

Modular Multiset Rewriting in Focused Linear Logic

Iliano Cervesato and Edmund S. L. Lam

July 2015
CMU-CS-15-117
CMU-CS-QTR-128

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Carnegie Mellon University, Qatar campus.
The author can be reached at iliano@cmu.edu or sllam@qatar.cmu.edu.

Abstract

Traditional multiset rewriting arises from a fragment of polarized linear logic, deriving its semantics from focused proof search. In recent years, languages based on multiset rewriting have been used for ever more ambitious applications. As program size grows however, so does the overhead of team-based development and the need for reusing components. Furthermore, there is a point where keeping a large flat collection of rules becomes impractical. In this report, we propose a module system for a small logically-motivated rule-based language that subsumes traditional multiset rewriting. The resulting modules are nothing more than rewrite rules of a specific form, and therefore are themselves just formulas in logic. Yet, they provide some of the same features found in advanced module systems such as that of Standard ML, in particular name space separation, support for abstract data types, parametrization by values and by other modules (functors in ML). Additionally, our modules offer features that are essential for concurrent programming, for example facilities of sharing private names. This work establishes a foundation for modularity in rule-based languages and is directly applicable to many languages based on multiset rewriting, most forward-chaining logic programming languages and many process algebras.

* This report was made possible by grant NPRP 09-667-1-100, *Effective Programming for Large Distributed Ensembles*, from the Qatar National Research Fund (a member of the Qatar Foundation). The statements made herein are solely the responsibility of the authors.

Keywords: Multiset Rewriting, Logic Programming, Modularity, Focused Search.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | Core Language | 3 |
| 2.1 | Multiset Rewriting with Existentials and Nested Rules | 3 |
| 2.1.1 | Syntax | 3 |
| 2.1.2 | Typing | 4 |
| 2.1.3 | Operational Semantics | 4 |
| 2.1.4 | Relationship to Other Languages | 8 |
| 2.2 | Logical Foundations | 9 |
| 2.3 | Mild Higher-Order Quantification | 10 |
| 3 | Adding Modularity | 13 |
| 3.1 | Name Space Separation | 13 |
| 3.2 | Modes | 15 |
| 3.3 | Abstract Data Types | 16 |
| 3.4 | Parametric Modules – I | 17 |
| 3.5 | Sharing Private Names | 19 |
| 3.6 | Parametric Modules – II | 20 |
| 4 | Multiset Rewriting with Modules | 23 |
| 5 | Related Work | 24 |
| 6 | Future Work and Conclusions | 25 |
| | References | 26 |
| A | Language Summary | 29 |
| A.1 | Syntax | 29 |
| A.2 | Typing | 31 |
| A.3 | Typing with Safety Checks | 33 |
| A.4 | First-Order Encoding | 36 |
| A.5 | Typechecking Modular Programs | 40 |
| A.6 | Congruence | 46 |
| A.7 | Unfocused Rewriting Semantics | 47 |
| A.8 | Focused Rewriting Semantics | 48 |
| B | Logical Interpretation | 50 |
| B.1 | Multiplicative-exponential Intuitionistic Linear Logic | 50 |
| B.2 | Focused MEILL | 52 |
| B.3 | Interpretation | 54 |
| B.3.1 | Translation | 54 |
| B.3.2 | Unfocused Transitions | 55 |
| B.3.3 | Focused Transitions | 56 |

List of Figures

| | | |
|-----|--|----|
| 2.1 | Typing Rules for \mathcal{L}^1 | 5 |
| A.1 | Grammar of $\mathcal{L}^{1.5}$ | 30 |
| A.2 | Syntax of \mathcal{L}^M | 45 |
| B.1 | Multiplicative Exponential Intuitionistic Linear Logic — MEILL | 51 |
| B.2 | Focused Presentation of MEILL | 53 |

1 Introduction

Rule-based programming, a model of computation by which rules modify a global state by concurrently rewriting disjoint portions of it, is having a renaissance as a number of domains are finding a use for its declarative and concise specifications, natural support for concurrency, and relative ease of reasoning. Indeed, special-purpose languages for security [23], networking [18, 24], robot coordination [2], multicore programming [11], graph programming [12] and mobile applications [21] have all recently been proposed as extensions of Datalog [16, 5], itself a rule-based database language founded on forward-chaining proof search. A few general-purpose languages based on this paradigm have also been proposed, for example the multiset rewriting language GAMMA [3] and more recently CHR [14]. The asynchronous state transformation model embodied by this paradigm has also been shown to subsume various models of concurrency [9], in particular multiset rewriting, Petri nets [30], process algebra [34], as well as the rare hybrid models [36, 8, 13].

As languages gain popularity, the need for modularity increases, since the overhead associated with writing code grows with program size. Yet, none of the languages just mentioned features a module system, even though some of them, for example CHR [14] and CoMingle [21], are starting to be used to write sizable applications. Modularity tames complexity. In traditional programming languages, it addresses the challenges of breaking a large program into a hierarchy of components with clear interfaces, implementation independence, team-based development, dependency management, code reuse, and separate compilation. These features alone justify extending popular rule-based languages with a module system.

Programming-in-the-large in a rule-based programming languages brings about additional challenges not typically found in imperative or functional languages. First, languages such as Datalog [16, 5], CHR [14] and GAMMA [3] have (like Prolog) a flat name space which gives no protections against accidentally reusing a name. Moreover, each rule in such a language adds to the definition of the names it contains rather than overriding them (as C does, for example). Second, these languages (like both C and Prolog) tend to have an open scope, meaning that there is no support for local definitions or private names. Finally, a rule in these languages can apply as soon as its prerequisites enter the global state, as opposed to when a procedure is called in a conventional language. This, together with the pitfalls of concurrency, makes writing correct code of even a moderate size difficult [7]. These challenges make enriching rule-based languages with a powerful module system all the more urgent if we want them to be used for large applications.

In this report, we develop a module system for a small rule-based programming language. This language, which we call \mathcal{L}^1 , is a fragment of the formalism in [9] and subsumes many languages founded on multiset rewriting, forward-chaining proof search and process algebra. Moreover, \mathcal{L}^1 is also a syntactic fragment of intuitionistic linear logic in that state transitions map to derivable sequents. In fact, the transition rules for each operator of \mathcal{L}^1 correspond exactly to the left sequent rules of this logic, and furthermore the notion of a whole-rule rewriting step originates in a focused presentation of proof search for it [35]. We extend \mathcal{L}^1 with a mild form of second-order quantification [15] to make our presentation more succinct (without adding expressiveness). We call this extension $\mathcal{L}^{1.5}$.

We engineer a module system for $\mathcal{L}^{1.5}$ by observing that certain programming patterns capture characteristic features of modularity such as hiding implementation details, providing functionalities to client code through a strictly defined interface, parametricity and the controlled sharing of names. We package these patterns into a handful of constructs that we provide to the programmer as a syntactic extension of $\mathcal{L}^{1.5}$ we call \mathcal{L}^M . The module system of \mathcal{L}^M support many of the facilities for modular programming found in Standard ML [25, 29], still considered by many an aspirational gold standard, in particular fine-grained name management and module parametricity (functors). Furthermore, \mathcal{L}^M naturally supports idioms such as higher-order functors and recursive modules, which are not found in [29]. Yet, because the modular constructs of \mathcal{L}^M are just programming patterns in $\mathcal{L}^{1.5}$, programs in \mathcal{L}^M can be faithfully compiled into $\mathcal{L}^{1.5}$ and thus into \mathcal{L}^1 . Moreover, since \mathcal{L}^1 subsumes the model of computation of a variety of rule-based languages (including those founded on forward-chaining, multiset rewriting and process algebra), it provides a blueprint for enriching these languages with a powerful, yet lightweight and declarative, module system.

With a few exceptions such as [26], research on modularity for rule-based languages has largely targeted backward-chaining logic programming [4]. Popular open-source and commercial implementations of Prolog (e.g., SWI Prolog and SICStus Prolog) do provide facilities for modular programming although not in a declarative fashion. The present

work is inspired by several attempts at understanding modularity in richer backward-chaining languages. In particular [27, 28] defines a module system for λ Prolog on the basis of this language’s support for embedded implication, while [1] achieves a form of modularization via a mild form of second-order quantification [15].

The main contributions of this report are thus threefold:

1. We define a language, \mathcal{L}^1 , that resides in an operational sweet spot between the stricture of traditional rule-based languages and the freedom of the rewriting reading of intuitionistic linear logic proposed in [9].
2. We engineer a powerful module system on top of this core language with support for name space separation, parametricity, and controlled name sharing.
3. We show that this module infrastructure is little more than syntactic sugar over the core language \mathcal{L}^1 , and can therefore be compiled away.

It is worth repeating that, because a large number of rule-based languages, including forward-chaining logic programming languages such as Datalog and its offshoots, multiset rewriting, as well as many process algebras can be viewed as linguistic fragments of our core language [9], this work defines a general mechanism for augmenting any of them with a powerful yet declarative module system. In fact, this work provides a logical foundation of modularity in rule-based languages in general.

The remainder of this report is organized as follows: in Section 2 we define the core multiset rewriting language on which we will build our module system, we relate it to more traditional presentations of multiset rewriting, and we discuss its logical underpinning. In Section 3, we introduce our modular infrastructure through a series of examples, incrementally developing syntax for it. In Section 4, we collect the language just developed, discuss additional syntactic checks, and explain how to compile it to the core language. We explore related work in Section 5 and outlines future developments in Section 6. Appendix A summarizes the syntax and semantics of our language and Appendix B lays down the details of its logical underpinning.

2 Core Language

This section develops a small, logically-motivated, rule-based language that will be used throughout the report. This formalism, a conservative extension of traditional multiset rewriting, will act as the core language in which we write (non-modular) programs. It is also the language our modular infrastructure will compile into. Section 2.1 defines this language, Section 2.2 highlights its foundations in linear logic, and Section 2.3 enriches it with a mild extension that will simplify writing programs in Section 3. As our development is incremental, we collect and organize the definition of the overall language in Appendix A.

2.1 Multiset Rewriting with Existentials and Nested Rules

Our core formalism is a first-order multiset rewriting language extended with dynamic name generation and support for nested rewrite rules. As such, it is a fragment of the logically-derived language of ω -multisets studied in [9]. Because we are interested in writing actual programs, we consider a simply-typed variant. We will refer to the language discussed in this section as \mathcal{L}^1 .

We now introduce the syntax of \mathcal{L}^1 in Section 2.1.1, give typing rules for it in Section 2.1.2, present its operational semantics in Section 2.1.3, and situate it with respect to other formalisms in Section 2.1.4.

2.1.1 Syntax

The syntax of the language \mathcal{L}^1 is specified by the following grammar:

$$\begin{array}{ll}
 \text{Types } \tau & ::= \iota \mid o \mid \tau \rightarrow \tau \mid \tau \times \tau \mid \top \\
 \text{Terms } t & ::= x \mid f t \mid (t, t) \mid () \\
 \text{Atoms } A & ::= p t \\
 \text{LHS } l & ::= \cdot \mid A, l \\
 \text{Rules } R & ::= l \multimap P \mid \forall x : \iota. R \\
 \text{Programs } P & ::= \cdot \mid P, P \mid A \mid !A \mid R \mid !R \mid \exists x : \tau \rightarrow \iota. P
 \end{array}$$

Our formalism is layered into a language of terms, used for describing entities of interest, and a language of formulas, used to specify computations.

Terms, written t , are built from other terms by pairing and by applying function symbols, f . The starting point is either the unit term $()$ or a term variable, generically written x . In examples, we abbreviate $f ()$ to just f — they correspond to the constants of a first-order term language. We classify terms by means of simple types, denoted τ . We consider two base types, the type of terms themselves, denoted ι , and the type of atoms, denoted o . The type constructors, products and function types, match the term constructors, with the type of $()$ written \top . While this minimal typing infrastructure is sufficient for the formal development in this report, it can be considerably enriched with additional type and term constructors — in fact, many of the examples in Section 4 will make use of the type nat of the unary natural numbers and the associated constructors.

The language of formulas consists of programs, themselves built out of rules and left-hand sides and ultimately atoms. An *atom* A is a predicate symbol p applied to a term. As with function symbols, we abbreviate $p ()$ in examples as just p — this is a propositional atom.

Atoms are used to build *rules*, R , which are the universal closures of rewrite directives of the form $l \multimap P$. The left-hand side l is a multiset of atoms, where we write “.” for the empty multiset and “ A, l ” for the extension of l with atom A . Considering the operator “,” commutative and associative with unit “.” will simplify our presentation, although this is not strictly necessary from an operational perspective. The right-hand side P of a rewrite directive is a multiset of either atoms A , reusable atoms $!A$, single-use rules R or reusable rules $!R$. A right-hand side can also have the form $\exists x : \tau \rightarrow \iota. P$, which, when executed, will have the effect of creating a new function symbol for x of type $\tau \rightarrow \iota$ for use in P . As rules consist of a rewrite directive embedded within a layer of universal quantifiers, generically

$$\forall x_1 : \tau_1. \dots \forall x_n : \tau_n. (l \multimap P)$$

(with τ_i equal to ι for the time being), we will occasionally use the notation $\forall \vec{x} : \vec{\tau}. (l \multimap P)$ where \vec{x} stands for x_1, \dots, x_n and $\vec{\tau}$ for τ_1, \dots, τ_n .

A *program* is what we just referred to as a right-hand side. A program is therefore a collection of single-use and reusable atoms and rules, and of existentially quantified programs.

The quantifiers $\forall x : \tau. R$ and $\exists x : \tau. P$ are the only binders in \mathcal{L}^1 (although richer languages could have more). We adopt the standard definition of free and bound variables (the above operators bind the variable x with scope R and P , respectively). A term or atom is *ground* if it does not contain free variables. A rule or program is *closed* if it does not contain free variables. Bound variables can be renamed freely as long as they do not induce capture. Given a syntactic entity O possibly containing a free variable x , we write $[t/x]O$ for the capture-avoiding substitution of term t for every free occurrence of x in O . As we perform a substitution, we will always be careful that the variable x and the term t have the same type (which for the time being can only be ι). Given sequences \vec{x} and \vec{t} of variables and terms of the same length, we denote the simultaneous substitution of every term t_i in \vec{t} for the corresponding variable x_i in \vec{x} in O as $[\vec{t}/\vec{x}]O$. We write θ for a generic substitution \vec{t}/\vec{x} and $O\theta$ for its application.

From the above grammar, it is apparent that traditional first-order multiset rewriting is the fragment of \mathcal{L}^1 where the right-hand side of a rule is a multiset of atoms (thereby disallowing rules and existentials) and where a program is a multiset of reusable rules. Maybe more surprisingly, this formalism also subsumes many process algebras: for example the asynchronous π -calculus corresponds to the restriction where a left-hand side always consists of exactly one atom (corresponding to a receive action) and where “;” is interpreted as parallel composition. This correspondence is explored further in Section 2.1.4 and at much greater depth in [9].

2.1.2 Typing

Function and predicate symbols have types $\tau \rightarrow \iota$ and $\tau \rightarrow o$ respectively. The symbols in use during execution together with their type are collected into a *signature*, denoted Σ . A signature induces a notion of type validity for the various entities of our language, which is captured by the typing judgments and rules in this section.

As we traverse universal quantifier, the typing rules collect the corresponding variables and their types into a *context*, denoted Γ . (Existential variables are treated as fresh function symbols and recorded in the signature during type checking.) Signatures and contexts are defined as follows:

$$\begin{array}{l} \text{Signatures } \Sigma ::= \cdot \mid \Sigma, f : \tau \rightarrow \iota \mid \Sigma, p : \tau \rightarrow o \\ \text{Contexts } \Gamma ::= \cdot \mid \Gamma, x : \iota \end{array}$$

Notice that context variables have type ι and therefore stand for terms.

The typing semantics of \mathcal{L}^1 is specified by the following judgments:

| | | |
|-----------------|---|--|
| Terms | $\Gamma \vdash_{\Sigma} t : \tau$ | <i>Term t has type τ in Γ and Σ</i> |
| Atoms | $\Gamma \vdash_{\Sigma} A \text{ atom}$ | <i>A is a well-typed atom in Γ and Σ</i> |
| Left-hand sides | $\Gamma \vdash_{\Sigma} l \text{ lhs}$ | <i>l is a well-typed left-hand side in Γ and Σ</i> |
| Rules | $\Gamma \vdash_{\Sigma} R \text{ rule}$ | <i>R is a well-typed rule in Γ and Σ</i> |
| Programs | $\Gamma \vdash_{\Sigma} P \text{ prog}$ | <i>P is a well-typed program in Γ and Σ</i> |

The fairly conventional rules defining these judgments are given in Figure 2.1.

Valid rewrite directives $l \multimap P$ are subject to the requirement that the free variables in P shall occur in l or in the left-hand side of an enclosing rule. While this *safety requirement* is not enforced by the rules in Figure 2.1, Appendix A.3 defines an extension that does enforce safety.

2.1.3 Operational Semantics

Computation in \mathcal{L}^1 takes the form of state transitions. A *state* is a triple $\Sigma. \langle \Omega ; \Pi \rangle$ consisting of a collection Ω of ground atoms and closed rules we call the *archive*, of a closed program Π , and a signature Σ that accounts for all the function and predicate symbols in Ω and Π . We emphasize that the program must be closed by writing it as Π rather

| | | | | |
|-----------|--|--|--|--|
| Terms: | $\frac{}{\Gamma, x : \tau \vdash_{\Sigma} x : \tau}$ | $\frac{\Gamma \vdash_{\Sigma, f: \tau \rightarrow \iota} t : \tau}{\Gamma \vdash_{\Sigma, f: \tau \rightarrow \iota} f t : \iota}$ | $\frac{\Gamma \vdash_{\Sigma} t_1 : \tau_1 \quad \Gamma \vdash_{\Sigma} t_2 : \tau_2}{\Gamma \vdash_{\Sigma} (t_1, t_2) : \tau_1 \times \tau_2}$ | $\frac{}{\Gamma \vdash_{\Sigma} () : \top}$ |
| Atoms: | $\frac{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} t : \tau}{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} p t \text{ atom}}$ | | | |
| LHS: | $\frac{}{\Gamma \vdash_{\Sigma} \cdot \text{lhs}}$ | $\frac{\Gamma \vdash_{\Sigma} A \text{ atom} \quad \Gamma \vdash_{\Sigma} l \text{ lhs}}{\Gamma \vdash_{\Sigma} A, l \text{ lhs}}$ | | |
| Rules: | $\frac{\Gamma \vdash_{\Sigma} l \text{ lhs} \quad \Gamma \vdash_{\Sigma} P \text{ prog}}{\Gamma \vdash_{\Sigma} l \multimap P \text{ rule}}$ | $\frac{\Gamma, x : \iota \vdash_{\Sigma} R \text{ rule}}{\Gamma \vdash_{\Sigma} \forall x : \iota. R \text{ rule}}$ | | |
| Programs: | $\frac{}{\Gamma \vdash_{\Sigma} \cdot \text{prog}}$ | $\frac{\Gamma \vdash_{\Sigma} P_1 \text{ prog} \quad \Gamma \vdash_{\Sigma} P_2 \text{ prog}}{\Gamma \vdash_{\Sigma} P_1, P_2 \text{ prog}}$ | | |
| | $\frac{\Gamma \vdash_{\Sigma} A \text{ atom}}{\Gamma \vdash_{\Sigma} A \text{ prog}}$ | $\frac{\Gamma \vdash_{\Sigma} A \text{ atom}}{\Gamma \vdash_{\Sigma} !A \text{ prog}}$ | $\frac{\Gamma \vdash_{\Sigma} R \text{ rule}}{\Gamma \vdash_{\Sigma} R \text{ prog}}$ | $\frac{\Gamma \vdash_{\Sigma} R \text{ rule}}{\Gamma \vdash_{\Sigma} !R \text{ prog}}$ |
| | $\frac{\Gamma \vdash_{\Sigma, x: \tau \rightarrow \iota} P \text{ prog}}{\Gamma \vdash_{\Sigma} \exists x : \tau \rightarrow \iota. P \text{ prog}}$ | | | |

Figure 2.1: Typing Rules for \mathcal{L}^1

than P . Archives will be used to store reusable components of the program Π as the computation unfolds in Π itself. We define archives next, and copy the definition of (closed) programs using this notation for emphasis:

$$\begin{aligned}
\text{Archives } \Omega &::= \cdot \mid \Omega, A \mid \Omega, R \\
\text{State programs } \Pi &::= \cdot \mid \Pi, \Pi \mid A \mid !A \mid R \mid !R \mid \exists x : \tau \rightarrow \iota. P \\
\text{State } \Sigma &::= \langle \Omega ; \Pi \rangle
\end{aligned}$$

We will occasionally refer to Π itself as a state when the other components are of secondary concern. When talking about states, it will be important that “;” be commutative and associative with unit “.”. This induces a notion of equivalence among states — a congruence indeed when applied to embedded programs. This congruence is examined in Appendix A.6. Note that all atoms A appearing in both the archive and the program component of a state must be ground. Furthermore, since rules R are closed and have the form $\forall \vec{x} : \vec{\tau}. (l \multimap P)$ we will occasionally use the notation $\forall(l \multimap P)$ with the understanding that the universal quantification is over all the free variables in $l \multimap P$.

The above definition combines elements of the notion of “state” found in both multiset rewriting and process algebras. In multiset rewriting, a state is traditionally understood a multiset of ground atoms that carry data (in the form of terms) on which a separate notion of immutable rules perform computation through rewriting. In process algebra, the “state” is a collection of processes that interact with each other as the computation takes place. In \mathcal{L}^1 , we have both: the ground atoms in Ω and Π act as data while non-atomic state items in them give rise to computation, like rules in multiset rewriting or processes in process algebra.

We extend the typing infrastructure discussed in the last section with the judgments

$$\begin{array}{ll}
\text{Archives} & \vdash_{\Sigma} \Omega \text{ archive} & \text{Archive } \Omega \text{ is well typed in } \Sigma \\
\text{States} & \vdash \Sigma. \langle \Omega ; \Pi \rangle \text{ state} & \text{State } \Sigma. \langle \Omega ; \Pi \rangle \text{ is well typed}
\end{array}$$

to describe typing over archives and valid states. A state $\Sigma.\langle\Omega ; \Pi\rangle$ is valid if both the archive Ω and the program Π are well typed in the empty context and Σ , as captured by the following rules:

$$\frac{}{\vdash_{\Sigma} \cdot \text{archive}} \quad \frac{\vdash_{\Sigma} \Omega \text{ archive} \quad \cdot \vdash_{\Sigma} A \text{ atom}}{\vdash_{\Sigma} \Omega, A \text{ archive}} \quad \frac{\vdash_{\Sigma} \Omega \text{ archive} \quad \cdot \vdash_{\Sigma} R \text{ rule}}{\vdash_{\Sigma} \Omega, R \text{ archive}}$$

$$\frac{\vdash_{\Sigma} \Omega \text{ archive} \quad \cdot \vdash_{\Sigma} \Pi \text{ prog}}{\vdash_{\Sigma} \Sigma.\langle\Omega ; \Pi\rangle \text{ state}}$$

We will give two characterizations of the rewriting semantics of \mathcal{L}^1 . The first interprets each operator in \mathcal{L}^1 as an independent state transformation directive. This makes it very close to traditional presentations of linear logic (see Section 2.2) but the resulting semantics is overly non-deterministic. The second presentation is instead focused on applying rules fully, and does so by interpreting them as meta-transformations, with the effect of giving a strong operational flavor to the language and an effective strategy for implementing an interpreter for it. This second characterization accounts for the standard semantics found in many languages based on multiset rewriting, forward chaining and process transformation. It also has its roots in linear logic, but this time in its focused presentation [22]. These two characterizations are not equivalent as the latter leads to fewer reachable states. It is however preferable as these states match the intuition of what is traditionally meant by applying a rewrite rule, leaving out states littered with half-applied rules. We will examine how this relates to the corresponding presentations of linear logic in Section 2.2.

Unfocused Rewriting Semantics The unfocused evaluation semantics makes use of a step judgment of the form

$$\Sigma.\langle\Omega ; \Pi\rangle \mapsto \Sigma'.\langle\Omega' ; \Pi'\rangle \quad \text{State } \Sigma.\langle\Omega ; \Pi\rangle \text{ transitions to state } \Sigma'.\langle\Omega' ; \Pi'\rangle \text{ in one step}$$

In this semantics, each operator in the language is understood as a directive to carry out one step of computation in a given state. Therefore, each operator yields one transition rule, given in the following table:

$$\begin{aligned} \Sigma.\langle\Omega, l_1 ; \Pi, l_2, (l_1, l_2) \multimap P\rangle &\mapsto \Sigma.\langle\Omega, l_1 ; \Pi, P\rangle \\ \Sigma.\langle\Omega ; \Pi, \forall x : \iota. R\rangle &\mapsto \Sigma.\langle\Omega ; \Pi, [t/x]R\rangle && \text{if } \cdot \vdash_{\Sigma} t : \iota \\ \Sigma.\langle\Omega ; \Pi, !A\rangle &\mapsto \Sigma.\langle\Omega, A ; \Pi\rangle \\ \Sigma.\langle\Omega ; \Pi, !R\rangle &\mapsto \Sigma.\langle\Omega, R ; \Pi\rangle \\ \Sigma.\langle\Omega ; \Pi, \exists x : \tau \rightarrow \iota. P\rangle &\mapsto (\Sigma, x : \tau \rightarrow \iota).\langle\Omega ; \Pi, P\rangle \\ \Sigma.\langle\Omega, R ; \Pi\rangle &\mapsto \Sigma.\langle\Omega, R ; \Pi, R\rangle \end{aligned}$$

In words, \multimap is a rewrite directive which has the effect of identifying its left-hand side atoms in the surrounding state and replacing them with the program in its right-hand side. It retrieves reusable atoms from the archive and single-use atoms from the program component of the state. Notice that the atom in the program side are consumed but atoms in the archive side are retained and can therefore be used over and over. The operator \forall is an instantiation directive: it picks a term of the appropriate type and replaces the bound variable with it. Instead, $!$ is to all effects a replication directive: it installs the entity it prefixes in the archive, enabling repeated accesses to atoms by the rewrite directive and to rules in the last transition, which copies a rule to the program side while retaining the master copy in the archive. Finally, \exists is a name generation directive which installs a new symbol of the appropriate type in the state's signature.

Evaluation steps transform valid states into valid states, a property captured by the following lemma. In particular, it maintains the invariant that a state is closed.

Lemma 1 (Type preservation) *If $\vdash_{\Sigma} \Sigma.\langle\Omega ; \Pi\rangle$ state and $\Sigma.\langle\Omega ; \Pi\rangle \mapsto \Sigma'.\langle\Omega' ; \Pi'\rangle$, then $\vdash_{\Sigma'} \Sigma'.\langle\Omega' ; \Pi'\rangle$ state.*

Proof The proof proceeds by cases on the step rule applied and then by inversion on the typing derivation. It relies on state equivalence and the corresponding equivalence it induces over typing derivation. \square

We write $\Sigma.\langle\Omega ; \Pi\rangle \mapsto^* \Sigma'.\langle\Omega' ; \Pi'\rangle$ for the reflexive and transitive closure of the above judgment, meaning that state $\Sigma.\langle\Omega ; \Pi\rangle$ transitions to $\Sigma'.\langle\Omega' ; \Pi'\rangle$ in zero or more steps. Naturally, type preservation holds for this iterated form as well.

The rules of this semantics are pleasantly simple as they tease out the specific behavior of each individual language construct. They are also highly non-deterministic as any of them is applicable whenever the corresponding operator appears as the top-level constructor of a state element. These transitions stem from the logic interpretation of our operators, and can in fact be read off directly from the rules of linear logic [9] — see Appendix B.3. This further supports the orthogonality of our constructs.

Focused Rewriting Semantics By considering each operator in isolation when applying a step, the above semantics falls short of the expected behavior of a rewriting language. For example, consider the state

$$\underbrace{(p : \iota \rightarrow o, q : \iota \rightarrow o, a : \iota, b : \iota)}_{\Sigma} . \langle \cdot ; \underbrace{p a, \forall x : \iota. p x \multimap q x}_{\Pi} \rangle$$

From it, the one transition sequence of interest is

$$\Pi \mapsto p a, (p a \multimap q a) \mapsto q a$$

where we have omitted the signature Σ and the empty archive (which do not change) for clarity. However, nothing prevents picking the “wrong” instance of x and taking the step

$$\Pi \mapsto p a, (p b \multimap q b)$$

from where we cannot proceed further. This second possibility is unsatisfactory as it does not apply the rule fully.

The focused operational semantics makes sure that rules are either fully applied, or not applied at all. It corresponds to the standard operational semantics of most languages based on multiset rewriting, forward chaining and process transformation. It also leverages the observation that some of the state transformations associated with individual operators have the potential to block other currently available transformations, while others do not. In particular, turning a top-level existential variable into a new name or moving a reusable entity to the archive never preempt other available transitions.

A closed program without top-level existential quantifiers or reusable items is called *stable*. We write Π for a state program Π that is stable. A stable program is therefore defined as follows:

$$\begin{aligned} \text{Stable state program } \Pi &::= \cdot && \text{empty state} \\ &| \Pi, A && \text{atom} \\ &| \Pi, R && \text{rule} \end{aligned}$$

A *stable state* has the form $\Sigma. \langle \Omega ; \Pi \rangle$.

The focused operational semantics is expressed by the following two judgments:

$$\begin{aligned} \Sigma. \langle \Omega ; \Pi \rangle &\Rightarrow \Sigma'. \langle \Omega' ; \Pi' \rangle && \text{Non-stable state } \Sigma. \langle \Omega ; \Pi \rangle \text{ transitions to state } \Sigma'. \langle \Omega' ; \Pi' \rangle \text{ in one step} \\ \Sigma. \langle \Omega ; \Pi \rangle &\Rightarrow \Sigma'. \langle \Omega' ; \Pi' \rangle && \text{Stable state } \Sigma. \langle \Omega ; \Pi \rangle \text{ transitions to state } \Sigma'. \langle \Omega' ; \Pi' \rangle \text{ in one step} \end{aligned}$$

The first judgment is realized by selecting an existential or reusable program component in Π and processing it by eliminating the quantifier and creating a new symbol and by moving its immediate subformula to the archive, respectively. By iterating this step, we will reach a stable state in finite time. At this point, the second judgment is applied. It selects a single-use or a reusable rule from the program or archive part of the state and fully applies it. To fully apply a rule $\forall(l \multimap P)$, the surrounding state must contain an instance $l\theta$ of the left-hand side l . Observe that some left-hand side atoms are retrieved from the archive and others from the program component of the state. The resulting state replaces it with the corresponding instance $P\theta$ of the right-hand side P . This state may not be stable as $P\theta$ could contain existentially quantified and reusable components. The following transitions formalize this insight.

$$\begin{aligned} \Sigma. \langle \Omega ; \Pi, !A \rangle &\Rightarrow \Sigma. \langle \Omega, A ; \Pi \rangle \\ \Sigma. \langle \Omega ; \Pi, !R \rangle &\Rightarrow \Sigma. \langle \Omega, R ; \Pi \rangle \\ \Sigma. \langle \Omega ; \Pi, \exists x : \tau \rightarrow \iota. P \rangle &\Rightarrow (\Sigma, x : \tau \rightarrow \iota). \langle \Omega ; \Pi, P \rangle \\ \Sigma. \langle \Omega, l_1\theta ; \Pi, l_2\theta, \forall(l_1, l_2 \multimap P) \rangle &\Rightarrow \Sigma. \langle \Omega, l_1\theta ; \Pi, P\theta \rangle \\ \underbrace{\Sigma. \langle \Omega, l_1\theta, \forall(l_1, l_2 \multimap P) ; \Pi, l_2\theta \rangle}_{\Omega^*} &\Rightarrow \Sigma. \langle \Omega^* ; \Pi, P\theta \rangle \end{aligned}$$

As usual, we write \Rightarrow^* and \Rightarrow^* for the reflexive and transitive closure of these relations.

Type preservation holds for the base relations and their transitive closure. Furthermore, any transition step achievable in the focused semantics is also achievable in the unfocused semantics, although possibly in more than one step in the case of rules.

Lemma 2

1. If $\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle$, then $\Sigma.\langle\Omega ; \Pi\rangle \mapsto \Sigma'.\langle\Omega' ; \Pi'\rangle$.
2. If $\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle$, then $\Sigma.\langle\Omega ; \Pi\rangle \mapsto^* \Sigma'.\langle\Omega' ; \Pi'\rangle$.

Proof The first part simply uses the isomorphic rule of the unfocused semantics. The second proceeds by induction on the structure of the selected rule. □

The reverse does not hold, as we saw earlier.

Because rules R are applicable only in stable states, it is typical to follow up such an application with zero or more transitions of the first kind to reach the next stable state. We write \Rightarrow for this relation: $\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle$ is defined as $\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma.\langle\Omega_1 ; \Pi_1\rangle \Rightarrow^* \Sigma'.\langle\Omega' ; \Pi'\rangle$.

2.1.4 Relationship to Other Languages

The language \mathcal{L}^1 is a syntactic fragment of the formalism of ω -multisets examined in [9] as a logical reconstruction of multiset rewriting and some forms of process algebra. It instantiates ω -multisets by fixing a language of terms as discussed earlier, while [9] kept it open-ended (requiring only predicativity). The main restriction concerns the left-hand side of rules, which in \mathcal{L}^1 is a multiset of atoms, while in an ω -multiset can be any formula in the language. This restriction makes implementing rule application in \mathcal{L}^1 much easier than in the general language, is in line with all forward logic programming languages we are aware of, and is much more directly supported by a focusing view of proof search (see Section 2.2). Therefore, \mathcal{L}^1 occupies a sweet spot between the freewheeling generality of ω -multiset rewriting and the implementation simplicity of many rule-based languages. Differently from ω -multisets, which are freely generated from the available connectives, \mathcal{L}^1 limits the usage of the ! operator to just rules and right-hand side atoms. These restrictions avoid expressions that are of little use in programming practice (for example doubly reusable rules $!!R$ or left-hand side atoms of the form $!A$). We also left out the choice operator of ω -multisets (written $\&$) because we did not need it in any of the examples in this report. Adding it back is straightforward.

The language \mathcal{L}^1 subsumes various rule-based formalisms founded on multiset rewriting, forward proof search and process algebra as a syntactic fragment. First-order multiset rewriting, as found for example in CHR [14], relies on rules whose right-hand side is a multiset of atoms and that can be used arbitrarily many times, and therefore correspond to \mathcal{L}^1 rules of the form $!\forall\vec{x}. (l_1 \multimap l_2)$. Languages such as MSR [6] additionally permit the creation of new symbols in the right-hand side of a rule, which is supported by \mathcal{L}^1 rules of the form $!\forall\vec{x}. (l_1 \multimap \exists\vec{y}. l_2)$. Datalog clauses [16, 5] are written in \mathcal{L}^1 as $!\forall\vec{x}. (l \multimap !A)$ while their facts are reusable program atoms as this language monotonically grows the set of known facts during computation.

As shown in [9], ω -multiset rewriting, and therefore \mathcal{L}^1 , also subsumes many formalisms based on process algebra. Key to doing so is the possibility to nest rules (thereby directly supporting the ability to sequentialize actions), a facility to create new symbols (which matches channel restriction), and the fact that multiset union is commutative and associative with unit the empty multiset (thereby accounting for the analogous properties of parallel composition and the inert process). The asynchronous π -calculus [34] is the fragment of \mathcal{L}^1 where rule left-hand sides consist of exactly one atom (corresponding to a receive action — send actions correspond to atoms on the right-hand side of a rule). Its synchronous variant can be simulated in \mathcal{L}^1 through a shallow encoding [8].

Our language, like ω -multiset rewriting itself, contains fragments that correspond to both the state transition approach to specifying concurrent computations (as multiset rewriting and Petri nets for example) and specifications in the process-algebraic style. It in fact supports hybrid specifications as well, as found in the Join calculus [13] and in CLF [8, 36].

2.2 Logical Foundations

The language \mathcal{L}^1 , like ω -multisets [9], corresponds exactly to a fragment of intuitionistic linear logic [17]. In fact, not only can we recognize the constructs of our language among the operators of this logic, but \mathcal{L}^1 's rewriting semantics stem directly from its proof theory. In this section, we outline this correspondence, with a much more detailed discussion relegated to Appendix B.

The operators “,” and “.”, \multimap , $!$, \forall and \exists of \mathcal{L}^1 correspond to the logical constructs \otimes , $\mathbf{1}$, \multimap , $!$, \forall and \exists , respectively, of multiplicative-exponential intuitionistic linear logic (MEILL). Because the left-hand side of a rule in \mathcal{L}^1 consists solely of atoms and only rules and atoms are reusable, the actual MEILL formulas our language maps to are similarly restricted — the exact correspondence is found in Appendix B.3. By comparison, ω -multiset rewriting also includes additive conjunction, written $\&$, and allows freely combining these operators, which matches exactly the corresponding logic.

The transitions of the unfocused rewriting semantics of \mathcal{L}^1 can be read off directly from the left sequent rules of the above connectives, for an appropriate presentation of the exponential $!$. We write the derivability judgment of MEILL as $\Gamma; \Delta \longrightarrow_{\Sigma} \varphi$, where φ is a formula, the linear context Δ is a multiset of formulas that can be used exactly once in a proof of φ , while the formulas in the persistent context Γ can be used arbitrarily many times, and Σ is a signature defined as in \mathcal{L}^1 — the rules for this judgment are given in Appendix B.1. Consider for example the transition for an existential \mathcal{L}^1 program and the left sequent rule $\exists\text{L}$ for the existential quantifier:

$$\Sigma.\langle \Omega; \Pi, \exists x : \tau \rightarrow \iota. P \rangle \mapsto (\Sigma, x : \tau \rightarrow \iota).\langle \Omega; \Pi, P \rangle \rightsquigarrow \frac{\Gamma; \Delta, \varphi \longrightarrow_{\Sigma, x : \tau \rightarrow \iota} \psi}{\Gamma; \Delta, \exists x : \tau \rightarrow \iota. \varphi \longrightarrow_{\Sigma} \psi} \exists\text{L}$$

The antecedent $\Omega; \Pi, \exists x : \tau \rightarrow \iota. P$ of the transition corresponds to the contexts $\Gamma; \Delta, \exists x : \tau \rightarrow \iota. \varphi$ of the rule conclusion, while its consequent $\Omega; (\Pi, P)$ matches the contexts of the premise $\Gamma; (\Delta, \varphi)$ and the signatures have been updated in the same way. A similar correspondence applies in all cases, once we account for shared structural properties of states and sequents. In particular, multiplicative conjunction \otimes and its unit $\mathbf{1}$ are the materialization of the formation operators for the linear context of a sequent, context union and the empty context. This means that linear contexts can be interpreted as the multiplicative conjunction of their formulas. The interplay between persistent formulas $!\varphi$ and the persistent context Γ matches the rules for reusable entities in \mathcal{L}^1 [9]. Altogether, this correspondence is captured by the following property, proved in [9] and discussed further in Appendix B.3, where $[\Omega]$ and $[\Pi]$ are MEILL contexts corresponding to archive Ω and state program Π , and the formula $\exists \Sigma'. \ulcorner \Omega'; \Pi' \urcorner$ is obtained by taking the multiplicative conjunction of the encodings of the entities in Π' and in Ω' (the latter are embedded within a $!$) and then prefixing the result with an existential quantification for each declaration in the signature Σ' .

Theorem 3 (Soundness) *If $\vdash \Sigma.\langle \Omega; \Pi \rangle$ state and $\Sigma.\langle \Omega; \Pi \rangle \mapsto \Sigma'.\langle \Omega'; \Pi' \rangle$, then $[\Omega]; [\Pi] \longrightarrow_{\Sigma} \exists \Sigma'. \ulcorner \Omega'; \Pi' \urcorner$.*

This correspondence between the left rules of linear logic and rewriting directives was observed in [9] for the larger formalism of ω -multiset rewriting.

The transitions of the focused rewriting semantics of \mathcal{L}^1 originate from the focused presentation of linear logic [22], and specifically of MEILL. Focusing is a proof search strategy that alternates two phases: an inversion phase where invertible sequent rules are applied exhaustively and a chaining phase where it selects a formula (the focus) and decomposes it maximally using non-invertible rules. Focusing is complete for many logics of interest, in particular for traditional intuitionistic logic [35] and for linear logic [22], and specifically for MEILL — see Appendix B.2 for a detailed definition. Transitions of the form $\Sigma.\langle \Omega; \Pi \rangle \mapsto \Sigma'.\langle \Omega'; \Pi' \rangle$ correspond to invertible rules and are handled as in the unfocused case. Instead, transitions of the form $\Sigma.\langle \Omega; \Pi \rangle \mapsto \Sigma'.\langle \Omega'; \Pi' \rangle$ correspond to the selection of a focus formula and the consequent chaining phase. Consider for example the transition for a single-use rule $\forall(l_1, l_2 \multimap P)$. The derivation snippet below selects a context formula $\forall(\varphi_{l_1} \otimes \varphi_{l_2} \multimap \varphi_P)$ of the corresponding form from a sequent where no invertible rule is applicable (generically written $\Gamma; \Delta \Longrightarrow_{\Sigma} \psi$), puts it into focus (indicated as a red box), and then applies non-invertible rules to it exhaustively. In the transcription of \mathcal{L}^1 into MEILL, the formulas $\varphi_{l_i} \theta$ are conjunctions of atomic formulas matching $l_i \theta$, which allows continuing the chaining phase in the left premise of rule $\multimap\text{L}$ into a complete derivation when Γ_1 and Δ_2 consist exactly of the atoms in l_1 and l_2 , respectively. The translation φ_P of the program P re-enables an invertible rule and therefore the derivation loses focus in the rightmost open

The typing rules for each of the new syntactic forms are given next:

$$\begin{array}{c}
\cdots \\
\frac{}{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} X : \tau \rightarrow o} \qquad \frac{}{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} p : \tau \rightarrow o} \\
\cdots \\
\frac{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} t : \tau}{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} X t \text{ atom}} \\
\cdots \\
\frac{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} R \text{ rule}}{\Gamma \vdash_{\Sigma} \forall X : \tau \rightarrow o. R \text{ rule}} \\
\cdots \\
\frac{\Gamma \vdash_{\Sigma, X: \tau \rightarrow o} P \text{ prog}}{\Gamma \vdash_{\Sigma} \exists X : \tau \rightarrow o. P \text{ prog}}
\end{array}$$

Just like for the first-order case, traversing a universal quantifier adds the bound variable to the context while type-checking an existential quantifier adds it to the signature. As in \mathcal{L}^1 , we treat existential variables as if they were names.

The notions of free and bound variables, ground atoms, closed rules and programs, and substitution carry over from the first-order case, and so does the safety requirement that a free universal variable in a program position must appear in an earlier left-hand side. Second-order universal variables are subject to an additional requirement: if a parametric atom $X t$ with X universally quantified occurs in the left-hand side of a rule, then X must occur in a term position either in the same left-hand side or in an outer left-hand side. This additional constraint is motivated by the need to avoid rules such as

$$!\forall X : \tau \rightarrow o. \forall x : \tau. X x \multimap \cdot$$

which would blindly delete any atom $p t$ for all p of type $\tau \rightarrow o$. Instead, we want a rule to be applicable only to atoms it “knows” something about, by requiring that their predicate name has been passed to it in a “known” atom. The following rule is instead acceptable:

$$!\forall X : \tau \rightarrow o. \forall x : \tau. \text{delete_all}(X), X x \multimap \cdot$$

as only the predicate names marked to be deleted can trigger this rule. Appendix A.3 presents a variant of the type system of $\mathcal{L}^{1.5}$ that checks this safety constraint on second-order variables as well as our original safety constraint.

The rewriting semantics of $\mathcal{L}^{1.5}$ extends that of \mathcal{L}^1 slightly by processing the second-order quantifiers. The notion of state remains unchanged. The unfocused rewriting semantics is extended with the following two transitions:

$$\begin{array}{l}
\Sigma. \langle \Omega ; \Pi, \forall X : \tau \rightarrow o. R \rangle \mapsto \Sigma. \langle \Omega ; \Pi, [p/X]R \rangle \qquad \text{if } p : \tau \rightarrow o \text{ in } \Sigma \\
\Sigma. \langle \Omega ; \Pi, \exists X : \tau \rightarrow o. P \rangle \mapsto (\Sigma, X : \tau \rightarrow o). \langle \Omega ; \Pi, P \rangle
\end{array}$$

Observe that the first rule replaces the second-order order variable X with a predicate symbol from the signature.

The focused transition semantics undergoes a similar extension. It adds a rule to process second-order existential programs in a non-stable state

$$\Sigma. \langle \Omega ; \Pi, \exists X : \tau \rightarrow o. P \rangle \mapsto (\Sigma, X : \tau \rightarrow o). \langle \Omega ; \Pi, P \rangle$$

while second-order universal quantifiers are handled in the transitions for rules in the archive and program portions of the state. These rules stay therefore as in Section 2.1.3.

The properties seen for \mathcal{L}^1 , in particular type preservation, also hold for $\mathcal{L}^{1.5}$.

The second-order artifacts of $\mathcal{L}^{1.5}$ are not essential for our development, although they will simplify the discussion and make it more intuitive. In fact, any program in $\mathcal{L}^{1.5}$ can be transformed into a program in \mathcal{L}^1 that performs the same computation. We will now give a (slightly simplified) intuition of the workings of this transformation — a formal definition can be found in Appendix A.4.

Let P be a program in $\mathcal{L}^{1.5}$. For every type τ such that the second-order quantifications $\forall X : \tau \rightarrow o$ or $\exists X : \tau \rightarrow o$ appears in P , define the new predicate symbol $p_{\tau} : \iota \times \tau \rightarrow o$ and collect these predicates into a signature $\Sigma^{1.5}$. Next, replace second-order quantifications and variable uses as follows:

- $\forall X : \tau \rightarrow o. \dots X t \dots (X) \dots$ becomes $\forall x_X : \iota. \dots p_\tau(x_X, t) \dots (x_X) \dots$
- $\exists X : \tau \rightarrow o. \dots X t \dots (X) \dots$ becomes $\exists x_X : \iota \rightarrow \top. \dots p_\tau(x_X (), t) \dots (x_X ()) \dots$

In both cases, all uses of the variable X as a predicate name are replaced by p_X as shown. Moreover, all occurrences of X in a term are replaced with the first-order variable x_X . Observe that universal and existential second-order variables are treated differently, due to the distinct typing patterns that universal and existential first-order variables have. For any predicate name p appearing in signature Σ of a state $\Sigma. \langle \Omega ; \Pi \rangle$ as $p : \tau \rightarrow o$, we define a new constant $x_p : \top \rightarrow \iota$ in $\Sigma^{1.5}$, replace every occurrence of p in a term position in Π with $x_p ()$ and every atom $p t$ with $p_\tau(x_p (), t)$.

This encoding transforms valid states in $\mathcal{L}^{1.5}$ into valid states in \mathcal{L}^1 and maps transition steps in $\mathcal{L}^{1.5}$ to transition steps in \mathcal{L}^1 . Appendix A.4 spells out the details of the transformation and a formal account of these results.

3 Adding Modularity

In this section, we synthesize a module system for $\mathcal{L}^{1.5}$ by examining a number of examples. In each case, we will see how to write rules in a way as to emulate some characteristic or other of a module system, and then develop syntax to abstract the resulting linguistic pattern. For readability, types will be grayed out throughout this section. In most cases, they can indeed be automatically inferred from the way the variables they annotate are used.

3.1 Name Space Separation

Our first goal will be to allow rules, possibly executing concurrently, to make use of the functionalities provided by a program component without interfering with each other via this component. Those rules will be the client code and the program component will be the module. Key to avoiding interferences will be to give access to the module functionalities to each client through a different name.

As our first case study, consider the problem of adding two numbers written in unary — we write z and $s(n)$ for zero and the successor of n , respectively, and refer to the type of such numerals as nat . Addition is then completely defined by the single rule

$$!\forall x: \text{nat}. \forall y: \text{nat}. \text{add}(s(x), y) \multimap \text{add}(x, s(y))$$

In fact, for any concrete values m and n , inserting the atom $\text{add}(m, n)$ in the state triggers a sequence of applications of this rule that will end in an atom of the form $\text{add}(z, r)$, with r the result of adding m and n .¹ The way this adding functionality will typically be used is by having one rule generate the request in its right-hand side and another retrieve the result in its left-hand side, as in the following code snippet:

$$\begin{array}{l} \forall m: \text{nat}. \forall n: \text{nat}. \quad \dots \quad \multimap \quad \dots, \text{add}(n, m) \\ \forall r: \text{nat}. \text{add}(z, r), \dots \quad \multimap \quad \dots \end{array}$$

Another common usage pattern, more in the flavor of process algebraic specifications, is to embed the second rule in the right-hand side of the first:

$$\forall m: \text{nat}. \forall n: \text{nat}. \dots \multimap \left[\begin{array}{l} \dots, \text{add}(n, m), \\ \forall r: \text{nat}. \text{add}(z, r), \dots \multimap \dots \end{array} \right]$$

While this code achieves the desired effect in this case, it is incorrect whenever there is the possibility of two clients performing an addition at the same time. In fact, any concurrent execution making use of add will cause an interference as the clients have no way to sort out which result (r in $\text{add}(z, r)$) is who's.

One way to obviate to this problem is to include an additional argument in the definition of add to disambiguate calls. The client code generates a unique identifier using existential quantification and retrieves the result that matches this identifier. The updated code for add and for the caller is as follows:

$$\begin{array}{l} !\forall x: \text{nat}. \forall y: \text{nat}. \forall id: \iota. \text{add}(id, s(x), y) \multimap \text{add}(id, x, s(y)) \\ \forall m: \text{nat}. \forall n: \text{nat}. \quad \dots \quad \multimap \quad \exists id: \iota. \left[\begin{array}{l} \dots, \text{add}(id, n, m), \\ \forall r: \text{nat}. \text{add}(id, z, r), \dots \multimap \dots \end{array} \right] \end{array}$$

While this intervention solves the above concern, it does not prevent a programmer from misusing the predicate add , either accidentally or intentionally. In fact, she can intercept results (or partial results) with rules of the form

$$\forall x: \iota. \forall m: \text{nat}. \forall n: \text{nat}. \text{add}(x, m, n) \multimap \text{add}(x, z, s(z))$$

¹This is not the only way to implement unary addition using rewrite rules and in fact not the best one. It is good for illustration however.

This returns the (typically incorrect) result $s(z)$ to any legitimate caller of this adder. The problems stems from two facts: first the predicate symbol used internally in the implementation to perform the addition is publicly known, the second is that terms, even when existentially generated, are not private as they can be matched as just shown.

Both issues can be solved by using a second-order existential quantifier to generate the predicate used to carry out the addition. The following code snippet fully address the possibility of interferences. Observe that it does so without polluting the definition of addition with a third argument.

$$\forall m: \text{nat}. \forall n: \text{nat}. \dots \multimap \exists \text{add}: \text{nat} \times \text{nat} \rightarrow o. \left[\begin{array}{l} !\forall x: \text{nat}. \forall y: \text{nat}. \text{add}(s(x), y) \multimap \text{add}(x, s(y)), \\ \dots, \text{add}(n, m), \\ \forall r: \text{nat}. \text{add}(z, r), \dots \multimap \dots \end{array} \right]$$

The safety constraint on second-order variables prevents our rogue programmer from intercepting the freshly generated predicate out of thin air. In fact, the rule

$$\forall X: \text{nat} \times \text{nat} \rightarrow o. \forall m: \text{nat}. \forall n: \text{nat}. X(m, n) \multimap X(z, s(z))$$

is invalid as X does not appear in a term position in the left-hand side.

While the above approach eliminates the possibility of interferences, it forces every rule that relies on addition to repeat the same piece of code (the inner reusable rule for addition above). Such code duplication goes against the spirit of modularity as it prevents code reuse and reduces maintainability. For definitions larger than this minimal case study, it also clutters the client code thereby diminishing readability.

One approach is to separate out the common code for addition and splice it back just before execution. This is essentially the solution advocated in [28, 27] as “elaboration”. Achieving it within the model of computation of $\mathcal{L}^{1.5}$ would require significant, and probably ad-hoc, extensions.

A better approach is to provide a public name for the addition functionality, for example through the predicate name `adder`, but pass to it the name of the private predicate used by the client code (`add`). Each client can then generate a fresh name for it. The definition of addition, triggered by a call to `adder`, can then be factored out from the client code. The resulting rules are as follows:

$$\begin{array}{l} !\forall \text{add}: \text{nat} \times \text{nat} \rightarrow o. \text{adder}(\text{add}) \multimap [!\forall x: \text{nat}. \forall y: \text{nat}. \text{add}(s(x), y) \multimap \text{add}(x, s(y))] \\ \forall m: \text{nat}. \forall n: \text{nat}. \dots \multimap \exists \text{add}: \text{nat} \times \text{nat} \rightarrow o. \left[\begin{array}{l} \text{adder}(\text{add}), \\ \dots, \text{add}(n, m), \\ \forall r: \text{nat}. \text{add}(z, r), \dots \multimap \dots \end{array} \right] \end{array}$$

Observe that, as before, the client generates a fresh name for its private adding predicate ($\exists \text{add}: \text{nat} \times \text{nat} \rightarrow o.$). It now passes it to `adder`, which has the effect of instantiating the rule for addition with the private name `add`. The client can then retrieve the result by putting $\text{add}(z, r)$ in the left-hand side of an embedded rule like before.

This idea will be the cornerstone of our approach to adding modules to $\mathcal{L}^{1.5}$. Because it requires that the module and client code abide by a rather strict format, it will be convenient to provide the programmer with derived syntax. We will write the exported module (the first line in our last example) as follows:

```

module adder
  provide   add : nat × nat → o
            !∀x: nat. ∀y: nat. add(s(x), y)  ⇨  add(x, s(y))
end

```

Here, the public name `adder` is used as the name of the module. The names of the exported operations are introduced by the keyword `provide`. By isolating the public predicate name `adder` in a special position (after the keyword

module), we can statically preempt one newly-introduced problem with the above definition: that a rogue client learn the private name via the atom $\text{adder}(X)$ in the left-hand side of a rule. We shall disallow predicates used as module names (here adder) from appearing in the left-hand side of other rules.

We also provide derived syntax for the client code to avoid the tedium of creating a fresh predicate name and using it properly. The second rule of our example is re-expressed as follows:

$$\forall m: \text{nat}. \forall n: \text{nat}. \dots \multimap A \text{ as } \text{adder}. \left[\dots, A.\text{add}(n, m), \forall r: \text{nat}. A.\text{add}(z, r), \dots \multimap \dots \right]$$

Here, A is a *reference name* introduced by the line “ A as adder .”. This syntactic artifact binds A in the right-hand side of the rule. The name A allows constructing *compound predicate names*, here $A.\text{add}$, which allow using the exact same names exported by a module in its **provide** stanza. Uses of compound names and the **as** construct are elaborated into our original code, as we will see in Section 4.

The **provide** stanza defines the functionalities exported by a module. We can collect these declarations and give them a name in an *interface* for the module, which corresponds to the notion of signature in Standard ML. Our ongoing example leads to the following interface declaration:

```
interface ADDER
  add : nat × nat → o
end
```

This gives us a way to provide a name for this interface, here ADDER . The module adder can then be explicitly annotated with the name of its interface:

```
module adder: ADDER
  provide   add : nat × nat → o
           !∀m: nat. ∀n: nat. add(s(m), n)  ∼ add(m, s(n))
end
```

As in SML, we could use such annotations to restrict or even replace the **provide** stanza in the module definition. We will not explore this possibility further. Interfaces will come handy later. Until then, we will not bother with them.

To summarize what we have achieved so far, we have devised a template to support basic modularity in $\mathcal{L}^{1.5}$ and we have developed some derived syntax for it. The same module can be used by multiple clients at the same time without the danger of (intentional or accidental) interferences. Multiple implementations for the same interface are supported, although our example is still too simple to demonstrate this.

3.2 Modes

In the face of it, our adder module is a bit strange as it requires client code to use an atom of the form $\text{add}(z, r)$ to retrieve the result r of an addition. A better approach is to split add into a predicate add_req for issuing an addition request and a predicate add_res for retrieving the result. We can write this directly in $\mathcal{L}^{1.5}$ as the follows:

$$\begin{array}{l}
! \forall \text{add_req} : \text{nat} \times \text{nat} \rightarrow o. \\
\forall \text{add_res} : \text{nat} \rightarrow o. \text{ adder}'(\text{add_req}, \text{add_res}) \multimap \\
\left[\begin{array}{l}
! \forall x : \text{nat}. \forall y : \text{nat}. \text{ add_req}(s(x), y) \multimap \text{ add_req}(x, s(y)) \\
\forall z : \text{nat}. \text{ add_req}(z, z) \multimap \text{ add_res}(z)
\end{array} \right] \\
\forall m : \text{nat}. \forall n : \text{nat}. \quad \dots \quad \multimap \\
\exists \text{add_req} : \text{nat} \times \text{nat} \rightarrow o. \left[\begin{array}{l}
\text{ adder}'(\text{add_req}, \text{add_res}), \\
\dots, \text{ add_req}(n, m), \\
\exists \text{add_res} : \text{nat} \rightarrow o. \left[\begin{array}{l}
\forall r : \text{nat}. \text{ add_res}(r), \dots \multimap \dots
\end{array} \right]
\end{array} \right]
\end{array}$$

Now, uses of *add_req* in the left-hand side of client rules would intercept requests, while occurrences of *add_res* on their right-hand side could inject forged results. We prevent such misuses by augmenting the syntax of modules (and interfaces) with *modes* that describe how exported predicates are to be used. The mode **in** in the **provide** stanza is used to check that a (compound) predicate is used only on the left-hand side of a rule, while the mode **out** forces usage on the right-hand side only. A declaration without such a marker, for example *add* in the last section, can be used freely. Our new *adder* module is written as follows in the concrete syntax:

```

module adder'
  provide out add_req : nat × nat → o in add_res : nat → o

  !∀x: nat. ∀y: nat. add_req(s(x), y)  ⇨  add_req(x, s(y))
    ∀z: nat. add_req(z, z)  ⇨  add_res(z)

end

```

and the client code assumes the following concrete form:

$$\forall m : \text{nat}. \forall n : \text{nat}. \dots \multimap A \text{ as } \text{adder}'. \left[\begin{array}{l}
\dots, A.\text{add_req}(n, m), \\
\forall r : \text{nat}. A.\text{add_res}(r), \dots \multimap \dots
\end{array} \right]$$

3.3 Abstract Data Types

The infrastructure we developed so far already allows us to write modules that implement abstract data types. For example, the following module (in concrete syntax) implements a simple dictionary with three operations: lookup (split into a request and a result predicate name), insert and delete. Keys are natural numbers (of type *nat*) and values are terms (of type *ι*).

```

module dictionary
  provide out lookup_req : nat → o in lookup_res : nat × ι → o
    out insert : nat × ι → o
    out delete : nat → o

  local data : nat × ι → o

  !∀k: nat. ∀v: ι. insert(k, v)  ⇨  data(k, v)
  !∀k: nat. ∀v: ι. lookup_req(k), data(k, v)  ⇨  lookup_res(k, v), data(k, v)
  !∀k: nat. ∀v: ι. delete(k), data(k, v)  ⇨  .

end

```

The **local** stanza introduces a name that is used locally within the module. It is elaborated into an existential surrounding the rules defining the module. The $\mathcal{L}^{1.5}$ program corresponding to this module is as follows:

$$\begin{array}{l}
\forall \text{lookup_req}: \text{nat} \rightarrow o. \forall \text{lookup_res}: \text{nat} \times \iota \rightarrow o. \forall \text{insert}: \text{nat} \times \iota \rightarrow o. \forall \text{delete}: \text{nat} \rightarrow o. \\
\text{dictionary}(\text{lookup_req}, \text{lookup_res}, \text{insert}, \text{delete}) \quad \dashv\!\!\dashv \quad \exists \text{data}: \text{nat} \times \iota \rightarrow o. \\
\left[\begin{array}{l}
! \forall k: \text{nat}. \forall v: \iota. \quad \text{insert}(k, v) \quad \dashv\!\!\dashv \quad \text{data}(k, v) \\
! \forall k: \text{nat}. \forall v: \iota. \quad \text{lookup_req}(k), \text{data}(k, v) \quad \dashv\!\!\dashv \quad \text{lookup_res}(k, v), \text{data}(k, v) \\
! \forall k: \text{nat}. \forall v: \iota. \quad \text{delete}(k), \text{data}(k, v) \quad \dashv\!\!\dashv \quad \cdot
\end{array} \right]
\end{array}$$

As a more complex example of abstract data type, the following module implements stacks of natural numbers. Internally, stacks are linked lists built from the local predicates $\text{head}(d'')$, $\text{elem}(n, d', d'')$ and $\text{empty}(d')$. Here, the terms d' and d'' implement the links of the linked list by identifying the predecessor and the successor element, respectively. Such a list will start with the atom $\text{head}(d_0)$, continue as series of atoms $\text{elem}(n_i, d_i, d_{i+1})$ and end with the atom $\text{empty}(d_n)$. The code maintains that invariant that links are used exactly this way.

$$\begin{array}{l}
\text{module stack} \\
\text{provide out } \text{push} : \text{nat} \rightarrow o \\
\text{out } \text{pop_req} : o \quad \text{in } \text{pop_res} : \text{nat} \rightarrow o \\
\text{out } \text{isempty} : o \\
\text{out } \text{size_req} : o \quad \text{in } \text{size_res} : \text{nat} \rightarrow o \\
\\
\text{local } \text{empty} : \iota \rightarrow o \\
\text{elem} : \text{nat} \times \iota \times \iota \rightarrow o \\
\text{head} : \text{nat} \times \iota \rightarrow o \\
\quad \cdot \quad \dashv\!\!\dashv \quad \exists d: \iota. \text{empty}(d), \text{head}(d) \\
! \forall e: \text{nat}. \quad \text{push}(e), \text{head}(d) \quad \dashv\!\!\dashv \quad \exists d': \iota. \text{elem}(e, d', d), \text{head}(d') \\
! \forall e: \text{nat}. \forall d: \iota. \forall d': \iota. \text{pop_req}, \text{head}(d), \text{elem}(e, d, d') \quad \dashv\!\!\dashv \quad \text{pop_res}(e), \text{head}(d') \\
! \forall d: \iota. \quad \text{isempty}, \text{head}(d), \text{empty}(d) \quad \dashv\!\!\dashv \quad \text{head}(d), \text{empty}(d) \\
! \forall d: \iota. \quad \text{size_req}, \text{head}(d) \quad \dashv\!\!\dashv \quad \exists \text{size}': \text{nat} \times \iota \rightarrow o. \\
\quad \text{head}(d), \text{size}'(z, d), \\
\left[\begin{array}{l}
! \forall n: \text{nat}. \forall e: \text{nat}. \forall d: \iota. \forall d': \iota. \text{size}'(n, d), \text{elem}(e, d, d') \quad \dashv\!\!\dashv \quad \text{elem}(e, d, d'), \text{size}'(s(n), d') \\
\forall n: \text{nat}. \forall d: \iota. \quad \text{size}'(n, d), \text{empty}(d) \quad \dashv\!\!\dashv \quad \text{empty}(d), \text{size_res}(n)
\end{array} \right] \\
\text{end}
\end{array}$$

3.4 Parametric Modules – I

Our next challenge will be to implement a module that provides a functionality to increment a number by n , where n is a parameter passed when the module is instantiated in the client code. This corresponds to a Standard ML functor that takes a value as an argument.

Such parametricity can be achieved by simply passing n as an argument to the public predicate used to call the module, in addition to the predicates for each declaration in its **provide** stanza. In concrete syntax, we can simply give n as an argument in our module declaration. Here is an implementation:

```

module increment (n: nat)
  provide out incr_req : nat → o in incr_res : nat → o

  local add : nat × nat → o
    ∀m: nat. incr_req(m)  → add(m, n)
    !∀m: nat. ∀n: nat. add(s(m), n)  → add(m, s(n))
    ∀r: nat. add(z, r)  → incr_res(r)

end

```

The value m passed as a parameter through the requester $incr_req(m)$ is added to n using the same code we wrote for our adder. The result is returned through the predicate $incr_res$. The above module is written as follows in $\mathcal{L}^{1.5}$:

$$\forall n: \text{nat}. \forall incr_req: \text{nat} \rightarrow o. \forall incr_res: \text{nat} \rightarrow o. \text{increment}(n, incr_req, incr_res) \rightarrow$$

$$\exists add: \text{nat} \times \text{nat} \rightarrow o. \left[\begin{array}{l} \forall m: \text{nat}. incr_req(m) \rightarrow add(m, n) \\ !\forall m: \text{nat}. \forall n: \text{nat}. add(s(m), n) \rightarrow add(m, s(n)) \\ \forall r: \text{nat}. add(z, r) \rightarrow incr_res(r) \end{array} \right]$$

Although plausible, this code is incorrect as two invocations of $incr_req$ within the same left-hand side may cause internal interferences as the local predicate symbol add is shared: the two calls would not be able to tell which result is which. The following code corrects this mistake by making add local to each use of $incr_req$.

```

module increment (n: nat)
  provide out incr_req : nat → o in incr_res : nat → o

  ∀m: nat. incr_req(m)  → ∃add: nat × nat → o.
    
$$\left[ \begin{array}{l} add(m, n) \\ !\forall m: \text{nat}. \forall n: \text{nat}. add(s(m), n) \rightarrow add(m, s(n)) \\ \forall r: \text{nat}. add(z, r) \rightarrow incr_res(r) \end{array} \right]$$


end

```

While correct, this code is a bit silly as it re-implements our adder. An even better solution is use the adder' module to perform the addition, as done here:

```

module increment (n: nat)
  provide out incr_req : nat → o in incr_res : nat → o

  ∀m: nat. incr_req(m)  → A as adder'. 
$$\left[ \begin{array}{l} A.add\_req(m, n), \\ \forall r: \text{nat}. A.add\_res(r) \rightarrow incr_res(r) \end{array} \right]$$


end

```

Modules can therefore call each other.

Now, client code would make use of our increment module as follows:

$$\forall m: \text{nat}. \forall n: \text{nat}. \dots \rightarrow I \text{ as increment}(n). \left[\begin{array}{l} \dots, I.incr_req(m), \\ \forall r: \text{nat}. I.incr_res(r), \dots \rightarrow \dots \end{array} \right]$$

3.5 Sharing Private Names

As our next example, we will implement — and use — a module that provides the functionalities of a reference cell in a stateful language. This cell is initialized as we instantiate the module (using the approach to parametrization seen in the last section) and provides private predicates to get the current value of the cell and to set it to a new value. At this stage, this is easily done by the following code, which relies on the private predicate *content* to store the content of the cell.

```

module cell (v: nat)
  provide out get : o           in got : nat → o
           out set : nat → o

  local content : nat → o

           ·                               ⇨ content(v)
           ∀v: nat. get, content(v)   ⇨ got(v), content(v)
           ∀v: nat. ∀v': nat. set(v'), content(v) ⇨ content(v')
end

```

Differently from the example in the last section, the fact that *content* is not local to each invocation of *get* and *set* is a feature as we want all uses of the cell to refer to the same content.

Reference cells are often shared by various subcomputations. One way to do so is to pass the exported predicates that manipulate them to the subcomputations, after instantiation. In the following example, the client code (first rule) creates a cell *C* initialized with the value *s*(*z*). It then passes the setter to the second rule through the predicate *p* and passes the getter to the third rule through the predicate *q*. The second rule can only write to the cell (it replaces its content with *z* in this example). The third rule can only read the content of the cell (here, it then outputs it through predicate *s*).

```

           ·                               ⇨ C as cell(s(z)). [ p(C.set),
           p(write: nat → o. p(write)           ⇨ write(z).
           [ ∀read_req: o. q(read_req, read)   ⇨ [ read_req,
           [ ∀read: nat → o. ]                   [ ∀r: nat. read(r)   ⇨ s(r) ] ]

```

This example shows how the private names obtained through a module invocation can be shared with other rules. Furthermore this sharing can be selective.

Note that reference cells can be implemented in a much simpler way by exporting the predicate that holds the content of the cell. This is reminiscent of what we did with our first adder. The code is as follows:

```

module cell' (v: nat)
  provide content : nat → o

           ·                               ⇨ content(v)
end

```

Now, however, there is no way for the client code to pass distinct capabilities to subcomputations. In the following code, the first rule passes the compound predicate name *C.content* to the second and the third rule, thereby giving both of them read and write access to the cell.


```

module queue': QUEUE
  provide out   enq : nat → o
           out deq_req : o           in deq : nat → o

  local   nil : ι
           cons : nat × ι → ι
           q : ι → o
           .
           . → q(nil)
           !∀e: nat. ∀tail: ι. enq(e), q(tail) → q(cons(e, tail))

           !∀e: nat. deq_req, q(tail) →
           [
             q(tail),
             q(last(tail)),
             ∀v: nat. last(cons(v, nil)) → deq(v)
             ! [
               ∀v: nat.
               ! [
                 ∀v': nat.
                 ! [
                   ∀tail: ι.
                   last(cons(v, cons(v', tail))) → last(cons(v', tail))
                 ]
               ]
             ]
           ]

end

```

Observe that both implementations ascribe to the same interface for queues, QUEUE.

Next, we define a producer-consumer module that exports predicates to produce a natural number or to consume it. In this module, we use a queue as the buffer where the producer deposits data and the consumer retrieves them from — in truth, our producer-consumer module does nothing more than map these operations to the underlying enqueueing and dequeuing facilities.

Now, rather than selecting a priori which implementation of queues to use, we can make producer-consumer module parametric with respect to the implementation of queues. This corresponds a functor parametrized by a structure in Standard ML. The resulting code is as follows.

```

module prodcons(Q: QUEUE)
  provide in   produce : nat → o
           in consume_req : o           out consume : nat → o

  . → B as Q. [
    !∀e: nat. produce(e) → B.enq(e)
    ! consume_req → [
      B.deq_req,
      ∀e: nat. B.deq(e) → consume(e)
    ]
  ]

end

```

The argument Q of `prodcons` is the name of a module with interface QUEUE, which gives use a choice between `queue` and `queue'`. The following code chooses `queue` (arbitrarily), and triggers two producers and three consumers by passing them the appropriate predicates exported by the module. We also show one such producer (triggered by predicate p_1) which pushes three integers in the buffer, and one consumer (c_3) that retrieves two elements.

4 Multiset Rewriting with Modules

The module infrastructure we developed in the previous section had two parts:

1. We introduced some convenience syntax for the programming patterns of $\mathcal{L}^{1.5}$ that realized module definitions (**module** ... **end**), module instantiations ($N \text{ as } p t. \dots$), and the use of exported names (e.g., $N.add$).
2. We imposed a series of restrictions on where and how predicates could be used (**in** for left-hand side only, **out** for right-hand side only), as well as a correspondence between the names exported by a module and the compound names used in client code.

We call the extension of $\mathcal{L}^{1.5}$ with both aspects \mathcal{L}^M . In this section, we describe how the added syntax in (1) can be compiled away, thereby showing that \mathcal{L}^M is just syntactic sugar for $\mathcal{L}^{1.5}$ — $\mathcal{L}^{1.5}$ has all the ingredients to write modular code already. We call this process *elaboration*. We handle the restrictions in (2) by typechecking \mathcal{L}^M code, as a user could violate them even if her code elaborated to a valid $\mathcal{L}^{1.5}$ program. Appendix A.5 describes an extension of the typing rules of Section 2.1.2 that checks that these restrictions are satisfied at the level of \mathcal{L}^M . It also summarizes the syntax of our module language (which is actually more flexible than what we saw in the examples of Section 3).

Once an \mathcal{L}^M program has been typechecked, it is elaborated into an $\mathcal{L}^{1.5}$ program by compiling the module-specific constructs into the native syntax of $\mathcal{L}^{1.5}$. We then execute this $\mathcal{L}^{1.5}$ program. We will now outline how to elaborate away the two added constructs, **module** and **end**. See Appendix A.5 for a precise description.

Recall the general syntax of modules definition has the form:

| | |
|----------------------------------|--|
| module $p(\Sigma_{par})$ | Module name and parameters (term and predicates) |
| provide Σ_{export} | Exported names |
| local Σ_{local} | Local predicates and constructors |
| P | Module definition — may use predicate names defined externally |
| end | |

What we called the module interface consists of the signature Σ_{export} on the **provide** stanza. The modes of the exported predicate names (**in** and **out** in Section 3) are irrelevant after type checking — we ignore them. Let Σ^* denote the tuple of the names declared in signature Σ . Then this construct is elaborated into the following $\mathcal{L}^{1.5}$ rule:

$$\forall \Sigma_{par}. \forall \Sigma_{export}. p(\Sigma_{par}^*, \Sigma_{export}^*) \multimap \exists \Sigma_{local}. P$$

where the notation $\forall \Sigma. R$ prefixes the rule R with one universal quantification for each declaration in Σ , and similarly for $\exists \Sigma. P$.

Next, we handle module instantiation, whose derived syntax, is

$$N \text{ as } p t. P$$

Let Σ_{export} be the interface exported by the module for p , defined as in the previous paragraph. It will be convenient to express Σ_{export} in the form $\vec{X} : \vec{\tau}$, where the i -th declaration in Σ_{export} is $X_i : \tau_i$. Then, we elaborate the above construct as

$$\exists \Sigma_{export}. \left[\begin{array}{l} p(t, \vec{X}) \\ [\vec{X}/N.\vec{X}]P \end{array} \right]$$

where $[\vec{X}/N.\vec{X}]P$ replaces each occurrence of $N.X_i$ in P with X_i . We implicitly assume that variables are renamed as to avoid capture.

Elaboration, as just described, transforms a valid \mathcal{L}^M program into a valid $\mathcal{L}^{1.5}$ program, ready to be executed. In particular, it removes all compound names of the form $N.X$.

5 Related Work

Programming languages equipped with a module system started appearing in the early 1970s, with Modula-2 [37] and Standard ML [29] still in use today. Modern languages, from Java to JavaScript, invariably feature a module system of some sort. The module system of ML [25, 29] is still considered by many the gold standard because of its support for parametric modules (functors) and fine control on exported functionalities. Extensions, in particular with higher-order modules are active areas of research. It should be noted however that large amounts of code are still written nowadays in non-modular languages, in particular in C.

Module research in logic programming has largely targeted languages based on backward search. See [4] for a comprehensive survey as of the mid-1990s. Several commercial and open-source implementations of Prolog feature a module system, for example SICStus Prolog and SWI Prolog. Declarative approaches to modularity have been investigated in relation to extensions of Prolog. In particular [27, 28] bases the module system for λ Prolog on this language's support for embedded implication and universal quantification, in a fashion that resembles our use of existential quantifiers. Modules are however dynamically spliced into the body of each clause that uses them, and therefore are not logically separate entities. Abdallah [1] proposes using second-order variables as a declarative support for procedure calls in a Prolog-like language, achieving a form of modularity along the way. This is closely related to our use of second-order quantification and to the logic recently investigated in [15]. Several proposals have been made for introducing modules in dependently-typed backward-chaining languages, for example [32] for LF.

Module systems for languages based on forward logic programming have been few and far between. Cervesato proposed an ad-hoc module system for the multiset rewriting language MSR [6]. Some languages based on general rewriting have been extended with a module system, for example Maude [26, 10]. Process algebras [34, 13] still tend to be used for specification rather than for general programming, yet some research has been done about modularizing them, especially for assurance purposes [33].

The relationship between logic and rule-based languages has been explored in great depth. Most relevant to the present work is [9] that shows how the left sequent rules of most connectives of linear logic yield a rewriting behavior that subsumes traditional multiset rewriting and established models of concurrency such as Petri nets [30] as well as many process algebras. The unfocused semantics of $\mathcal{L}^{1.5}$ discussed in this report is largely based on a subset of these operators. The connection between focusing and forward chaining has also been the subject of recent study, for example in [35].

Forward-chaining logic programming gained prominence in the 1980s with Datalog [16, 5] in the database area. This paradigm has been having a renaissance of late and has been the basis for languages in the domains of security and trust management [23], networking [18, 24], robot coordination [2], multicore programming [11], graph programming [12] and concurrent specifications [36] just to cite a few. It is also a foundation for the general purpose programming language CHR [14] and earlier GAMMA [3] as well as CoMingle, a language for mobile application development [21].

6 Future Work and Conclusions

In this report, we developed an advanced module system for \mathcal{L}^1 , a small rule-based programming language. This language corresponds to a large fragment of polarized intuitionistic linear logic under a derivation strategy based on focusing. Modules are rewrite rules of a specific form, which means that not only do they too have a natural interpretation in logic, but they are also first-class entities in our language. They make use of a mild form of second-order quantification, which can be compiled to first-order quantification but at the cost of complicating somewhat the static semantics of our language. Although our module infrastructure is nothing more than normal rules used in a specific way, they share many of the features found in advanced module systems such as that of Standard ML [25]. In particular, they enable private names, which in turn support abstract data types, name space separation and other forms of information hiding. Our modules can be invoked with parameters, both values supplied by a client rule as well as other modules. This corresponds directly to the notion of functor in Standard ML. Being first-class, higher-order modules and recursive functors are natively supported. The underlying rewriting model also enables sharing private names across threads of computation, which is critical for the kind of concurrent settings multiset rewriting is applied most frequently nowadays [20].

The rewriting language underlying our modular scaffolding is of interest in its own right. Similarly to the language of ω -multisets [9], it can be viewed as a specific reading of the traditional inference rules for a fragment of linear logic. It subsumes most multiset rewriting formalisms we are aware of (e.g., [14, 36]), forward logic programming languages such as Datalog [16, 5] and recent refinements [18, 24, 2, 11, 12], and many process algebras [34, 13]. This entails that the module infrastructure developed in this report can be applied directly to any of these formalisms.

Our immediate next step will be to implement our module system in the CoMingle system [21]. CoMingle is a logic programming framework aimed at simplifying the development of applications distributed over multiple Android devices. It currently relies on the fragment of the language discussed in this report deprived of nested rules. It however extends this language with multiset comprehensions, which allow atomically manipulating arbitrarily many atoms that match a pattern, and provides a declarative interface to the underlying Android runtime system. CoMingle programs are compiled, and one of the most interesting challenges of extending it with modules will be to support separate compilation, one of the key feature of advanced module systems à la Standard ML.

We also intend to further extend the core language we defined in this report. Here, second-order variables stood for predicate names only, making them more like the quantifiers of Gabbay and Mathijssen’s “one-and-a-halfth-order logic” [15]. Although matching in second-order logic is decidable, it is expensive in its most general form [19]. However, second-order matching also supports features such as reflection and the runtime construction of rules, which are valuable programming tools. We want to investigate relaxing our current constraint while preserving performance. Our core language also possesses numerous features found in process algebras, in particular nested actions and the dynamic creation of names. We expect many of the verification techniques developed for process algebras, for example bisimilarity and congruence in the π -calculus [34], can be adapted to our language. This would be particularly interesting as reasoning about multiset rewriting is usually based on a different set of techniques, often reachability arguments.

References

- [1] M.A. Nait Abdallah. Procedures in Horn-Clause Programming. In Ehud Shapiro, editor, *Third International Conference on Logic Programming — ICLP'86*, pages 433–447. Springer LNCS 225, 1986.
- [2] Michael P. Ashley-Rollman, Peter Lee, Seth Copen Goldstein, Padmanabhan Pillai, and Jason D. Campbell. A Language for Large Ensembles of Independently Executing Nodes. In *25th International Conference on Logic Programming — ICLP'09*, pages 265–280. Springer LNCS 5649, 2009.
- [3] Jean-Pierre Banâtre and Daniel Le Métayer. Programming by Multiset Transformation. *Communications of the ACM*, 36(1):98–111, 1993.
- [4] Michele Bugliesi, Evelina Lamma, and Paola Mello. Modularity in Logic Programming. *The Journal of Logic Programming*, 19–20(1):443–502, 1994.
- [5] Stefano Ceri, Georg Gottlob, and Letizia Tanca. What You Always Wanted to Know About Datalog (And Never Dared to Ask). *IEEE Transactions on Knowledge and Data Engineering*, 1(1):146–166, March 1989.
- [6] Iliano Cervesato. MSR 2.0: Language Definition and Programming Environment. Technical Report CMU-CS-11-141, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, November 2011.
- [7] Iliano Cervesato. Typed Multiset Rewriting Specifications of Security Protocols. Technical Report CMU-CS-11-140, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, October 2011.
- [8] Iliano Cervesato, Frank Pfenning, David Walker, and Kevin Watkins. A Concurrent Logical Framework II: Examples and Applications. Technical Report CMU-CS-02-102, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 2003.
- [9] Iliano Cervesato and Andre Scedrov. Relating State-Based and Process-Based Concurrency through Linear Logic. *Information & Computation*, 207(10):1044–1077, 2009.
- [10] Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Carolyn Talcott. System Modules. In *All About Maude — A High-Performance Logical Framework*, pages 131–157. Springer LNCS 4350, 2007.
- [11] Flávio Cruz, Michael P. Ashley-Rollman, Seth Copen Goldstein Goldstein, Ricardo Rocha, and Frank Pfenning. Bottom-Up Logic Programming for Multicores. In Vítor Santos Costa, editor, *7th International Workshop on Declarative Aspects and Applications of Multicore Programming — DAMP'12*. ACM Digital Library, January 2012. Short paper.
- [12] Flávio Cruz, Ricardo Rocha, Seth Copen Goldstein Goldstein, and Frank Pfenning. A Linear Logic Programming Language for Concurrent Programming over Graph Structures. In *30th International Conference on Logic Programming — ICLP'14*, Vienna, Austria, 2014.
- [13] Cédric Fournet. *The Join-Calculus: a Calculus for Distributed Mobile Programming*. PhD thesis, École Polytechnique, Palaiseau, 1998. INRIA TU-0556. Also available from <http://research.microsoft.com/~fournet>.
- [14] Thom Frühwirth. *Constraint Handling Rules*. Cambridge University Press, 2009.
- [15] Murdoch J. Gabbay and Aad Mathijssen. One-and-a-halfth-order Logic. *Journal of Logic and Computation*, 18(4):521–562, August 2008.
- [16] Hervé Gallaire, Jack Minker, and Jean-Marie Nicolas. Logic and Databases: A Deductive Approach. *ACM Computing Survey*, 16(2):153–185, June 1984.
- [17] Jean-Yves Girard. Linear Logic. *Theoretical Computer Science*, 50:1–102, 1987.

- [18] Stéphane Grumbach and Fang Wang. Netlog, a Rule-based Language for Distributed Programming. In *12th International Conference on Practical Aspects of Declarative Languages — PADL'10*, pages 88–103. Springer LNCS 5937, 2010.
- [19] Kouichi Hirata, Keizo Yamada, and Masateru Harao. Tractable and Intractable Second-order Matching Problems. *Journal of Symbolic Computation*, 37(5):611–628, 2004.
- [20] Edmund S.L. Lam and Iliano Cervesato. Optimized Compilation of Multiset Rewriting with Comprehensions. In Jacque Garrigue, editor, *12th Asian Symposium on Programming Languages and Systems — APLAS'14*, pages 19–38, Singapore, 2014. Springer LNCS 8858.
- [21] Edmund S.L. Lam, Iliano Cervesato, and Nabeeha Fatima Haque. Comingle: Distributed Logic Programming for Decentralized Mobile Ensembles. In Tom Holvoet and Mirko Viroli, editors, *17th IFIP International Conference on Coordination Models and Languages — COORDINATION'15*, pages 51–66, Grenoble, France, 2–4 June 2015. Inria, Springer LNCS 9037.
- [22] Olivier Laurent, Myriam Quatrini, and Lorenzo Tortora de Falco. Polarized and Focalized Linear and Classical Proofs. *Annals of Pure and Applied Logic*, 134(2–3):217–264, July 2005.
- [23] Ninghui Li and John C. Mitchell. Datalog with Constraints: A Foundation for Trust-management Languages. In *5th International Conference on Practical Aspects of Declarative Languages — PADL'03*, pages 58–73. Springer LNCS 2562, 2003.
- [24] Boon Thau Loo, Tyson Condie, Minos Garofalakis, David E. Gay, Joseph M. Hellerstein, Petros Maniatis, Raghu Ramakrishnan, Timothy Roscoe, and Ion Stoica. Declarative Networking: Language, Execution and Optimization. In *2006 ACM SIGMOD International Conference on Management Of Data — SIGMOD'06*, pages 97–108. ACM, 2006.
- [25] David MacQueen. Modules for Standard ML. In *1984 ACM Symposium on LISP and Functional Programming Languages*, pages 198–207, 1984.
- [26] José Meseguer and Christiano Braga. Modular Rewriting Semantics of Programming Languages. In Charles Rattray, Savitri Maharaj, and Carron Shankland, editors, *10th International Conference on Algebraic Methodology and Software Technology — AMAST'04*, pages 364–378. Springer LNCS 3116, 2004.
- [27] Dale Miller. A Logical Analysis of Modules in Logic Programming. *Journal of Logic Programming*, 6(1-2):79–108, January 1989.
- [28] Dale Miller. A Proposal for Modules in λ Prolog. In Roy Dyckhoff, editor, *4th Workshop on Extensions to Logic Programming — ELP'94*, pages 206–221. Springer LNCS 798, 1994.
- [29] Robin Milner, Robert Harper, David MacQueen, and Mads Tofte. *The Definition of Standard ML – Revised*. MIT Press, 1997.
- [30] Carl Adam Petri. Fundamentals of a Theory of Asynchronous Information Flow. In *Proceedings of IFIP Congress 62 on Information Processing*, pages 386–390, Munich, Germany, 1963. North Holland.
- [31] Frank Pfenning. Structural Cut Elimination I. Intuitionistic and Classical Logic. *Information and Computation*, 157(1/2):84–141, 2000.
- [32] Florian Rabe and Carsten Schürmann. A Practical Module System for LF. In James Cheney and Amy Felty, editors, *Proceedings of the Fourth International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice — LFMTP'09*, pages 40–48, Montreal, Canada, 2009.
- [33] Sriram Rajamani and Jakob Rehof. A Behavioral Module System for the Pi-Calculus. In Patrick Cousot, editor, *8th International Symposium on Static Analysis — SAS'01*, pages 375–394. Springer LNCS 2126, 2001.

- [34] Davide Sangiorgi and David Walker. *The Pi-Calculus — a Theory of Mobile Processes*. Cambridge University Press, 2001.
- [35] Robert J. Simmons. Structural Focalization. *ACM Transaction in Computational Logic*, 15(3):1–33, September 2014.
- [36] Kevin Watkins, Frank Pfenning, David Walker, and Iliano Cervesato. Specifying Properties of Concurrent Computations in CLF. In Carsten Schürmann, editor, *Fourth Workshop on Logical Frameworks and Meta-languages — LFM'04*, pages 67–87, Cork, Ireland, 2007. ENTCS 199.
- [37] Niklaus Wirth. *Programming in Modula-2*. Springer, 4th edition edition, 1988.

A Language Summary

This appendix summarizes the syntax (Appendix A.1), typing and operational semantics of the language $\mathcal{L}^{1.5}$ developed in this report. We collect the base typing rules of the language, deprived of safety checks, in Appendix A.2 and enrich them with such checks in Appendix A.3. We show how to translate $\mathcal{L}^{1.5}$ programs back to \mathcal{L}^1 in Appendix A.4. The syntax of modular language \mathcal{L}^M is collected in Appendix A.5 together with typing rules for it. The summary of the operational semantics of our language starts in Appendix A.6 with a discussion of the underlying notion of congruence, continues in Appendix A.7 with an abstract presentation of its abstract rewriting semantics read off directly from the rules of the logic operator associated with each operator of $\mathcal{L}^{1.5}$ (see Appendix B), and ends with an operational semantics which is based on a focused presentation of this logic. Such an operational semantics is an essential step towards a practical implementation of this rewriting language.

A.1 Syntax

This appendix summarizes the syntax of the our core language $\mathcal{L}^{1.5}$. The addition that $\mathcal{L}^{1.5}$ brings about with respect to \mathcal{L}^1 are highlighted in [blue](#).

The grammar below refers to the following syntactic categories, which are assumed given and distinguishable:

| | |
|----------------------------|-----|
| <i>Term variables</i> | x |
| <i>Predicate variables</i> | X |
| <i>Function symbols</i> | f |
| <i>Predicate symbols</i> | p |

In the typing and evaluation semantics below, we blur the distinction between (term and predicate) variables introduced by an existential quantifier and (function and predicate) symbols. A more precise definition would instantiate such variables with new symbols when they are used, but at the cost of notational clutter — we prefer a simpler albeit slightly imprecise notation to a fully formal description.

The other entities of our language are defined by the grammar in Figure A.1.

In examples, we abbreviated the type $\top \rightarrow \iota$ as simply ι . This is the type of a function symbol that does not take any (proper) argument — a constant. Any signature declaration $f : \iota$ is therefore understood as $f : \top \rightarrow \iota$ and every term written f stands for $f()$. Similarly, we write the type $\top \rightarrow o$ as o — this is the type predicate names and variables without (proper) argument. A declaration $p : o$ stands therefore for $p : \top \rightarrow o$ and an atom p is shorthand for $p()$ — these are the propositional atoms.

| | | | | |
|-----------------|--------|-------|---|--------------------------|
| <i>Types</i> | τ | $::=$ | ι | individuals |
| | | | o | state objects |
| | | | $\tau \rightarrow \tau$ | function type |
| | | | $\tau \times \tau$ | product type |
| | | | \top | unit type |
| <i>Terms</i> | t | $::=$ | x | term variables |
| | | | $f t$ | term constructors |
| | | | (t, t) | products |
| | | | $()$ | unit |
| | | | X | predicate variables |
| | | | p | predicate names |
| <i>Atoms</i> | A | $::=$ | $p t$ | atoms |
| | | | $X t$ | parametric atoms |
| <i>LHS</i> | l | $::=$ | \cdot | empty left-hand side |
| | | | A, l | left-hand side extension |
| <i>Rules</i> | R | $::=$ | $l \multimap P$ | rewrite |
| | | | $\forall x : \iota. R$ | term abstraction |
| | | | $\forall X : \tau \rightarrow o. R$ | predicate abstraction |
| <i>Programs</i> | P | $::=$ | \cdot | empty program |
| | | | P, P | program union |
| | | | A | atom |
| | | | $!A$ | reusable atom |
| | | | R | rule |
| | | | $!R$ | reusable rule |
| | | | $\exists x : \tau \rightarrow \iota. P$ | new term variable |
| | | | $\exists X : \tau \rightarrow o. P$ | new predicate variable |

Figure A.1: Grammar of $\mathcal{L}^{1.5}$

A.2 Typing

This appendix summarizes the typing rules of our core language, without regard for safety constraint checks (which are the subject of Appendix A.3) or the module syntax (which is given in Appendix A.5).

Typing makes use of signatures and contexts, defined next.

| | | | | |
|-------------------|----------|-------|---|-------------------------|
| <i>Signatures</i> | Σ | $::=$ | \cdot | empty signature |
| | | | $ \ \Sigma, f : \tau \rightarrow \iota$ | constructor declaration |
| | | | $ \ \Sigma, p : \tau \rightarrow o$ | predicate declaration |
| <i>Contexts</i> | Γ | $::=$ | \cdot | empty context |
| | | | $ \ \Gamma, x : \iota$ | term assumption |
| | | | $ \ \Gamma, X : \tau \rightarrow o$ | predicate assumption |

The typing rules for our language are very conventional. They are as follows.

Terms $\boxed{\Gamma \vdash_{\Sigma} t : \tau}$

$$\frac{}{\Gamma, x : \tau \vdash_{\Sigma} x : \tau} \quad \frac{\Gamma \vdash_{\Sigma, f: \tau \rightarrow \iota} t : \tau}{\Gamma \vdash_{\Sigma, f: \tau \rightarrow \iota} f t : \iota} \quad \frac{\Gamma \vdash_{\Sigma} t_1 : \tau_1 \quad \Gamma \vdash_{\Sigma} t_2 : \tau_2}{\Gamma \vdash_{\Sigma} (t_1, t_2) : \tau_1 \times \tau_2} \quad \frac{}{\Gamma \vdash_{\Sigma} () : \bar{\tau}}$$

$$\frac{}{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} X : \tau \rightarrow o} \quad \frac{}{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} p : \tau \rightarrow o}$$

Atoms $\boxed{\Gamma \vdash_{\Sigma} A \text{ atom}}$

$$\frac{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} t : \tau}{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} p t \text{ atom}} \quad \frac{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} t : \tau}{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} X t \text{ atom}}$$

Rule left-hand-sides $\boxed{\Gamma \vdash_{\Sigma} l \text{ lhs}}$

$$\frac{}{\Gamma \vdash_{\Sigma} \cdot \text{lhs}} \quad \frac{\Gamma \vdash_{\Sigma} A \text{ atom} \quad \Gamma \vdash_{\Sigma} l \text{ lhs}}{\Gamma \vdash_{\Sigma} A, l \text{ lhs}}$$

Rules $\boxed{\Gamma \vdash_{\Sigma} R \text{ rule}}$

$$\frac{\Gamma \vdash_{\Sigma} l \text{ lhs} \quad \Gamma \vdash_{\Sigma} P \text{ prog}}{\Gamma \vdash_{\Sigma} l \multimap P \text{ rule}} \quad \frac{\Gamma, x : \iota \vdash_{\Sigma} R \text{ rule}}{\Gamma \vdash_{\Sigma} \forall x : \iota. R \text{ rule}} \quad \frac{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} R \text{ rule}}{\Gamma \vdash_{\Sigma} \forall X : \tau \rightarrow o. R \text{ rule}}$$

Programs $\boxed{\Gamma \vdash_{\Sigma} P \text{ prog}}$

$$\frac{}{\Gamma \vdash_{\Sigma} \cdot \text{prog}} \quad \frac{\Gamma \vdash_{\Sigma} P_1 \text{ prog} \quad \Gamma \vdash_{\Sigma} P_2 \text{ prog}}{\Gamma \vdash_{\Sigma} P_1, P_2 \text{ prog}}$$

$$\frac{\Gamma \vdash_{\Sigma} A \text{ atom}}{\Gamma \vdash_{\Sigma} A \text{ prog}} \quad \frac{\Gamma \vdash_{\Sigma} A \text{ atom}}{\Gamma \vdash_{\Sigma} !A \text{ prog}} \quad \frac{\Gamma \vdash_{\Sigma} R \text{ rule}}{\Gamma \vdash_{\Sigma} R \text{ prog}} \quad \frac{\Gamma \vdash_{\Sigma} R \text{ rule}}{\Gamma \vdash_{\Sigma} !R \text{ prog}}$$

$$\frac{\Gamma \vdash_{\Sigma, x: \tau \rightarrow \iota} P \text{ prog}}{\Gamma \vdash_{\Sigma} \exists x : \tau \rightarrow \iota. P \text{ prog}} \quad \frac{\Gamma \vdash_{\Sigma, X: \tau \rightarrow o} P \text{ prog}}{\Gamma \vdash_{\Sigma} \exists X : \tau \rightarrow o. P \text{ prog}}$$

Although evaluation level entities, we also define archives and states and give their typing rules.

$$\begin{array}{lcl}
 \text{Archives } \Omega & ::= & \cdot \quad \text{empty archive} \\
 & & | \quad \Omega, A \quad \text{reusable atom} \\
 & & | \quad \Omega, R \quad \text{reusable rule} \\
 \text{States } \Sigma.\langle \Omega ; \Pi \rangle & &
 \end{array}$$

The typing rules for states amount to requiring that the contained archive and program be valid and closed.

Archives $\boxed{\vdash_{\Sigma} \Omega \text{ archive}}$

$$\frac{}{\vdash_{\Sigma} \cdot \text{archive}} \qquad \frac{\vdash_{\Sigma} \Omega \text{ archive} \quad \cdot \vdash_{\Sigma} A \text{ atom}}{\vdash_{\Sigma} \Omega, A \text{ archive}} \qquad \frac{\vdash_{\Sigma} \Omega \text{ archive} \quad \cdot \vdash_{\Sigma} R \text{ rule}}{\vdash_{\Sigma} \Omega, R \text{ archive}}$$

States $\boxed{\vdash_{\Sigma} \langle \Omega ; \Pi \rangle \text{ state}}$

$$\frac{\vdash_{\Sigma} \Omega \text{ archive} \quad \cdot \vdash_{\Sigma} \Pi \text{ prog}}{\vdash_{\Sigma} \langle \Omega ; \Pi \rangle \text{ state}}$$

A.3 Typing with Safety Checks

In this appendix, we refine the typing rules seen in Appendix A.2 to capture the safety constraints on variable occurrences in a program.

An atom A in a program is a *left-hand side atom* (or lhs-atom) if it is introduced by the production $l ::= A, l$ in the grammar of Appendix A.1. It is a *right-hand side atom* (a rhs-atom) if it is instead introduced by the productions $P ::= A$ and $P ::= !A$. We write A_L and A_R for a generic lhs- and rhs-atom, respectively. Terms occurring in a lhs-atom (rhs-atom) are called lhs-terms (rhs-terms, respectively). We write $A_L[x]$ to indicate that the variable x occurs in a term in lhs-atom A_L , and similarly for rhs-atoms. We adopt the same notation for predicate variables.

The safety constraints can then be expressed as follows:

1. If a (term or predicate) variable introduced by universal quantification occurs in a rhs-atom, this same variable must also occur in an earlier lhs-atom.

Pictorially, every occurrence of a rhs-atom $A_R[x]$ within the scope of a universal declaration for x shall adhere to the following template

$$\forall x. \dots A_L[x] \dots \multimap (\dots, A_R[x], \dots) \dots$$

The case for a predicate variable X is analogous.

2. If a universal predicate variable X leads a parametric lhs-atom $X t$, then X must occur in a term position either in the same left-hand side or in a the left-hand side of an enclosing rule. If the first case, it must be possible to rearrange this left-hand side so that the term occurrence appears to the left of the $X t$ to avoid circularities such as $X(Y), Y(X)$.

Pictorially, given any program, it should be possible to reorder its lhs-atoms so that, for every lhs-atom $X t$ where X is universal, there is another lhs-atom $A_L[X]$ so that either $A_L[X]$ occurs in the left-hand side of an outer rule, as in the following template,

$$\forall X. \dots A_L[X] \dots \multimap (\dots, X t, \dots \multimap \dots)$$

or $A_L[X]$ occurs to the left of $X t$ in the same left-hand side, as depicted by the following template:

$$\forall X. \dots (\dots, A_L[X], \dots, X t, \dots) \multimap \dots$$

We enforce the safety constraints on variables by augmenting the typing judgments in Appendix A.2 with either one or two *constraint sets*. A constraint set annotates each universal variable in scope with a marker that indicates whether this variable has been seen in a lhs-term before. Specifically, x^0 indicates that variable x has been declared (i.e., introduced by means of a universal quantifier) but has not yet appeared in a term in a lhs-atom. Instead, x^1 designates x as having been seen in an earlier lhs-atom. The same convention applies for predicate variables.

Constraint sets are therefore defined by the following grammar:

| | | | |
|----------------------------------|-------|---|--|
| <i>Constraints</i> \mathcal{X} | $::=$ | . | no constraints |
| | | | \mathcal{X}, x^0 unseen term variable |
| | | | \mathcal{X}, x^1 seen term variable |
| | | | \mathcal{X}, X^0 unseen predicate variable |
| | | | \mathcal{X}, X^1 seen predicate variable |

The typing judgments seen in Appendix A.2 acquire a new component, a constraint set that tracks the use of universal variables. The judgments for syntactic objects that appear in a program position (including rule right-hand sides) are updated differently from the judgments for objects that appear in a left-hand side. Since atoms and terms can appear in both, they are refined into two versions, one for when they occur in program position (rhs-atoms and rhs-terms) and a distinct one for when they appear in a left-hand side (lhs-atoms and lhs-terms).

Program-side judgments are extended with a single constraint set, which we indicate in red. These judgments are

| | |
|--|-------------|
| $\Gamma \vdash_{\Sigma} t : \tau \mid \mathcal{X}$ | Terms (rhs) |
| $\Gamma \vdash_{\Sigma} A \text{ atom} \mid \mathcal{X}$ | Atoms (rhs) |
| $\Gamma \vdash_{\Sigma} R \text{ rule} \mid \mathcal{X}$ | Rules |
| $\Gamma \vdash_{\Sigma} P \text{ prog} \mid \mathcal{X}$ | Programs |

In each rule, the constraint set \mathcal{X} is extended as the corresponding syntactic object is traversed, going from the conclusion to the premises of the rules. The markers of the variables in \mathcal{X} are checked in rules where variables are used (in the judgments for terms and atoms).

Left-hand side judgments are enriched by two constraint sets, which we write as $\mathcal{X} > \mathcal{X}'$ and color in blue. These judgments are

| | |
|---|-----------------|
| $\Gamma \vdash_{\Sigma} t : \tau \mid \mathcal{X} > \mathcal{X}'$ | Terms (lhs) |
| $\Gamma \vdash_{\Sigma} A \text{ atom} \mid \mathcal{X} > \mathcal{X}'$ | Atoms (lhs) |
| $\Gamma \vdash_{\Sigma} l \text{ lhs} \mid \mathcal{X} > \mathcal{X}'$ | Left-hand sides |

The constraint set \mathcal{X} is understood as an input to the judgment while \mathcal{X}' is its output: \mathcal{X} is propagated from the conclusion to one of the premises of a rule while \mathcal{X}' flows the other direction, from a premise to the conclusion. The set \mathcal{X}' differs from \mathcal{X} by the fact that the marker of some variables has been updated from 0 to 1 . This takes place when a variable is encountered within a (left-hand side) term

The updated rules are displayed next, with commentary when needed to clarify the refinement. Here, we take advantage of the commutativity of “,” and implicitly reorder atoms in a left-hand side as needed to obtain a derivation. It is not necessary to do so in program positions. In all cases, we treat existential variables as if they were constants.

Terms (lhs) $\boxed{\Gamma \vdash_{\Sigma} t : \tau \mid \mathcal{X} > \mathcal{X}'}$

$$\frac{}{\Gamma, x : \tau \vdash_{\Sigma} x : \tau \mid \mathcal{X}, x^? > \mathcal{X}, x^1} \qquad \frac{\Gamma \vdash_{\Sigma, f: \tau \rightarrow \iota} t : \tau \mid \mathcal{X} > \mathcal{X}'}{\Gamma \vdash_{\Sigma, f: \tau \rightarrow \iota} f t : \iota \mid \mathcal{X} > \mathcal{X}'}$$

$$\frac{\Gamma \vdash_{\Sigma} t_1 : \tau_1 \mid \mathcal{X} > \mathcal{X}' \quad \Gamma \vdash_{\Sigma} t_2 : \tau_2 \mid \mathcal{X}' > \mathcal{X}''}{\Gamma \vdash_{\Sigma} (t_1, t_2) : \tau_1 \times \tau_2 \mid \mathcal{X} > \mathcal{X}''} \qquad \frac{}{\Gamma \vdash_{\Sigma} () : \top \mid \mathcal{X} > \mathcal{X}}$$

$$\frac{}{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} X : \tau \rightarrow o \mid \mathcal{X}, X^? > \mathcal{X}, X^1} \qquad \frac{}{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} p : \tau \rightarrow o \mid \mathcal{X} > \mathcal{X}}$$

The rules for term and predicate variables update the marker for that variable to 1 , whatever it was before — we write $x^?$ for either x^0 or x^1 and similarly for predicate variables. The other rules thread the variables sets through. We arbitrarily proceed from left to right in the case of pairs.

Terms (rhs) $\boxed{\Gamma \vdash_{\Sigma} t : \tau \mid \mathcal{X}}$

$$\frac{}{\Gamma, x : \tau \vdash_{\Sigma} x : \tau \mid \mathcal{X}, x^1} \qquad \frac{\Gamma \vdash_{\Sigma, f: \tau \rightarrow \iota} t : \tau \mid \mathcal{X}}{\Gamma \vdash_{\Sigma, f: \tau \rightarrow \iota} f t : \iota \mid \mathcal{X}} \qquad \frac{\Gamma \vdash_{\Sigma} t_1 : \tau_1 \mid \mathcal{X} \quad \Gamma \vdash_{\Sigma} t_2 : \tau_2 \mid \mathcal{X}}{\Gamma \vdash_{\Sigma} (t_1, t_2) : \tau_1 \times \tau_2 \mid \mathcal{X}} \qquad \frac{}{\Gamma \vdash_{\Sigma} () : \top \mid \mathcal{X}}$$

$$\frac{}{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} X : \tau \rightarrow o \mid \mathcal{X}, X^1} \qquad \frac{}{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} p : \tau \rightarrow o \mid \mathcal{X}}$$

The right-hand side rules for term and predicate variables require that these variables have been seen earlier in a left-hand side, which is witnessed by the marker 1 . All other rules thread the constraint set through.

Atoms (lhs) $\boxed{\Gamma \vdash_{\Sigma} A \text{ atom} \mid \mathcal{X} > \mathcal{X}'}$

$$\frac{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} t : \tau \mid \mathcal{X} > \mathcal{X}'}{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} p t \text{ atom} \mid \mathcal{X} > \mathcal{X}'}$$

$$\frac{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} t : \tau \mid \mathcal{X}, X^1 > \mathcal{X}'}{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} X t \text{ atom} \mid \mathcal{X}, X^1 > \mathcal{X}'}$$

A predicate variable X introducing a parametric atom on the left-hand side must have been seen earlier within a left-hand side term, as indicated by the marker X^1 . Other variables are threaded through.

Atoms (rhs) $\boxed{\Gamma \vdash_{\Sigma} A \text{ atom} \mid \mathcal{X}}$

$$\frac{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} t : \tau \mid \mathcal{X}}{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o} p t \text{ atom} \mid \mathcal{X}} \qquad \frac{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} t : \tau \mid \mathcal{X}, X^1}{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} X t \text{ atom} \mid \mathcal{X}, X^1}$$

Like any right-hand side variable, a predicate variable introducing a parametric atom must have been seen previously within a left-hand side term.

Rule left-hand-sides $\boxed{\Gamma \vdash_{\Sigma} l \text{ lhs} \mid \mathcal{X} > \mathcal{X}'}$

$$\frac{}{\Gamma \vdash_{\Sigma} \cdot \text{lhs} \mid \mathcal{X} > \mathcal{X}} \qquad \frac{\Gamma \vdash_{\Sigma} A \text{ atom} \mid \mathcal{X} > \mathcal{X}' \quad \Gamma \vdash_{\Sigma} l \text{ lhs} \mid \mathcal{X}' > \mathcal{X}''}{\Gamma \vdash_{\Sigma} A, l \text{ lhs} \mid \mathcal{X} > \mathcal{X}''}$$

Constraint variables are refined (i.e., some markers ⁰ are upgraded to ¹) going from left to right in a left-hand side.

Rules $\boxed{\Gamma \vdash_{\Sigma} R \text{ rule} \mid \mathcal{X}}$

$$\frac{\Gamma \vdash_{\Sigma} l \text{ lhs} \mid \mathcal{X} > \mathcal{X}' \quad \Gamma \vdash_{\Sigma} P \text{ prog} \mid \mathcal{X}'}{\Gamma \vdash_{\Sigma} l \multimap P \text{ rule} \mid \mathcal{X}} \qquad \frac{\Gamma, x : \iota \vdash_{\Sigma} R \text{ rule} \mid \mathcal{X}, x^0}{\Gamma \vdash_{\Sigma} \forall x : \iota. R \text{ rule} \mid \mathcal{X}} \qquad \frac{\Gamma, X : \tau \rightarrow o \vdash_{\Sigma} R \text{ rule} \mid \mathcal{X}, X^0}{\Gamma \vdash_{\Sigma} \forall X : \tau \rightarrow o. R \text{ rule} \mid \mathcal{X}}$$

The typing judgment for rules is where a lot of the action takes place. Traversing a universal quantifier installs its variable in the constraint set in the premise with marker ⁰. A bare rule of the form $l \multimap P$ is handled by passing the current constraint set \mathcal{X} to the left-hand side l , which will refine it into the constraint set \mathcal{X}' which is what is passed to the program part P .

Programs $\boxed{\Gamma \vdash_{\Sigma} P \text{ prog} \mid \mathcal{X}}$

$$\frac{}{\Gamma \vdash_{\Sigma} \cdot \text{prog} \mid \mathcal{X}} \qquad \frac{\Gamma \vdash_{\Sigma} P_1 \text{ prog} \mid \mathcal{X} \quad \Gamma \vdash_{\Sigma} P_2 \text{ prog} \mid \mathcal{X}}{\Gamma \vdash_{\Sigma} P_1, P_2 \text{ prog} \mid \mathcal{X}} \qquad \frac{\Gamma \vdash_{\Sigma} A \text{ atom} \mid \mathcal{X}}{\Gamma \vdash_{\Sigma} A \text{ prog} \mid \mathcal{X}} \qquad \frac{\Gamma \vdash_{\Sigma} A \text{ atom} \mid \mathcal{X}}{\Gamma \vdash_{\Sigma} !A \text{ prog} \mid \mathcal{X}}$$

$$\frac{\Gamma \vdash_{\Sigma} R \text{ rule} \mid \mathcal{X}}{\Gamma \vdash_{\Sigma} R \text{ prog} \mid \mathcal{X}} \qquad \frac{\Gamma \vdash_{\Sigma} R \text{ rule} \mid \mathcal{X}}{\Gamma \vdash_{\Sigma} !R \text{ prog} \mid \mathcal{X}} \qquad \frac{\Gamma \vdash_{\Sigma, x: \tau \rightarrow \iota} P \text{ prog} \mid \mathcal{X}}{\Gamma \vdash_{\Sigma} \exists x : \tau \rightarrow \iota. P \text{ prog} \mid \mathcal{X}} \qquad \frac{\Gamma \vdash_{\Sigma, X: \tau \rightarrow o} P \text{ prog} \mid \mathcal{X}}{\Gamma \vdash_{\Sigma} \exists X : \tau \rightarrow o. P \text{ prog} \mid \mathcal{X}}$$

The typing rules for programs simply pass the constraint set in their conclusion to their premise. Note in particular that existential variables are ignored.

The typing judgment for archives and states remains unchanged, but the rules need to be upgraded to seed their constituents with an empty constraint set.

Archives $\boxed{\vdash_{\Sigma} \Omega \text{ archive}}$

$$\frac{}{\vdash_{\Sigma} \cdot \text{archive}} \qquad \frac{\vdash_{\Sigma} \Omega \text{ archive} \quad \cdot \vdash_{\Sigma} A \text{ atom} \mid \cdot}{\vdash_{\Sigma} \Omega, A \text{ archive}} \qquad \frac{\vdash_{\Sigma} \Omega \text{ archive} \quad \cdot \vdash_{\Sigma} R \text{ rule} \mid \cdot}{\vdash_{\Sigma} \Omega, R \text{ archive}}$$

States $\boxed{\vdash_{\Sigma} \langle \Omega ; \Pi \rangle \text{ state}}$

$$\frac{\vdash_{\Sigma} \Omega \text{ archive} \quad \cdot \vdash_{\Sigma} \Pi \text{ prog} \mid \cdot}{\vdash_{\Sigma} \langle \Omega ; \Pi \rangle \text{ state}}$$

A.4 First-Order Encoding

In this appendix, we formalize the encoding that transforms an $\mathcal{L}^{1.5}$ state into a state in \mathcal{L}^1 with the same typing and operational behavior. To do so, we need to extend this encoding to all entities of $\mathcal{L}^{1.5}$. Intuitively, we will map every universal second-order variable $X : \tau \rightarrow o$ to a first-order variables $x_X : \iota$ and replace each parametric atom $X t$ with the atom $p_\tau(x_X, t)$, where p_τ is a predicate name associated with type τ . Occurrences of X in a term will be translated to just x_X . Existential second-order variables X are treated slightly differently because existential first-order variables must have a type of the form $\tau \rightarrow \iota$. Therefore, we will associate to each X a first-order variable $x_X : \top \rightarrow \iota$ and proceed as in the universal case, except that wherever we had x_X we now have $x_X ()$.

The encoding below assumes that predicates occurring within terms are not themselves applied to terms, i.e., that they have the bare form p rather than $p t$. This will simplify the presentation. We can always rewrite a program that makes use of applied predicates by turning the application into a pair. For example, the term $f(p t)$ would be rewritten as $f(p, t)$. This transformation can be done automatically. A consequence of this simplification is that the type o cannot appear immediately underneath a product or as the antecedent of a function type, but only in the form $\tau \rightarrow o$.

As we encounter second-order variables in terms and atoms, we need to know whether they are universal or existential. We do so by maintaining a set \mathcal{X} which marks each variable declaration in scope as universal or existential. Its structure is defined by means of the following grammar:

$$\begin{array}{lcl} \text{Declaration sets } \mathcal{X} & ::= & \cdot \quad \text{empty set} \\ & | & \mathcal{X}, \forall X : \tau \rightarrow o \quad \text{universal declaration} \\ & | & \mathcal{X}, \exists X : \tau \rightarrow o \quad \text{existential declaration} \end{array}$$

We now describe the encoding $[O]_{\mathcal{X}}^\#$ of each syntactic entity O in $\mathcal{L}^{1.5}$, providing commentary as needed.

Terms $[t]_{\mathcal{X}}^\# = t'$

$$\left\{ \begin{array}{l} [x]_{\mathcal{X}}^\# = x \\ [f t]_{\mathcal{X}}^\# = f [t]_{\mathcal{X}}^\# \\ [(t_1, t_2)]_{\mathcal{X}}^\# = ([t_1]_{\mathcal{X}}^\#, [t_2]_{\mathcal{X}}^\#) \\ [()]_{\mathcal{X}}^\# = () \\ [X t]_{\mathcal{X}, \forall X : \tau \rightarrow o}^\# = x_X \\ [X t]_{\mathcal{X}, \exists X : \tau \rightarrow o}^\# = x_X () \\ [p]_{\mathcal{X}}^\# = x_p () \end{array} \right.$$

Terms are translated homomorphically except for second-order variables and predicate names. A variable X is mapped to x_X if it is universal and to $x_X ()$ if it is existential. Predicate names can occur in $\mathcal{L}^{1.5}$ terms and are treated as if they were introduced by an existential quantifier. In fact, we will later associate a function symbol $x_p : \top \rightarrow \iota$ to each predicate name p in the signature.

Atoms $[A]_{\mathcal{X}}^\# = A'$

$$\left\{ \begin{array}{l} [p t]_{\mathcal{X}}^\# = p [t]_{\mathcal{X}}^\# \\ [X t]_{\mathcal{X}, \forall X : \tau \rightarrow o}^\# = p_\tau(x_X, [t]_{\mathcal{X}, \forall X : \tau \rightarrow o}^\#) \\ [X t]_{\mathcal{X}, \exists X : \tau \rightarrow o}^\# = p_\tau(x_X (), [t]_{\mathcal{X}, \exists X : \tau \rightarrow o}^\#) \end{array} \right.$$

A second-order variable $X : \tau \rightarrow o$ heading a parametric predicate $X t$ is replaced by the predicate p_τ associated with the type τ — we will see how this is done momentarily. We remember X by passing either x_X or $x_X ()$ as the first argument of p_τ .

Rule left-hand sides $[l]_{\mathcal{X}}^{\#} = l'$

$$\begin{cases} [\cdot]_{\mathcal{X}}^{\#} = \cdot \\ [A, l]_{\mathcal{X}}^{\#} = [A]_{\mathcal{X}}^{\#}, [l]_{\mathcal{X}}^{\#} \end{cases}$$

Left-hand sides simply propagates the translation to their constituent atoms.

Rules $[R]_{\mathcal{X}}^{\#} = \langle \Sigma'; R' \rangle$

$$\begin{cases} [l \multimap P]_{\mathcal{X}}^{\#} = \langle \Sigma'; [l]_{\mathcal{X}}^{\#} \multimap P' \rangle & \text{where } [P]_{\mathcal{X}}^{\#} = \langle \Sigma'; P' \rangle \\ [\forall x : \iota. R]_{\mathcal{X}}^{\#} = \langle \Sigma'; \forall x : \iota. R' \rangle & \text{where } [R]_{\mathcal{X}}^{\#} = \langle \Sigma'; R' \rangle \\ [\forall X : \tau \rightarrow o. R]_{\mathcal{X}}^{\#} = \langle (\Sigma' \cup p_{\tau} : \iota \times \tau' \rightarrow o); \forall x_X : \iota. R' \rangle & \text{where } [R]_{\mathcal{X}, \forall X : \tau \rightarrow o}^{\#} = \langle \Sigma'; R' \rangle \end{cases}$$

Translating a rule diverges from the above pattern because we may need to account for new second-order variables. The interesting case is that of second-order universal rules $\forall X : \tau \rightarrow o. R$. The matrix R may make use of X , which means that while translating it we need to know that X was introduced universally. Therefore, we extend the declaration set with a universal marker $\forall X : \tau \rightarrow o$ while doing so. Now, if X heads a parametric predicate $X t$ inside R , it will get replaced by the predicate p_{τ} associated to τ . We remember which such predicates are needed by returning a signature Σ' as an additional output of the translation function. Because the translation of t may replace some predicate names and variables with terms, the type of p_{τ} is $\iota \times \tau' \rightarrow o$ where τ' is the translation of the type τ , defined later. Note that the predicate name p_{τ} is fixed. Therefore, were the translation of R to turn out another copy, we shall retain just one. We write $\Sigma \cup \Sigma'$ for the set union of signatures Σ and Σ' . It keep just one copy of any common declaration. This differs from the notation “ $\Sigma, p_{\tau} : \iota \times \tau' \rightarrow o$ ” which α -renames p_{τ} apart from Σ .

Programs $[P]_{\mathcal{X}}^{\#} = \langle \Sigma'; P' \rangle$

$$\begin{cases} [\cdot]_{\mathcal{X}}^{\#} = \langle \cdot; \cdot \rangle \\ [P_1, P_2]_{\mathcal{X}}^{\#} = \langle \Sigma'_1 \cup \Sigma'_2; (P'_1, P'_2) \rangle & \text{where } [P_1]_{\mathcal{X}}^{\#} = \langle \Sigma'_1; P'_1 \rangle \text{ and } [P_2]_{\mathcal{X}}^{\#} = \langle \Sigma'_2; P'_2 \rangle \\ [A]_{\mathcal{X}}^{\#} = \langle \cdot; [A]_{\mathcal{X}}^{\#} \rangle \\ [!A]_{\mathcal{X}}^{\#} = \langle \cdot; ![A]_{\mathcal{X}}^{\#} \rangle \\ [R]_{\mathcal{X}}^{\#} = [R]_{\mathcal{X}}^{\#} \\ [!R]_{\mathcal{X}}^{\#} = \langle \Sigma'; !R' \rangle & \text{where } [R]_{\mathcal{X}}^{\#} = \langle \Sigma'; R' \rangle \\ [\exists x : \tau \rightarrow \iota. P]_{\mathcal{X}}^{\#} = \langle \Sigma'; \exists x : \tau \rightarrow \iota. P' \rangle & \text{where } [P]_{\mathcal{X}}^{\#} = \langle \Sigma'; P' \rangle \\ [\exists X : \tau \rightarrow o. P]_{\mathcal{X}}^{\#} = \langle (\Sigma' \cup p_{\tau} : \iota \times \tau' \rightarrow o); \exists x_X : \top \rightarrow \iota. P' \rangle & \text{where } [P]_{\mathcal{X}, \exists X : \tau \rightarrow o}^{\#} = \langle \Sigma'; P' \rangle \end{cases}$$

The translation of programs follows the template we just saw for rules. This time, the variables are existential.

We now give a translation for types, which is used in the encoding of rules and programs as we saw. We also give translations for signatures, contexts and archives.

Types $\tau^{\#} = \tau'$

$$\begin{cases} \iota^{\#} = \iota \\ o^{\#} = \iota & \text{(unused)} \\ (\tau_1 \rightarrow \iota)^{\#} = \tau_1^{\#} \rightarrow \iota \\ (\tau_1 \rightarrow o)^{\#} = \top \rightarrow \iota \\ (\tau_1 \rightarrow \tau_2)^{\#} = \iota & \text{(unused for } \tau \neq \iota \text{ and } \tau \neq o) \\ (\tau_1 \times \tau_2)^{\#} = \tau_1^{\#} \times \tau_2^{\#} \\ \top^{\#} = \top \end{cases}$$

The type τ translated by τ^\sharp will always occur on the left-hand side of function type $\tau \rightarrow \tau'$. In fact τ' can be only ι or o in $\mathcal{L}^{1.5}$. Because we restricted predicate name occurrences within terms to be bare, the type o can only appear as the target of a function type, of the form $\tau' \rightarrow o$. In particular isolated occurrences of o will never be encountered. We arbitrarily map both unused forms to ι .

Signatures $\boxed{[\Sigma]^\sharp = \Sigma'}$

$$\left\{ \begin{array}{l} [\cdot]^\sharp = \cdot \\ [\Sigma, f : \tau \rightarrow \iota]^\sharp = [\Sigma]^\sharp, f : \tau^\sharp \rightarrow \iota \\ [\Sigma, p : \tau \rightarrow o]^\sharp = [\Sigma]^\sharp, p : \tau^\sharp \rightarrow o, x_p : \top \rightarrow \iota \end{array} \right.$$

The translation of signatures bears much in common with what we did for existential variables. Specifically, we associate to each predicate name $p : \tau \rightarrow o$ not only its natural encoding $p : \tau^\sharp \rightarrow o$ but also the function symbol $x_p : \top \rightarrow \iota$ in case this predicate name is used in a term (see our encoding of terms above).

Contexts $\boxed{[\Gamma]^\sharp = \langle \Sigma'; \Gamma' \rangle}$

$$\left\{ \begin{array}{l} [\cdot]^\sharp = \langle \cdot; \cdot \rangle \\ [\Gamma, x : \iota]^\sharp = \langle \Sigma'; \Gamma', x : \iota \rangle \\ [\Gamma, X : \tau \rightarrow o]^\sharp = \langle (\Sigma' \cup p_\tau : \iota \times \tau^\sharp \rightarrow o); \Gamma', x_X : \iota \rangle \end{array} \right. \quad \begin{array}{l} \text{where } [\Gamma]^\sharp = \langle \Sigma'; \Gamma' \rangle \\ \text{where } [\Gamma]^\sharp = \langle \Sigma'; \Gamma' \rangle \end{array}$$

We encode a predicate variable $X : \tau \rightarrow o$ in a context into a first-order variable x_X of type ι , but also define the signature declaration for p_τ , which will account for parametric atoms $X t$ that may make use of X . Therefore, the translation of a context returns both a signature and another context.

Archives $\boxed{[\Omega]^\sharp = \langle \Sigma'; \Omega' \rangle}$

$$\left\{ \begin{array}{l} [\cdot]^\sharp = \langle \cdot; \cdot \rangle \\ [\Omega, A]^\sharp = \langle \Sigma'; (\Omega', A^\sharp) \rangle \\ [\Omega, R]^\sharp = \langle \Sigma' \cup \Sigma''; (\Omega', R') \rangle \end{array} \right. \quad \begin{array}{l} \text{where } [\Omega]^\sharp = \langle \Sigma'; \Omega' \rangle \\ \text{where } [\Omega]^\sharp = \langle \Sigma'; \Omega' \rangle \text{ and } [R]^\sharp = \langle \Sigma''; R' \rangle \end{array}$$

A archive is encoded by collecting the encoding of each constituent. Since the encoding of rules output a signature, so does the encoding of archives.

The above translation maps a well-typed entity in $\mathcal{L}^{1.5}$ to a similarly well-typed object in \mathcal{L}^1 . This is formally expressed by the following property.

Lemma 5 For $\mathcal{L}^{1.5}$ signature Σ and context Γ , let $[\Sigma]^\sharp = \Sigma'$ and $[\Gamma]^\sharp = \langle \Sigma''; \Gamma' \rangle$.

1. If $\Gamma \vdash_\Sigma t : \tau$ in $\mathcal{L}^{1.5}$, then $\Gamma' \vdash_{\Sigma' \cup \Sigma''} [t : \tau^\sharp]^\sharp$ in \mathcal{L}^1 .
2. If $\Gamma \vdash_\Sigma A$ atom in $\mathcal{L}^{1.5}$, then $\Gamma' \vdash_{\Sigma' \cup \Sigma''} [A]^\sharp$ atom in \mathcal{L}^1 .
3. If $\Gamma \vdash_\Sigma l$ lhs in $\mathcal{L}^{1.5}$, then $\Gamma' \vdash_{\Sigma' \cup \Sigma''} [l]^\sharp$ lhs in \mathcal{L}^1 .
4. If $\Gamma \vdash_\Sigma R$ rule in $\mathcal{L}^{1.5}$, then $\Gamma' \vdash_{\Sigma' \cup \Sigma'' \cup \Sigma'''} R'$ rule in \mathcal{L}^1 where $[R]^\sharp = \langle \Sigma'''; R' \rangle$.
5. If $\Gamma \vdash_\Sigma P$ prog in $\mathcal{L}^{1.5}$, then $\Gamma' \vdash_{\Sigma' \cup \Sigma'' \cup \Sigma'''} P'$ prog in \mathcal{L}^1 where $[P]^\sharp = \langle \Sigma'''; P' \rangle$.
6. If $\vdash_\Sigma \Omega$ archive in $\mathcal{L}^{1.5}$, then $\vdash_{\Sigma' \cup \Sigma''} \Omega'$ archive in \mathcal{L}^1 where $[\Omega]^\sharp = \langle \Sigma'''; \Omega' \rangle$.

7. If $\vdash \Sigma.\langle \Omega ; \Pi \rangle$ state in $\mathcal{L}^{1.5}$, then $\vdash (\Sigma' \cup \Sigma''' \cup \Sigma'''').\langle \Omega' ; \Pi' \rangle$ state in \mathcal{L}^1 where $[\Omega]^\sharp = \langle \Sigma'''; \Omega' \rangle$ and $[\Pi]^\sharp = \langle \Sigma''''; \Pi' \rangle$.

Proof The proof proceeds by mutual induction on the given typing derivations. □

The encoding also maps transitions in $\mathcal{L}^{1.5}$ to transitions in \mathcal{L}^1 .

Lemma 6 For $\mathcal{L}^{1.5}$ state $\Sigma_A.\langle \Omega_A ; \Pi_A \rangle$, let $[\Sigma_A]^\sharp = \Sigma'_A$ and $[\Omega_A]^\sharp = \langle \Sigma''_A; \Omega'_A \rangle$ and $[\Pi_A]^\sharp = \langle \Sigma'''_A; \Pi'_A \rangle$.

Similarly, for $\mathcal{L}^{1.5}$ state $\Sigma_B.\langle \Omega_B ; \Pi_B \rangle$, let $[\Sigma_B]^\sharp = \Sigma'_B$ and $[\Omega_B]^\sharp = \langle \Sigma''_B; \Omega'_B \rangle$ and $[\Pi_B]^\sharp = \langle \Sigma'''_B; \Pi'_B \rangle$.

Let $\Sigma_A^* = \Sigma'_A \cup \Sigma''_A \cup \Sigma'''_A$ and $\Sigma_B^* = \Sigma'_B \cup \Sigma''_B \cup \Sigma'''_B$.

If $\vdash \Sigma_A.\langle \Omega_A ; \Pi_A \rangle$ state and $\Sigma_A.\langle \Omega_A ; \Pi_A \rangle \mapsto \Sigma_B.\langle \Omega_B ; \Pi_B \rangle$ in $\mathcal{L}^{1.5}$, then $\Sigma_A^*.\langle \Omega'_A ; \Pi'_A \rangle \mapsto \Sigma_B^*.\langle \Omega'_B ; \Pi'_B \rangle$ in \mathcal{L}^1 .

Proof The proof proceeds by cases on the given $\mathcal{L}^{1.5}$ transitions. □

A.5 Typechecking Modular Programs

This appendix defines the extension \mathcal{L}^M of the language $\mathcal{L}^{1.5}$ with the module infrastructure devised in Section 3. We describe both the syntax of \mathcal{L}^M and its extended typing semantics, inclusive of mode declarations and checks. Extensions are shown in blue. Differently from Section 4 and Appendix A.4, we do not compile \mathcal{L}^M into $\mathcal{L}^{1.5}$ but handle it head on. Furthermore, we allow the programmer to mix and match the module syntax that characterizes \mathcal{L}^M with more basic constructs of $\mathcal{L}^{1.5}$. For example, she will be able to define a module using the **module** syntax but use it without relying on **as**, and vice versa. We conclude by summarizing the elaboration of \mathcal{L}^M into $\mathcal{L}^{1.5}$.

The basic syntactic categories of \mathcal{L}^M extend those of $\mathcal{L}^{1.5}$ with module reference names N . The overall lexicon is as follows, where we make the same provisos as for $\mathcal{L}^{1.5}$ in Appendix A.1.

| | |
|----------------------------|-----|
| <i>Term variables</i> | x |
| <i>Predicate variables</i> | X |
| <i>Function symbols</i> | f |
| <i>Predicate symbols</i> | p |
| <i>Module references</i> | N |

The other entities of \mathcal{L}^M are defined by the grammar in Figure A.2, where all modifications (not just extensions) are highlighted in blue. These modifications are as follows:

- The type of state objects has been refined into the form o^ξ to incorporate mode ascriptions for predicate names, as used in the **provide** stanza of the examples in Section 3. The type o^l refers to the predicate names whose atoms can only appear in the left-hand side of a rule (except at their definition point) — that’s the marker **in**, o^r stands for predicate names that can occur only on the right-hand side (again, except when being defined) — that’s the marker **out**, and o indicates predicate names that can occur anywhere — they correspond to the type o of $\mathcal{L}^{1.5}$. We allow mode restrictions on any predicate name, not just the symbols exported by a module.

Given an atom type o^ξ , we write $\overline{o^\xi}$ for the type that flips the restrictions on use. Specifically,

$$\begin{cases} \overline{o^l} & = & o^r \\ \overline{o^r} & = & o^l \\ \overline{o} & = & o \end{cases}$$

This is useful as module definitions provide predicates to be used on the left-hand side in client code by producing them on their right-hand side, and vice versa.

- A predicate variable V is either one of the predicate variable X of $\mathcal{L}^{1.5}$ or a compound predicate variable $N.X$ where N is a module reference. Either form can be existentially or universally quantified. Similarly, existential term variables (and later term-level signature declaration) can be either simple or compound names.
- An interface Υ is a multiset of declarations for function and predicate symbols. It is what is exported by a module. It also corresponds exactly to the signatures of $\mathcal{L}^{1.5}$. Although they did not appear in any of the examples in Section 3, term-level symbols can be exported by a module. As we will see, interfaces in \mathcal{L}^M are richer.
- Programs are extended with the construct for using a module and the construct for defining one.

The resulting grammar is given in Figure A.2.

We proceed by giving typing rules for \mathcal{L}^M , as just defined. We update the definition of signature to allow compound term and predicate names and the definition of context to allow moded, possibly compound, predicate declarations.

Compound term variables cannot enter a context.

| | | | |
|-------------------|--------------|---|----------------------------------|
| <i>Signatures</i> | $\Sigma ::=$ | \cdot | empty signature |
| | | $ \ \Sigma, f : \tau \rightarrow \iota$ | constructor declaration |
| | | $ \ \Sigma, N.f : \tau \rightarrow \iota$ | compound constructor declaration |
| | | $ \ \Sigma, p : \tau \rightarrow o^\xi$ | predicate declaration |
| | | $ \ \Sigma, N.p : \tau \rightarrow o^\xi$ | compound predicate declaration |
| <i>Contexts</i> | $\Gamma ::=$ | \cdot | empty context |
| | | $ \ \Gamma, x : \iota$ | term assumption |
| | | $ \ \Gamma, V : \tau \rightarrow o^\xi$ | predicate assumption |

The definition of archives and states remains unchanged.

Two extensions to the typing semantics of $\mathcal{L}^{1.5}$ are required. First we need to check that predicate variables are used according to their mode declaration. This entails providing different rules for atoms occurring on the left- and right-hand side of a rule. Second, we need to keep track of the modules available for instantiation. To do so, we define *module sets* M , which associate the name and interface of each module in scope.

| | | | |
|-------------------|---------|----------------------|------------------|
| <i>Module set</i> | $M ::=$ | \cdot | empty module set |
| | | $ \ M, p : \Upsilon$ | available module |

The judgments for terms and states remain unchanged. We update the judgments for atoms, rules and programs as follows:

| | | |
|---------------------------------|------------------------|-------------------------|
| $\Gamma \vdash_\Sigma t : \tau$ | <i>lhs atom</i> | Atoms (left-hand side) |
| $\Gamma \vdash_\Sigma A$ | <i>rhs atom</i> | Atoms (right-hand side) |
| $\Gamma \vdash_\Sigma R$ | <i>rule</i> $ M$ | Rules |
| $\Gamma \vdash_\Sigma P$ | <i>prog</i> $ M > M'$ | Programs |

When analyzing a program P , the set M represents known modules before examining P while M' is the extended module set after having gone through P . Scope considerations prevent a rule from exporting more modules than it knows about.

We now give the updated typing rules, highlighting the changes in blue. We comment on changes as required.

Terms $\Gamma \vdash_\Sigma t : \tau$

| | | |
|--|--|--|
| $\frac{}{\Gamma, x : \tau \vdash_\Sigma x : \tau}$ | $\frac{\Gamma \vdash_{\Sigma, f : \tau \rightarrow \iota} t : \tau}{\Gamma \vdash_{\Sigma, f : \tau \rightarrow \iota} f t : \iota}$ | $\frac{\Gamma \vdash_{\Sigma, N.f : \tau \rightarrow \iota} t : \tau}{\Gamma \vdash_{\Sigma, N.f : \tau \rightarrow \iota} N.f t : \iota}$ |
| $\frac{\Gamma \vdash_\Sigma t_1 : \tau_1 \quad \Gamma \vdash_\Sigma t_2 : \tau_2}{\Gamma \vdash_\Sigma (t_1, t_2) : \tau_1 \times \tau_2}$ | $\frac{}{\Gamma \vdash_\Sigma () : \top}$ | $\frac{}{\Gamma, V : \tau \rightarrow o^\xi \vdash_\Sigma V : \tau \rightarrow o^\xi}$ |
| | | $\frac{}{\Gamma \vdash_{\Sigma, p : \tau \rightarrow o^\xi} p : \tau \rightarrow o^\xi}$ |

The rules for terms change only in as much as predicate variables may use compound names and carry modes, and that we could have compound term constructors in the signature.

Atoms (lhs) $\Gamma \vdash_\Sigma A$ *lhs atom*

| | |
|---|---|
| $\frac{\Gamma \vdash_{\Sigma, p : \tau \rightarrow o^{l?}} t : \tau}{\Gamma \vdash_{\Sigma, p : \tau \rightarrow o^{l?}} p t$ <i>lhs atom</i> | $\frac{\Gamma, V : \tau \rightarrow o^{l?} \vdash_\Sigma t : \tau}{\Gamma, V : \tau \rightarrow o^{l?} \vdash_\Sigma V t$ <i>lhs atom</i> |
|---|---|

Here, we write $o^{l?}$ for either type o or o^l . Predicate names in a left-hand side atom can be either unrestricted (o) or defined for left-hand side use only (o^l).

Atoms (rhs) $\boxed{\Gamma \vdash_{\Sigma} A \text{ rhs atom}}$

$$\frac{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o^{r?}} t : \tau}{\Gamma \vdash_{\Sigma, p: \tau \rightarrow o^{r?}} p t \text{ rhs atom}} \quad \frac{\Gamma, V : \tau \rightarrow o^{r?} \vdash_{\Sigma} t : \tau}{\Gamma, V : \tau \rightarrow o^{r?} \vdash_{\Sigma} V t \text{ rhs atom}}$$

Similar considerations apply to right-hand side atoms. Here, $o^{r?}$ stands for either o or o^r .

Rule left-hand-sides $\boxed{\Gamma \vdash_{\Sigma} l \text{ lhs}}$

$$\frac{}{\Gamma \vdash_{\Sigma} \cdot \text{lhs}} \quad \frac{\Gamma \vdash_{\Sigma} A \text{ lhs atom} \quad \Gamma \vdash_{\Sigma} l \text{ lhs}}{\Gamma \vdash_{\Sigma} A, l \text{ lhs}}$$

Rule left-hand sides are updated only in so far as their constituents are sent to the left-hand side atom judgment.

Rules $\boxed{\Gamma \vdash_{\Sigma} R \text{ rule} \mid M}$

$$\frac{\Gamma \vdash_{\Sigma} l \text{ lhs} \quad \Gamma \vdash_{\Sigma} P \text{ prog} \mid M > M'}{\Gamma \vdash_{\Sigma} l \multimap P \text{ rule} \mid M} \quad \frac{\Gamma, x : \iota \vdash_{\Sigma} R \text{ rule} \mid M}{\Gamma \vdash_{\Sigma} \forall x : \iota. R \text{ rule} \mid M} \quad \frac{\Gamma, V : \tau \rightarrow o^{\xi} \vdash_{\Sigma} R \text{ rule} \mid M}{\Gamma \vdash_{\Sigma} \forall V : \tau \rightarrow o^{\xi}. R \text{ rule} \mid M}$$

Rules are where we first encounter module sets. All they do is to propagate them to the program on the right-hand side P of a rewriting directive $l \multimap P$. Note that the module set M' that comes back is discarded as the modules defined in P are not in scope outside of P itself.

Programs $\boxed{\Gamma \vdash_{\Sigma} P \text{ prog} \mid M > M'}$

$$\frac{}{\Gamma \vdash_{\Sigma} \cdot \text{prog} \mid M > M} \quad \frac{\Gamma \vdash_{\Sigma} P_1 \text{ prog} \mid M > M' \quad \Gamma \vdash_{\Sigma} P_2 \text{ prog} \mid M' > M''}{\Gamma \vdash_{\Sigma} P_1, P_2 \text{ prog} \mid M > M''}$$

$$\frac{\Gamma \vdash_{\Sigma} A \text{ rhs atom}}{\Gamma \vdash_{\Sigma} A \text{ prog} \mid M > M} \quad \frac{\Gamma \vdash_{\Sigma} A \text{ rhs atom}}{\Gamma \vdash_{\Sigma} !A \text{ prog} \mid M > M} \quad \frac{\Gamma \vdash_{\Sigma} R \text{ rule} \mid M}{\Gamma \vdash_{\Sigma} R \text{ prog} \mid M > M} \quad \frac{\Gamma \vdash_{\Sigma} R \text{ rule} \mid M}{\Gamma \vdash_{\Sigma} !R \text{ prog} \mid M > M}$$

$$\frac{\Gamma \vdash_{\Sigma, v: \tau \rightarrow \iota} P \text{ prog} \mid M > M'}{\Gamma \vdash_{\Sigma} \exists v : \tau \rightarrow \iota. P \text{ prog} \mid M > M'} \quad \frac{\Gamma \vdash_{\Sigma, V: \tau \rightarrow o^{\xi}} P \text{ prog} \mid M > M'}{\Gamma \vdash_{\Sigma} \exists V : \tau \rightarrow o^{\xi}. P \text{ prog} \mid M > M'}$$

$$\frac{\Gamma \vdash_{\Sigma} p t \text{ lhs atom} \quad \Gamma \vdash_{\Sigma, N, \Upsilon} P \text{ prog} \mid M, p : \Upsilon > M'}{\Gamma \vdash_{\Sigma} N \text{ as } p t. P \text{ prog} \mid M, p : \Upsilon > M'}$$

$$\frac{p : \Upsilon_{par} \times \Upsilon_{export} \rightarrow o^r \text{ in } \Sigma \quad p \text{ not in } M \quad \Gamma, \Upsilon_{par}, \overline{\Upsilon_{export}} \vdash_{\Sigma, \Upsilon_{local}} P \text{ prog} \mid M, p : \Upsilon_{export} > M'}{\Gamma \vdash_{\Sigma} \text{module } p(\Upsilon_{par}) \text{ provide } \Upsilon_{export} \text{ local } \Upsilon_{local} P \text{ end prog} \mid M > M'}$$

The typing rules for programs are where most of the action takes place. Notice first that the union P_1, P_2 of two programs accumulates modules from left to right. As we understand “,” as being commutative and associative, this does not force a module to be defined before being used in a strictly lexical sense. Because some ordering of a program constituents must be picked, it however prevents circularity. Naturally, the empty program and right-hand side atoms simply return the module set they were handed. The updated version of the remaining rules of $\mathcal{L}^{1.5}$ simply propagate the sets to and from their premises.

The rules for the new constructs are more interesting. The rule for a module definition

$$\text{module } p(\Upsilon_{par}) \text{ provide } \Upsilon_{export} \text{ local } \Upsilon_{local} P \text{ end}$$

first checks that the predicate name p (the name of the module) is defined in the signature Σ and that it takes as input both the parameters Υ_{par} and the module's interface Υ_{export} . Abusing notation, we write $\Upsilon_{par} \times \Upsilon_{export}$ for the overall product type — for notational simplicity we keep the parameters on the left and the exported symbols on the right although there is no reason for this. Since a hand-crafted use of this module (i.e., without relying on `as`) would make use of it on the right-hand side of a rule, we give p mode o^r as using it on the left-hand side of a rule (other than the above module definition) would have the effect of intercepting the instantiation request. Therefore, we require that p appears in the signature Σ with type $\Upsilon_{par} \times \Upsilon_{export} \rightarrow o^r$. The second premise ensures that module p is defined no more than once.

The last premise of this rule extends the context with both the parameter declarations Υ_{par} and the inverted exported declarations Υ_{export} , and the signature with the local declarations Υ_{local} . We invert the modes of the exported declarations as the module uses these symbols in the opposite way as a client. The module set that can be used inside P is extended with $p : \Upsilon_{export}$, thereby allowing recursive instantiations. Note that this extension will be returned through M' .

The rule for the instantiation $N \text{ as } p \text{ t. } P$ of a module p with interface Υ , as reported in the incoming module set, starts by checking the validity of $p \text{ t}$ as a right-hand side atom. It then extends the signature in which to typecheck P with the exported symbols Υ prefixed by the module reference name N . This is expressed by the notation $N.\Upsilon$.

Archives $\boxed{\vdash_{\Sigma} \Omega \text{ archive } | M}$

$$\frac{}{\vdash_{\Sigma} \cdot \text{archive } | M} \quad \frac{\vdash_{\Sigma} \Omega \text{ archive } | M \quad \cdot \vdash_{\Sigma} A \text{ rhs atom}}{\vdash_{\Sigma} \Omega, A \text{ archive } | M} \quad \frac{\vdash_{\Sigma} \Omega \text{ archive } | M \quad \cdot \vdash_{\Sigma} R \text{ rule } | M}{\vdash_{\Sigma} \Omega, R \text{ archive } | M}$$

The rules for archives pass the module set they are given to any contained \mathcal{L}^M rule. As we will see next when examining states, this module set will be empty in this report.

States $\boxed{\vdash_{\Sigma} \langle \Omega ; \Pi \rangle \text{ state}}$

$$\frac{\vdash_{\Sigma} \Omega \text{ archive } | \cdot \quad \cdot \vdash_{\Sigma} \Pi \text{ prog } | \cdot > M'}{\vdash_{\Sigma} \langle \Omega ; \Pi \rangle \text{ state}}$$

The rule for states seeds the embedded archive Ω and program Π with an empty module set and discards what the analysis of Π reports. Note that, were an implementation to provide system modules, this is where they would be plugged in.

A program in \mathcal{L}^M is elaborated into an $\mathcal{L}^{1.5}$ program before execution. This means that we can rely on the operational semantics for this language, as discussed in Section 2.1.3, for running \mathcal{L}^M programs, without a need to account for any of the modular infrastructure. The elaboration, adapted from Section 4, is defined as follows:

- A module declaration

```

module  $p(\Upsilon_{par})$ 
  provide  $\Upsilon_{export}$ 
  local  $\Upsilon_{local}$ 
   $P$ 
end

```

is compiled into the rule

$$\forall \Upsilon_{par}. \forall \Upsilon_{export}. p(\Upsilon_{par}^*, \Upsilon_{export}^*) \multimap \exists \Upsilon_{local}. P$$

where the predicate name p is called on the names declared in Υ_{par} and Υ_{export} , dropping the types, something we express as $p(\Upsilon_{par}^*, \Upsilon_{export}^*)$.

- Let p be a module with interface Υ , which we write as $\vec{X} : \vec{\tau}$ for convenience — note that \vec{X} may contain term variables as well as predicate variables. Then, a module instantiation

$$N \text{ as } p \text{ t. } P$$

is compiled into the program

$$\exists \Upsilon. \left[\begin{array}{l} p(t, \vec{X}), \\ [\vec{X}/N.\vec{X}]/P \end{array} \right]$$

where $[\vec{X}/N.\vec{X}]/P$ replaces each occurrence of the compound name $N.X_i$ in P with X_i , and similarly for term variables. We implicitly rename variables as required to avoid capture.

Starting with a well-typed \mathcal{L}^M program without compound predicate names in its signature, this transformation results in a valid $\mathcal{L}^{1.5}$ program without compound predicate names at all. Mode annotations are dropped altogether.

| | | | | |
|----------------------------|------------|-------|---|--|
| <i>Atom types</i> | o^ξ | $::=$ | o o^l o^r | unrestricted atom lhs-only atom rhs-only atom |
| <i>Types</i> | τ | $::=$ | ι o^ξ $\tau \rightarrow \tau$ $\tau \times \tau$ \top | individuals state objects function type product type unit type |
| <i>Interface</i> | Υ | $::=$ | \cdot $\Upsilon, f : \tau \rightarrow \iota$ $\Upsilon, p : \tau \rightarrow o^\xi$ | empty interface function symbol predicate symbol |
| <i>Term variables</i> | v | $::=$ | x $N.x$ | simple term variable compound term variable |
| <i>Predicate variables</i> | V | $::=$ | X $N.X$ | simple predicate variable compound predicate variable |
| <i>Terms</i> | t | $::=$ | v $f t$ (t, t) $()$ V p | term variables term constructors products unit predicate variables predicate names |
| <i>Atoms</i> | A | $::=$ | $p t$ $V t$ | atoms parametric atoms |
| <i>LHS</i> | l | $::=$ | \cdot A, l | empty left-hand side left-hand side extension |
| <i>Rules</i> | R | $::=$ | $l \multimap P$ $\forall x : \iota. R$ $\forall V : \tau \rightarrow o^\xi. R$ | rewrite term abstraction predicate abstraction |
| <i>Programs</i> | P | $::=$ | \cdot P, P A $!A$ R $!R$ $\exists v : \tau \rightarrow \iota. P$ $\exists V : \tau \rightarrow o^\xi. P$ $N \text{ as } p t. P$ module $p(\Upsilon)$ provide Υ local Υ P end | empty program program union atom reusable atom rule reusable rule new term variable new predicate variable module instantiation module definition |

Figure A.2: Syntax of \mathcal{L}^M

A.6 Congruence

Throughout this report, we implicitly viewed the operator “;” used to combine left-hand side atoms and (right-hand side) programs as associative with unit “.” and commutative. This is logically justified as the corresponding operators in linear logic, \otimes and $\mathbf{1}$, have isomorphic properties relative to derivability — see Appendix B.

For the sake of completeness, we spell out these properties in the case of left-hand sides:

$$\begin{aligned} l_1, (l_2, l_3) &= (l_1, l_2), l_3 && \text{associativity} \\ l, \cdot &= l && \text{unit} \\ l_1, l_2 &= l_2, l_1 && \text{commutativity} \end{aligned}$$

Analogous properties apply to programs. These properties form an equivalence relation. Moreover, they apply arbitrarily deep within a term and therefore define a congruence.

We also wrote “;” for the operation of extending (and joining) signatures, contexts and state components and “.” for the empty signature, context, archive and state. Also in these case, “;” is commutative and associative with unit “.” — this corresponds directly to the often-used conventions for the homonymous formation operators in logic.

The use of “;” as a pairing operator in terms is neither associative nor commutative.

A.7 Unfocused Rewriting Semantics

The full set of transitions in the unfocused evaluation semantics is as follows.

$$\begin{array}{l}
 \text{Transitions } \boxed{\Sigma.\langle\Omega ; \Pi\rangle \mapsto \Sigma'.\langle\Omega' ; \Pi'\rangle} \\
 \\
 \Sigma.\langle\Omega, l_1 ; \Pi, l_2, (l_1, l_2) \multimap P\rangle \mapsto \Sigma.\langle\Omega, l_1 ; \Pi, P\rangle \\
 \Sigma.\langle\Omega ; \Pi, \forall x : \iota. R\rangle \mapsto \Sigma.\langle\Omega ; \Pi, [t/x]R\rangle \quad \text{if } \cdot \vdash_{\Sigma} t : \iota \\
 \Sigma.\langle\Omega ; \Pi, \forall X : \tau \rightarrow o. R\rangle \mapsto \Sigma.\langle\Omega ; \Pi, [p/X]R\rangle \quad \text{if } p : \tau \rightarrow o \text{ in } \Sigma \\
 \\
 \Sigma.\langle\Omega ; \Pi, !A\rangle \mapsto \Sigma.\langle\Omega, A ; \Pi\rangle \\
 \Sigma.\langle\Omega ; \Pi, !R\rangle \mapsto \Sigma.\langle\Omega, R ; \Pi\rangle \\
 \Sigma.\langle\Omega ; \Pi, \exists x : \tau \rightarrow \iota. P\rangle \mapsto (\Sigma, x : \tau \rightarrow \iota).\langle\Omega ; \Pi, P\rangle \\
 \Sigma.\langle\Omega ; \Pi, \exists X : \tau \rightarrow o. P\rangle \mapsto (\Sigma, X : \tau \rightarrow o).\langle\Omega ; \Pi, P\rangle \\
 \\
 \Sigma.\langle\Omega, R ; \Pi\rangle \mapsto \Sigma.\langle\Omega, R ; \Pi, R\rangle
 \end{array}$$

The first three transitions refer to the rules of $\mathcal{L}^{1.5}$ according to the definition in Appendix A.1. The next four handles the remaining program forms. The last rule transfers a copy of an archived rule to the program part of the state for further processing. Archived atoms are used in the first rule, which processes rewriting directives. As we will see, each transition corresponds to the left sequent rule of the corresponding logical operator, as given in Appendix B.1. Observe that we are implicitly making use of the congruences in Appendix A.6 in the way we organize the starting state of each transition.

With the partial exception of $!$ which we decomposed over several rules, these transitions view each operator in the language as a local state transformation directive. We can make this correspondence tighter by doing away with the archive component of a state and merging the rules that process $!$ with the rules that use the resulting program components. The following presentation does precisely this:

$$\begin{array}{l}
 \text{Transitions } \boxed{\Sigma.\langle\Pi\rangle \mapsto \Sigma'.\langle\Pi'\rangle} \\
 \\
 \Sigma.\langle\Pi, !(l_1), l_2, (l_1, l_2) \multimap P\rangle \mapsto \Sigma.\langle\Pi, !(l_1), P\rangle \\
 \Sigma.\langle\Pi, \forall x : \iota. R\rangle \mapsto \Sigma.\langle\Pi, [t/x]R\rangle \quad \text{if } \cdot \vdash_{\Sigma} t : \iota \\
 \Sigma.\langle\Pi, \forall X : \tau \rightarrow o. R\rangle \mapsto \Sigma.\langle\Pi, [p/X]R\rangle \quad \text{if } p : \tau \rightarrow o \text{ in } \Sigma \\
 \\
 \Sigma.\langle\Pi, !R\rangle \mapsto \Sigma.\langle\Pi, !R, R\rangle \\
 \Sigma.\langle\Pi, \exists x : \tau \rightarrow \iota. P\rangle \mapsto (\Sigma, x : \tau \rightarrow \iota).\langle\Pi, P\rangle \\
 \Sigma.\langle\Pi, \exists X : \tau \rightarrow o. P\rangle \mapsto (\Sigma, X : \tau \rightarrow o).\langle\Pi, P\rangle
 \end{array}$$

where $!(l)$ denotes a portion of the state that consists of zero or more reusable atoms. This presentation is equivalent to the transitions at the beginning of this section and it too has its roots in linear logic. It is however somewhat harder to work with. We will not pursue this presentation further.

A.8 Focused Rewriting Semantics

The focused operational semantics relies on the notion of a stable state $\Sigma.\langle\Omega ; \Pi\rangle$, whose program component Π is defined as follows:

$$\begin{array}{lcl} \text{Stable state program } \Pi & ::= & \cdot \quad \text{empty state} \\ & | & \Pi, A \quad \text{atom} \\ & | & \Pi, R \quad \text{rule} \end{array}$$

The full set of focused transitions is as given next.

Rule transitions $\boxed{\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$

$$\begin{array}{l} \Sigma.\langle\Omega, l_1\theta ; \Pi, l_2\theta, \forall(l_1, l_2 \multimap P)\rangle \Rightarrow \Sigma.\langle\Omega, l_1\theta ; \Pi, P\theta\rangle \\ \Sigma.\langle\Omega, l_1\theta, \forall(l_1, l_2 \multimap P) ; \Pi, l_2\theta\rangle \Rightarrow \Sigma.\langle\Omega, l_1\theta, \forall(l_1, l_2 \multimap P) ; \Pi, P\theta\rangle \end{array}$$

Stabilizing transitions $\boxed{\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega ; \Pi'\rangle}$

$$\begin{array}{l} \Sigma.\langle\Omega ; \Pi, !A\rangle \Rightarrow \Sigma.\langle\Omega, A ; \Pi\rangle \\ \Sigma.\langle\Omega ; \Pi, !R\rangle \Rightarrow \Sigma.\langle\Omega, R ; \Pi\rangle \\ \Sigma.\langle\Omega ; \Pi, \exists x : \tau \rightarrow \iota. P\rangle \Rightarrow (\Sigma, x : \tau \rightarrow \iota).\langle\Omega ; \Pi, P\rangle \\ \Sigma.\langle\Omega ; \Pi, \exists X : \tau \rightarrow o. P\rangle \Rightarrow (\Sigma, X : \tau \rightarrow o).\langle\Omega ; \Pi, P\rangle \end{array}$$

As we will see in Appendix B.2, these transitions stem from the interpretation of $\mathcal{L}^{1.5}$ in focused intuitionistic logic. As for the unfocused rules, we have implicitly used the congruences in Appendix A.6 in this presentation.

To facilitate the parallel with focused linear logic, it is useful to give a second presentation of the focused operational semantics of $\mathcal{L}^{1.5}$, and specifically of the rule transition judgment $\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle$. We introduce an additional rule processing judgment $\Sigma.\langle\Omega ; \Pi, \boxed{R}\rangle \mapsto \Sigma'.\langle\Omega' ; \Pi'\rangle$ and define both by inference rules as follows:

Rule selection $\boxed{\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$

$$\frac{\Sigma.\langle\Omega ; \Pi, \boxed{R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega ; \Pi, R\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle} \qquad \frac{\Sigma.\langle\Omega ; \Pi, \boxed{R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega, R ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

Rule processing $\boxed{\Sigma.\langle\Omega ; \Pi, \boxed{R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$

$$\frac{\Sigma.\langle\Omega, l_1 ; \Pi, l_2, \boxed{l_1, l_2 \multimap P}\rangle \Rightarrow \Sigma.\langle\Omega, l_1 ; \Pi, P\rangle}{\cdot \vdash_{\Sigma} t : \iota \quad \Sigma.\langle\Omega ; \Pi, \boxed{[t/x]R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle} \qquad \frac{(\Sigma, p : \tau \rightarrow o).\langle\Omega ; \Pi, \boxed{[p/X]R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{(\Sigma, p : \tau \rightarrow o).\langle\Omega ; \Pi, \boxed{\forall X : \tau \rightarrow o. R}\rangle \Rightarrow \Sigma'.\langle\Omega ; \Pi'\rangle}$$

It is easy to show that the transition $\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle$ at the beginning of this section is achievable if and only if there is a derivation of $\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle$ according to the above rules. The rule selection rules matches closely the focus rules of focused linear logic, while the rule processing rules correspond to the chaining of appropriate non-invertible rules. This is discussed in further detail in Appendix B.2.

We can go one step further and directly define the judgment $\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle$ mentioned in Section 2.1.3. Recall that this judgment maps stable states to stable states by selecting one rule to apply and then stabilizing the resulting state. It is realized by the following rules:

Stable transition $\boxed{\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$

$$\frac{\Sigma.\langle\Omega ; \Pi, \boxed{R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega ; \Pi, R\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

$$\frac{\Sigma.\langle\Omega, R ; \Pi, \boxed{R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega, R ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

Rule processing $\boxed{\Sigma.\langle\Omega ; \Pi, \boxed{R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$

$$\frac{\Sigma.\langle\Omega, l_1 ; \Pi, P\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega, l_1 ; \Pi, l_2, \boxed{l_1, l_2 \multimap P}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

$$\frac{\cdot \vdash_{\Sigma} t : \iota \quad \Sigma.\langle\Omega ; \Pi, \boxed{[t/x]R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega ; \Pi, \boxed{\forall x : \iota. R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

$$\frac{(\Sigma, p : \tau \rightarrow o).\langle\Omega ; \Pi, \boxed{[p/X]R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{(\Sigma, p : \tau \rightarrow o).\langle\Omega ; \Pi, \boxed{\forall X : \tau \rightarrow o. R}\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

Stabilization $\boxed{\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$

$$\overline{\Sigma.\langle\Omega ; \Pi\rangle \Rightarrow \Sigma.\langle\Omega ; \Pi\rangle}$$

$$\frac{\Sigma.\langle\Omega, A ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega ; \Pi, !A\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

$$\frac{\Sigma.\langle\Omega, R ; \Pi\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega ; \Pi, !R\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

$$\frac{(\Sigma, x : \tau \rightarrow \iota).\langle\Omega ; \Pi, P\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega ; \Pi, \exists x : \tau \rightarrow \iota. P\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

$$\frac{(\Sigma, X : \tau \rightarrow o).\langle\Omega ; \Pi, P\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}{\Sigma.\langle\Omega ; \Pi, \exists X : \tau \rightarrow o. P\rangle \Rightarrow \Sigma'.\langle\Omega' ; \Pi'\rangle}$$

The last four rules corresponds to invertible rules in focused linear logic.

We could simplify the operational semantics of $\mathcal{L}^{1.5}$ even further by observing that the existentials in a program P can always be brought outward in prenex position (possibly after renaming bound variables). Therefore, every program P is equivalent to a program of the form $\exists\Sigma. P_{\Pi}, !(P_{\Omega})$, where P_{Π} is a collection of atoms A and rules R , and $!(P_{\Omega})$ denotes a similar collection of reusable atoms and rules. Notice that, when closed, P_{Π} has the form of a stable state program and P_{Ω} of an archive. This all implies that a rule R matches the template $\forall(l \multimap \exists\Sigma. P_{\Pi}, !(P_{\Omega}))$.

Then, the above rules collapse into the following two rules, each of which carries out a single step of execution that transforms a stable state into another stable state.

$$\overline{\Sigma.\langle\Omega, l_1\theta ; \Pi, l_2\theta, \forall(l_1, l_2 \multimap \exists\Sigma'. P_{\Pi}, !(P_{\Omega}))\rangle \Rightarrow (\Sigma, \Sigma').\langle\Omega, l_1\theta, P_{\Omega}\theta ; \Pi, P_{\Pi}\theta\rangle}$$

$$\overline{\Sigma.\langle\underbrace{\Omega, l_1\theta, \forall(l_1, l_2 \multimap \exists\Sigma'. P_{\Pi}, !(P_{\Omega}))}_{\Omega^*} ; \Pi, l_2\theta\rangle \Rightarrow (\Sigma, \Sigma').\langle\Omega^*, P_{\Omega}\theta ; \Pi, P_{\Pi}\theta\rangle}$$

We can rewrite them as state transitions over stable states as follows:

$$\begin{aligned} \Sigma.\langle\Omega, l_1\theta ; \Pi, l_2\theta, \forall(l_1, l_2 \multimap \exists\Sigma'. P_{\Pi}, !(P_{\Omega}))\rangle &\Rightarrow (\Sigma, \Sigma').\langle\Omega, l_1\theta, P_{\Omega}\theta ; \Pi, P_{\Pi}\theta\rangle \\ \Sigma.\langle\underbrace{\Omega, l_1\theta, \forall(l_1, l_2 \multimap \exists\Sigma'. P_{\Pi}, !(P_{\Omega}))}_{\Omega^*} ; \Pi, l_2\theta\rangle &\Rightarrow (\Sigma, \Sigma').\langle\Omega^*, P_{\Omega}\theta ; \Pi, P_{\Pi}\theta\rangle \end{aligned}$$

Although possible in this language, this last simplification may not be easily expressible in a larger language, for example were we to include the choice operator $\&$ of [9].

B Logical Interpretation

This appendix discusses a fragment of intuitionistic linear logic that is related to the language examined in this report. Section B.1 defines this fragment, multiplicative-exponential intuitionistic linear logic or MEILL, and shows its traditional derivation rules. Section B.2 gives a focused presentation. Section B.3 identifies the precise fragment of MEILL our core language corresponds to and relates its rewriting semantics both to traditional and focused derivability. For the sake of brevity, we do not distinguish between first- and second-order variables or quantification at the logical level.

B.1 Multiplicative-exponential Intuitionistic Linear Logic

Multiplicative-exponential intuitionistic linear logic, MEILL for short, is the fragment of intuitionistic linear logic [17] deprived of additive connectives and units.

Its formulas, denoted φ and ψ possibly with subscripts and superscripts, are defined by the following grammar:

$$\text{Formulas } \varphi ::= A \mid \mathbf{1} \mid \varphi \otimes \varphi \mid \varphi \multimap \varphi \mid !\varphi \mid \forall x : \tau. \varphi \mid \exists x : \tau. \varphi$$

Atomic formulas A are constructed as in $\mathcal{L}^{1.5}$. In particular, symbols wherein are drawn from a signature Σ .

We describe derivability for MEILL in the sequent calculus style using a two-context judgment of the form

$$\Gamma; \Delta \longrightarrow_{\Sigma} \varphi$$

where Γ and Δ are multisets of formulas called the persistent and linear context respectively. A context is defined as follows:

$$\begin{array}{l} \text{Contexts } \Delta ::= \cdot \quad \text{Empty context} \\ \quad \quad \quad | \Delta, \varphi \quad \text{Context extension} \end{array}$$

We will often use “,” as a context union operator and consider it commutative and associative with unit “.”. In a sequent $\Gamma; \Delta \longrightarrow_{\Sigma} \varphi$, the linear context Δ holds formulas that can be used exactly once in a derivation of φ , while the formulas in the persistent context Γ can be used arbitrarily many times.

The rules defining the above judgment are given in Figure B.1. The cut rules for this language are admissible [31] and have therefore been omitted. We have grayed out rules for the fragment of MEILL that do not correspond to any operator in our rewriting language (see Appendix B.3 for details).

Observe that, aside for \otimes and $\mathbf{1}$ which play a structural role, only left rules are retained.

$$\begin{array}{c}
\frac{}{\Gamma; \varphi \rightarrow_{\Sigma} \varphi} \text{init} \\
\frac{\Gamma; \Delta \rightarrow_{\Sigma} \psi}{\Gamma; \Delta, \mathbf{1} \rightarrow_{\Sigma} \psi} \mathbf{1L} \\
\frac{\Gamma; \Delta, \varphi_1, \varphi_2 \rightarrow_{\Sigma} \psi}{\Gamma; \Delta, \varphi_1 \otimes \varphi_2 \rightarrow_{\Sigma} \psi} \otimes L \\
\frac{\Gamma; \Delta_1 \rightarrow_{\Sigma} \varphi_1 \quad \Gamma; \Delta_2, \varphi_2 \rightarrow_{\Sigma} \psi}{\Gamma; \Delta_1, \Delta_2, \varphi_1 \multimap \varphi_2 \rightarrow_{\Sigma} \psi} \multimap L \\
\frac{\Gamma, \varphi; \Delta \rightarrow_{\Sigma} \psi}{\Gamma; \Delta, !\varphi \rightarrow_{\Sigma} \psi} !L \\
\frac{\cdot \vdash_{\Sigma} t : \tau \quad \Gamma; \Delta, [t/x]\varphi \rightarrow_{\Sigma} \psi}{\Gamma; \Delta, \forall x : \tau. \varphi \rightarrow_{\Sigma} \psi} \forall L \\
\frac{\Gamma; \Delta, \varphi \rightarrow_{\Sigma, x: \tau} \psi}{\Gamma; \Delta, \exists x : \tau. \varphi \rightarrow_{\Sigma} \psi} \exists L
\end{array}
\qquad
\begin{array}{c}
\frac{\Gamma, \varphi; \Delta, \varphi \rightarrow_{\Sigma} \psi}{\Gamma, \varphi; \Delta \rightarrow_{\Sigma} \psi} \text{clone} \\
\frac{}{\Gamma; \cdot \rightarrow_{\Sigma} \mathbf{1}} \mathbf{1R} \\
\frac{\Gamma; \Delta_1 \rightarrow_{\Sigma} \varphi_1 \quad \Gamma; \Delta_2 \rightarrow_{\Sigma} \varphi_2}{\Gamma; \Delta_1, \Delta_2 \rightarrow_{\Sigma} \varphi_1 \otimes \varphi_2} \otimes R \\
\frac{\Gamma; \Delta, \varphi_1 \rightarrow_{\Sigma} \varphi_2}{\Gamma; \Delta \rightarrow_{\Sigma} \varphi_1 \multimap \varphi_2} \multimap R \\
\frac{\Gamma; \cdot \rightarrow_{\Sigma} \varphi}{\Delta; \cdot \rightarrow_{\Sigma} !\varphi} !R \\
\frac{\Gamma; \Delta \rightarrow_{\Sigma, x: \tau} \varphi}{\Gamma; \Delta \rightarrow_{\Sigma} \forall x : \tau. \varphi} \forall R \\
\frac{\cdot \vdash_{\Sigma} t : \tau \quad \Gamma; \Delta \rightarrow_{\Sigma} [t/x]\varphi}{\Gamma; \Delta \rightarrow_{\Sigma} \exists x : \tau. \varphi} \exists R
\end{array}$$

Figure B.1: Multiplicative Exponential Intuitionistic Linear Logic — MEILL

B.2 Focused MEILL

A focusing presentation of logic hinges on the observation that some sequent rules are invertible, i.e., their application during proof search can always be “undone”, while other rules are not. Invertible rules can therefore be applied eagerly and exhaustively until a *stable sequent* emerges — a sequent where no such rules are applicable. In well-behaved logics, in particular traditional and linear logic, non-invertible rules can be chained without compromising completeness. Chaining means that, when applying a non-invertible rule, the next non-invertible rules target just the by-products in the rule’s premises of the decomposition of the principal formula in its conclusion. Completeness ensures that this can be carried out for as long as non-invertible rules can be applied to these by-products.

This suggests a proof-search strategy that alternates two phases: an *inversion phase* where invertible rules are applied exhaustively. A formula in the resulting stable sequent is then selected as the focus of the search. A *chaining phase* applies non-invertible rules to the formula in focus exclusively, putting its immediate subformulas in focus in its premises. Chaining proceeds until no non-invertible rule is applicable to the formula in focus. At this point, a new inversion phase begins, and so on until the derivation is completed or a failure is reported. In the latter case, proof search backtracks to the beginning of the last chaining phase, where it selects another formula and repeats the process. This approach to proof search is called *focusing*, and a logic for which it is complete is called focused (every logic for which chaining is complete is focused).

A focused presentation of linear logic relies on the observation that either the left rules or the right rules of each logical operator is invertible. This suggests classifying formulas accordingly. Positive formulas φ^+ have invertible left rules and negative formulas φ^- have invertible right rules. One of the most remarkable facts about focusing is that we can arbitrarily classify atomic formulas as positive or negative without impacting derivability, although the choice determines the shape of derivations. This classification effectively assigns polarities to each logical operator and to atoms. This leads to a polarized presentation of logic. A polarized grammar for MEILL is as follows:

$$\begin{aligned} \text{Positive formulas } \varphi^+ &::= A^+ \mid \mathbf{1} \mid \varphi \otimes \varphi \mid !\varphi \mid \exists x : \tau. \varphi \\ \text{Negative formulas } \varphi^- &::= A^- \mid \varphi \multimap \varphi \mid \forall x : \tau. \varphi \\ \text{Formulas } \varphi &::= \varphi^+ \mid \varphi^- \end{aligned}$$

Forward-chaining search, which is what $\mathcal{L}^{1.5}$ relies on, is obtained by making all atoms positive. This is why we have grayed out negative atoms in this grammar.

A stable sequent is one where the linear context consists solely of negative formulas or positive atoms, and the right-hand side is either a positive formula or a negative atom (this latter case does not apply in our setting).

$$\begin{aligned} \text{Stable linear context } \Delta &::= \cdot \mid \Delta, \varphi^- \mid \Delta, A^+ \\ \text{Stable right-hand side } \psi &::= \varphi^+ \mid A^- \end{aligned}$$

A focused presentation of MEILL (and linear logic in general) relies on the following four judgments:

$$\begin{aligned} \Gamma; \Delta &\Longrightarrow_{\Sigma} \varphi && \text{Generic sequent} \\ \Gamma; \Delta &\Longrightarrow_{\Sigma} \psi && \text{Stable sequent} \\ \Gamma; \Delta, \boxed{\varphi} &\Longrightarrow_{\Sigma} \psi && \text{Left-focused sequent} \\ \Gamma; \Delta &\Longrightarrow_{\Sigma} \boxed{\psi} && \text{Right-focused sequent} \end{aligned}$$

A generic sequent is subject to the application of invertible rules until a stable sequent is produced. Then, a formula on either the left or the right is selected as the focus and non-invertible rules are applied to it and its descendant for as long as possible. The focus is highlighted in red.

These judgments are realized by the rules in Figure B.2, where again we have grayed out rules that do not apply in our setting. Notice that again, aside from \otimes and $\mathbf{1}$, only left rules are retained. Once more, we omit the cut rules, which are admissible.

The lines marked “focus” and “blur” are where phase alternation takes place. Specifically each of the “focus” rules selects a formula from a stable sequent to start the chaining phase. The “blur” rules terminate the chaining phase once no non-invertible rule is applicable to the formula in focus, thereby starting a new inversion phase. Observe that the above rules maintain the invariant that each focused sequent has exactly one formula in focus.

Focus:

$$\begin{array}{c}
\frac{\Gamma; \Delta, \boxed{\varphi^-} \Longrightarrow_{\Sigma} \psi}{\Gamma; \Delta, \varphi^- \Longrightarrow_{\Sigma} \psi} \text{focusR} \quad \frac{\Gamma, \varphi; \Delta, \boxed{\varphi} \Longrightarrow_{\Sigma} \psi \quad (\varphi \text{ not } A^+)}{\Gamma, \varphi; \Delta \Longrightarrow_{\Sigma} \psi} \text{focus!R} \quad \frac{\Gamma; \Delta \Longrightarrow_{\Sigma} \boxed{\varphi^+}}{\Gamma; \Delta \Longrightarrow_{\Sigma} \varphi^+} \text{focusR} \\
\\
\frac{}{\Gamma; A^+ \Longrightarrow_{\Sigma} \boxed{A^+}} \text{atmL} \quad \frac{}{\Gamma, A^+; \cdot \Longrightarrow_{\Sigma} \boxed{A^+}} \text{atm!L} \quad \frac{}{\Gamma; \boxed{A^-} \Longrightarrow_{\Sigma} A^-} \text{atmR} \\
\\
\frac{\Gamma; \Delta \Longrightarrow_{\Sigma} \psi}{\Gamma; \Delta, \mathbf{1} \Longrightarrow_{\Sigma} \psi} \mathbf{1L} \quad \frac{}{\Gamma; \cdot \Longrightarrow_{\Sigma} \boxed{\mathbf{1}}} \mathbf{1R} \\
\\
\frac{\Gamma; \Delta, \varphi_1, \varphi_2 \Longrightarrow_{\Sigma} \psi}{\Gamma; \Delta, \varphi_1 \otimes \varphi_2 \Longrightarrow_{\Sigma} \psi} \otimes L \quad \frac{\Gamma; \Delta_1 \Longrightarrow_{\Sigma} \boxed{\varphi_1} \quad \Gamma; \Delta_2 \Longrightarrow_{\Sigma} \boxed{\varphi_2}}{\Gamma; \Delta_1, \Delta_2 \Longrightarrow_{\Sigma} \boxed{\varphi_1 \otimes \varphi_2}} \otimes R \\
\\
\frac{\Gamma; \Delta_1 \Longrightarrow_{\Sigma} \boxed{\varphi_1} \quad \Gamma; \Delta_2, \boxed{\varphi_2} \Longrightarrow_{\Sigma} \psi}{\Gamma; \Delta_1, \Delta_2, \boxed{\varphi_1 \multimap \varphi_2} \Longrightarrow_{\Sigma} \psi} \multimap L \quad \frac{\Gamma; \Delta, \varphi_1 \Longrightarrow_{\Sigma} \varphi_2}{\Gamma; \Delta \Longrightarrow_{\Sigma} \varphi_1 \multimap \varphi_2} \multimap R \\
\\
\frac{\Gamma, \varphi; \Delta \Longrightarrow_{\Sigma} \psi}{\Gamma; \Delta, !\varphi \Longrightarrow_{\Sigma} \psi} !L \quad \frac{\Gamma; \cdot \Longrightarrow_{\Sigma} \varphi}{\Delta; \cdot \Longrightarrow_{\Sigma} \boxed{!\varphi}} !R \\
\\
\frac{\cdot \vdash_{\Sigma} t : \tau \quad \Gamma; \Delta, \boxed{[t/x]\varphi} \Longrightarrow_{\Sigma} \psi}{\Gamma; \Delta, \boxed{\forall x : \tau. \varphi} \Longrightarrow_{\Sigma} \psi} \forall L \quad \frac{\Gamma; \Delta \Longrightarrow_{\Sigma, x:\tau} \varphi}{\Gamma; \Delta \Longrightarrow_{\Sigma} \forall x : \tau. \varphi} \forall R \\
\\
\frac{\Gamma; \Delta, \varphi \Longrightarrow_{\Sigma, x:\tau} \psi}{\Gamma; \Delta, \exists x : \tau. \varphi \Longrightarrow_{\Sigma} \psi} \exists L \quad \frac{\cdot \vdash_{\Sigma} t : \tau \quad \Gamma; \Delta \Longrightarrow_{\Sigma} \boxed{[t/x]\varphi}}{\Gamma; \Delta \Longrightarrow_{\Sigma} \boxed{\exists x : \tau. \varphi}} \exists R \\
\\
\text{Blur:} \quad \frac{\Gamma; \Delta, \varphi^+ \Longrightarrow_{\Sigma} \psi}{\Gamma; \Delta, \boxed{\varphi^+} \Longrightarrow_{\Sigma} \psi} \text{blurL} \quad \frac{\Gamma; \Delta \Longrightarrow_{\Sigma} \varphi^-}{\Gamma; \Delta \Longrightarrow_{\Sigma} \boxed{\varphi^-}} \text{blurR}
\end{array}$$

Figure B.2: Focused Presentation of MEILL

Focusing is sound and complete for intuitionistic linear logic as shown in [22]. Worth mentioning is also the elegant structural proof in [35] — it applies to traditional rather than linear logic although nothing prevents porting it to the linear setting.

Theorem 7 (Soundness and Completeness of Focusing) $\Gamma; \Delta \longrightarrow_{\Sigma} \psi$ if and only if $\Gamma; \Delta \Longrightarrow_{\Sigma} \psi$.

Proof See [22, 35]. □

Formula templates that, when focused upon, are processed through non-invertible rules only (i.e., that have derivation stubs consisting of a chain of non-invertible rules) can be viewed as *synthetic connectives*. This derivation stub coalesces into derived (non-invertible) rule, a right rule if the chain starts on the right and a left rule if it starts on the left. It turns out that this formula template can be decomposed by means of invertible rules only when put on the other side of the sequent. The corresponding derived rule is therefore itself invertible.

B.3 Interpretation

In this section, we identify $\mathcal{L}^{1.5}$ as a fragment of MEILL, show that its unfocused rewriting semantics stems from the traditional derivation rules of MEILL for that fragment, and that its focused transitions are explained by the corresponding focused rules.

B.3.1 Translation

The rewriting language $\mathcal{L}^{1.5}$ examined in this report corresponds to a fragment of MEILL. Atoms A of $\mathcal{L}^{1.5}$ are exactly the atomic formulas A of MEILL. The quantifiers, ! and the rewrite directive \multimap are mapped to the identical operators of linear logic. Multiset union “,” and the empty multiset “.” within a program correspond instead to the tensor \otimes and its unit $\mathbf{1}$.

This correspondence is formalized as follows, where the MEILL formula for program P is denoted $\ulcorner P \urcorner$, and similarly for left-hand sides and rules.

| | Rewriting | Logic |
|-----------------|---|--------------|
| <i>LHS</i> | $\ulcorner \cdot \urcorner = \mathbf{1}$ | |
| | $\ulcorner A, l \urcorner = A \otimes \ulcorner l \urcorner$ | |
| <i>Rules</i> | $\ulcorner l \multimap P \urcorner = \ulcorner l \urcorner \multimap \ulcorner P \urcorner$ | |
| | $\ulcorner \forall x : \iota. P \urcorner = \forall x : \iota. \ulcorner P \urcorner$ | |
| | $\ulcorner \forall X : \tau \rightarrow o. P \urcorner = \forall X : \tau \rightarrow o. \ulcorner P \urcorner$ | |
| <i>Programs</i> | $\ulcorner \cdot \urcorner = \mathbf{1}$ | |
| | $\ulcorner P_1, P_2 \urcorner = \ulcorner P_1 \urcorner \otimes \ulcorner P_2 \urcorner$ | |
| | $\ulcorner A \urcorner = A$ | |
| | $\ulcorner !A \urcorner = !\ulcorner A \urcorner$ | |
| | $\ulcorner R \urcorner = \ulcorner R \urcorner$ | |
| | $\ulcorner !R \urcorner = !\ulcorner R \urcorner$ | |
| | $\ulcorner \exists x : \tau \rightarrow \iota. P \urcorner = \exists x : \tau \rightarrow \iota. \ulcorner P \urcorner$ | |
| | $\ulcorner \exists X : \tau \rightarrow o. P \urcorner = \exists X : \tau \rightarrow o. \ulcorner P \urcorner$ | |

Here, we have made explicit the distinction between first- and second-order entities in MEILL since we distinguished them in $\mathcal{L}^{1.5}$.

This translation identifies a linguistic fragment of MEILL that is given by the following grammar:

| | |
|-----------------|---|
| <i>LHS</i> | $\varphi_l ::= \cdot \mid A \otimes \varphi_l$ |
| <i>Rules</i> | $\varphi_R ::= \varphi_l \multimap \varphi_P \mid \forall x : \iota. \varphi_R$ $\mid \forall X : \tau \rightarrow o. \varphi_R$ |
| <i>Programs</i> | $\varphi_P ::= \cdot \mid \varphi_P \otimes \varphi_P \mid A \mid !A \mid \varphi_R \mid !\varphi_R \mid \exists x : \tau \rightarrow \iota. \varphi_P$ $\mid \exists X : \tau \rightarrow o. \varphi$ |

This is nothing more than the grammar of left-hand sides, rules and programs from Appendix A.1 expressed using MEILL formulas.

We also give an encoding of archives to linear formulas consisting of the conjunction of their constituents, as we rely on these entities to discuss the correspondence between $\mathcal{L}^{1.5}$ and MEILL.

| | Rewriting | Logic |
|-----------------|---|--------------|
| <i>Archives</i> | $\ulcorner \cdot \urcorner = \mathbf{1}$ | |
| | $\ulcorner \Omega, A \urcorner = \ulcorner \Omega \urcorner \otimes !A$ | |
| | $\ulcorner \Omega, R \urcorner = \ulcorner \Omega \urcorner \otimes !\ulcorner R \urcorner$ | |

Then, given a state $\Sigma. \langle \Omega ; \Pi \rangle$, we write $\exists \Sigma. \ulcorner !\Omega ; \Pi \urcorner$ for the MEILL formula $\exists \Sigma. \ulcorner \Omega \urcorner \otimes \ulcorner \Pi \urcorner$ obtained by prefixing the combined translation $\ulcorner \Omega \urcorner \otimes \ulcorner \Pi \urcorner$ of Ω and Π with an existential quantifier for every declaration in the signature Σ .

Below, we will make use of the following homomorphic mapping of an archive Ω to (persistent) contexts Γ in MEILL:

$$\begin{array}{ccc}
 & \textbf{Rewriting} & \textbf{Logic} \\
 \textit{Archives} & [\cdot] = \cdot & \\
 & [\Omega, A] = [\Omega], A & \\
 & [\Omega, R] = [\Omega], [R] &
 \end{array}$$

We will also rely on the following non-deterministic encoding of the program component Π of a state $\Sigma.\langle\Omega ; \Pi\rangle$ to a (linear) context Δ :

$$\begin{array}{ccc}
 & \textbf{Rewriting} & \textbf{Logic} \\
 \textit{State programs} & [P_1, P_2] = [P_1], [P_2] & \\
 & [P] = \ulcorner P \urcorner &
 \end{array}$$

This encoding is non-deterministic as the union (P_1, P_2) of a two $\mathcal{L}^{1.5}$ programs can be translated either a context that contains a single formula $\ulcorner P_1, P_2 \urcorner$ or as a context which contains possible more formulas obtained by encoding P_1 and P_2 in the same way recursively.

B.3.2 Unfocused Transitions

State transition rules in $\mathcal{L}^{1.5}$ correspond to derivable sequents of MEILL, a result that has been proved in a more general setting in [9]. In fact, each of our transition rules stems from one of the left rules of the operators in Appendix B.1, possibly helped by some other rules that play a structural role. The transitions for the rewriting quantifiers match exactly the left rules for the logic quantifiers when read from conclusion to premise. The transition for \multimap stems from rule \multimap L, relying on the rules for \otimes and $\mathbf{1}$ to break the corresponding elements in the state program and rules init and clone to match each atom within with a corresponding left-hand side atom.

Given a state $\Sigma.\langle\Omega ; \Pi\rangle$, the following property is proved as in [9].

Theorem 8 (Soundness) *If $\vdash \Sigma.\langle\Omega ; \Pi\rangle$ state and $\Sigma.\langle\Omega ; \Pi\rangle \mapsto \Sigma'.\langle\Omega' ; \Pi'\rangle$, then $[\Omega]; [\Pi] \longrightarrow_{\Sigma} \exists \Sigma'. \ulcorner \Omega'; \Pi' \urcorner$.*

While a formal proof can be found in [9], we now show the derivation snippet that corresponds to each transition. As we do so, we make implicit use of the left derivation rules for \otimes and $\mathbf{1}$ to break $[\Pi]$ into context formulas that do not have these connectives as their topmost operator. The mapping between unfocused transitions in $\mathcal{L}^{1.5}$ and derivation snippets in linear logic is as follows, where the formula ψ in each case is as in the above theorem [9] and we are

making implicit uses of the definition of the encodings $\lceil - \rceil$ and $\lfloor - \rfloor$.

$$\begin{array}{lcl}
\Sigma.\langle \Omega, l_1 ; \Pi, l_2, (l_1, l_2) \multimap P \rangle \mapsto \Sigma.\langle \Omega, l_1 ; \Pi, P \rangle & \rightsquigarrow & \frac{\dots \overline{\lceil \Omega, l_1 \rceil}; A \longrightarrow_{\Sigma} A \dots}{\frac{[\Omega, l_1]; [l_2] \longrightarrow_{\Sigma} \lceil l_1, l_2 \rceil \quad [\Omega, l_1]; [\Pi, P] \longrightarrow_{\Sigma} \psi}{[\Omega, l_1]; [\Pi, l_2, (l_1, l_2) \multimap P] \longrightarrow_{\Sigma} \psi} \multimap L} \\
\Sigma.\langle \Omega ; \Pi, \forall x : \iota. R \rangle \mapsto \Sigma.\langle \Omega ; \Pi, [t/x]R \rangle & \rightsquigarrow & \frac{\cdot \vdash_{\Sigma} t : \iota \quad [\Omega]; [\Pi, [t/x]R] \longrightarrow_{\Sigma} \psi}{[\Omega]; [\Pi, \forall x : \iota. R] \longrightarrow_{\Sigma} \psi} \forall L \\
\text{if } \cdot \vdash_{\Sigma} t : \iota & & \\
\Sigma.\langle \Omega ; \Pi, \forall X : \tau \rightarrow o. R \rangle \mapsto \Sigma.\langle \Omega ; \Pi, [p/X]R \rangle & \rightsquigarrow & \frac{p : \tau \rightarrow o \text{ in } \Sigma \quad [\Omega]; [\Pi, [p/X]R] \longrightarrow_{\Sigma} \psi}{[\Omega]; [\Pi, \forall X : \tau \rightarrow o. R] \longrightarrow_{\Sigma} \psi} \forall L \\
\text{if } p : \tau \rightarrow o \text{ in } \Sigma & & \\
\Sigma.\langle \Omega ; \Pi, !A \rangle \mapsto \Sigma.\langle \Omega, A ; \Pi \rangle & \rightsquigarrow & \frac{[\Omega, A]; [\Pi] \longrightarrow_{\Sigma} \psi}{[\Omega]; [\Pi, !A] \longrightarrow_{\Sigma} \psi} !L \\
\Sigma.\langle \Omega ; \Pi, !R \rangle \mapsto \Sigma.\langle \Omega, R ; \Pi \rangle & \rightsquigarrow & \frac{[\Omega, R]; [\Pi] \longrightarrow_{\Sigma} \psi}{[\Omega]; [\Pi, !R] \longrightarrow_{\Sigma} \psi} !L \\
\Sigma.\langle \Omega ; \Pi, \exists x : \tau \rightarrow \iota. P \rangle \mapsto (\Sigma, x : \tau \rightarrow \iota).\langle \Omega ; \Pi, P \rangle & \rightsquigarrow & \frac{[\Omega]; [\Pi, P] \longrightarrow_{\Sigma, x: \tau \rightarrow \iota} \psi}{[\Omega]; [\Pi, \exists x : \tau \rightarrow \iota. P] \longrightarrow_{\Sigma} \psi} \exists L \\
\Sigma.\langle \Omega ; \Pi, \exists X : \tau \rightarrow o. P \rangle \mapsto (\Sigma, X : \tau \rightarrow o).\langle \Omega ; \Pi, P \rangle & \rightsquigarrow & \frac{[\Omega]; [\Pi, P] \longrightarrow_{\Sigma, X: \tau \rightarrow o} \psi}{[\Omega]; [\Pi, \exists X : \tau \rightarrow o. P] \longrightarrow_{\Sigma} \psi} \exists L \\
\Sigma.\langle \Omega, R ; \Pi \rangle \mapsto \Sigma.\langle \Omega, R ; \Pi, R \rangle & \rightsquigarrow & \frac{[\Omega, R]; [\Pi, R] \longrightarrow_{\Sigma} \psi}{[\Omega, R]; [\Pi] \longrightarrow_{\Sigma} \psi} \text{clone}
\end{array}$$

The first correspondence, for \multimap , requires some explanation: the derivation on the left branch is completed by using rules $\otimes L$ and $!L$ to decompose $\lceil l_2 \rceil$ into individual atomic formulas in the linear context. Then, rules $\otimes R$ and $!R$ are applied to break the formula $\lceil l_1, l_2 \rceil$ into atoms. Doing so exhaustively leaves sequents of either the form $[\Omega, l_1]; A \longrightarrow_{\Sigma} A$ or $[\Omega, l'_1, A]; \cdot \longrightarrow_{\Sigma} A$ where $l_1 = l'_1, A$. The first form is immediately resolved by a use of rule init , the second prepares for such an application by means of rule clone .

It is worth noting that this proof does not makes use of the grayed out rules from Appendix B.1. However, were we to replace the initial rule

$$\overline{\Gamma; \varphi \longrightarrow_{\Sigma} \varphi}^{\text{init}} \quad \text{with the more canonical} \quad \overline{\Gamma; A \longrightarrow_{\Sigma} A}^{\text{atm}}$$

the resulting derivation would need the grayed out rules, in general.

The above soundness theorem holds for the iterated transition relation \mapsto^* .

B.3.3 Focused Transitions

Accounting for polarities in the translation of $\mathcal{L}^{1.5}$ we just saw, left-hand sides and rules yield negative formulas while programs map to positive formulas:

$$\begin{array}{lcl}
\text{LHS} & \varphi_l^+ & ::= \cdot \mid A^+ \otimes \varphi_l^+ \\
\text{Rules} & \varphi_R^- & ::= \varphi_l \multimap \varphi_P^+ \mid \forall x : \iota. \varphi_R^- \\
& & \mid \forall X : \tau \rightarrow o. \varphi_R^- \\
\text{Programs} & \varphi_P^+ & ::= \cdot \mid \varphi_P^+ \otimes \varphi_P^+ \mid A^+ \mid !A^+ \mid \varphi_R^- \mid !\varphi_R^- \mid \exists x : \tau \rightarrow \iota. \varphi_P^+ \\
& & \mid \exists X : \tau \rightarrow o. \varphi
\end{array}$$

The lone occurrence of φ_R^- in the production for φ_P^+ can be accounted for by rewriting it as $\varphi_R^- \otimes \mathbf{1}$, for example

Observe that a program state, which we wrote Π , consists of atoms A^+ , negative conjunctions $\varphi_P^+ \otimes \varphi_P^+$ and units $\mathbf{1}$ (which can be simplified immediately as the left rules of these operators are invertible), existentials (whose left rules are also invertible), exponentials (whose left rule is invertible too), and formulas corresponding to either rules φ_R^- or atoms A^+ . A stable state corresponds to a stable program: it contains only atoms and the representation of single-use rules only. This is also the form the persistent context has — it is populated from an archive.

The focused transition rules discussed in Appendix A.8 match focused derivation rules of MEILL in a more sophisticated way than their unfocused counterpart. The transitions for the existential programs and reusable entities work in the same way: existential and exponential formulas are positive since the left rule of their operator is invertible. Rules $\forall(l \multimap P)$ in $\mathcal{L}^{1.5}$ correspond however to negative formulas according to our translations. Therefore they start a chaining phase which ends when the (translation of) the embedded program P (a positive formula) is processed. The rewriting transition for $\mathcal{L}^{1.5}$ rules captures exactly this chaining phase. Reusable $\mathcal{L}^{1.5}$ rules are handled similarly, but their use is introduced by rule `focus!L`.

The program component Π of a stable state is transformed into a linear context obtained by mapping all occurrences of “;” and “.” to their identically-denoted context-level operators. Therefore, a stable state is mapped to a context consisting solely of single-use atoms and rules. It is defined as follows.

| | Rewriting | Logic |
|------------------------------|------------------|--------------------------------|
| <i>Stable state programs</i> | $[\cdot]$ | \cdot |
| | $[\Pi, A]$ | $[\Pi], A^+$ |
| | $[\Pi, R]$ | $[\Pi], \ulcorner R \urcorner$ |

The above argument is summarized by the following soundness theorem, which distinguishes cases based on whether the starting state is stable or not (after applying rules $\otimes L$ and $\mathbf{1}L$ exhaustively).

Theorem 9 (Soundness) *Assume that $\vdash \Sigma.\langle \Omega ; \Pi \rangle$ state.*

1. *If Π is stable (i.e., Π has the form Π) and $\Sigma.\langle \Omega ; \Pi \rangle \Leftrightarrow \Sigma'.\langle \Omega' ; \Pi' \rangle$, then $[\Omega]; [\Pi] \Longrightarrow_{\Sigma} \exists \Sigma'. \ulcorner \Omega' ; \Pi' \urcorner$.*
2. *If Π is not stable and $\Sigma.\langle \Omega ; \Pi \rangle \Leftrightarrow \Sigma'.\langle \Omega' ; \Pi' \rangle$, then $[\Omega]; [\Pi] \Longrightarrow_{\Sigma} \exists \Sigma'. \ulcorner \Omega' ; \Pi' \urcorner$.*

Proof The proof proceeds as in the unfocused case, except for an additional induction on the number of universal quantifiers appearing in an $\mathcal{L}^{1.5}$ rule. □

It should be observed that the constructed derivation can make use of all the rules of focused linear logic, including the ones that have been grayed out. This is the reason why focused and unfocused transitions in $\mathcal{L}^{1.5}$ are not equivalent, despite Theorem 7.

As in the unfocused case, we distill the key elements of the proof by showing how each focused transition in $\mathcal{L}^{1.5}$

maps to a focused derivation snippet in MEILL.

$$\begin{array}{c}
\Sigma.\langle \Omega, l_1\theta ; \Pi, l_2\theta, \forall(l_1, l_2 \multimap P) \rangle \Rightarrow \Sigma.\langle \Omega, l_1\theta ; \Pi, P\theta \rangle \rightsquigarrow \\
\frac{\dots \frac{[\Omega, l_1\theta]; [A\theta] \Rightarrow_{\Sigma} [A^+\theta] \text{ atmL} \quad \dots \quad [\Omega, l'_1\theta, A\theta]; \cdot \Rightarrow_{\Sigma} [A^+\theta] \text{ atm!L} \quad \dots \quad \frac{[\Omega, l_1\theta]; [\Pi, P\theta] \Rightarrow_{\Sigma} \psi}{[\Omega, l_1\theta]; [\Pi], [P\theta^{\neg}] \Rightarrow_{\Sigma} \psi} \text{ blurL}}{[\Omega, l_1\theta]; [l_2\theta] \Rightarrow_{\Sigma} [l_1\theta, l_2\theta^{\neg}]} \quad \frac{[\Omega, l_1\theta]; [\Pi, P\theta] \Rightarrow_{\Sigma} \psi}{[\Omega, l_1\theta]; [\Pi], [P\theta^{\neg}] \Rightarrow_{\Sigma} \psi} \text{ blurL}}{[\Omega, l_1\theta]; [\Pi, l_2\theta], [l_1, l_2 \multimap P^{\neg}] \Rightarrow_{\Sigma} \psi} \text{ } \multimap\text{L}}{\frac{[\Omega, l_1\theta]; [\Pi, l_2\theta], [l_1, l_2 \multimap P^{\neg}] \Rightarrow_{\Sigma} \psi}{[\Omega, l_1\theta]; [\Pi, l_2\theta], [\forall(l_1, l_2 \multimap P)^{\neg}] \Rightarrow_{\Sigma} \psi} \text{ } \forall\text{L(repeated)}}{\frac{[\Omega, l_1\theta]; [\Pi, l_2\theta], [\forall(l_1, l_2 \multimap P)^{\neg}] \Rightarrow_{\Sigma} \psi}{[\Omega, l_1\theta]; [\Pi, l_2\theta, \forall(l_1, l_2 \multimap P)] \Rightarrow_{\Sigma} \psi} \text{ focusL}} \text{ } \multimap\text{L}} \\
\Sigma.\langle \underbrace{\Omega, l_1\theta, \forall(l_1, l_2 \multimap P)}_{\Omega^*}; \Pi, l_2\theta \rangle \Rightarrow \Sigma.\langle \Omega^*; \Pi, P\theta \rangle \rightsquigarrow \frac{[\Omega^*]; [\Pi, l_2\theta], [\forall(l_1, l_2 \multimap P)^{\neg}] \Rightarrow_{\Sigma} \psi}{[\Omega^*]; [\Pi, l_2\theta] \Rightarrow_{\Sigma} \psi} \text{ focusL} \\
\Sigma.\langle \Omega ; \Pi, !A \rangle \Rightarrow \Sigma.\langle \Omega, A ; \Pi \rangle \rightsquigarrow \frac{[\Omega, A]; [\Pi] \Rightarrow_{\Sigma} \psi}{[\Omega]; [\Pi, !A] \Rightarrow_{\Sigma} \psi} \text{ clone} \\
\Sigma.\langle \Omega ; \Pi, !R \rangle \Rightarrow \Sigma.\langle \Omega, R ; \Pi \rangle \rightsquigarrow \frac{[\Omega, R]; [\Pi] \Rightarrow_{\Sigma} \psi}{[\Omega]; [\Pi, !R] \Rightarrow_{\Sigma} \psi} \text{ clone} \\
\Sigma.\langle \Omega ; \Pi, \exists x : \tau \rightarrow \iota. P \rangle \Rightarrow (\Sigma, x : \tau \rightarrow \iota).\langle \Omega ; \Pi, P \rangle \rightsquigarrow \frac{[\Omega]; [\Pi, P] \Rightarrow_{\Sigma, x: \tau \rightarrow \iota} \psi}{[\Omega]; [\Pi, \exists x : \tau \rightarrow \iota. P] \Rightarrow_{\Sigma} \psi} \exists\text{L} \\
\Sigma.\langle \Omega ; \Pi, \exists X : \tau \rightarrow o. P \rangle \Rightarrow (\Sigma, X : \tau \rightarrow o).\langle \Omega ; \Pi, P \rangle \rightsquigarrow \frac{[\Omega]; [\Pi, P] \Rightarrow_{\Sigma, X: \tau \rightarrow o} \psi}{[\Omega]; [\Pi, \exists X : \tau \rightarrow o. P] \Rightarrow_{\Sigma} \psi} \exists\text{L}
\end{array}$$

Observe how chaining tracks exactly the structure of a $\mathcal{L}^{1.5}$ rule $\forall(l \multimap P)$. Because this pattern is fully handled by invertible rule when translated to logic, it gives rise to a synthetic connective, whose left rule has the form

$$\frac{[\Omega, l_1\theta]; [\Pi], [P\theta^{\neg}] \Rightarrow_{\Sigma} \psi}{[\Omega, l_1\theta]; [\Pi, l_2\theta], [\forall(l_1, l_2 \multimap P)^{\neg}] \Rightarrow_{\Sigma} \psi}$$

which corresponds closely to the focused transition for $\mathcal{L}^{1.5}$ rules. An exact correspondence is obtained by bracketing it between uses of rules focusL and blurL:

$$\frac{[\Omega, l_1\theta]; [\Pi, P\theta] \Rightarrow_{\Sigma} \psi}{[\Omega, l_1\theta]; [\Pi, l_2\theta, \forall(l_1, l_2 \multimap P)] \Rightarrow_{\Sigma} \psi}$$

The same happens to a reusable rule $!(l \multimap P)$ once it has been archived (i.e., inserted in the persistent context using rule clone on the logic side). Bracketing the above synthetic rule between focus!L and blurL yields the derived rule

$$\frac{[\Omega, l_1\theta, \forall(l_1, l_2 \multimap P)]; [\Pi, P\theta] \Rightarrow_{\Sigma} \psi}{[\Omega, l_1\theta, \forall(l_1, l_2 \multimap P)]; [\Pi, l_2\theta] \Rightarrow_{\Sigma} \psi}$$

which corresponds exactly to the focused $\mathcal{L}^{1.5}$ transition for reusable rules.

Some of the same comments we made when discussing the unfocused correspondence apply here also. Additionally, the substitution term used for each instance of rule $\forall\text{L}$ is exactly what the substitution θ associates to the variable processed by this rule. The proof of the sequent $[\Omega, l_1\theta]; [l_2\theta] \Rightarrow_{\Sigma} [l_1\theta, l_2\theta^{\neg}]$ is again interesting. Because the linear context $[l_1\theta^{\neg}]$ is stable, it does not have occurrences of \otimes nor $\mathbf{1}$ (it is a multiset of bare atoms). As a consequence, this derivation is chained all the way to the applications of rules atmL and atm!L (or 1R if l_2 is empty). As earlier, we wrote l_1 as l'_1, A in the displayed use of rule atm!L.

The same argument can be made starting from the other formulations of the focused transition semantics of $\mathcal{L}^{1.5}$ seen in Appendix A.8. In particular, our second presentation, which decomposed the transitions for single-use and reusable rules into discrete steps, is essentially isomorphic to the focused left sequent rules for the logical operators corresponding to each case. Our third presentation, which went from stable state to stable state, combines the chaining phase of a step with the immediately following inversion phase.