \text{locs}(M \ N)$  s.t.  $\mathfrak{I}_{2.1} \ T \ J \ \ell_1 \rightsquigarrow J \ \ell_2$  and  $M \ N|_{\ell_2} \sqsubseteq M' \ N|_{\ell_1}$ .

**Lemma 6.** If  $h \Vdash L \triangleright R \ S$ , then  $h \Vdash L \circ \mathfrak{I}_{2.1}^{\text{Loc}} \triangleright R$ .

**Proof of Theorem 2** Suppose  $W \vdash M \xrightarrow{T} N$ : then for all labellings  $L$  and  $\ell \in \text{locs}(N)$ , if  $T L \ell \xrightarrow{*} \sim a$ , there exists  $\ell'$  such that  $L \ell' \rightsquigarrow \sim a$  and  $M|_{\ell'} \sqsubseteq N|_{\ell}$ .

*Proof.* By induction on the derivation of  $W \vdash M \xrightarrow{T} N$ . We consider only some of the rules, but the full proof can be obtained following the same approach.

- $W \vdash (f(x).R) S \xrightarrow{W_{\beta} (\beta \rho(f,x.R))} R[f(x).R, S/f, x]$   
we have

$$W_{\beta} (\beta \rho(f, x.R)) J \ell \xrightarrow{*}$$

case  $\ell$  of

$$\{\epsilon \mapsto w_{\beta} (J \epsilon) (J [1]) (J (\beta \rho(f, x.R) \epsilon))\}$$

$$; - \mapsto J (\beta \rho(f, x.R) \ell)\}$$

and by cases on  $\ell$  we prove that this reduces to  $J (\beta \rho(f, x.R) \ell)$ ; by Lemma 3 we obtain  $\ell'$  such that  $J (\beta \rho(f, x.R) \ell) \xrightarrow{*} J \ell'$  and  $(f(x).R) S|_{\ell'} \sqsubseteq R[f(x).R, S/f, x]|_{\ell'}$ , which proves the thesis.

- $W \vdash \oplus(\vec{V}_n) \xrightarrow{W_{\oplus} \vec{V}_n} \hat{\oplus}(\vec{V}_n)$   
 $\ell \in \text{locs}(\hat{\oplus}(\vec{V}_n))$  means  $\ell = \epsilon$ , thus we have  $W_{\oplus} \vec{V}_n J \epsilon \xrightarrow{*} w_{\oplus} \vec{V}_n (J [i])_{i=1\dots n} (J \epsilon)$ . Since this term only reduces to  $\bullet$ , we have nothing to prove.
- $W \vdash \text{fst}(\langle R, S \rangle) \xrightarrow{W_{\text{fst}}} R$   
by a simple case analysis on  $\ell$  we obtain  $W_{\text{fst}} J \ell \xrightarrow{*} J (1 \# 1 \# \ell)$ ; then we take  $\ell' = 1 \# 1 \# \ell$  and prove  $\text{fst}(\langle R, S \rangle)|_{1 \# 1 \# \ell} = R|_{\ell'}$ , thus the thesis holds.
- $W \vdash \text{case} (R_1 \# R_2) \text{ of } m \xrightarrow{W_{\#} (\beta_{\#} \rho(x,y.S))} S[R_1, R_2/x, y]$ , where  $(x, y \mapsto S) \in m$   
by cases on  $\ell$  we show  $W_{\#} (\beta_{\#} \rho(x, y.S)) J \ell \xrightarrow{*} J (\beta_{\#} \rho(x, y.S) \ell)$ ; we use Lemma 4 to obtain  $\ell'$  such that  $\beta_{\#} \rho(x, y.S) \ell \xrightarrow{*} \ell'$  and  $\text{case} (R_1 \# R_2) \text{ of } m|_{\ell'} \sqsubseteq S[R_1, R_2/x, y]|_{\ell'}$ , which proves the thesis.
- $W \vdash R S \xrightarrow{\mathcal{I}_{2.1} T} R' S$   
assume by IH that for all  $J$  and all  $\ell \in \text{locs}(R')$  there exists  $\ell' \in \text{locs}(R)$  s.t.  $T J \ell \xrightarrow{*} J \ell'$  and  $R|_{\ell'} \sqsubseteq R'|_{\ell}$ .  
By Lemma 5 we prove that for all  $J$  and all  $\ell \in \text{locs}(R' S)$  there exists  $\ell' \in \text{locs}(R S)$  s.t.  $\mathcal{I}_{2.1} T I \ell \xrightarrow{*} \ell'$ ,  $R S|_{\ell'} \sqsubseteq R' S|_{\ell}$ .

□

**Proof of Theorem 3** If  $h \Vdash L \triangleright M$  and  $E \vdash M \xrightarrow{T} N$ , then  $h \Vdash T L \triangleright N$ .

*Proof.* We proceed by induction on the derivation of  $E \vdash M \xrightarrow{T} N$ . We consider only some of the rules, but the full proof can be obtained following the same approach.

- $E \vdash (f(x).R) S \xrightarrow{E_{\beta} \beta(f(x).R,S)} R[f(x).R, S/f, x]$   
assume  $\ell \in \text{locs}(R[f(x).R, S/f, x])$ : we have

$$E_{\beta} (\beta \rho(f, x.R)) L \ell \xrightarrow{*}$$

case  $\ell$  of

$$\{\epsilon \mapsto e_{\beta} (L \epsilon) (L [1]) (L (\beta \rho(f, x.R) \epsilon))\}$$

$$; - \mapsto L (\beta \rho(f, x.R) \ell)\}$$

by cases on  $\ell$  we prove that this reduces to  $L (\beta \rho(f, x.R) \ell)$ ; by Lemma 3 we obtain  $\ell'$  such that  $L (\beta \rho(f, x.R) \ell) \rightsquigarrow L \ell'$  and  $(f(x).R) S|_{\ell'} \sqsubseteq R[f(x).R, S/f, x]|_{\ell'}$ ; since by hypothesis  $h$  is a consistent map for the function application, we obtain  $V$  such that  $\hat{h}(V) \sqsubseteq (f(x).R) S|_{\ell'}$ ; by transitivity of  $\sqsubseteq$ , we obtain  $\hat{h}(V) \sqsubseteq R[f(x).R, S/f, x]|_{\ell'}$ , which proves the thesis.

- $E \vdash \oplus(\overrightarrow{U}_n) \xrightarrow{E_{\oplus} \overrightarrow{U}_n} \hat{\oplus}(\overrightarrow{U}_n)$   
the only location in  $\hat{\oplus}(\overrightarrow{U}_n)$  is  $\epsilon$ , so we have to find  $V$  such that  $E_{\oplus} \overrightarrow{U}_n L \in \overset{*}{\hookrightarrow} V$  and  $\hat{h}(V) \sqsubseteq \hat{\oplus}(\overrightarrow{U}_n)$ . We show that  $E_{\oplus} \overrightarrow{U}_n L \in \rightsquigarrow k_{\oplus}(\overrightarrow{(L [i] * U_i)_{i=1 \dots n}})$  and choose this as  $V$ ; we also compute  $\hat{h}(V) = \hat{h}(k_{\oplus}(\overrightarrow{(L [i] * U_i)_{i=1 \dots n}})) = \oplus(\hat{h}(\overrightarrow{(L [i] * U_i)_{i=1 \dots n}}))$ ; for each  $i = 1 \dots n$  we prove  $\hat{h}((L [i] * U_i) \sqsubseteq U_i$ : in fact, the argument of the function is either a symbol  $k_{U_i}$  representing  $U_i$  (which must be a constant), and  $\hat{h}(k_{U_i}) = U_i$ , or some value  $U^*$  such that  $L [i] \rightsquigarrow U^*$ , and this is enough to prove  $\hat{h}(U^*) \sqsubseteq U_i$  by the hypothesis on  $h$ ; this implies  $\hat{h}(V) \sqsubseteq \oplus(\overrightarrow{U}_n) \rightsquigarrow \hat{\oplus}(\overrightarrow{U}_n)$ , which proves the thesis.
- $E \vdash \mathbf{fst}(\langle R, S \rangle) \xrightarrow{E_{\mathbf{fst}}} R$   
by a simple case analysis on  $\ell$  we obtain  $E_{\mathbf{fst}} L \ell \rightsquigarrow L (1 \# 1 \# \ell)$ , where  $1 \# 1 \# \ell \in \text{locs}(\mathbf{fst}(\langle R, S \rangle))$ ; since  $h \Vdash L \triangleright \mathbf{fst}(\langle R, S \rangle)$ , we obtain  $V$  such that  $\hat{h}(V) \sqsubseteq \mathbf{fst}(\langle R, S \rangle)|_{1 \# 1 \# \ell} = R|_{\ell}$ , which proves the thesis.
- $E \vdash R S \xrightarrow{\mathcal{J}_{2.1} T} R' S$   
assume by IH that, for all  $L'$ , if  $h \Vdash L' \triangleright R$ , then  $h \Vdash T L' \triangleright R'$ ; we assume  $\ell \in \text{locs}(R' S)$  and proceed by cases:
  - $\ell = \epsilon$ : then we have  $\mathcal{J}_{2.1} T L \in \overset{*}{\hookrightarrow} L \epsilon$ ; since  $\epsilon \in \text{locs}(R S)$  and  $h \Vdash L \triangleright R S$ , we obtain  $V$  such that  $\hat{h}(V) \sqsubseteq R S|_{\epsilon} = R S$ ; since  $R S \rightsquigarrow R' S$ , this also implies  $\hat{h}(V) \sqsubseteq R' S$ .
  - $\ell = 2 \# \ell_1$ , where  $\ell_1 \in \text{locs}(S)$ : then we have  $\mathcal{J}_{2.1} T L (2 \# \ell_1) \overset{*}{\hookrightarrow} L (2 \# \ell_1)$ : similarly to the previous case, we obtain  $V$  such that  $\hat{h}(V) \sqsubseteq R S|_{2 \# \ell_1} = S|_{\ell_1} = R' S|_{2 \# \ell_1}$ , which proves the thesis.
  - $\ell = 1 \# \ell_1$ , where  $\ell_1 \in \text{locs}(R')$ : then we have  $\mathcal{J}_{2.1} T L (1 \# \ell_1) \overset{*}{\hookrightarrow} T (L \circ \mathcal{J}_{2.1}^{\text{loc}} \ell_1)$ ; by Lemma 6 we prove  $h \Vdash L \circ \mathcal{J}_{2.1}^{\text{loc}} \triangleright R$ , and consequently, by IH, we obtain  $V$  such that  $\hat{h}(V) \sqsubseteq R'|_{\ell_1} = R' S|_{1 \# \ell_1}$ , which proves the thesis.

□

**Definition 3.** Given a relation  $\mathcal{R}$  on values, we define its extension to expressions  $\mathcal{R}^{\rightarrow}$  as follows:

$$M_1 \mathcal{R}^{\rightarrow} M_2 \iff (\forall V_1. M_1 \overset{*}{\hookrightarrow} V_1 \implies \exists V_2. M_2 \overset{*}{\hookrightarrow} V_2 \wedge V_1 \mathcal{R} V_2)$$

To prove the soundness theorem for dependency provenance, we introduce the following notations:

- quasi-equality  $L_M \triangleright M \approx_a L_N \triangleright N$  (Fig. 8)
- substitution  $\xrightarrow{\quad}$  for provenance-labelled expressions  $(L \triangleright M)\{\overrightarrow{L'} \triangleright \overrightarrow{N} / \overrightarrow{x}\}$  (Fig. 9)
- guarded reduction  $D \vdash (L \triangleright M) \xrightarrow{T}_a N$  (Fig. ??)
- non-guarded reduction  $D \vdash (L \triangleright M) \xrightarrow{T}_{\bar{a}} N$  (Fig. ??)

**Lemma 7.** Assume  $L_M \triangleright M \approx_a L_N \triangleright N$ . Let  $L'_M, L'_N$  such that for all  $\ell \in \text{locs}(M) \cup \text{locs}(N)$  we have  $L_M \ell \subseteq^{\rightarrow} L'_M \ell$  and  $L_N \ell \subseteq^{\rightarrow} L'_N \ell$ . Then  $L'_M \triangleright M \approx_a L'_N \triangleright N$

**Lemma 8.** Suppose  $L_M \triangleright M \approx_a L'_M \triangleright M'$  and  $\overrightarrow{L_N \triangleright N} \approx_a \overrightarrow{L'_N \triangleright N'}$  (pointwise). Then  $(L_M \triangleright M)\{\overrightarrow{L_N \triangleright N} / \overrightarrow{x}\} \approx_a (L'_M \triangleright M')\{\overrightarrow{L'_N \triangleright N'} / \overrightarrow{x}\}$ .

**Lemma 9.** Suppose:

$$\begin{aligned} \mathfrak{E}_{2.1}^{\text{Prov}} L \triangleright f(x).M_1 &\approx_a \mathfrak{E}_{2.1}^{\text{Prov}} L' \triangleright f(x).M'_1 \\ \mathfrak{E}_{1.1}^{\text{Prov}} (\mathfrak{E}_{2.1}^{\text{Prov}} L) \triangleright M_1 &\approx_a \mathfrak{E}_{1.1}^{\text{Prov}} (\mathfrak{E}_{2.1}^{\text{Prov}} L') \triangleright M'_1 \\ \mathfrak{E}_{2.2}^{\text{Prov}} L \triangleright M_2 &\approx_a \mathfrak{E}_{2.2}^{\text{Prov}} L' \triangleright M'_2 \end{aligned}$$

Then:

$$\begin{aligned} D_{\beta} \beta(f(x).M_1, M_2) L \triangleright M_1[f(x).M_1, M_2/f, x] \\ \approx_a D_{\beta} \beta(f(x).M'_1, M'_2) L' \triangleright M'_1[f(x).M'_1, M'_2/f, x] \end{aligned}$$

$$\begin{array}{c}
\frac{a \in \rightarrow L \in \cup L' \in}{L \triangleright M \approx_a L' \triangleright M'} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{L \triangleright c \approx_a L' \triangleright c} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{L \triangleright x \approx_a L' \triangleright x} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{L \triangleright \bullet \approx_a L' \triangleright \bullet} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{L \triangleright \epsilon \approx_a L' \triangleright \epsilon} \\
\\
\frac{a \notin \rightarrow L \in \cup L' \in}{(\mathfrak{E}_{n.i} L \triangleright M_i \approx_a \mathfrak{E}^{n.i} L' \triangleright M'_i)_{i=1 \dots n}} \frac{}{L \triangleright \oplus(\overrightarrow{M_n}) \approx_a L' \triangleright \oplus(\overrightarrow{M'_n})} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{\mathfrak{E}_{1.1} L \triangleright M \approx_a \mathfrak{E}^{1.1} L' \triangleright M'} \frac{}{L \triangleright f(x).M \approx_a L' \triangleright f(x).M'} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{\mathfrak{E}_{2.1} L \triangleright M_1 \approx_a \mathfrak{E}^{2.1} L' \triangleright M'_1} \frac{\mathfrak{E}_{2.2} L \triangleright M_2 \approx_a \mathfrak{E}^{2.2} L' \triangleright M'_2}{L \triangleright M_1 M_2 \approx_a L' \triangleright M'_1 M'_2} \\
\\
\frac{a \notin \rightarrow L \in \cup L' \in}{\mathfrak{E}_{1.1} L \triangleright M \approx_a \mathfrak{E}^{1.1} L' \triangleright M'} \frac{}{L \triangleright \text{fst}(M) \approx_a L' \triangleright \text{fst}(M')} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{\mathfrak{E}_{1.1} L \triangleright M \approx_a \mathfrak{E}^{1.1} L' \triangleright M'} \frac{}{L \triangleright \text{snd}(M) \approx_a L' \triangleright \text{snd}(M')} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{\mathfrak{E}_{2.1} L \triangleright M_1 \approx_a \mathfrak{E}^{2.1} L' \triangleright M'_1} \frac{\mathfrak{E}_{2.2} L \triangleright M_2 \approx_a \mathfrak{E}^{2.2} L' \triangleright M'_2}{L \triangleright \langle M_1, M_2 \rangle \approx_a L' \triangleright \langle M'_1, M'_2 \rangle} \\
\\
\frac{a \notin \rightarrow L \in \cup L' \in}{\mathfrak{E}_{1.1} L \triangleright M \approx_a \mathfrak{E}^{1.1} L' \triangleright M'} \frac{}{L \triangleright \sim M \approx_a L' \triangleright \sim M'} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{\mathfrak{E}_{2.1} L \triangleright M_1 \approx_a \mathfrak{E}^{2.1} L' \triangleright M'_1} \frac{\mathfrak{E}_{2.2} L \triangleright M_2 \approx_a \mathfrak{E}^{2.2} L' \triangleright M'_2}{L \triangleright M_1 \# M_2 \approx_a L' \triangleright M'_1 \# M'_2} \quad \frac{a \notin \rightarrow L \in \cup L' \in}{\mathfrak{E}_{3.1} L \triangleright M_1 \approx_a \mathfrak{E}^{3.1} L' \triangleright M'_1} \frac{\mathfrak{E}_{3.2} L \triangleright M_2 \approx_a \mathfrak{E}^{3.2} L' \triangleright M'_2}{\mathfrak{E}_{3.3} L \triangleright M_3 \approx_a \mathfrak{E}^{3.3} L' \triangleright M'_3} \frac{}{L \triangleright \text{case } M_1 \text{ of } \{- \mapsto M_2; - \mapsto M_3\} \approx_a L' \triangleright \text{case } M'_1 \text{ of } \{- \mapsto M'_2; - \mapsto M'_3\}} \\
\\
\frac{a \notin \rightarrow L \in \cup L' \in}{L \triangleright \rho(\overrightarrow{x}.M) \approx_a L' \triangleright \rho(\overrightarrow{x}.M)}
\end{array}$$

Figure 8: Quasi-equality

$$\begin{array}{c}
\frac{x_i \in \overrightarrow{x_n}}{(L \triangleright x_i)\{L'_n \triangleright N_n / \overrightarrow{x_n}\} = L'_i \triangleright N_i} \quad \frac{y \notin \overrightarrow{x_n}}{(L \triangleright y)\{L'_n \triangleright N_n / \overrightarrow{x_n}\} = L \triangleright y} \\
\\
\frac{\{f, x\} \cap \overrightarrow{x_n} = \emptyset}{(\mathfrak{E}_{1.1}^{\text{Prov}} L \triangleright M)\{L'_n \triangleright N_n / \overrightarrow{x_n}\} = L'' \triangleright M''} \quad \frac{(\mathfrak{E}_{2.1}^{\text{Prov}} L \triangleright M_1)\{L'_n \triangleright N_n / \overrightarrow{x_n}\} = L'_1 \triangleright M'_1}{(\mathfrak{E}_{2.2}^{\text{Prov}} L \triangleright M_2)\{L'_n \triangleright N_n / \overrightarrow{x_n}\} = L'_2 \triangleright M'_2} \\
\frac{}{(L \triangleright f(x).M)\{L'_n \triangleright N_n / \overrightarrow{x_n}\} = \mathfrak{J}_1^{\text{Prov}} L'' (L \in) \triangleright f(x).M''} \quad \frac{}{(L \triangleright M_1 M_2)\{L'_n \triangleright N_n / \overrightarrow{x_n}\} = \mathfrak{J}_2^{\text{Prov}} L'_1 L'_2 (L \in) \triangleright M'_1 M'_2}
\end{array}$$

Figure 9: Substitution for provenance-labelled expressions (selected cases)

$$\begin{array}{c}
\frac{a \in \rightarrow L \in \quad D \vdash M \xrightarrow{T} N}{D \vdash L \triangleright M \xrightarrow{T}_a N} \\
\\
\frac{D \vdash \mathfrak{E}_{n.i}^{\text{Prov}} L \triangleright M_i \xrightarrow{T}_a M'_i}{D \vdash L \triangleright \oplus(\overrightarrow{M_n}) \xrightarrow{\mathfrak{J}_{n.i}^T}_a \oplus(M_1, \dots, M'_i, \dots, M_n)} \quad \frac{D \vdash \mathfrak{E}_{2.1}^{\text{Prov}} L \triangleright M \xrightarrow{T}_a M'}{D \vdash L \triangleright M N \xrightarrow{\mathfrak{J}_{2.1}^T}_a M' N} \quad \frac{D \vdash \mathfrak{E}_{2.2}^{\text{Prov}} L \triangleright N \xrightarrow{T}_a N'}{D \vdash L \triangleright M N \xrightarrow{\mathfrak{J}_{2.2}^T}_a M N'} \\
\\
\frac{D \vdash \mathfrak{E}_{2.1}^{\text{Prov}} L \triangleright M \xrightarrow{T}_a M'}{D \vdash L \triangleright \langle M, N \rangle \xrightarrow{\mathfrak{J}_{2.1}^T}_a \langle M', N \rangle} \quad \frac{D \vdash \mathfrak{E}_{2.2}^{\text{Prov}} L \triangleright N \xrightarrow{T}_a N'}{D \vdash L \triangleright \langle M, N \rangle \xrightarrow{\mathfrak{J}_{2.2}^T}_a \langle M, N' \rangle} \\
\\
\frac{D \vdash \mathfrak{E}_{1.1}^{\text{Prov}} L \triangleright M \xrightarrow{T}_a M'}{D \vdash L \triangleright \text{fst}(M) \xrightarrow{\mathfrak{J}_{1.1}^T}_a \text{fst}(M')} \quad \frac{D \vdash \mathfrak{E}_{1.1}^{\text{Prov}} L \triangleright M \xrightarrow{T}_a M'}{D \vdash L \triangleright \text{snd}(M) \xrightarrow{\mathfrak{J}_{1.1}^T}_a \text{snd}(M')} \\
\\
\frac{D \vdash \mathfrak{E}_{3.1}^{\text{Prov}} L \triangleright M \xrightarrow{T}_a M'}{D \vdash L \triangleright \text{case } M \text{ of } m \xrightarrow{\mathfrak{J}_{3.1}^T}_a \text{case } M' \text{ of } m}
\end{array}$$

Figure 10: Guarded reduction

$$\begin{array}{c}
\frac{a \not\rightarrow L \in \cup L [1]}{D \vdash L \triangleright (f(x).M) N \xrightarrow{D_\beta (\beta \rho(f,x.M))} \bar{a} M[f(x).M, N/f, x]} \quad \frac{a \not\rightarrow L \in}{F \vdash \oplus(\vec{V}_n) \xrightarrow{D_\oplus \vec{V}_n} \hat{\oplus}(\vec{V}_n)} \\
\frac{a \not\rightarrow L \in \cup L [1]}{D \vdash L \triangleright \mathbf{fst}(\langle M, N \rangle) \xrightarrow{D_{\mathbf{fst}}} \bar{a} M} \quad \frac{a \not\rightarrow L \in \cup L [1]}{D \vdash L \triangleright \mathbf{snd}(\langle M, N \rangle) \xrightarrow{D_{\mathbf{snd}}} \bar{a} N} \\
\frac{(\mathbf{tt} \mapsto N) \in m \quad a \not\rightarrow L \in \cup L [1]}{D \vdash L \triangleright \mathbf{case \ ff \ of \ } m \xrightarrow{D_{\mathbf{tt}}} \bar{a} N} \quad \frac{(\mathbf{ff} \mapsto N) \in m \quad a \not\rightarrow L \in \cup L [1]}{D \vdash L \triangleright \mathbf{case \ ff \ of \ } m \xrightarrow{D_{\mathbf{ff}}} \bar{a} N} \\
\frac{(\bullet \mapsto N) \in m \quad a \not\rightarrow L \in \cup L [1]}{D \vdash L \triangleright \mathbf{case \ } \bullet \mathbf{ \ of \ } m \xrightarrow{D_\bullet} \bar{a} N} \quad \frac{(\sim x \mapsto N) \in m \quad a \not\rightarrow L \in \cup L [1]}{D \vdash L \triangleright \mathbf{case \ } \sim M \mathbf{ \ of \ } m \xrightarrow{D_\sim (\beta_\sim \rho(x.N))} \bar{a} N[M/x]} \\
\frac{(\epsilon \mapsto N) \in m \quad a \not\rightarrow L \in \cup L [1]}{D \vdash L \triangleright \mathbf{case \ } \epsilon \mathbf{ \ of \ } m \xrightarrow{D_\epsilon} \bar{a} N} \quad \frac{(h \# t \mapsto N) \in m \quad a \not\rightarrow L \in \cup L [1]}{D \vdash L \triangleright \mathbf{case \ } M_1 \# M_2 \mathbf{ \ of \ } m \xrightarrow{D_\# (\beta_\# \rho(h,t.N))} \bar{a} N[M_1, M_2/h, t]} \\
\frac{M|_\ell = M', \epsilon \quad M' \notin \vec{x}_n \quad a \not\rightarrow L \in}{D \vdash L \triangleright \rho(\vec{x}_n.M) \ell \xrightarrow{\lambda.\perp} \bar{a} (n+1) \# \ell} \quad \frac{M|_\ell = x_i, \ell' \quad 1 \leq i \leq n \quad a \not\rightarrow L \in}{D \vdash L \triangleright \rho(\vec{x}_n.M) \ell \xrightarrow{\lambda.\perp} \bar{a} i \# \ell'} \\
\frac{\ell' \neq \epsilon \quad a \not\rightarrow L \in \quad M|_\ell = N, \ell' \quad \text{for all } i \text{ s.t. } 1 \leq i \leq n : N \neq x_i}{D \vdash L \triangleright \rho(\vec{x}_n.M) \ell \xrightarrow{\lambda.\perp} \bar{a} \epsilon} \\
\frac{D \vdash \mathfrak{E}_{n,i}^{\text{Prov}} L \triangleright M_i \xrightarrow{T} \bar{a} M'_i}{D \vdash L \triangleright \oplus(\vec{M}_n) \xrightarrow{\mathfrak{J}_{n,i} T} \bar{a} \oplus(M_1, \dots, M'_i, \dots, M_n)} \quad \frac{D \vdash \mathfrak{E}_{2,1}^{\text{Prov}} L \triangleright M \xrightarrow{T} \bar{a} M'}{D \vdash L \triangleright M N \xrightarrow{\mathfrak{J}_{2,1} T} \bar{a} M' N} \quad \frac{D \vdash \mathfrak{E}_{2,2}^{\text{Prov}} L \triangleright N \xrightarrow{T} \bar{a} N'}{D \vdash L \triangleright M N \xrightarrow{\mathfrak{J}_{2,2} T} \bar{a} M N'} \\
\frac{D \vdash \mathfrak{E}_{2,1}^{\text{Prov}} L \triangleright M \xrightarrow{T} \bar{a} M'}{D \vdash L \triangleright \langle M, N \rangle \xrightarrow{\mathfrak{J}_{2,1} T} \bar{a} \langle M', N \rangle} \quad \frac{D \vdash \mathfrak{E}_{2,2}^{\text{Prov}} L \triangleright N \xrightarrow{T} \bar{a} N'}{D \vdash L \triangleright \langle M, N \rangle \xrightarrow{\mathfrak{J}_{2,2} T} \bar{a} \langle M, N' \rangle} \\
\frac{D \vdash \mathfrak{E}_{1,1}^{\text{Prov}} L \triangleright M \xrightarrow{T} \bar{a} M'}{D \vdash L \triangleright \mathbf{fst}(M) \xrightarrow{\mathfrak{J}_{1,1} T} \bar{a} \mathbf{fst}(M')} \quad \frac{D \vdash \mathfrak{E}_{1,1}^{\text{Prov}} L \triangleright M \xrightarrow{T} \bar{a} M'}{D \vdash L \triangleright \mathbf{snd}(M) \xrightarrow{\mathfrak{J}_{1,1} T} \bar{a} \mathbf{snd}(M')} \\
\frac{D \vdash \mathfrak{E}_{3,1}^{\text{Prov}} L \triangleright M \xrightarrow{T} \bar{a} M'}{D \vdash L \triangleright \mathbf{case \ } M \mathbf{ \ of \ } m \xrightarrow{\mathfrak{J}_{3,1} T} \bar{a} \mathbf{case \ } M' \mathbf{ \ of \ } m}
\end{array}$$

Figure 11: Non-guarded reduction

*Proof.* Let

$$\begin{aligned} L_N \triangleright N &= (\mathfrak{E}_{1.1}^{\text{Prov}} (\mathfrak{E}_{2.1}^{\text{Prov}} L) \triangleright M_1) \\ &\quad \{\mathfrak{E}_{2.1}^{\text{Prov}} L \triangleright f(x).M_1, \mathfrak{E}_{2.2}^{\text{Prov}} L \triangleright M_2/f, x\} \\ L'_N \triangleright N' &= (\mathfrak{E}_{1.1}^{\text{Prov}} (\mathfrak{E}_{2.1}^{\text{Prov}} L') \triangleright M'_1) \\ &\quad \{\mathfrak{E}_{2.1}^{\text{Prov}} L' \triangleright f(x).M'_1, \mathfrak{E}_{2.2}^{\text{Prov}} L' \triangleright M'_2/f, x\} \end{aligned}$$

Let us use the following notation:

$$L \preceq_{\mathcal{S}} L' \iff (\forall \ell \in \mathcal{S} : L \ell \subseteq \rightarrow L' \ell)$$

We prove:

$$\begin{aligned} L_N \preceq_{\text{locs}(N) \cup \text{locs}(N')} D_{\beta} \beta(f(x).M_1, M_2) L \\ L'_N \preceq_{\text{locs}(N) \cup \text{locs}(N')} D_{\beta} \beta(f(x).M'_1, M'_2) L' \end{aligned}$$

So, by Lemma 7, to prove the lemma we only need to show that

$$L_N \triangleright N \approx_a L'_N \triangleright N'$$

which is a consequence of Lemma 8. □

#### Proof of Theorem 4

Suppose  $L \triangleright M \approx_a L' \triangleright M'$ . Then:

1. If  $D \vdash (L \triangleright M) \xrightarrow{T}_a N$ , we have  $T L \triangleright N \approx_a L' \triangleright M'$ .
2. If  $D \vdash (L \triangleright M) \xrightarrow{T}_{\bar{a}} N$ , there exist  $T', N'$  such that  $D \vdash M' \xrightarrow{T'} N'$  and  $T L \triangleright N \approx_a T' L' \triangleright N'$ .

*Proof.* We proceed by induction on the derivation of  $L \triangleright M \approx_a L' \triangleright M'$ . Selected cases:

- $a \in \rightarrow (L \epsilon) \cup (L' \epsilon)$ 
  1. Assume  $D \vdash (L \triangleright M) \xrightarrow{T}_a N$ ; we prove that  $a \in \rightarrow T L \epsilon$  (under the dependency view, annotations are monotonically increasing), thus  $a \in \rightarrow (T L \epsilon) \cup (L' \epsilon)$  and thus  $T L \triangleright N \approx_a L' \triangleright M'$ .
  2. Assume  $D \vdash (L \triangleright M) \xrightarrow{T}_{\bar{a}} N$ ; by case analysis on the last rule of this derivation we show that  $a \notin \rightarrow L \epsilon$ , reaching a contradiction
- $a \notin \rightarrow (L \epsilon) \cup (L' \epsilon)$ ,  $M = M_1 M_2$   
 $M' = M'_1 M'_2$   
 $\mathfrak{E}_{2.i}^{\text{Prov}} L \triangleright M_i \approx_{\mathfrak{E}_{2.i}^{\text{Prov}}} L' \triangleright M'_i$  for  $i = 1, 2$ 
  1. Assume  $D \vdash (L \triangleright M_1 M_2) \xrightarrow{T}_a N$ ; since  $a \notin \rightarrow L \epsilon$ , by a case analysis on the derivation of reduction we see that the reduction is happening entirely within  $M_i$  (for  $i = 1, 2$ ); if  $i = 1$  (the case for  $i = 2$  is similar, for some  $T_1, N_1$  we have  $T = \mathcal{J}_{2.1} T_1$ ,  $N = N_1 M_2$ , and

$$D \vdash (\mathfrak{E}_{2.1}^{\text{Prov}} L \triangleright M_1) \xrightarrow{T_1}_{\bar{a}} N_1$$

By IH we have  $T_1 (\mathfrak{E}_{2.1}^{\text{Prov}} L) \triangleright N_1 \approx_a \mathfrak{E}_{2.1}^{\text{Prov}} L' \triangleright M'_1$ , thus also  $T L \triangleright N_1 M_2 \approx_a L' \triangleright M'_1 M'_2$ .

2. Assume  $D \vdash (L \triangleright M_1 M_2) \xrightarrow{T}_{\bar{a}} N$ ; by case analysis on the last rule of this derivation we show that either the reduction happened entirely within  $M_1$  or  $M_2$  (in which case the proof is similar to the previous subcase), or  $M_1 = f(x).M_0$  and the reduction is a beta reduction

$$\begin{aligned} D \vdash L \triangleright (f(x).M_0) M_2 &\xrightarrow{D_{\beta} (\beta \rho(f(x).M_0))}_{\bar{a}} \\ &M_0[f(x).M_0, M_2/f, x] \end{aligned}$$

Additionally,  $a \notin \rightarrow L [1]$ . By case analysis on  $\mathfrak{E}_{2.1}^{\text{Prov}} L \triangleright f(x).M_0 \approx_{\mathfrak{E}_{2.1}^{\text{Prov}}} L' \triangleright M'_1$  we prove  $M'_1 = f(x).M'_0$ , with  $\mathfrak{E}_{1.1}^{\text{Prov}} (\mathfrak{E}_{2.1}^{\text{Prov}} L) \triangleright M_0 \approx_a \mathfrak{E}_{1.1}^{\text{Prov}} (\mathfrak{E}_{2.1}^{\text{Prov}} L') \triangleright M'_0$ . Then we prove:

$$\begin{aligned} D \vdash (f(x).M'_0) M'_2 &\xrightarrow{D_{\beta} (\beta \rho(f(x).M'_0))}_{\bar{a}} \\ &M'_0[f(x).M'_0, M'_2/f, x] \end{aligned}$$

Finally, by Lemma 9, we have:

$$\begin{aligned} D_\beta (\beta \rho(f, x.M_0)) L \triangleright M_0[f(x).M_0, M_2/f, x] \approx_a \\ D_\beta (\beta \rho(f, x.M'_0)) L' \triangleright M'_0[f(x).M'_0, M'_2/f, x] \end{aligned}$$

which proves the thesis

□