

Type safety for substructural specifications: preliminary results

Robert J. Simmons

July 15, 2010
CMU-CS-10-134

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Substructural logics, such as linear logic and ordered logic, have an inherent notion of state and state change. This makes them a natural choice for developing logical frameworks that specify evolving stateful systems. Our previous work has shown that the so-called *forward reasoning* fragment of ordered linear logic can be used to give clear, concise, and modular specifications of stateful and concurrent features of programming languages in a style called Substructural Operational Semantics (SSOS).

This report discusses two new ideas. First, I discuss a re-orientation of existing logics and logical frameworks based on a presentation of logic as a state transition system. Second, utilizing this transition-based interpretation of logic, I present a safety proof for the SSOS specification of a simple sequential programming language, and discuss how that proof can be updated as the language is extended with parallel evaluation and exception handling.

This report is derived from my thesis proposal [35], which was certified on July 7, 2010.

This material is based upon work supported under a National Science Foundation Graduate Research Fellowship. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Keywords: logical frameworks, ordered logic, preservation, progress, safety, substructural logic, specification, thesis proposal

1 Introduction

Robust and flexible frameworks for specifying and reasoning about programming languages form the bedrock of programming languages research. In part, this is because they provide a common convention: historically, the most successful specification framework, Structural Operational Semantics (i.e. SOS), is the *lingua franca* of PL research. Even more critical is the ability of a framework to permit reasoning about properties of programming languages; in the case of SOS, the traditionally interrelated notion of “Safety = Progress + Preservation” has essentially come to define what it means for a programming language to make sense.

One problem with SOS specifications of programming languages is that they are *non-modular*: especially when dealing with imperative and concurrent programming language features, the natural specification of language *parts* cannot, in general, be combined to give a natural specification of the language *whole*. This makes it difficult to straightforwardly extend a pure “toy” programming language with a feature like imperative references without a complete overhaul of the language’s definition and type safety proof; a similar overhaul is necessary to extend the definition of a large sequential language to handle concurrency. This has troubling practical consequences: it means that it is more difficult to scale from a small language formalization to a large one, and it means that existing language specifications are difficult to augment with features that were not anticipated in advance.

In this preliminary report, which is based on the content of my thesis proposal, I discuss a solution that involves two interconnected ideas. The first idea is a *specification style*, Substructural Operational Semantics (i.e. SSOS). SSOS specifications generalize abstract machine specifications and permit clear and modular specification of programming language features that involve state and concurrency — the specification style that I consider is a generalization of the ordered SSOS specifications as presented in existing work [28].

The second idea is a *logical framework* based on forward-chaining in ordered linear logic. The specifics of this framework are broadly similar to the ordered logical framework discussed in the aforementioned paper on SSOS specifications, and the framework also shares broad similarities with an existing family of logical frameworks such as LF [11], LLF [1], HLF [30], OLF [29], and (in particular) CLF [40]. The interesting difference between existing work and this framework is that this framework is based on an interpretation of logic that has a first-class notion of *state transitions*; as a result, a “step” of an SSOS specification can be represented by a state transition in the logical framework.

The ability to talk about individual transitions and potentially non-terminating computations is essential for giving a full account of type safety, since type safety must also account for the behavior of non-terminating programs. The primary novel contribution in this report is a proof of type safety, by way of progress and preservation lemmas, for the SSOS specification of a simple programming language. I then discuss how that proof can be straightforwardly extended to handle program extensions which introduce concurrency, non-local transfer of control, and state.

In Section 2 I motivate substructural operational semantics specifications by discussing modular and non-modular extensions to programming language specifications, and in Section 3 I motivate the treatment of logic as a state transition system. Section 4 presents a logical framework based on the ideas in Section 3; this section is rather technical and can probably be skipped on a first reading. In Section 5 I present a variety of substructural operational semantics specifications and proofs of their type safety. Both the dynamic and static semantics of the programming languages are presented as ordered logical specifications.

2 Modular and non-modular specification

In the introduction, the phrase *modular specification* was used but not defined. Modularity means different things in different settings; in this section, we will motivate the problem by considering modular and non-modular extensions to a very simple programming language. We start by assuming that we have natural numbers n and some operations (such as addition) on those natural numbers; expressions e can then be defined by giving a BNF specification:

$$e ::= \text{num}(n) \mid e_1 + e_2$$

Call this language of numbers and addition \mathcal{L}_0 ; we next give a small-step structural operational semantics for \mathcal{L}_0 that specifies a left-to-right evaluation order. We also adopt the standard convention of writing expressions that are known to be values as v , not e :

$$\frac{}{\text{num}(n_1) \text{ value}} \quad \frac{e_1 \mapsto e'_1}{e_1 + e_2 \mapsto e'_1 + e_2} \quad \frac{v_1 \text{ value} \quad e_2 \mapsto e'_2}{v_1 + e_2 \mapsto v_1 + e'_2} \quad \frac{n_1 + n_2 = n_3}{\text{num}(n_1) + \text{num}(n_2) \mapsto \text{num}(n_3)}$$

The extension of \mathcal{L}_0 with eager pairs ($e ::= \dots \mid \pi_1 e \mid \pi_2 e \mid \langle e_1, e_2 \rangle$) is an example of a modular extension. We can take the following natural small-step SOS specification of eager pairs and concatenate it together with the \mathcal{L}_0 specification to form a coherent (and, with a suitable type system, provably type-safe) specification of a language that has both features.

$$\frac{v_1 \text{ value} \quad v_2 \text{ value}}{\langle v_1, v_2 \rangle \text{ value}} \quad \frac{e_1 \mapsto e'_1}{\langle e_1, e_2 \rangle \mapsto \langle e'_1, e_2 \rangle} \quad \frac{v_1 \text{ value} \quad e_2 \mapsto e'_2}{\langle v_1, e_2 \rangle \mapsto \langle v_1, e'_2 \rangle}$$

$$\frac{e \mapsto e'}{\pi_1 e \mapsto \pi_1 e'} \quad \frac{\langle v_1, v_2 \rangle \text{ value}}{\pi_1 \langle v_1, v_2 \rangle \mapsto v_1} \quad \frac{e \mapsto e'}{\pi_2 e \mapsto \pi_2 e'} \quad \frac{\langle v_1, v_2 \rangle \text{ value}}{\pi_2 \langle v_1, v_2 \rangle \mapsto v_2}$$

In contrast, the extension of \mathcal{L}_0 with ML-style imperative references is an obvious example of a non-modular extension. The extended syntax is ($e ::= \dots \mid \text{ref } e \mid \text{loc}[l] \mid !e \mid e_1 \text{ gets } e_2$) — the expression $\text{ref } e$ creates a reference, $!e$ dereferences a reference, $e_1 \text{ gets } e_2$ assigns the value of e_2 to the location represented by e_1 , and $\text{loc}[l]$ is a value, which only exists at runtime, referencing the abstract heap location l where some value is stored. In order to give a small-step SOS specification for the extended language, we must mention a store σ in every rule:

$$\frac{}{\text{loc}[l] \text{ value}} \quad \frac{(\sigma, e) \mapsto (\sigma', e')}{(\sigma, \text{ref } e) \mapsto (\sigma', \text{ref } e')} \quad \frac{v \text{ value} \quad l \notin \text{dom}(\sigma)}{(\sigma, \text{ref } v) \mapsto (\sigma[l := v], \text{loc}[l])}$$

$$\frac{(\sigma, e) \mapsto (\sigma', e')}{(\sigma, !e) \mapsto (\sigma', !e')} \quad \frac{\sigma(l) = v}{(\sigma, !\text{loc}[l]) \mapsto (\sigma, v)}$$

$$\frac{(\sigma, e_1) \mapsto (\sigma', e'_1)}{(\sigma, e_1 \text{ gets } e_2) \mapsto (\sigma', e'_1 \text{ gets } e_2)} \quad \frac{v_1 \text{ value} \quad (\sigma, e_2) \mapsto (\sigma', e'_2)}{(\sigma, v_1 \text{ gets } e_2) \mapsto (\sigma', v_1 \text{ gets } e'_2)}$$

$$\frac{v_2 \text{ value}}{(\sigma, \text{loc}[l] \text{ gets } v_2) \mapsto (\sigma[l := v_2], v_2)}$$

The non-modularity of this extension comes from the fact that it forces us to rewrite the rules describing addition as well:

$$\frac{(\sigma, e_1) \mapsto (\sigma', e'_1)}{(\sigma, e_1 + e_2) \mapsto (\sigma', e'_1 + e_2)}$$

$$\frac{v_1 \text{ value} \quad (\sigma, e_2) \mapsto (\sigma, e'_2)}{(\sigma, v_1 + e_2) \mapsto (\sigma', v_1 + e'_2)} \quad \frac{n_1 + n_2 = n_3}{(\sigma, \text{num}(n_1) + \text{num}(n_2)) \mapsto (\sigma, \text{num}(n_3))}$$

2.1 Modularity and specification styles

Another interesting loss of modularity in SOS specifications appears when we introduce exceptions and exception handling ($e ::= \dots \mid \text{error} \mid \text{try } e_1 \text{ ow } e_2$).

$$\frac{e_1 \mapsto e_2}{\text{try } e_1 \text{ ow } e_2 \mapsto \text{try } e'_1 \text{ ow } e_2} \quad \frac{v_1 \text{ value}}{\text{try } v_1 \text{ ow } e_2 \mapsto v_1} \quad \frac{}{\text{try error ow } e_2 \mapsto e_2}$$

Our rules for addition do not need to be revised to be compatible with this extension, but in order to preserve type safety, we must provide a way for errors to “bubble up” through additions.

$$\frac{}{\text{error} + e_2 \mapsto \text{error}} \quad \frac{v_1 \text{ value}}{v_1 + \text{error} \mapsto \text{error}}$$

We can avoid this particular form of non-modularity while technically staying within the SOS framework. This is possible if we use an *abstract machine* style of specification. Abstract machine-style specifications have an explicit control stack (or continuation), usually denoted k , that represents unfinished parts of the computation; a control stack is a list of *continuation frames*; we introduce two continuation frames f in order to specify \mathcal{L}_0 :

- $(\square + e_2)$ is a frame waiting on the first argument to be evaluated to a value so that the evaluation of e_2 can begin, and
- $(v_1 + \square)$ is a frame holding the already-evaluated value v_1 waiting on the second argument to be evaluated to a value.

Basic abstract machine specifications have two kinds of states: an expression being evaluated on a stack ($k \triangleright e$) and a value being returned to the stack ($k \triangleleft v$). The abstract machine specification of the operational semantics of \mathcal{L}_0 has four rules:

$$\begin{aligned} k \triangleright \text{num}(n) &\mapsto k \triangleleft \text{num}(n) \\ k \triangleright (e_1 + e_2) &\mapsto (k, \square + e_2) \triangleright e_1 \\ (k, \square + e_2) \triangleleft v_1 &\mapsto (k, v_1 + \square) \triangleright e_2 \\ (k, \text{num}(n_1) + \square) \triangleleft \text{num}(n_2) &\mapsto k \triangleleft \text{num}(n_3) \quad (\text{if } n_1 + n_2 = n_3) \end{aligned}$$

Given this specification, we can define the behavior of exceptions by adding a new frame ($\text{try } \square \text{ ow } e_2$) and a new kind of state — ($k \blacktriangleleft$), a stack dealing with an error.

$$\begin{aligned} k \triangleright \text{try } e_1 \text{ ow } e_2 &\mapsto (k, \text{try } \square \text{ ow } e_2) \triangleright e_1 \\ (k, \text{try } \square \text{ ow } e_2) \triangleleft v_1 &\mapsto k \triangleleft v_1 \\ k \triangleright \text{error} &\mapsto k \blacktriangleleft \\ (k, \text{try } \square \text{ ow } e_2) \blacktriangleleft &\mapsto k \triangleright e_2 \\ (k, f) \blacktriangleleft &\mapsto k \blacktriangleleft \quad (\text{if } f \neq \text{try } \square \text{ ow } e_2) \end{aligned}$$

When it comes to control features like exceptions, abstract machine specifications are more modular. We did not have to specifically consider or modify the stack frames for addition (or for multiplication, or function application...) in order to introduce exceptions and exception handling; one rule (the last one above) interacts modularly with essentially any “pure” language features. Both the SOS specification and the abstract machine specification were reasonable ways of specifying pure features like addition, but the abstract machine specification better anticipates the addition of control features.

Observation 1. *Different styles of specification can allow different varieties of programming language features to be specified in a modular style.*

Another way of stating this observation is that we must *precisely* anticipate the language features dealing with state and control if SOS-style specifications are our only option. When we pull back and look at a choice between these two specification styles, we only need to *generally* anticipate whether we might need control features like exceptions or first-class continuations; if so, we should use an abstract machine specification.

2.2 Modularity and ambient state

Of course, an abstract-machine-style presentation does not solve the problem with mutable references that we discussed at first! Various attempts have been made to address this problem. The Definition of Standard ML used an ad-hoc convention that, applied to our example, would allow us to write the original \mathcal{L}_0 rules while actually meaning that we were writing the state-annotated version of the rules [21, p. 40]. In this case, the convention is less about modular specification and more about not writing hundreds and hundreds of basically-obvious symbols in the language definition.

Mosses’s Modular Structural Operational Semantics (MSOS) attempts to formalize and standardize this sort of convention [22]. A transition in MSOS is written as $(e_1 \xrightarrow{X} e_2)$, where X is an open-ended description of the ambient state. By annotating congruence rules with $X = \{\dots\}$, the MSOS specification of \mathcal{L}_0 specifies that the congruence rules simply pass on whatever effects happen in subexpressions. The reduction rule is annotated with $X = \{\leftarrow\}$ to indicate that it has no effect on the ambient state.

$$\frac{v_1 \text{ value} \quad e_2 \xrightarrow{\{\dots\}} e'_2}{v_1 + e_2 \xrightarrow{\{\dots\}} v_1 + e'_2} \quad \frac{n_1 + n_2 = n_3}{\text{num}(n_1) + \text{num}(n_2) \xrightarrow{\{\leftarrow\}} \text{num}(n_3)}$$

Reduction rules that actually access some part of the ambient state, such as reduction rules for referencing the store, can mention and manipulate the store as necessary:

$$\frac{\sigma(l) = v}{! \text{loc}(l) \xrightarrow{\{\sigma, \leftarrow\}} v} \quad \frac{v_2 \text{ value}}{\text{loc}(l) \text{ gets } v_2 \xrightarrow{\{\sigma, \sigma' = \sigma[l := v_2], \leftarrow\}} v_2}$$

There is a good bit of related work along the lines of MSOS; for example, Implicitly-Modular Structural Operational Semantics (I-MSOS) is an extension of MSOS that partially removes the need for the explicit annotations X [23]. In all cases, however, the goal is to provide a regular notion of ambient state that allows each part of a language specification to leave irrelevant parts of the state implicit and open-ended.

Observation 2. *A framework that provides an open-ended notion of ambient state assists in writing modular programming language specifications.*

2.3 Modularity and the LF context

Logical frameworks in the LF family (such as LF [11], LLF [1], OLF [29], CLF [40], and HLF [30]) have an inherent notion of ambient information provided by the framework’s *context*. As an example, consider the following LF specification of the static semantics of \mathcal{L}_0 (we use the usual convention in which capital letters are understood to be universally quantified):

of : exp \rightarrow tp \rightarrow type.
of/num : of (num N) int.
of/plus : of E₁ int \rightarrow of E₂ int \rightarrow of (plus E₁ E₂) int.

The contents of the LF context are specified independently from the specification of the inference rules. If we declare the LF context to be empty, our specification corresponds to an “on-paper” specification of the language’s static semantics that looks like this:

$$\frac{}{\text{num}(n) : \text{int}} \quad \frac{e_1 : \text{int} \quad e_2 : \text{int}}{e_1 + e_2 : \text{int}}$$

This assumption that there is nothing interesting in the context, frequently called a *closed world assumption*, is not the only option. Another option is to declare that the LF context may include free expression variables x in tandem with an assumption that $x:\tau$ for some type τ . In this case, the LF specification written above corresponds to the following on-paper specification:

$$\frac{}{\Gamma, x:\tau \vdash x : \tau} \quad \frac{}{\Gamma \vdash \text{num}(n) : \text{int}} \quad \frac{\Gamma \vdash e_1 : \text{int} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash e_1 + e_2 : \text{int}}$$

Language features that incorporate binding require a variable context in the specification of their static semantics; as a result, our original on-paper static semantics could not be composed with an on-paper static semantics of language features like let-expressions and first-class functions. But because the contents of the LF context are open-ended, we *can* combine the LF specifications of the static semantics. The specification we wrote down in LF is open-ended in precisely the way we need it to be, even though the on-paper version is not.

Observation 3. *The LF context provides an open-ended notion of ambient information that enhances the modularity of certain specifications.*

Unfortunately, the open-endedness we describe here is useful for capturing ambient *information* in a static semantics; it is not useful for capturing ambient *state* in a dynamic semantics. The LF context only supports the addition of new assumptions that are persistent and unchanging, but capturing the dynamic semantics of stateful features like mutable references requires adding, modifying, and removing information about the state of the specified system. The solution is, in principle, well understood: logical frameworks based on substructural logics (especially linear logic) can capture stateful change — it is possible to remove facts from a linear logic context. Therefore, assumptions in a linear logic context can be thought of as representing *ephemeral* facts about the state of a system, facts which might change, as opposed to persistent assumptions which can never be removed or modified. Examples of logical frameworks that provide an open-ended notion of ambient state through the use of a linear context include Forum [18], Lolli [14], LLF [1], and CLF [40]. Furthermore, all of these frameworks have been given an operational interpretation as logic programming languages, so specifications are, at least in theory, executable.

2.4 Modularity in substructural operational semantics

These observations motivate the design of substructural operational semantics and, in turn, the design of the logical framework. We can give a initial picture of SSOS specifications at this point by thinking of them as collections of string rewriting rules where, for instance, the rewrite rule $a \cdot b \rightarrow c$ allows the string $a \cdot a \cdot b \cdot b$ to transition to the string $a \cdot c \cdot b$.

SSOS specifications resemble the abstract machine-style specifications above, but instead of capturing the entire control stack k as a single piece of syntax, each continuation frame f is captured by a token $\text{cont}(f)$ in the string. The token $\text{eval}(e)$ represents an evaluating expression e , and is analogous to the state $(k \triangleright e)$ in the abstract machine specification, whereas the token $\text{retn}(v)$ represents a returned value and is

analogous to the state $(k \triangleleft v)$ in the abstract machine specification. The resulting rules are quite similar to the abstract machine rules:

$$\begin{aligned}
& \text{eval}(\text{num}(n)) \rightarrow \text{retn}(\text{num}(n)) \\
& \text{eval}(e_1 + e_2) \rightarrow \text{cont}(\square + e_2) \cdot \text{eval}(e_1) \\
& \text{cont}(\square + e_2) \cdot \text{retn}(v_1) \rightarrow \text{cont}(v_1 + \square) \cdot \text{eval}(e_2) \\
& \text{cont}(\text{num}(n_1) + \square) \cdot \text{retn}(\text{num}(n_2)) \rightarrow \text{retn}(\text{num}(n_3)) \quad (\text{if } n_1 + n_2 = n_3)
\end{aligned}$$

We can then give modular specifications of exceptions and exception handling much as we did in the abstract machine for control, and the non-local nature of the specifications means that we can add state. To give a very simple example that does not generalize well, we can modularly extend the language with a single counter; evaluating the expression `get` gets the current value of the counter, and evaluating the expression `inc` gets the current value of the counter and then increments it by one. The value of the counter is stored in a token `counter(n)` that appears to the right of the `eval(e)` or `retn(v)` tokens (whereas the stack grows off to the left).

$$\begin{aligned}
& \text{eval}(\text{get}) \cdot \text{counter}(n) \rightarrow \text{retn}(\text{num}(n)) \cdot \text{counter}(n) \\
& \text{eval}(\text{inc}) \cdot \text{counter}(n) \rightarrow \text{retn}(\text{num}(n)) \cdot \text{counter}(n + 1)
\end{aligned}$$

Just as the rules need only reference the parts of the control stack that they modify, the rules only need to mention the counter if they access or modify that counter in some way. Again, this is a terrifically non-generalizable way of extending the system with state, because the use of ambient state is not “open-ended” — where would we put the *second* counter? Nevertheless, it is a specification of a stateful feature that can be composed modularly with the specification of addition.

After we motivate and present the logical framework we plan to use in the next two sections, we will return to this style of specification and see how it can be used to give modular specifications of exceptions and parallel evaluation.

3 Logics of deduction and transition

The aforementioned logical frameworks that provide an open-ended notion of ambient state through the use of substructural logics — Forum, CLF, etc. — can be difficult to use for certain types of reasoning about properties of systems specified in the framework. One source of difficulty is that many properties of specified systems are naturally described as properties of the *partial* proofs that arise during proof search, but in all of these frameworks it is more natural to reason about properties of *complete* proofs — put another way, SSOS-like specifications are properly understood *small-step* specifications, but the logical frameworks only make it easy to reason about *big-step* properties of those specifications. In this section, we discuss a solution to this mismatch: we define linear logic as a state transition system, at which point it is natural to talk about sequences of transitions (which do, in fact, correspond to partial proofs in a sequent calculus) as first-class entities. The only other solutions that I am aware of are based on Ludics, a presentation of logic that treats proofs with missing “leaves” as objects of consideration [10].

Underlying logical frameworks like LF and LLF is a particular justification of logic, dating back to Martin-Löf’s 1983 Siena Lectures [15], which holds that the meaning of a proposition is captured by what counts as a verification of that proposition. In Section 3.1 I briefly revisit Martin-Löf’s verificationist meaning-theory and how it relates to the use of canonical forms to adequately capture deductive systems.

There is another presentation of logic — first outlined by Pfenning [25] and developed further in Section 3.2 — which underlies a line of work on using forward-chaining in substructural logics for logic programming and logical specification [36, 37, 28]. Substructural contexts are treated as descriptions of the ephemeral state of a system, and the meaning of a proposition is captured by the state transitions it generates.

The integration of these two ways of understanding logic requires care, as a naïve approach for combining the two paradigms will destroy desirable properties of both the canonical-forms-based framework and the state-transition-based framework. In Section 3.3, I explain my proposed approach, which is based on the observations of adjoint logic [31].

3.1 A canonical-forms-based view of logic

Existing work on the proof theory of logical frameworks, including the LF family of logical frameworks [11], is based in part on a *verificationist* philosophy described in Martin-Löf’s Siena Lectures [15]. A verificationist meaning-theory provides that the meaning of a proposition is precisely captured by its introduction rules — that is, by the ways in which it may be verified. A proposition $A \wedge B$ is true, then, precisely if A is true and B is true, and a proposition $A \supset B$ is true precisely if we can verify B given an additional assumption that A is true.

The elimination rules, on the other hand, must be *justified* with respect to the introduction rules. We justify the elimination rules for a connective in part by checking their *local soundness* and *local completeness* [26]. Local soundness ensures that the elimination rules are not too strong: if an introduction rule is immediately followed by an elimination rule, the premise gives us all the necessary evidence for the conclusion.

$$\frac{\frac{\mathcal{D}_1 \quad \mathcal{D}_2}{A \text{ true} \quad B \text{ true}} \wedge I}{\frac{A \wedge B \text{ true}}{A \text{ true}} \wedge E_1} \Rightarrow_R \quad \mathcal{D}_1 \quad A \text{ true}$$

Local completeness, on the other hand, ensures that the elimination rules are not too weak: given a derivation of the compound connective, we can use elimination rules to get all the evidence necessary to reapply the introduction rule:

$$A \wedge B \text{ true} \stackrel{\mathcal{D}}{\Rightarrow_R} \frac{\frac{A \wedge B \text{ true}}{A \text{ true}} \wedge E_1 \quad \frac{A \wedge B \text{ true}}{B \text{ true}} \wedge E_2}{A \wedge B \text{ true}} \wedge I$$

3.1.1 Verifications and canonical forms

An examination of the requirement that introduction rules *precisely* define the connectives leads to a restricted set of proofs called *verifications*. If there is a proof of the truth of $A \wedge B$, then we don’t know a great deal about its proof. Maybe it looks like this!

$$\frac{\frac{\frac{\vdots}{D \wedge (A \wedge B) \text{ true}}{A \wedge B \text{ true}} \wedge E_2 \quad \frac{\vdots}{C \text{ true}} \wedge I}{(A \wedge B) \wedge C \text{ true}} \wedge I}{A \wedge B \text{ true}} \wedge E_1$$

On the other hand, if we have a *verification* of $A \wedge B$ then we know that the last step in its proof combines two (smaller) verifications, one of which is a verification of A and the other of which is a verification of B . We consider the verifications to be the canonical proofs of a proposition — when we introduce proof terms, verifications correspond to a restricted set of proofs called the *canonical forms*.

3.1.2 Definitions, atomic propositions, and adequacy

The canonical forms are interesting from the point of view of logical frameworks because they allow us to capture a representative of a derivation as a proof term within the logic. To use a standard example, we can introduce a new atomic proposition ($\text{add } N M P$) that is defined by the two constants add/z and add/s . We continue to use the convention that identifiers beginning with capital letters represent variables that are implicitly universally quantified.

```

z : nat.
s : nat → nat.

add : nat → nat → nat → type.
add/z : add z N N.
add/s : add N M P → add (s N) M (s P).

```

Given this signature, we would like to say that we also know what counts as a verification of the proposition $\text{add } (s z) (s z) (s(s z))$ — in particular, that only one thing will do: a use of add/s and a verification of $\text{add } z (s z) (s z)$. But this is not the case! We can also verify a proposition $\text{add } (s z) (s z) (s(s z))$ by using an *assumption* of the same proposition, or by using an assumption of the form $\text{add } N M P \rightarrow \text{add } N M P$ and a verification of $\text{add } (s z) (s z) (s(s z))$, or by using an assumption of the form $\text{add } N (s M) P \rightarrow \text{add } N M P$ and a verification of $\text{add } (s z) (s(s z)) (s(s z))$. . . The point is that the very open-endedness of the LF context that we pointed out in Observation 3 is preventing us from making the claim that we know something specific about what counts as a verification of the proposition $\text{add } (s z) (s z) (s(s z))$.

However, it is intuitively reasonable enough to require that our current assumptions include neither assumptions of the form $\text{add } N M P$ nor assumptions that can be immediately used to prove $\text{add } N M P$ (which is effectively a *closed world assumption*). Under this requirement, our intuition for what counts as a verification of the proposition $\text{add } (s z) (s z) (s(s z))$ is correct. Considerations such as these lead to a formal notion of *adequate encodings*, which in this case means that there is a bijection between terms of the type $\text{add } N M P$ and derivations of $N + M = P$ as defined by the following deductive system:

$$\frac{N + M = P}{s N + M = s P} \text{ add}/s \quad \frac{}{z + N = N} \text{ add}/z$$

There is, as the example of a static semantics from Section 2 illustrates, tension between considerations of modularity (which dictate that the form of the context should be as open-ended as possible) and adequacy (which dictate that the form of the context should be as fully specified as possible). However, experience shows that, by incorporating concepts such as subordination [39], this tension is unproblematic in practice.

The above discussion is informal, and adequacy is given a much more careful treatment elsewhere [11, 12]. The point of the preceding discussion is to introduce the use of canonical forms for specifying inductively defined systems and to underscore the necessity of imposing some sort of regularity constraint on the context.

3.2 A state-transition-based view of logic

The presentation of logic as verifications and canonical forms treats atomic propositions and their proofs as the primary objects of consideration. We are interested in treating the *state of a system* and *transitions between states* as the primary objects of consideration. The result has strong similarities to sequent calculus presentations of linear logic, which is unsurprising: previous work has universally started with a sequent presentation of logic and then “read off” a state transition system. Cervesato and Scedrov’s multiset rewriting language ω is a prime example of this process; an introduction to ω and a review of related work in the area can be found in their article [3].

We will treat the state of a specified system as being defined by a set of *ephemeral facts* about that state; we think of these ephemeral facts as *resources*. Therefore, states are multisets of ephemeral facts ($A \text{ res}$); the empty multiset is written as (\cdot) and multiset union is written as (Δ_1, Δ_2) . We define the meaning of propositions by the effect they have on resources. The proposition $A \& B$ represents a resource that can provide *either* a resource A or a resource B , and \top is an unusable resource, so the resource $(\top \text{ res})$ will always stick around uselessly. The proposition $\forall x:\tau.A(x)$ represents a resource that can provide the resource $A(t)$ for any term t with type τ (we assume there is some signature Σ defining base types and constants).

$$\begin{aligned}
 (\&_{T1}) \quad & A \& B \text{ res} \rightsquigarrow A \text{ res} \\
 (\&_{T2}) \quad & A \& B \text{ res} \rightsquigarrow B \text{ res} \\
 (\text{no rule for } \top) \quad & \\
 (\forall_T) \quad & \forall x:\tau.A(x) \text{ res} \rightsquigarrow A(t) \text{ res} \quad \text{if} \quad \Sigma \vdash t : \tau
 \end{aligned}$$

Linear implication represents a sort of test. It simplifies matters if, initially, we only consider linear implication $(Q \multimap A)$ where Q is an atomic proposition; it means that we can consume Q to produce A .

$$(\multimap_{T1}) \quad Q \text{ res}, Q \multimap A \text{ res} \rightsquigarrow A \text{ res}$$

The transition relation is, in general, $\Delta_1 \rightsquigarrow \Delta_2$, and there is a congruence or frame property that says a transition can happen to any part of the state.

$$\frac{\Delta_1 \rightsquigarrow \Delta'_1}{\Delta_1, \Delta_2 \rightsquigarrow \Delta'_1, \Delta_2} \text{ frame}$$

We also may allow the signature Σ to mention certain resources as being valid or unrestrictedly available; in this case we can spontaneously generate new copies of the resource:

$$(\text{copy}) \quad \cdot \rightsquigarrow A \text{ res} \quad \text{if} \quad (A \text{ valid}) \in \Sigma$$

This is a very brief version of this story, and in its current telling it can be understood as a subset of Cervesato and Scedrov’s ω multiset rewriting language [3]. However, that language and other similar languages take the notion of derivability in a linear logic sequent calculus as primary and derive a transition semantics from open proofs. The story we have told here starts with transitions and generalizes to all of first-order linear logic save for additive disjunction $(A \oplus B)$ and its unit $\mathbf{0}$.¹ I believe there is an inherent value in reducing the gap between the specification of linear logic and the specification of systems within that logic, and having transitions as a first-class notion within the logical framework appears to clarify many issues related to reasoning about specifications.

¹Furthermore, I have made substantial, though unfinished, progress on incorporating these connectives into a transition-based presentations as well.

3.2.1 Canonical transitions

Our motivation for describing logic as a transition system is to make it more natural to use logic for the specification of transition systems. An example of a very simple transition system that we would hope to represent is the following bare-bones asynchronous pi-calculus represented by the usual structural congruences and the following reduction rule:

$$c(x).p \parallel c\langle v \rangle \mapsto p[v/x]$$

To represent this system in logic, we can show that the above transition system corresponds exactly to the following signature where there are two sorts of terms, channels `chan` and processes `process`. A process can be either asynchronous send along a channel (send `CV`, i.e. $c\langle v \rangle$) or a receive along a channel (recv `C` ($\lambda x. P x$), i.e. $c(x).p$). The resource (proc `P`) represents the active process `P`.

`send` : `chan` \rightarrow `chan` \rightarrow `process`.

`recv` : `chan` \rightarrow (`chan` \rightarrow `process`) \rightarrow `process`.

`proc` : `process` \rightarrow `type`.

`synch` : `proc(recv C` ($\lambda x. P x$)) \multimap `proc(send C V)` \multimap `proc(P V)`.

What do we mean by *corresponds exactly*? Perhaps something like weak bisimulation or strong bisimulation would do, but we really want something akin to the adequacy properties discussed earlier. Adequacy in this setting means that there should be a bijective correspondence between states in linear logic and states in the simple process calculus; furthermore, transitions in both systems must respect this correspondence — any single transition that is made in the logic can be mapped onto a single transition in the process calculus, and vice-versa.

This, certainly, is not the case in system we have described so far. Take the synchronization transition ($a(x).x\langle x \rangle \parallel a\langle b \rangle$) \mapsto $b\langle b \rangle$ in the simple process calculus. We can say that $a(x).x\langle x \rangle$ is represented as the atomic proposition `proc(recv a` $\lambda x. \text{send } x x$) and that $a\langle b \rangle$ is represented as the atomic proposition `proc(send a b)`. However, because the `synch` rule has two premises, fully applying it will take at least two transitions in linear logic (five if we count the three transitions necessary due to the implicit quantification of `C`, `P`, and `V` in the `synch` rule).

The analogue to canonical forms and verifications in this setting addresses this problem. A *focused* presentation of our linear logic allows us to restrict attention to transitions that atomically copy a proposition from the signature, instantiate all of its variables, and satisfy all of its premises. Then we can say that the `synch` rule defined in the signature corresponds to the following *synthetic transition*:

$$\text{(synch)} \quad \text{proc}(\text{send } C V) \text{ res}, \text{proc}(\text{recv } C \lambda x. P x) \text{ res} \rightsquigarrow \text{proc}(P V) \text{ res}$$

The synthetic transition associated with `synch` precisely captures the transition $c(x).p \parallel c\langle v \rangle \mapsto p[v/x]$ in the process calculus, so transitions in the encoded system are adequately represented by transitions in the logical framework. In order to ensure that states are adequately represented in the framework, we again must impose a constraint on the regular structure of the context, namely that it only consists of atomic resources of the form (proc `P`).

3.3 Combining transitions and canonical forms

In Section 3.1, we saw that deductive systems, such as the addition of unary natural numbers, can be adequately represented in a logical framework based on verifications and canonical forms. In Section 3.2 we

saw that transition systems can be characterized by an interpretation of linear logic that treats the logic as a transition system. However, a transition system frequently needs to refer to some notion of inductive definition as (essentially) a side condition. As a slightly contrived example, consider a transition system where each resource contains a natural number and where two distinct numbers can be added together:

$$\text{num } N_1 \text{ res}, \text{ num } N_2 \text{ res} \rightsquigarrow \text{num}(N_3) \text{ res} \quad (\text{if add } N_1 \ N_2 \ N_3 \text{ can be derived})$$

Recall that, in Section 2 when we presented the psedo-SSOS specification of a language with addition, we also used a side condition of this form, but now we have made it explicit that the *content* of the side condition is an inductively defined judgment.

3.3.1 Introduction to adjoint logic

In the next section, we will take a fragment of LF — a canonical forms-based logic for representing syntax and derivations (which are used to express side conditions) — with a transition-based presentation of ordered linear logic that is used to represent evolving systems. The basis for this combination is adopted from Reed’s adjoint logic [31], in which two syntactically distinct logics can be connected through unary connectives U and F . In the logical framework presented in the next section, these two logics will be a fragment of LF and a transition-based ordered linear logic, but in this section, we will consider a logic that combines persistent and linear logic:

$$\begin{array}{ll} \text{Persistent propositions} & P, Q, R ::= p \mid UA \mid P \supset Q \\ \text{Linear/ephemeral propositions} & A, B, C ::= a \mid FP \mid A \multimap B \end{array}$$

Following Reed (and in order to match up with the naturally sequent-calculus-oriented state transition presentation of logic) we will consider adjoint logic as a sequent calculus. Persistent contexts Γ hold persistent facts and linear contexts Δ hold ephemeral resources; the sequent $\Gamma \vdash P$ characterizes proofs of persistent propositions and the sequent $\Gamma; \Delta \vdash A$ characterizes proofs of linear (ephemeral) propositions. Reading the rules in a bottom-up fashion, the linear goal FP is only true if there are no linear hypotheses and the persistent proposition P is provable; the linear hypothesis FP can be replaced by the persistent hypothesis P .

$$\frac{}{\Gamma; a \vdash a} \quad \frac{\Gamma \vdash P}{\Gamma; \cdot \vdash FP} \quad \frac{\Gamma, P; \Delta \vdash C}{\Gamma; \Delta, FP \vdash C} \quad \frac{\Gamma; \Delta, A \vdash B}{\Gamma; \Delta \vdash A \multimap B} \quad \frac{\Gamma; \Delta_A \vdash A \quad \Gamma; \Delta, B \vdash C}{\Gamma; \Delta, \Delta_A, A \multimap B \vdash C}$$

The description of persistent connectives below has one quirk: there are two left rules for $(P \supset Q)$, one where the ultimate goal is to prove a persistent proposition and one where the ultimate goal is to prove a linear proposition. This is not a quirk that is unique to adjoint logic; it is the same issue that requires Pfenning and Davies’ judgmental S4 to have two elimination rules for $\Box A$ [26], for instance.²

$$\frac{\frac{p \in \Gamma}{\Gamma \vdash p} \quad \frac{\Gamma; \cdot \vdash A}{\Gamma \vdash UA} \quad \frac{(UA) \in \Gamma \quad \Gamma; \Delta, A \vdash C}{\Gamma; \Delta \vdash C} \quad \frac{\Gamma, P \vdash Q}{\Gamma \vdash P \supset Q}}{\frac{(P \supset Q) \in \Gamma \quad \Gamma \vdash P \quad \Gamma, Q \vdash R}{\Gamma \vdash R} \quad \frac{(P \supset Q) \in \Gamma \quad \Gamma \vdash P \quad \Gamma, Q; \Delta \vdash C}{\Gamma; \Delta \vdash C}}$$

As Reed notes, if we erase every persistent proposition save for UA , the resulting logic is equivalent to linear logic with exponentials, where $!A \equiv FUA$. Similarly, if we erase every linear proposition save for FP , the resulting logic is equivalent to the Pfenning-Davies reconstruction of lax logic where $\circ P \equiv UFP$.

²Another way to look at this “quirk” is that the *real* rule is parametric in the conclusion, able to prove either $A \text{ res}$ or $P \text{ valid}$, and we just give the two instances of the real rule here.

3.3.2 Combining transitions and canonical forms with adjoint logic

My approach in the following section is to use adjoint logic to present a framework that connects a persistent logical framework of canonical forms and a substructural logical framework of transitions. The result, as we have just seen, combines elements of substructural logic with elements of lax logic, so it is not surprising that the end result turns out to be very similar to the CLF logical framework that (in a rather different way) also segregates transitions and deductions using a combination of lax and substructural logic.³ It is even closer to the *semantic effects* fragment of CLF, which requires the canonical forms-based fragment of the language to forgo any reference to the state-transition-based fragment [5].

The framework I present in the next section has no way of incorporating transitions within canonical forms; syntactically, this is achieved by removing the proposition UA from the framework. This exclusion decreases the expressiveness of the framework relative to CLF, but brings it closer in line to the semantic effects fragment. The connective F is then the only connection between the two systems, and since it retains much of the character of the linear logic exponential, it is written as $!P$ even though it actually acts only as the “first half” of the exponential.

4 A logical framework for evolving systems

In the previous section, we considered the use of logic to describe both deductive systems and state-transition systems, and our immediate goal is to use a logical framework based on these principles to specify the operational semantics of programming languages with state and concurrency. In this section we will present that logical framework, which based on a substructural logic that includes both ordered and linear propositions.

In addition to CLF, the design synthesizes ideas from Polakow et al.’s ordered logic framework [29], Reed’s adjoint logic [31], and Zeilberger’s polarized higher-order focusing [41]. There is no fundamental reason why the canonical forms fragment of the framework described here cannot be the full dependently typed logical framework LF [11, 12]; however, it is convenient for our current purposes to just consider a restricted fragment of LF.

4.1 Representing terms

We have already been using LF-like representations of terms: both natural numbers and process-calculus terms can be adequately represented as canonical forms in a simply-typed lambda calculus under the following signature.

$\text{nat} : \text{type}.$	$\text{process} : \text{type}.$
$z : \text{nat}.$	$\text{chan} : \text{type}.$
$s : \text{nat} \rightarrow \text{nat}.$	$\text{send} : \text{chan} \rightarrow \text{chan} \rightarrow \text{process}.$
	$\text{recv} : \text{chan} \rightarrow (\text{chan} \rightarrow \text{process}) \rightarrow \text{process}.$

A term type τ is either an atomic term type a defined by the signature (like nat , process , and chan) or an implication $\tau \rightarrow \tau$. The signature also defines term constants (like z , s , send , and recv) which are defined to have some type τ . Canonical terms are the η -long, β -normal terms of the simply-typed lambda calculus

³I believe that adjoint logic provides a better explanation than lax logic for how CLF is actually used and thought about, but I should re-emphasize that the critical difference between CLF and the framework presented in the next section — the reason I cannot use CLF directly — is that CLF does not treat individual state transitions and sequences of state transitions as objects of consideration.

using constants drawn from the signature: the natural number 3 is represented as $s(s(sz))$ and the process $a(x).x(x)$ is represented as the canonical term $(\text{recv } a \lambda x. \text{ send } x x)$.

4.2 Judgments as types

An organizing principle of LF-like frameworks is the representation of *judgments as types* — in the case of the aforementioned example of addition, the three-place judgment $n_1 + n_2 = n_3$ is represented by the type family add of kind $\text{nat} \rightarrow \text{nat} \rightarrow \text{nat} \rightarrow \text{type}$, and for any specific n_1, n_2 , and n_3 we represent derivations of $n_1 + n_2 = n_3$ as canonical proof terms of type $(\text{add } n_1 n_2 n_3)$ in the following signature:

$\text{add} : \text{nat} \rightarrow \text{nat} \rightarrow \text{nat} \rightarrow \text{type}.$
 $\text{add}/z : \prod N:\text{nat}. \text{add } z N N.$
 $\text{add}/s : \prod N:\text{nat}. \prod M:\text{nat}. \prod P:\text{nat}. \text{add } N M P \rightarrow \text{add } (s N) M (s P).$

In this signature, the proof term $(\text{add}/s z (sz) (sz) (\text{add}/z (sz)))$ acts as a representative of this derivation:

$$\frac{\frac{}{z + sz = sz} \text{add}/z}{sz + sz = s(sz)} \text{add}/s$$

More generally, type families are declared as $(a : \tau_1 \rightarrow \dots \rightarrow \tau_n \rightarrow \text{type})$ in the signature; atomic types in the type family a are written as $a t_1 \dots t_n$, where each of the arguments t_i has type τ_i . Types A are either atomic types, implications $A \rightarrow B$, or universal quantifications over terms $\prod x:\tau. A$. We usually leave universal quantification implicit in the signature, which allows us to write rules as we did earlier:

$\text{add}/z : \text{add } z N N.$
 $\text{add}/s : \text{add } N M P \rightarrow \text{add } (s N) M (s P).$

Furthermore, we leave the instantiation of implicit quantifiers implicit as well, which allows us to write the proof term corresponding to the above derivation more concisely as $(\text{add}/s \text{add}/z)$.

4.3 States as contexts

In the transition-based fragment of the logic, the states of the systems we intend to specify are represented as collections of atomic propositions. In sequent calculus-based logical frameworks, these collections correspond to *contexts*; we adopt this terminology.

Ephemeral contexts Linear logic and other substructural logics are useful for describing systems that change over time; part of the reason for this is that a linear resource is effectively an *ephemeral* fact about the state of a system — it describes something that is true in the current state of the system but which may not be true in a future state of the system. Our framework considers two kinds of ephemeral facts: *linear* facts (insensitive to their position relative to other facts) and *ordered* facts (sensitive to their position relative to other ordered facts).

We incorporate ephemeral facts into our framework by introducing two different kinds of type family: linear type families (where we write type_l — l for linear — instead of type in the type family declaration) and ordered type families (where we write type_o — o for ordered — instead of type in the type family declaration). This effectively enforces a syntactic separation between those persistent, linear, and ordered

atomic propositions; the need for this syntactic separation has been noted in previous work [36, 37] but was only given a logical interpretation in [28].

An ephemeral context (written Δ) is syntactically a list of linear and ordered types, but we consider contexts to be equivalent if they can be made syntactically equal by reordering linear propositions. As an example, if a_l and b_l are linear atomic propositions and c_o and d_o are ordered atomic propositions, then the contexts (a_l, b_l, c_o, d_o) , (c_o, d_o, b_l, a_l) , and (c_o, a_l, d_o, b_l) are all equivalent, but they are not equivalent to (a_l, b_l, d_o, c_o) because the ordered propositions c_o and d_o appear in a different order. We write this equivalence on contexts as $\Delta \approx \Delta'$.

We can use ordered resources to represent the *concrete* syntax of the language \mathcal{L}_0 from Section 2 by defining an ordered atomic proposition for each syntactic token:

$+ : \text{type}_o$.
 $\text{num} : \text{nat} \rightarrow \text{type}_o$.

The proposition $+$ represents the addition token, and $\text{num } N$ (where N is a term of type nat) represents a token for the natural number N . While we cannot yet actually describe parsing, we can intuitively say that the ephemeral context $(\text{num}(s z), +, (\text{num } z))$ corresponds to the \mathcal{L}_0 program $1 + 0$, the ephemeral context $(\text{num}(s z), +, (\text{num } z), +, (\text{num}(s(s z))))$ is ambiguous and corresponds to either the \mathcal{L}_0 program $(1 + 0) + 2$ or the \mathcal{L}_0 program $1 + (0 + 2)$, and the ephemeral context $+, +, (\text{num}(s z))$ is not syntactically well-formed and does not correspond to any \mathcal{L}_0 program.

As in sequent calculus presentations of a logic, we annotate the propositions in a context with unique variable names. Therefore, the context corresponding to the first \mathcal{L}_0 program above could also have been written as $(x_1 : \text{num}(s z), x_2 : +, x_3 : \text{num } z)$. We further follow the convention of sequent calculus presentations of logic by often not mentioning these variable names when it is possible to reconstruct them.

Persistent contexts In addition to ephemeral state, is important to allow systems to have *persistent* state, which includes both dynamically introduced parameters and facts that, once true, necessarily stay true.⁴ We write free term parameters as $x : \tau$ and persistent assumptions as $x : A$, where A is a persistent type. A collection of these assumptions is a persistent context, written as Γ , and we treat all contexts containing the same terms and assumptions as equivalent, writing this equivalence as $\Gamma \approx \Gamma'$.

4.4 Substitutions

The connection point between contexts and terms is the definition of *substitutions*. A substitution σ is a term which gives evidence that, given the persistent facts Γ and ephemeral resources Δ , we can model or represent the state described by the persistent facts Γ' and ephemeral facts Δ' . Alternatively, we can think of Γ' and Δ' as “context-shaped holes,” in which case a substitution is evidence that Γ and Δ “fit in the holes.”

The way a substitution gives this evidence is by providing a resource (designated by its label y) for every for every resource in Δ' , providing a proof term \mathcal{D} that is defined in the context Γ for every fact $x : A$ in Γ' , and providing a term t that is defined in the context Γ for every $x : \tau$ in Γ' .⁵ Syntactically, a substitution is described by the following grammar:

$$\sigma ::= [] \mid \sigma, t/x \mid \sigma, \mathcal{D}/x \mid \sigma, y/x$$

⁴Parameters can, for instance, be used to model the channels in the process calculus from the previous section, or to model abstract heap locations in a specification of ML-style references. We need to be able to dynamically introduce new parameters in order to represent producing a new channel or allocating a new location on the heap.

⁵In an appropriate dependently-typed setting, term types τ and types A are members of the same syntactic class, as are terms t and proof terms \mathcal{D} .

The definition of substitution typing $\Gamma; \Delta \vdash_{\Sigma} \sigma : \Gamma'; \Delta'$ captures the judgment that σ is a witness to the fact that the state $\Gamma; \Delta$ can model or represent the state Γ', Δ' (under the signature Σ). The definition relies on two other judgments, $\Gamma \vdash_{\Sigma} t : \tau$ and $\Gamma \vdash_{\Sigma} \mathcal{D} : A$, which describe well-typed terms and well-typed proof terms, respectively. It also relies on the operation $A[\sigma]$, the application of a substitution to a type. Because we are using only a restricted fragment of LF, only the term components (t/x) of substitutions actually matter for this substitution.

$$\frac{}{\Gamma; \cdot \vdash_{\Sigma} [] : \cdot; \cdot} \quad \frac{\Gamma; \cdot \vdash_{\Sigma} \sigma : \Gamma'; \cdot \quad \Gamma \vdash_{\Sigma} t : \tau}{\Gamma; \cdot \vdash_{\Sigma} (\sigma, t/x) : (\Gamma', x : \tau); \cdot} \quad \frac{\Gamma; \cdot \vdash_{\Sigma} \sigma : \Gamma'; \cdot \quad \Gamma \vdash_{\Sigma} \mathcal{D} : A[\sigma]}{\Gamma; \cdot \vdash_{\Sigma} (\sigma, \mathcal{D}/x) : (\Gamma', x : A); \cdot}$$

$$\frac{\Delta \approx (\Delta'', y : Q') \quad Q' = Q[\sigma] \quad \Gamma; \Delta'' \vdash_{\Sigma} \sigma : \Gamma'; \Delta'}{\Gamma; \Delta \vdash_{\Sigma} (\sigma, y/x) : \Gamma'; (\Delta', x : Q)}$$

4.5 Positive types as patterns

In the discussion of a transition-based view of logic in Section 3.2, the only connectives we discussed were those that were defined in terms of their effect on the context; these are the connectives that make up the so-called *negative* propositions. Another family of connectives, which make up the so-called *positive* connectives, are more naturally defined by the contexts they describe. The pattern judgment $\Gamma; \Delta \Vdash p :: A^+$ expresses that p is the evidence that A^+ describes the context $\Gamma; \Delta$. The most natural way to think about Γ and Δ , which should appropriately be seen as an *output* of the judgment, is that it is describing a context-shaped hole that will be filled by a substitution in order to prove A^+ .

Ordered conjunction ($A^+ \cdot B^+$) describes a context containing, to the left, the context described by A^+ and then, to the right, the context described by B^+ .

$$\frac{\Gamma_1; \Delta_1 \Vdash p_1 :: A^+ \quad \Gamma_2; \Delta_2 \Vdash p_2 :: B^+}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \Vdash (p_1 \cdot p_2) :: (A^+ \cdot B^+)}$$

Existential quantification ($\exists x:\tau. A^+$) describes a context described by A^+ , plus a variable $x : \tau$. If we are thinking of the output contexts as context-shaped holes, then x is a hole that must be filled by a term of type τ . The rule below should make it clear that pattern judgments are fundamentally different from typing judgments, where we would expect the x to appear in the premise, not in the conclusion.

$$\frac{\Gamma; \Delta \Vdash p :: A^+}{x : \tau, \Gamma; \Delta \Vdash (x.p) :: (\exists x:\tau. A^+)}$$

Finally, $\mathbf{1}$ describes a context containing no ephemeral propositions, and we treat atomic propositions as positive as well: Q (where Q is an ordered atomic proposition) describes a context with just one ordered atomic proposition Q , $!Q$ (where Q is a linear atomic proposition) describes a context with just one linear atomic proposition Q , and $!A$ (where A is a persistent proposition) describes a context with no ephemeral propositions where A is true.

$$\frac{}{\cdot; \cdot \Vdash () :: \mathbf{1}} \quad \frac{}{x : A; \cdot \Vdash x :: !A} \quad \frac{}{\cdot; x : Q \Vdash x :: !Q} \quad \frac{}{\cdot; x : Q \Vdash x :: Q}$$

Example. Because we may think of the pattern judgment $\Gamma; \Delta \Vdash p :: A^+$ as producing context-shape holes Γ and Δ , looking at a pattern (which produces context-shape holes) together with a substitution (which provides evidence that those holes can be filled) is illustrative. If we want to show that the ephemeral

context $(x_1 : \text{num}(sz), x_2 : +, x_3 : \text{num } z)$ is described by the positive proposition $\exists n. (\text{num } n \cdot +) \cdot (\text{num } z \cdot !\text{add } n n (s(sz)))$, then we first find the pattern associated with that proposition (we write n instead of num for brevity):

$$\frac{\frac{\frac{\frac{\overline{; y_1 : n n \Vdash y_1 :: n n} \quad \overline{; y_2 : + \Vdash y_2 :: +}}{\cdot; y_1 : n n, y_2 : + \Vdash y_1 \cdot y_2 :: (n n \cdot +)}}{\quad} \quad \frac{\frac{\overline{; y_3 : n z \Vdash y_3 :: n z} \quad \overline{y_4 : \text{add } n n (s(sz)); \cdot \Vdash y_4 :: !\text{add } n n (s(sz))}}{\quad}}{\quad} \quad \frac{\overline{y_4 : \text{add } n n (s(sz)); y_3 : n z \Vdash (y_3 \cdot y_4) :: (n z \cdot !\text{add } n n (s(sz)))}}{\quad}}{\quad} \quad \frac{\overline{y_4 : \text{add } n n (s(sz)); y_1 : n n, y_2 : +, y_3 : n z \Vdash (y_1 \cdot y_2) \cdot (y_3 \cdot y_4) :: (n n \cdot +) \cdot (n z \cdot !\text{add } n n (s(sz)))}}{\quad}}{\quad} \quad \frac{\overline{n : \text{nat}, y_4 : \text{add } n n (s(sz)); y_1 : n n, y_2 : +, y_3 : n z \Vdash (n \cdot (y_1 \cdot y_2) \cdot (y_3 \cdot y_4)) :: \exists n. (n n \cdot +) \cdot (n z \cdot !\text{add } n n (s(sz)))}}{\quad}}{\quad}$$

Then, we find a substitution σ such that we can derive the following:

$$\cdot; (x_1 : \text{num}(sz), x_2 : +, x_3 : \text{num } z) \vdash \sigma : (n : \text{nat}, y_4 : \text{add } n n (s(sz))); (y_1 : \text{num } n, y_2 : +, y_3 : \text{num } z)$$

The substitution $\sigma = \square, (sz)/n, (\text{add}/s \text{ add}/z)/y_4, x_1/y_1, x_2/y_2, x_3/y_3$ works in this case.

In future examples, we will use a derived notation of a *filled pattern*, a substitution applied to a pattern. For example, rather than first deriving a pattern and then a substitution in the preceding example, we will simply say that the filled pattern $((sz) \cdot (x_1 \cdot x_2) \cdot (x_3 \cdot (\text{add}/s \text{ add}/z)))$ is a proof term showing that the ephemeral context $(x_1 : \text{num}(sz), x_2 : +, x_3 : \text{num } z)$ is described by the positive proposition $\exists n. (\text{num } n \cdot +) \cdot (\text{num } z \cdot !\text{add } n n (s(sz)))$.

4.6 Negative types as transition rules

Negative types were introduced in Section 3.2 as describing ways in which contexts can change. The proposition $b \cdot c \rightarrow \uparrow e$, for example, is essentially a rewriting rule that allows an ephemeral context such as $(a \ b \ c \ d)$ to be rewritten as an ephemeral context $(a \ e \ d)$.

The judgment associated with negative propositions is $(\Gamma; \Delta \Vdash g :: A^- \gg \Gamma'; \Delta' \Vdash p')$. This judgment expresses that, if the premises of A^- — the context-shaped holes Γ and Δ — can be filled by consuming some number of ephemeral resources, then the consumed ephemeral resources can be replaced with the conclusion of A^- — the resources described by Δ' and the new facts described by Γ' . The *goal* g captures the premises, and the pattern p describes the conclusion; together, a goal g and a pattern p are the proof term associated with a negative proposition.

The simplest negative proposition is $\uparrow A^+$, which consumes no resources and produces the resources described by A^+ :

$$\frac{\Gamma'; \Delta' \Vdash p' :: A^+}{\cdot; \cdot \Vdash () :: \uparrow A^+ \gg \Gamma'; \Delta' \Vdash p'}$$

The proposition $A^- \ \& \ B^-$ can represent either the transition described by A^- or the transition described by B^- :

$$\frac{\Gamma; \Delta \Vdash p :: A^- \gg \Gamma'; \Delta' \Vdash p'}{\Gamma; \Delta \Vdash (\pi_1 p) :: (A^- \ \& \ B^-) \gg \Gamma'; \Delta' \Vdash p'} \quad \frac{\Gamma; \Delta \Vdash p :: B^- \gg \Gamma'; \Delta' \Vdash p'}{\Gamma; \Delta \Vdash (\pi_2 p) :: (A^- \ \& \ B^-) \gg \Gamma'; \Delta' \Vdash p'}$$

Transitions have to be located at a particular point in the ephemeral context; ordered implication $(A^+ \rightarrow B^-)$ is a transition that consumes a part of the context described by A^+ to the right of its initial position and

then proceeds with the transition described by B^- .⁶

$$\frac{\Gamma_1; \Delta_1 \Vdash p :: A^+ \quad \Gamma_2; \Delta_2 \Vdash g :: B^- \gg \Gamma'; \Delta' \Vdash p'}{\Gamma_1, \Gamma_2; \Delta_1, \Delta_2 \Vdash (p; g) :: (A^+ \multimap B^-) \gg \Gamma'; \Delta' \Vdash p'}$$

Finally, universal quantification $\forall x:\tau. A^-$ can behave like $A^-[t/x]$ for any $t:\tau$; therefore, its pattern introduces a new

$$\frac{\Gamma; \Delta \Vdash g :: A^- \gg \Gamma'; \Delta' \Vdash p'}{x:\tau, \Gamma; \Delta \Vdash (x.g) :: (\forall x:\tau. A^-) \gg \Gamma'; \Delta' \Vdash p'}$$

When the placement of the shift operator \uparrow can be inferred from the context, we will leave it implicit, writing $A^+ \multimap B^+$ instead of $A \multimap \uparrow B^+$.

4.7 Transitions and expressions

We now can actually describe parsing the concrete syntax of \mathcal{L}_0 using transition rules. First, we describe the abstract syntax of \mathcal{L}_0 as canonical terms of type `exp`; then we describe an ordered type family `parsed E` and two rewriting rules describing how the concrete syntax can be parsed.

`exp` : type.
`n` : `nat` \rightarrow `exp`.
`plus` : `exp` \rightarrow `exp` \rightarrow `exp`.
`parsed` : `exp` \rightarrow `typeo`.
`parse/num` : `num N` \multimap `parsed(n N)`.
`parse/plus` : `parsed E1` \cdot `+` \cdot `parsed E2` \multimap `parsed(plus E1 E2)`.

Intuitively, `parse/num` is a rewriting rule that allows us to transition from the ephemeral context $(\text{num}(sz)), +, (\text{num } z)$ to the ephemeral context $(\text{parsed}(n(sz))), +, (\text{num } z)$. To describe the proof term capturing this transition, we first discover the goal and pattern associated with the type of `parse/num` (note that the implicit quantification and the implicit shift operator \uparrow are both made explicit):

$$\frac{\frac{\frac{(\cdot; y_1 : \text{parsed}(n n)) \Vdash y_1 :: \text{parsed}(n n)}{(\cdot; y : \text{num } n) \Vdash y :: \text{num } n} \quad (\cdot; \cdot) \Vdash () :: \uparrow \text{parsed}(n n) \gg (\cdot; y_1 : \text{parsed}(n n)) \Vdash y_1}{(\cdot; y : \text{num } n) \Vdash (y; ()) :: (\text{num } n \multimap \uparrow \text{parsed}(n n)) \gg (\cdot; y_1 : \text{parsed}(n n)) \Vdash y_1}}{(n : \text{nat}; y : \text{num } n) \Vdash (n.y; ()) :: (\forall n. \text{num } n \multimap \uparrow \text{parsed}(n n)) \gg (\cdot; y_1 : \text{parsed}(n n)) \Vdash y_1}}$$

If x_1 is the variable marking the ordered resource $(\text{parsed}(n(sz)))$, we can fill this goal with the substitution $([], sz/n, x_1/y)$; the resulting filled pattern is $(sz. x_1; ())$. The proof term corresponding to the transition can be written in a CLF-like notation as the let-expression $(\text{let } (y_1) = \text{parse/num}(sz. x_1; ()) \text{ in})$. However, I prefer a more “directional” notation for transitions that looks like this: $(\text{parse/num}(sz. x_1; ()) \triangleright (y_1))$. In addition, we will leave universal quantification implicit and omit the trailing $()$ in a goal; the final result is a proof term for the aforementioned transition that looks like this:

⁶In ordered logic there are actually two distinct implications, right ordered implication $A^+ \multimap B^-$ and left ordered implication $A^+ \multimap B^-$. Either one or the other is usually sufficient for the logical fragment we are considering, and I conjecture that as long as the ephemeral context only contains atomic propositions, the addition of left ordered implication does not actually add any expressiveness.

$$\begin{aligned} \text{parse/num}(x_1) \triangleright (y_1) \\ : (x_1 : \text{num}(\text{s z}), x_2 : +, x_3 : \text{num z}) \rightsquigarrow_{\Sigma \text{parse}} (y_1 : \text{parsed}(\text{n}(\text{s z})), x_2 : +, x_3 : \text{num z}) \end{aligned}$$

We'll return to the Σparse annotation in the next section; in the meantime, the typing rule for a transition $\mathcal{T} : (\Gamma; \Delta \rightsquigarrow_{\Sigma} \Gamma'; \Delta')$ is as follows:

$$\frac{c : A^- \in \Sigma \quad \Gamma_{in}; \Delta_{in} \Vdash g :: A^- \gg \Gamma_{out}; \Delta_{out} \Vdash p \quad \Delta \approx \Delta_L, \Delta', \Delta_R \quad \Gamma; \Delta' \vdash_{\Sigma} \sigma : \Gamma_{in}; \Delta_{in}}{c(g[\sigma]) \triangleright p : \Gamma; \Delta \rightsquigarrow_{\Sigma} \Gamma, \Gamma_{out}; \Delta_L, \Delta_{out}, \Delta_R}$$

This transition can be read like this:

- Starting from the state $\Gamma; \Delta$,
- Pick a rule A^- from the signature Σ (this is the constant c),
- Determine the input context $\Gamma_{in}; \Delta_{in}$ (this is the goal g) and the output context $\Gamma_{out}; \Delta_{out}$ (this is the pattern p) associated with A^- ,
- Split the ephemeral context Δ into three parts, Δ_L , Δ' , and Δ_R ,
- Show that, using the persistent facts Γ and the ephemeral resources Δ' you can fulfill the demands represented by the input pattern (this is the substitution σ), and
- Extend the persistent context with Γ_{out} and replace Δ' in the context with Δ_{out} to get the result of the transition.

4.8 Expressions

A sequence of one or more transitions \mathcal{T} is an expression $\mathcal{E} : (\Gamma; \Delta \rightsquigarrow_{\Sigma}^* \Gamma'; \Delta')$. An expression is either a single transition \mathcal{T} , the sequential composition of two expressions $\mathcal{E}_1; \mathcal{E}_2$, or the empty expression \diamond . We treat sequential composition as an associative operation with unit \diamond .

The following expression represents a complete parse of an \mathcal{L}_0 program:

$$\begin{aligned} \text{parse/num}(x_1) \triangleright (y_1); \text{parse/num}(x_3) \triangleright (y_3); \text{parse/plus}(y_1 \cdot x_2 \cdot y_3) \triangleright (y) \\ : (x_1 : \text{num}(\text{s z}), x_2 : +, x_3 : \text{num z}) \rightsquigarrow_{\Sigma \text{parse}}^* (y : \text{parsed}(\text{plus}(\text{n}(\text{s z}))(\text{n z}))) \end{aligned}$$

The annotation of a signature Σparse on the transition $\mathcal{T} : (\Gamma; \Delta \rightsquigarrow_{\Sigma} \Gamma'; \Delta')$ will be critical in practice. We frequently want to define multiple families of transitions that can act on a single proposition. For instance, consider the following rules which describe the direct calculation of a number from the concrete syntax of \mathcal{L}_0 (note, in particular, the use of $\text{!add } N_1 N_2 N_3$ to capture the “side condition” in calc/plus):

$$\begin{aligned} \text{calc} : \text{nat} \rightarrow \text{type}_o. \\ \text{calc/num} : \text{num } N \rightarrow \text{calc } N. \\ \text{calc/plus} : \text{calc } N_1 \cdot + \cdot \text{calc } N_2 \cdot \text{!add } N_1 N_2 N_3 \rightarrow \text{calc } N_3. \end{aligned}$$

Then we can talk about the transitions and expressions where we refer only to the transition rules beginning with calc .

$$\text{calc/num}(x_1) \triangleright (y_1) : (x_1 : \text{num}(\text{s z}), x_2 : +, x_3 : \text{num z}) \rightsquigarrow_{\Sigma \text{calc}} (y_1 : \text{calc}(\text{s z}), x_2 : +, x_3 : \text{num z})$$

$$\begin{aligned} \text{calc/num}(x_1) \triangleright (y_1); \text{calc/num}(x_3) \triangleright (y_3); \text{calc/plus}(y_1 \cdot x_2 \cdot y_3 \cdot (\text{add/s add/z})) \triangleright (y) \\ : (x_1 : \text{num}(s z), x_2 : +, x_3 : \text{num } z) \rightsquigarrow_{\Sigma_{\text{calc}}}^* (y : \text{calc}(s z)) \end{aligned}$$

We use these annotations to specify theorems about the relationship between two different transitions or expressions. For instance, if Σ_e describes the (not-yet-specified) SSOS specification of \mathcal{L}_0 , then we can write the following conjecture about the relationship between the rules starting with calc/\dots and the rules starting with parse/\dots .


Conjecture 1. *If $(\Delta \rightsquigarrow_{\Sigma_{\text{calc}}}^* y : \text{calc } N)$ and such that $(\Delta \rightsquigarrow_{\Sigma_{\text{parse}}}^* y : \text{parsed } E)$, then $(x : \text{eval } E \rightsquigarrow_{\Sigma_e}^* y : \text{retn}(n N))$.*

4.9 Concurrent equivalence

The two expressions below both represent a complete parse of the \mathcal{L}_0 program that we have been using in our example:

$$\begin{aligned} \text{parse/num}(x_3) \triangleright (y_3); \text{parse/num}(x_1) \triangleright (y_1); \text{parse/plus}(y_1 \cdot x_2 \cdot y_3) \triangleright (y) \\ \text{parse/num}(x_1) \triangleright (y_1); \text{parse/num}(x_3) \triangleright (y_3); \text{parse/plus}(y_1 \cdot x_2 \cdot y_3) \triangleright (y) \end{aligned}$$

In fact, these two expressions represent the *same* parse if we look at the parse not as a list of transitions but as a tree of string rewrites:



The notion of concurrent (or permutative) equivalence $\mathcal{E}_1 \approx \mathcal{E}_2$ captures the notion that, while we represent expressions as sequences, they are perhaps more appropriately thought of as directed acyclic graphs. Concurrent equivalence is the least equivalence relation such that

$$(\mathcal{E}_1; c(g[\sigma]) \triangleright p; c'(g'[\sigma']) \triangleright p'; \mathcal{E}_2) \approx (\mathcal{E}_1; c'(g'[\sigma']) \triangleright p'; c(g[\sigma]) \triangleright p; \mathcal{E}_2)$$

whenever σ does not include any variables bound by p' and σ' likewise does not mention any variables bound by p .

4.10 Notational conventions

We have discussed a number of notational conventions, such as omitting the shift operator \uparrow and simplifying the way that goals are written down. In discussions of SSOS in Section 5, we will use a more concise shorthand notation for describing transitions and expressions; this section will introduce this shorthand notation.

Consider the last transition in the parse of our example \mathcal{L}_0 program:

$$\begin{aligned} \text{parse/plus}(y_1 \cdot x_2 \cdot y_3) \triangleright (y) \\ : (y_1 : \text{parsed}(n(s z)), x_2 : +, y_3 : \text{parsed}(n z)) \rightsquigarrow_{\Sigma_{\text{parse}}} (y : \text{parsed}(\text{plus}(n(s z)) (n z))) \end{aligned}$$

The variables that describe what parts of the context changed are not strictly necessary here; we can convey the exact same information in a more concise form by omitting the mention of any variables:

$$\text{parse/plus} : \text{parsed}(n(sz)), +, \text{parsed}(nz) \rightsquigarrow_{\Sigma_{\text{parse}}} \text{parsed}(\text{plus}(n(sz))(nz))$$

This notation must be used carefully: for instance, if we say that the complete parse is represented by the expression $(\text{parse/num}; \text{parse/num}; \text{parse/plus})$, it is unclear whether the first token or the last token was rewritten first (although in this case, both of the traces are concurrently equivalent!)

It is not always the case that the name of a rule alone is sufficient. The exception is when one of the premises to a rule is a canonical form, such as the last step in the calculation of our example \mathcal{L}_0 program:

$$\begin{aligned} \text{calc/plus}(y_1 \cdot x_2 \cdot y_3 \cdot (\text{add/s add/z})) \triangleright (y) \\ : (y_1 : \text{calc}(sz), x_2 : +, y_3 : \text{calc } z) \rightsquigarrow_{\Sigma_{\text{calc}}} (y : \text{calc}(sz)) \end{aligned}$$

The shorthand version of this rule only mentions the canonical form $(\text{add/s add/z} : \text{add}(sz)z(sz))$ — unlike variable names, that part of the original transition conveys more information than just a position.

$$\text{calc/plus}(\text{add/s add/z}) : \text{calc}(sz), +, \text{calc } z \rightsquigarrow_{\Sigma_{\text{calc}}} \text{calc}(sz)$$

Finally, in the following section I will refer to states $\Gamma; \Delta$ generally as \mathcal{S} , and will use the notation $\mathcal{S}[\Delta]$ to describe a state that includes somewhere inside of it the ephemeral context Δ — a “one-hole context” over contexts, in other words. For instance, we can write the following:

$$\begin{aligned} \text{parse/num} : \mathcal{S}[\text{num}(s(sz))] \rightsquigarrow_{\Sigma_{\text{parse}}} \mathcal{S}[\text{parsed}(s(sz))] \\ \text{calc/num} : \mathcal{S}[\text{num}(s(sz))] \rightsquigarrow_{\Sigma_{\text{calc}}} \mathcal{S}[\text{calc}(s(sz))] \\ \text{calc/plus}(\text{plus/z}) : \mathcal{S}[\text{calc } z, +, \text{calc}(s(sz))] \rightsquigarrow_{\Sigma_{\text{calc}}} \mathcal{S}[\text{calc}(s(sz))] \end{aligned}$$

5 Substructural operational semantics

Now that we have discussed the outlines of our logical framework, we can finally return to the discussion of substructural operational semantics specifications that we motivated in the introduction and in Section 2. In this section, we consider a series of relatively simple substructural operational semantics specifications, mostly taken from [28]. Substructural operational semantics (or SSOS) is a style of specifying the operational semantics of programming languages as transitions in substructural logics. Its origin lies in a set of examples illustrating the expressiveness of the logical framework CLF [2], though elements were anticipated by Chirimar’s Ph.D. thesis [4] and by the continuation-based specification of ML with references in LLF [1]. A methodology of SSOS specifications was refined and generalized to include other substructural logics in subsequent work [24, 28].

In this section, we will develop a base language \mathcal{L}_1 that extends \mathcal{L}_0 with functions and function application (Section 5.1) and then extend that language with parallel evaluation of pairs (Section 5.3) and with exceptions and exception handling (Section 5.4). The novel content of this section is a presentation of the *static semantics* of this language and proof of language safety via progress and preservation lemmas that it enables. We give a proof of safety for \mathcal{L}_1 in Section 5.2, and then briefly discuss how that safety proof can be extended when we consider the static semantics of parallel evaluation and exceptions.

Fundamental components of SSOS specifications Many of the elements of substructural operational semantics specifications are immediately familiar to anyone who has written or read traditional SOS-style operational semantics formalized in Twelf. The syntax, for instance, consists of expressions E , types T , stack frames F , and a judgment (value V) capturing the subsort of expressions that are also values (we continue to write V for those expressions which are known to be values).

exp : type.
 tp : type.
 frame : type.
 value : exp \rightarrow type.

Particular to SSOS specifications are three groups of propositions. The *active* propositions are those that can always eagerly participate in a transition; the prototypical active proposition is (eval E), an ordered atomic proposition containing an expression E that we are trying to evaluate to a value.

eval : exp \rightarrow type_o.

The *latent* propositions represent suspended computations. The primary latent proposition we will consider is (cont F), which stores a part of a continuation (in the form of a stack frame) waiting for a value to be returned.⁷

cont : frame \rightarrow type_o.

Finally, *passive* propositions are those that do not drive transitions on their own, but which may take part in transitions when combined with latent propositions. The only passive proposition in the SSOS specification of an effect-free languages is (retn V), which holds a value being returned from the evaluation of an expression. On its own, it is passive, representing a completed computation. If there is a latent stack frame to its left, a transition should occur by returning the value to the stack frame.

retn : exp \rightarrow type_o.

5.1 Specifying \mathcal{L}_1

We will introduce SSOS specifications by encoding the operational semantics of the very simple language \mathcal{L}_0 that has been used as running example; we extend \mathcal{L}_0 with call-by-value functions to make the language less boring, and call the result \mathcal{L}_1 .

5.1.1 Syntax

The syntax of expressions can be given by the following BNF specification:

$$e ::= x \mid \lambda x.e \mid e_1(e_2) \mid n \mid e_1 + e_2$$

The encoding of that BNF is standard, including the use of higher-order abstract syntax to represent binding. The judgments v/n and v/lam indicate that functions and numbers are values.

⁷Latent propositions can also be seen as artifacts of kind of defunctionalization that is performed on “higher-order” SSOS specifications to make them more amenable to extension and analysis. (See Section ?? for a further discussion.)

$\text{lam} : (\text{exp} \rightarrow \text{exp}) \rightarrow \text{exp}.$
 $\text{app} : \text{exp} \rightarrow \text{exp} \rightarrow \text{exp}.$
 $n : \text{nat} \rightarrow \text{exp}.$
 $\text{plus} : \text{exp} \rightarrow \text{exp} \rightarrow \text{exp}.$
 $v/\text{lam} : \text{value}(\text{lam } \lambda x. E x).$
 $v/n : \text{value}(n N).$

We also need to define the syntax of frames.

$$f ::= \square(e_2) \mid v_1(\square) \mid \square + e \mid v + \square$$

$\text{app}_1 : \text{exp} \rightarrow \text{frame}. \quad \text{— } \textit{this is } \square(e_2)$
 $\text{app}_2 : \text{exp} \rightarrow \text{frame}. \quad \text{— } \textit{this is } v_1(\square)$
 $\text{plus}_1 : \text{exp} \rightarrow \text{frame}. \quad \text{— } \textit{this is } \square + e_2$
 $\text{plus}_2 : \text{exp} \rightarrow \text{frame}. \quad \text{— } \textit{this is } v_1 + \square$

5.1.2 Dynamic semantics

The dynamic semantics for \mathcal{L}_1 is straightforward by analogy with an abstract machine for control. A function is a value, and therefore it is always returned immediately, and a function application first evaluates the function part to a value, then evaluates the argument part to a value, and then finally substitutes the argument into the function. In the usual style of higher-order abstract syntax, substitution is performed by application in the rule e/app_2 .

$e/\text{lam} : \text{eval}(\text{lam } \lambda x. E x) \rightarrow \text{retn}(\text{lam } \lambda x. E x).$
 $e/\text{app} : \text{eval}(\text{app } E_1 E_2) \rightarrow \text{cont}(\text{app}_1 E_2) \cdot \text{eval } E_1.$
 $e/\text{app}_1 : \text{cont}(\text{app}_1 E_2) \cdot \text{retn } V_1 \rightarrow \text{cont}(\text{app}_2 V_1) \cdot \text{eval } E_2.$
 $e/\text{app}_2 : \text{cont}(\text{app}_2(\text{lam } \lambda x. E_0 x)) \cdot \text{retn } V_2 \rightarrow \text{eval}(E_0 V_2).$

With one exception, the dynamic semantics of numbers and addition are just as straightforward. As mentioned before, when both subterms have been evaluated to numerical values, there must be a primitive operation that actually adds the two numbers together. As discussed in the context of the calc/\dots rules in the previous section, the solution is to rely on the persistent type family $(\text{add } N M P)$ that adequately encodes the inductive definition of addition for natural numbers. Given this definition, the dynamic semantics of addition are also straightforward: a number is a value, and to evaluate $\text{plus } E_1 E_2$ we evaluate E_1 , then evaluate E_2 , and then add the resulting numbers.

$e/n : \text{eval}(\text{nat } N) \rightarrow \text{retn}(\text{nat } N).$
 $e/\text{plus} : \text{eval}(\text{plus } E_1 E_2) \rightarrow \text{cont}(\text{plus}_1 E_2) \cdot \text{eval } E_1.$
 $e/\text{plus}_1 : \text{cont}(\text{plus}_1 E_2) \cdot \text{retn } V_1 \rightarrow \text{cont}(\text{plus}_2 V_1) \cdot \text{eval } E_2.$
 $e/\text{plus}_2 : \text{cont}(\text{plus}_2(n N_1)) \cdot \text{retn}(n N_2) \cdot \text{!add } N_1 N_2 N_3 \rightarrow \text{retn}(n N_3).$

5.1.3 Example trace

As an example, Figure 1 shows a complete evaluation of an \mathcal{L}_1 expression that we would write on paper as $((\lambda x. 2 + x)(1 + 5))$. The expression represented by the series of transitions on the right has the following type:

$$\text{eval}(\text{app}(\text{lam}(\lambda x. \text{plus}(\text{n}(\text{s}(\text{s}z)))x))(\text{plus}(\text{n}(\text{s}z))(\text{n}(\text{s}(\text{s}(\text{s}(\text{s}z))))))) \rightsquigarrow_{\Sigma_e}^* \text{retn}(\text{n}(\text{s}(\text{s}(\text{s}(\text{s}(\text{s}z))))))$$

For brevity and clarity, the on-paper notation is used in the figure.

$$\begin{array}{ll}
\text{eval}((\lambda x. 2 + x)(1 + 5)) & \\
\rightsquigarrow_{\Sigma_e} \text{cont}(\square(1 + 5)), \text{eval}(\lambda x. 2 + x) & \text{by e/app} \\
\rightsquigarrow_{\Sigma_e} \text{cont}(\square(1 + 5)), \text{retn}(\lambda x. 2 + x) & \text{by e/lam} \\
\rightsquigarrow_{\Sigma_e} \text{cont}((\lambda x. 2 + x)(\square)), \text{eval}(1 + 5) & \text{by e/app}_1 \\
\rightsquigarrow_{\Sigma_e} \text{cont}((\lambda x. 2 + x)(\square)), \text{cont}(\square + 5), \text{eval}(1) & \text{by e/plus} \\
\rightsquigarrow_{\Sigma_e} \text{cont}((\lambda x. 2 + x)(\square)), \text{cont}(\square + 5), \text{retn}(1) & \text{by e/n} \\
\rightsquigarrow_{\Sigma_e} \text{cont}((\lambda x. 2 + x)(\square)), \text{cont}(1 + \square), \text{eval}(5) & \text{by e/plus}_1 \\
\rightsquigarrow_{\Sigma_e} \text{cont}((\lambda x. 2 + x)(\square)), \text{cont}(1 + \square), \text{retn}(5) & \text{by e/n} \\
\rightsquigarrow_{\Sigma_e} \text{cont}((\lambda x. 2 + x)(\square)), \text{retn}(6) & \text{by e/plus}_2(\text{add/s add/z}) \\
\rightsquigarrow_{\Sigma_e} \text{eval}(2 + 6) & \text{by e/app}_2 \\
\rightsquigarrow_{\Sigma_e} \text{cont}(\square + 6), \text{eval}(2) & \text{by e/plus} \\
\rightsquigarrow_{\Sigma_e} \text{cont}(\square + 6), \text{retn}(2) & \text{by e/n} \\
\rightsquigarrow_{\Sigma_e} \text{cont}(2 + \square), \text{eval}(6) & \text{by e/plus}_1 \\
\rightsquigarrow_{\Sigma_e} \text{cont}(2 + \square), \text{retn}(6) & \text{by e/n} \\
\rightsquigarrow_{\Sigma_e} \text{retn}(8) & \text{by e/plus}_2(\text{add/s}(\text{add/s add/z}))
\end{array}$$

Figure 1: Example trace in \mathcal{L}_1 .

5.2 Static semantics, progress, and preservation

The description of the dynamic semantics of \mathcal{L}_1 is not enough on its own; there's nothing we have said so far that prevents us from attempting to evaluate $5 + \lambda x.x$ or $3(\lambda x.x + 2)$, both of which would get *stuck*. A state is stuck if it has not evaluated all the way to the single atomic proposition ($\text{retn } E$) but which nevertheless cannot take a step. A safe expression is one that cannot lead to a stuck state:

Definition 1 (Safety). *The expression E is **safe** if $(\text{eval } E \rightsquigarrow_{\Sigma_e}^* \mathcal{S})$ implies that either $\mathcal{S} \rightsquigarrow_{\Sigma_e} \mathcal{S}'$ for some \mathcal{S}' or else $\mathcal{S} = \text{retn } V$ for some V .*

While this definition of safety is appropriate for our current purposes, it is not necessarily the only definition of safety we might want to consider. For instance, a system that might deadlock is usually considered to be safe, but would not be safe according to this definition.

Integral to establishing safety will be the static semantics, a new set of rewriting rules annotated with \mathfrak{t}/\dots (so we use $\Sigma \mathfrak{t}$ to refer to them collectively). The goal of these rewriting rules is to rewrite the entire state \mathcal{S} to a single atomic proposition $\text{abs } T$. The atomic proposition $\text{abs } T$ is intended to represent an abstraction of states that will produce a value of type T if they produce any value at all.⁸ These rewriting rules, which

⁸This characterization of $\text{abs } T$ is actually a corollary of the preservation lemma.

we will refer to as the static semantics, rely on the not-yet-defined judgments of $E \vdash T$ (the expression E has type T) and $\text{off } F \vdash T \vdash T'$ (the frame F takes values of type T to states of type T'). The rules themselves are straightforward:

$$\begin{aligned} \text{abs} &: \text{tp} \rightarrow \text{type}_o. \\ \text{t/eval} &: \text{eval } E \cdot \text{!of } E \vdash T \Rightarrow \text{abs } T. \\ \text{t/retn} &: \text{retn } E \cdot \text{!value } E \cdot \text{!of } E \vdash T \Rightarrow \text{abs } T. \\ \text{t/cont} &: \text{cont } F \cdot \text{abs } T \cdot \text{!off } F \vdash T \vdash T' \Rightarrow \text{abs } T'. \end{aligned}$$

These static semantics (together with the not-yet-specified typing rules) allow us to prove progress and preservation theorems, which in turn allow us to establish type safety for the language.

Progress If $\mathcal{S} \rightsquigarrow_{\Sigma_t}^* \text{abs } T$ and \mathcal{S} contains no resources $\text{abs } T$, then either $\mathcal{S} \rightsquigarrow_{\Sigma_e} \mathcal{S}'$ for some \mathcal{S}' or else $\mathcal{S} = \text{retn } V$.

$$\begin{array}{c} \mathcal{S} \xrightarrow{\Sigma_e} \mathcal{S}' \\ \text{or} \\ \mathcal{S} = \text{retn } V \\ \left. \begin{array}{c} \mathcal{S} \\ \downarrow \Sigma_t \\ \text{abs } T \end{array} \right\} \downarrow^* \end{array}$$

Preservation If $\mathcal{S} \rightsquigarrow_{\Sigma_e} \mathcal{S}'$ and $\mathcal{S} \rightsquigarrow_{\Sigma_t}^* \text{abs } T$, then $\mathcal{S}' \rightsquigarrow_{\Sigma_t}^* \text{abs } T$.

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{\Sigma_e} & \mathcal{S}' \\ \left. \begin{array}{c} \mathcal{S} \\ \downarrow \Sigma_t \\ \text{abs } T \end{array} \right\} \downarrow^* & & \left. \begin{array}{c} \mathcal{S}' \\ \downarrow \Sigma_t \\ \text{abs } T \end{array} \right\} \downarrow^* \end{array}$$

Given progress and preservation as stated above, we can prove the desired type safety theorem:

Theorem 1. *If $\vdash \text{of } E \vdash T$, then E is safe.*

Proof. The overall picture for safety is this one:

$$\begin{array}{ccccccc} \text{eval } E & \xrightarrow{\Sigma_e} & \mathcal{S}_1 & \dots & \mathcal{S}_n & \xrightarrow{\Sigma_e} & \mathcal{S} \xrightarrow{\Sigma_e} \mathcal{S}' \\ \left. \begin{array}{c} \text{eval } E \\ \downarrow \Sigma_t \\ \text{abs } T \end{array} \right\} \downarrow^* & & \left. \begin{array}{c} \mathcal{S}_1 \\ \downarrow \Sigma_t \\ \text{abs } T \end{array} \right\} \downarrow^* & & \left. \begin{array}{c} \mathcal{S}_n \\ \downarrow \Sigma_t \\ \text{abs } T \end{array} \right\} \downarrow^* & & \left. \begin{array}{c} \mathcal{S} \\ \downarrow \Sigma_t \\ \text{abs } T \end{array} \right\} \downarrow^* \\ & & & & & & \text{or} \\ & & & & & & \text{retn } V \end{array}$$

We are given an expression ($\text{eval } E \rightsquigarrow_{\Sigma_e}^* \mathcal{S}$), and we must show that either $\mathcal{S} \rightsquigarrow_{\Sigma_e} \mathcal{S}'$ for some \mathcal{S}' or else $\mathcal{S} = \text{retn } V$ for some V . By rule t/eval and the premise, ($\text{eval } E \rightsquigarrow_{\Sigma_t}^* \text{abs } T$), and by induction on the length of the trace and progress lemma, ($\mathcal{S} \rightsquigarrow_{\Sigma_t}^* \text{abs } T$). By a similar induction on the trace and the observation that $\text{abs } T$ appears nowhere in any of the rules of the form e/\dots , \mathcal{S} contains no resources $\text{abs } T$. The result then follows by the progress lemma. \square

Note that both the proof of type safety and the statements of the progress and preservation theorems talk about individual transition steps and about taking a series of transitions and extending them with an additional transition. Both of these notions are naturally represented by the transitions and expressions of our logical framework, but neither of these are easy to represent in the CLF framework. The CLF framework has proof terms that express a completed derivations that has ended by proving something, but a series of transitions in our framework translates to a *partial* derivation in CLF, not a complete derivation. It is for this reason that I previously described the focus on a state-transition-based presentation of logic as the most critical difference between CLF and the framework presented here.

In the remainder of this section, we will present the typing rules and static semantics for \mathcal{L}_1 , and then give proofs of the progress and preservation lemmas above.

5.2.1 Static semantics

Typing rules are completely conventional. The syntax of types can be written as a BNF specification as $\tau ::= \text{nat} \mid \tau \rightarrow \tau$, and the signature for types is the following:

$\text{tnat} : \text{tp}.$
 $\text{arrow} : \text{tp} \rightarrow \text{tp} \rightarrow \text{tp}.$

The typing rules are the familiar ones:

$\text{of} : \text{exp} \rightarrow \text{tp} \rightarrow \text{type}.$
 $\text{of}/\text{lam} : (\Pi x. \text{of } x \text{ T} \rightarrow \text{of } (E x) \text{ T}') \rightarrow \text{of } (\text{lam } \lambda x. E x)(\text{arrow } \text{T T}').$
 $\text{of}/\text{app} : \text{of } E_1 (\text{arrow } \text{T}' \text{ T}) \rightarrow \text{of } E_2 \text{ T}' \rightarrow \text{of } (\text{app } E_1 E_2) \text{ T}.$
 $\text{of}/\text{n} : \text{of } (\text{nat } N) \text{tnat}.$
 $\text{of}/\text{plus} : \text{of } E_1 \text{tnat} \rightarrow \text{of } E_2 \text{tnat} \rightarrow \text{of } (\text{plus } E_1 E_2) \text{tnat}.$

Frame typing rules are also straightforward:

$\text{off}/\text{app}_1 : \text{of } E_2 \text{ T}' \rightarrow \text{off } (\text{app}_1 E_2) (\text{arrow } \text{T}' \text{ T}) \text{ T}.$
 $\text{off}/\text{app}_2 : \text{value } E_1 \rightarrow \text{of } E_1 (\text{arrow } \text{T}' \text{ T}) \rightarrow \text{off } (\text{app}_2 E_2) \text{ T}' \text{ T}.$
 $\text{off}/\text{plus}_1 : \text{of } E_2 \text{tnat} \rightarrow \text{off } (\text{plus}_1 E_2) \text{tnat} \text{tnat}.$
 $\text{off}/\text{plus}_2 : \text{value } E_1 \rightarrow \text{of } E_1 \text{tnat} \rightarrow \text{off } (\text{plus}_2 E_1) \text{tnat} \text{tnat}.$

5.2.2 Preservation

The preservation lemma establishes that our purported invariant is actually an invariant. The form of the preservation lemma should be very familiar: if the invariant described by the static semantics holds of a state, the invariant still holds of the state after any transition made under the dynamic semantics.

Lemma 1 (Preservation). *If $S \rightsquigarrow_{\Sigma_e} S'$ and $S \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f$, then $S' \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f$.*

Preservation proofs tend to rely on typing inversion lemmas, and we use traditional typing inversion lemmas as needed without proof. However, there is another lemma that is critical for preservation proofs for SSOS specifications. It captures the intuition that if $S \rightsquigarrow_{\Sigma_t}^* \text{abs } T$, either S is already a substructural context containing a single resource ($\text{abs } T$) or else everything in S must eventually be consumed in the service of producing a single resource ($\text{abs } T$). We call this the *consumption lemma* because it represents the obligation

that the static semantics consume everything in the context. Recall that $\mathcal{S}[\Delta]$ is notation representing a state \mathcal{S} containing the ephemeral state Δ inside of it.

Lemma 2 (Consumption).

- If $\mathcal{E} : \mathcal{S}[\text{eval } E] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f$,
then $\mathcal{E} \approx (t/\text{eval}(\mathcal{D}_t); \mathcal{E}')$, where $\mathcal{D}_t : \text{of } E \ T$ and $\mathcal{E}' : \mathcal{S}[\text{abs } T] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f$ for some T .
- If $\mathcal{E} : \mathcal{S}[\text{retn } V] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f$,
then $\mathcal{E} \approx (t/\text{retn}(\mathcal{D}_v \cdot \mathcal{D}_t); \mathcal{E}')$, where $\mathcal{D}_v : \text{value } V$, $\mathcal{D}_t : \text{of } V \ T$, and $\mathcal{E}' : \mathcal{S}[\text{abs } T] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f$ for some T .
- If $\mathcal{E} : \mathcal{S}[\text{cont } F, \text{abs } T] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f$,
then $\mathcal{E} \approx (t/\text{cont}(\mathcal{D}_f); \mathcal{E}')$, where $\mathcal{D}_f : \text{off } F \ T \ T'$ and $\mathcal{E}' : \mathcal{S}[\text{abs } T'] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f$ for some T' .

Proof. Each of the statements can be proved independently by induction on the structure of \mathcal{E} . The eval and retn cases are quite similar to each other and are both simpler versions of the cont case, which we give below.

Case: $\mathcal{E} = \diamond$

This case cannot occur, as there is no way for $\mathcal{S}[\text{cont } F, \text{abs } T]$ to be equivalent to $(\text{abs } T_f)$.

Case: $\mathcal{E} = t/\text{eval}(\mathcal{D}'_t); \mathcal{E}'$

$\mathcal{D}'_t : \text{of } E \ T_e$,
 $\mathcal{E}' : (\mathcal{S}'[\text{abs } T_e] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f)$, and
 $\mathcal{S}[\text{cont } F, \text{abs } T] \approx \mathcal{S}'[\text{eval } E]$

Because $\mathcal{S}[\text{cont } F, \text{abs } T] \approx \mathcal{S}'[\text{eval } E]$, it must be the case that both of these are also equivalent to $\mathcal{S}''[\text{cont } F, \text{abs } T][\text{eval } E]$ for some \mathcal{S}' and $\mathcal{E}' : (\mathcal{S}''[\text{cont } F, \text{abs } T][\text{abs } T_e] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f)$.

$\mathcal{E}' \approx (t/\text{cont}(\mathcal{D}_f); \mathcal{E}'')$	By i.h.
$\mathcal{D}_f : \text{off } E \ T \ T'$	"
$\mathcal{E}'' : (\mathcal{S}''[\text{abs } T'][\text{abs } T_e] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f)$	"
$\mathcal{E} \approx (t/\text{eval}(\mathcal{D}'_t); t/\text{cont}(\mathcal{D}_f); \mathcal{E}'')$	By construction
$\mathcal{E} \approx (t/\text{cont}(\mathcal{D}_f); t/\text{eval}(\mathcal{D}'_t); \mathcal{E}'')$	By concurrent equivalence

Case: $\mathcal{E} = t/\text{retn}(\mathcal{D}'_v \cdot \mathcal{D}'_t); \mathcal{E}'$

Similar to previous case.

Case: $\mathcal{E} = t/\text{cont}(\mathcal{D}'_f); \mathcal{E}'$

$\mathcal{D}'_f : \text{off } F_1 \ T_1 \ T'_1$,
 $\mathcal{E}' : (\mathcal{S}'[\text{abs } T'_1] \rightsquigarrow_{\Sigma_t}^* \text{abs } T_f)$, and
 $\mathcal{S}[\text{cont } F, \text{abs } T] \approx \mathcal{S}'[\text{cont } F_1, \text{abs } T_1]$

Because $\mathcal{S}[\text{cont } F, \text{abs } T] \approx \mathcal{S}'[\text{cont } F_1, \text{abs } T_1]$, there are two possibilities. The first possibility is that $\mathcal{S}[\text{abs } T'] \approx \mathcal{S}'[\text{abs } T']$, $F = F_1$, and $T = T_1$, in which case we are done.

The other possibility is $\mathcal{S}[\text{cont } F, \text{abs } T] \approx \mathcal{S}'[\text{cont } F_1, \text{abs } T_1] \approx \mathcal{S}''[\text{cont } F, \text{abs } T][\text{cont } F_1, \text{abs } T_1]$ for some \mathcal{S}'' . In that case, $\mathcal{E}' : (\mathcal{S}''[\text{cont } F, \text{abs } T][\text{abs } T'_1])$ and we proceed with the following reasoning:

$\mathcal{E}' \approx (\text{t}/\text{cont}(\mathcal{D}_f); \mathcal{E}'')$	By i.h.
$\mathcal{D}_f : \text{off E T T}'$	”
$\mathcal{E}'' \approx (\mathcal{S}''[\text{abs T}'][\text{abs T}'_1])$	”
$\mathcal{E} \approx (\text{t}/\text{cont}(\mathcal{D}'_f); \text{t}/\text{cont}(\mathcal{D}_f); \mathcal{E}'')$	By construction
$\mathcal{E} \approx (\text{t}/\text{cont}(\mathcal{D}_f); \text{t}/\text{cont}(\mathcal{D}'_f); \mathcal{E}'')$	By concurrent equivalence

This completes the proof. □

It's worth emphasizing what may be obvious: the proof of the consumption lemma is a bit more general than necessary. In the static semantics we have defined so far, $\mathcal{S} \rightsquigarrow_{\Sigma_t}^* \text{abs T}$ means that \mathcal{S} contains *precisely one* atomic proposition of the form (eval E), (retn E), or (abs T), which means that most of the cases we consider are actually impossible. However, this impossibility is *also* a fact that must be proven by induction on the structure of expressions. The proof we have given is both straightforward in its own right and allows the proof of preservation to “scale” to specifications with parallel evaluation (see Section 5.3). Furthermore, the theorem *statement* is not at all complicated as a result of handling the general case. If we wanted to write a version of the consumption lemma for the special case of sequential SSOS specifications, the only difference is that we would write “If $\mathcal{E} : (\mathcal{S}, \text{eval E} \rightsquigarrow_{\Sigma_t}^* \text{abs T}_f)$, then. . .” instead of “If $\mathcal{E} : (\mathcal{S}[\text{eval E}] \rightsquigarrow_{\Sigma_t}^* \text{abs T}_f)$, then. . .” because we know that **eval E** appears only on the right-hand edge of the ephemeral context.

Now we can consider the actual proof of preservation:

Proof of the Preservation Lemma (Lemma 1). The proof proceeds by case analysis on the transition $\mathcal{T} : (\mathcal{S} \rightsquigarrow_{\Sigma_e} \mathcal{S}')$, followed in each case by applying the consumption lemma to the expression $\mathcal{E} : (\mathcal{S} \rightsquigarrow_{\Sigma_t}^* \text{abs T}_f)$.

Case: $\mathcal{T} = e/\text{lam} : \mathcal{S}[\text{eval}(\text{lam } \lambda x. E x)] \rightsquigarrow_{\Sigma_e} \mathcal{S}[\text{retn}(\text{lam } \lambda x. E x)]$

$\mathcal{E} \approx (\text{t}/\text{eval}(\mathcal{D}_t); \mathcal{E}')$	By consumption lemma
$\mathcal{D}_t : \text{of}(\text{lam } \lambda x. E x) \text{ T}$	”
$\mathcal{E}' : (\mathcal{S}[\text{abs T}] \rightsquigarrow_{\Sigma_t}^* \text{abs T}_f)$	”
$(\text{t}/\text{retn}(v/\text{lam} \cdot \mathcal{D}_t); \mathcal{E}') : (\mathcal{S}[\text{retn}(\text{lam } \lambda x. E x)] \rightsquigarrow_{\Sigma_t}^* \text{abs T}_f)$	By construction

Case: $\mathcal{T} = e/\text{app} : \mathcal{S}[\text{eval}(\text{app } E_1 E_2)] \rightsquigarrow_{\Sigma_e} \mathcal{S}[\text{cont}(\text{app}_1 E_2), \text{eval } E_1]$

$\mathcal{E} \approx (\text{t}/\text{eval}(\mathcal{D}_t); \mathcal{E}')$	By consumption lemma
$\mathcal{D}_t : \text{of}(\text{app } E_1 E_2) \text{ T}$	”
$\mathcal{E}' : \mathcal{S}[\text{abs T}] \rightsquigarrow_{\Sigma_t}^* \text{abs T}_f$	”
$\mathcal{D}_t = \text{of}/\text{app } \mathcal{D}_1 \mathcal{D}_2$	By inversion on \mathcal{D}_t
$\mathcal{D}_1 : \text{of } E_1 (\text{arrow } \text{T}' \text{ T})$	”
$\mathcal{D}_2 : \text{of } E_2 \text{ T}'$	”
$(\text{t}/\text{cont}(\text{off}/\text{app}_1 \mathcal{D}_2); \mathcal{E}') : (\mathcal{S}[\text{cont}(\text{app}_1 E_2), \text{abs}(\text{arrow } \text{T}' \text{ T})] \rightsquigarrow_{\Sigma_t}^* \text{abs T}_f)$	By construction
$(\text{t}/\text{eval}(\mathcal{D}_1); \text{t}/\text{cont}(\text{off}/\text{app}_1 \mathcal{D}_2); \mathcal{E}') : (\mathcal{S}[\text{cont}(\text{app}_1 E_2), \text{eval } E_1] \rightsquigarrow_{\Sigma_t}^* \text{abs T}_f)$	By construction

Case: $\mathcal{T} = e/\text{app}_1 : \mathcal{S}[\text{cont}(\text{app}_1 E_2), \text{retn } V_1] \rightsquigarrow_{\Sigma_e} \mathcal{S}[\text{cont}(\text{app}_2 V_1), \text{eval } E_2]$

$\mathcal{E} \approx (\text{t}/\text{retn}(\mathcal{D}_v \cdot \mathcal{D}_t); \mathcal{E}')$	By consumption lemma
$\mathcal{D}_v : \text{value } V_1$	”
$\mathcal{D}_t : \text{of } V_1 \text{ T}$	”
$\mathcal{E}' : (\mathcal{S}[\text{cont}(\text{app}_1 E_2), \text{abs T}] \rightsquigarrow_{\Sigma \text{t}}^* \text{abs T}_f)$	”
$\mathcal{E}' \approx \text{t}/\text{cont}(\mathcal{D}_f); \mathcal{E}''$	By consumption lemma
$\mathcal{D}_f : \text{off}(\text{app}_1 E_2) \text{ T T}'$	”
$\mathcal{E}'' : (\mathcal{S}[\text{abs T}'] \rightsquigarrow_{\Sigma \text{t}}^* \text{abs T}_f)$	”
$\mathcal{D}_f = \text{off}/\text{app}_1(\mathcal{D}'_t)$	By inversion on \mathcal{D}_f
$\text{T} = \text{arrow } T_0 \text{ T}'$	”
$\mathcal{D}'_t : \text{of } E_2 T_0$	”
$(\text{t}/\text{cont}(\text{off}/\text{app}_2 \mathcal{D}_v \mathcal{D}_t); \mathcal{E}'') : (\mathcal{S}[\text{cont}(\text{app}_2 V_1), \text{abs T}_0] \rightsquigarrow_{\Sigma \text{t}}^* \text{abs T}_f)$	By construction
$(\text{t}/\text{eval}(\mathcal{D}'_t); \text{t}/\text{cont}(\text{off}/\text{app}_2 \mathcal{D}_v \mathcal{D}_t); \mathcal{E}'') : (\mathcal{S}[\text{cont}(\text{app}_2 V_1), \text{eval } E_2] \rightsquigarrow_{\Sigma \text{t}}^* \text{abs T}_f)$	By construction

Case: $\mathcal{T} = \text{e}/\text{app}_2 : \mathcal{S}[\text{cont}(\text{app}_2(\text{lam } \lambda x. E x)), \text{retn } V_2] \rightsquigarrow_{\Sigma \text{e}} \mathcal{S}[\text{eval}(E V_2)]$

$\mathcal{E} \approx \text{t}/\text{retn}(\mathcal{D}_v \cdot \mathcal{D}_t); \mathcal{E}'$	By consumption lemma
$\mathcal{D}_t : \text{of } V_2 \text{ T}$	”
$\mathcal{E}' : (\mathcal{S}[\text{cont}(\text{app}_2(\text{lam } \lambda x. E x)), \text{abs T}] \rightsquigarrow_{\Sigma \text{t}}^* \text{abs T}_f)$	”
$\mathcal{E}' \approx \text{t}/\text{cont}(\mathcal{D}_f); \mathcal{E}''$	By consumption lemma
$\mathcal{D}_f : \text{off}(\text{app}_2(\text{lam } \lambda x. E x)) \text{ T T}'$	”
$\mathcal{E}'' : (\mathcal{S}[\text{abs T}'] \rightsquigarrow_{\Sigma \text{t}}^* \text{abs T}_f)$	”
$\mathcal{D}_f = \text{off}/\text{app}_2 \mathcal{D}'_v (\lambda x. \lambda d. \mathcal{D}'_t x d)$	By inversion on \mathcal{D}_f
$(\lambda x. \lambda d. \mathcal{D}'_t x d) : \text{II}x. \text{of } x \text{ T} \rightarrow \text{of}(E x) \text{ T}'$	”
$\mathcal{D}'_t V_2 \mathcal{D}_t : \text{of}(E V_2) \text{ T}'$	By substitution
$(\text{t}/\text{eval}(\mathcal{D}'_t V_2 \mathcal{D}_t); \mathcal{E}'') : (\mathcal{S}[\text{eval}(E V_2)] \rightsquigarrow_{\Sigma \text{t}}^* \text{abs T}_f)$	By construction

The remaining cases for e/nat , e/plus , e/plus_1 , and e/plus_2 are similar. □

5.2.3 Progress

Preservation theorems, and in particular big-step preservation theorems, have traditionally been the primary correctness criteria for language specifications in substructural logics [1]. However, preservation alone is insufficient to ensure language safety as specified by Definition 1. The preservation lemma ensures that the invariant is maintained; the progress lemma ensures that the invariant actually establishes safety.

The critical lemma for progress, the analogue to the consumption lemma in the preservation lemma, is actually quite a bit more general. It relies on a general property of the static semantics, namely that it is *contractive*, because each rule has at most one conclusion.

Definition 2 (Contractive signature). *If every rule in Σ has at most one atomic proposition in the conclusion, and that conclusion is an ordered or linear atomic proposition, then Σ is contractive.*

When a signature is contractive, then different pieces of the context at the end of an expression can be traced backwards to different pieces at the beginning of the expression.

Lemma 3 (Splitting). *If Σ is contractive and $\mathcal{E} : (\mathcal{S} \rightsquigarrow_{\Sigma}^* \mathcal{S}'_1, \mathcal{S}'_2)$, then $\mathcal{S} \approx \mathcal{S}_1, \mathcal{S}_2$, and there exist expressions \mathcal{E}_1 and \mathcal{E}_2 such that $\mathcal{E}_1 : (\mathcal{S}_1 \rightsquigarrow_{\Sigma}^* \mathcal{S}'_1)$ and $\mathcal{E}_2 : (\mathcal{S}_2 \rightsquigarrow_{\Sigma}^* \mathcal{S}'_2)$, where \mathcal{E} has the same number of transitions as \mathcal{E}_1 and \mathcal{E}_2 combined.*

Proof. By induction on the structure of \mathcal{E} . If $\mathcal{E} = \diamond$, the result is immediate, and if $\mathcal{E} = (\mathcal{E}'; \mathcal{T})$ then we apply the induction hypothesis on \mathcal{E}' and add \mathcal{T} to the end of the appropriate smaller expression: \mathcal{E}_1 if the conclusion of \mathcal{T} was bound in \mathcal{S}_1 , \mathcal{E}_2 otherwise. \square

A stronger version of the splitting lemma would also establish that $\mathcal{E} \approx (\mathcal{E}_1; \mathcal{E}_2) \approx (\mathcal{E}_2; \mathcal{E}_1)$, but this lemma is sufficient to allow us to prove the progress lemma.

Lemma 4 (Progress). *If $\mathcal{S} \rightsquigarrow_{\Sigma t}^* \text{abs } T$ and \mathcal{S} contains no resources $\text{abs } T$, then either $\mathcal{S} \rightsquigarrow_{\Sigma e} \mathcal{S}'$ for some \mathcal{S}' or else $\mathcal{S} = \text{retn } V$ for some value V with type T .*

At a high level, this proof works by showing that any state \mathcal{S} that satisfies the typing invariant must be one of the following:

- $\mathcal{S} = \text{retn } V$, a safe state that takes no step,
- $\mathcal{S} = \text{eval } E$, in which case we can always take a step because $\text{eval } E$ is an active proposition,
- $\mathcal{S} = (\text{comp } F, \mathcal{S}')$ where $\mathcal{S}' \rightsquigarrow_{\Sigma e} \mathcal{S}''$; in this case, the larger state takes a step as well, or
- $\mathcal{S} = (\text{comp } F, \text{retn } V)$, in which case we use canonical forms lemmas (which are standard, and which we therefore use without proof) to ensure that we can perform a reduction.

Proof. By induction on the number of transitions in \mathcal{E} . We note that $\mathcal{E} \neq \diamond$, because that would mean that $\mathcal{S} = \text{abs } T$, contradicting the premise that \mathcal{S} does not contain any atomic propositions of the form $\text{abs } T$. So we can assume $\mathcal{E} = (\mathcal{E}'; \mathcal{T})$ and do case analysis on the form of the transition \mathcal{T} .

Case: $\mathcal{T} = \text{t/retn}(\mathcal{D}_v \cdot \mathcal{D}_t)$ and $\mathcal{E}' : \mathcal{S} \rightsquigarrow_{\Sigma t}^* \text{retn } V$
where $\mathcal{D}_v : \text{value } V$ and $\mathcal{D}_t : \text{of } E T$.

There are no transitions in Σt that can appear as the last transition in \mathcal{E}' , so $\mathcal{E}' = \diamond$ and $\mathcal{S} = \text{retn } V$; therefore, we are done.

Case: $\mathcal{T} = \text{t/eval}(\mathcal{D}_t)$ and $\mathcal{E}' : \mathcal{S} \rightsquigarrow_{\Sigma t}^* \text{eval } E$.

There are no transitions in Σt that can appear as the last transition in \mathcal{E}' , so $\mathcal{E}' = \diamond$ and $\mathcal{S} = \text{eval } E$. We proceed by case analysis on the structure of E .

Subcase: $E = \text{lam } \lambda x. E x$ — $e/\text{lam} : \text{eval}(\text{lam } \lambda x. E x) \rightsquigarrow_{\Sigma e} \text{retn}(\text{lam } \lambda x. E x)$

Subcase: $E = \text{app } E_1 E_2$ — $e/\text{app} : \text{eval}(\text{app } E_1 E_2) \rightsquigarrow_{\Sigma e} (\text{cont}(\text{app}_1 E_2), \text{eval } E_1)$

Subcase: $E = n N$ — $e/n : \text{eval}(n N) \rightsquigarrow_{\Sigma e} \text{retn}(n N)$

Subcase: $E = \text{plus } E_1 E_2$ — $e/\text{plus} : \text{eval}(\text{plus } E_1 E_2) \rightsquigarrow_{\Sigma e} (\text{cont}(\text{plus}_1 E_2), y_2 : \text{eval } E_1)$

Case: $\mathcal{T} = \text{t/cont}(\mathcal{D}_f)$ and $\mathcal{E}' : (\mathcal{S} \rightsquigarrow_{\Sigma t}^* \text{cont } F, \text{abs } T')$, where $\mathcal{D}_f : \text{off } F T' T$.

By the splitting lemma, $\mathcal{S} \approx \mathcal{S}_1, \mathcal{S}_2$ and there exist two expressions $\mathcal{E}_1 : (\mathcal{S}_1 \rightsquigarrow_{\Sigma t}^* \text{cont } F)$ and $\mathcal{E}_2 : (\mathcal{S}_2 \rightsquigarrow_{\Sigma t}^* \text{abs } T')$. There are no transitions in Σt that can appear as the last transition in \mathcal{E}_1 , so $\mathcal{E}_1 = \diamond$ and $\mathcal{S}_1 = \text{cont } F$.

We then apply the induction hypothesis on \mathcal{E}_2 (which has one transition less than \mathcal{E} , justifying the call to the induction hypothesis). If $\mathcal{T} : \mathcal{S}_2 \rightsquigarrow_{\Sigma e} \mathcal{S}'_2$, then $\mathcal{T} : (\text{cont } F, \mathcal{S}_2 \rightsquigarrow_{\Sigma e} \text{cont } F, \mathcal{S}'_2)$ by the rules for typing transitions,⁹ and we are done.

⁹Generically we could call this a *frame property* of the framework.

Otherwise, we have $\mathcal{S}_2 = \text{retn } V$, $\mathcal{D}_v : \text{value } V$, and $\mathcal{D}_t : \text{of } V \ T'$. We proceed by case analysis on the proof term \mathcal{D}_f establishing that the frame F is well typed.

Subcase: $\mathcal{D}_f = \text{off}/\text{app}_1 \ \mathcal{D}_{t_2}$, so $F = \text{app}_1 \ E_2$.

$$e/\text{app}_1 : (\text{cont}(\text{app}_1 \ E_2), \text{retn } V) \rightsquigarrow_{\Sigma_e} (\text{cont}(\text{app}_2 \ V), \text{eval } E_2).$$

Subcase: $\mathcal{D}_f = \text{off}/\text{app}_2 \ \mathcal{D}_{v_1} \ \mathcal{D}_{t_1}$, so $F = \text{app}_2 \ V_1$,

$\mathcal{D}_{v_1} : \text{value } V_1$, and $\mathcal{D}_{t_1} : \text{of } E_1 \ (\text{arrow } T' \ T)$.

By the canonical forms lemma on \mathcal{D}_{v_1} and \mathcal{D}_{t_1} , $V_1 = \text{lam } \lambda x. E_0 \ x$.

$$e/\text{app}_2 : (\text{cont}(\text{app}_2(\text{lam } \lambda x. E_0 \ x)), \text{retn } V) \rightsquigarrow_{\Sigma_e} \text{eval}(E_0 \ V).$$

Subcase: $\mathcal{D}_f = \text{off}/\text{plus}_1 \ \mathcal{D}_{t_2}$, so $F = \text{plus}_1 \ E_2$.

$$e/\text{plus}_1 : (\text{cont}(\text{plus}_1 \ E_2), \text{retn } V) \rightsquigarrow_{\Sigma_e} (\text{cont}(\text{plus}_2 \ V), \text{eval } E_2).$$

Subcase: $\mathcal{D}_f = \text{off}/\text{plus}_2 \ \mathcal{D}_{v_1} \ \mathcal{D}_{t_1}$, so $F = \text{plus}_2 \ V_1$ and $T = T' = \text{tnat}$

$\mathcal{D}_{v_1} : \text{value } E_1$, and $\mathcal{D}_{t_1} : \text{of } E_1 \ \text{tnat}$.

By the canonical forms lemma on \mathcal{D}_{v_1} and \mathcal{D}_{t_1} , $V_1 = n \ N_1$.

By the canonical forms lemma on \mathcal{D}_v and \mathcal{D}_t , $V = n \ N_2$.

By the effectiveness of addition on N_1 and N_2 , there exists a natural number N_3 and a proof term $\mathcal{D}_a : \text{add } N_1 \ N_2 \ N_3$.

$$e/\text{plus}_2(\mathcal{D}_a) : (\text{cont}(\text{app}_2(n \ N_1)), \text{retn}(n \ N_2)) \rightsquigarrow_{\Sigma_e} \text{retn}(n \ N_3).$$

This completes the proof; we have assumed standard canonical forms lemmas and the effectiveness of addition, provable by induction on the structure of N_1 . \square

5.3 Parallel evaluation

One way in which parallel evaluation has been incorporated into functional programming languages is by allowing pairs (or, more generally, tuples) to evaluate in parallel [9]. In this section we will consider \mathcal{L}_2 , the first modular extension to the language \mathcal{L}_1 with parallel pairs. The syntax and typing rules pairs are a straightforward and standard addition:

$\text{pair} : \text{exp} \rightarrow \text{exp} \rightarrow \text{exp}$.

$\text{fst} : \text{exp} \rightarrow \text{exp}$.

$\text{snd} : \text{exp} \rightarrow \text{exp}$.

$v/\text{pair} : \text{value } E_1 \rightarrow \text{value } E_2 \rightarrow \text{value}(\text{pair } E_1 \ E_2)$.

$\text{pairtp} : \text{tp} \rightarrow \text{tp} \rightarrow \text{tp}$.

$\text{of}/\text{pair} : \text{of } E_1 \ T_1 \rightarrow \text{of } E_2 \ T_2 \rightarrow \text{of}(\text{pair } E_1 \ E_2) (\text{pairtp } T_1 \ T_2)$.

$\text{of}/\text{fst} : \text{of } E (\text{pairtp } T_1 \ T_2) \rightarrow \text{of}(\text{fst } E) \ T_1$.

$\text{of}/\text{snd} : \text{of } E (\text{pairtp } T_1 \ T_2) \rightarrow \text{of}(\text{snd } E) \ T_2$.

5.3.1 Dynamic semantics

Nothing in the previous signature indicates that pairs are any more interesting than the language features we have already presented. It is in the specification of the dynamic semantics that we add a new capability to the language: *parallel evaluation*. Our dynamic semantics will evaluate both parts of a pair in parallel. First,

we add a new type of latent ordered proposition, ($\text{cont}_2 F$), representing a frame waiting on *two* values to be returned to it.

$$\begin{aligned} \text{cont}_2 &: \text{frame} \rightarrow \text{type}_o. \\ \text{fst}_0 &: \text{frame}. \\ \text{snd}_0 &: \text{frame}. \\ \text{pair}_0 &: \text{frame}. \\ \text{e}/\text{fst} &: \text{eval}(\text{fst } E) \rightarrow \text{cont } \text{fst}_0 \cdot \text{eval } E. \\ \text{e}/\text{fst}_0 &: \text{cont } \text{fst}_0 \cdot \text{retn}(\text{pair } V_1 V_2) \rightarrow \text{retn } V_1. \\ \text{e}/\text{snd} &: \text{eval}(\text{snd } E) \rightarrow \text{cont } \text{snd}_0 \cdot \text{eval } E. \\ \text{e}/\text{snd}_0 &: \text{cont } \text{snd}_0 \cdot \text{retn}(\text{pair } V_1 V_2) \rightarrow \text{retn } V_2. \\ \text{e}/\text{pair} &: \text{eval}(\text{pair } E_1 E_2) \rightarrow \text{cont}_2 \text{pair}_0 \cdot \text{eval } E_1 \cdot \text{eval } E_2. \\ \text{e}/\text{pair}_0 &: \text{cont}_2 \text{pair}_0 \cdot \text{retn } V_1 \cdot \text{retn } V_2 \rightarrow \text{retn}(\text{pair } V_1 V_2). \end{aligned}$$

Our specification no longer corresponds to an on-paper description of a stack machine with one expression evaluating on top of the stack: with this change, we have made the substructural context a treelike structure where multiple independent groups of propositions may be able to transition at any given time.

5.3.2 Static semantics

Beyond the typing rules for pairs given at the beginning of this section, we need a new typing judgment for continuation frames waiting on two values:

$$\text{off2} : \text{frame} \rightarrow \text{tp} \rightarrow \text{tp} \rightarrow \text{tp} \rightarrow \text{type}.$$

The fst_0 and snd_0 frames lead to new cases for the regular frame typing rule, and we give a parallel frame typing rule for the pair_0 frame:

$$\begin{aligned} \text{off}/\text{fst}_0 &: \text{off } \text{fst}_0 (\text{pairtp } T_1 T_2) T_1. \\ \text{off}/\text{snd}_0 &: \text{off } \text{fst}_0 (\text{pairtp } T_1 T_2) T_2. \\ \text{off2}/\text{pair}_0 &: \text{off2 } \text{pair}_0 T_1 T_2 (\text{pairtp } T_1 T_2). \end{aligned}$$

And a rule in the static semantics explaining how parallel frames interact with ($\text{abs } T$):

$$\begin{aligned} \text{t}/\text{cont}_2 &: \text{cont}_2 F \cdot \text{abs } T_1 \cdot \text{abs } T_2 \cdot \text{!off2 } F T_1 T_2 T' \\ &\rightarrow \text{abs } T'. \end{aligned}$$

Those three declarations entirely describe the static semantics of parallel evaluation. The typing rule $\text{off2}/\text{pair}_0$ and the rewriting rule t/cont_2 could easily be merged, but by separating them it is possible to incorporate additional parallel features into the language in a modular way.

5.3.3 Safety

We have just seen that it is possible to modularly extend both the static and dynamic semantics of \mathcal{L}_1 to obtain the language \mathcal{L}_2 with parallel evaluation of pairs. We are also able to straightforwardly extend the safety proof of \mathcal{L}_1 to incorporate parallel evaluation. We have to do the following things to extend the existing proof:

- Add a new consumption lemma:
 If $\mathcal{E} : \mathcal{S}[\text{cont}_2 F, \text{abs } T_1, \text{abs } T_2] \rightsquigarrow_{\Sigma t}^* \text{abs } T_f$,
 then $\mathcal{E} \approx (t/\text{cont}(\mathcal{D}_t); \mathcal{E}')$, where $\mathcal{D}_f : \text{off}_2 F T_1 T_2 T'$ and $\mathcal{E}' : \mathcal{S}[\text{abs } T'] \rightsquigarrow_{\Sigma t}^* \text{abs } T_f$ for some T' .
- Add a case for t/cont_2 to the proof of each of the other consumption lemmas.
- Add cases for e/pair and e/pair_0 , e/fst , etc. to the proof of preservation lemma.
- Add new subcases for $E = \text{pair } E_1 E_2$, $E = \text{fst } E$, and $E = \text{snd } E$ to the second case of the progress lemma where it is the case that $\mathcal{S} = \text{eval } E$.
- Add new subcases for $\mathcal{D}_f = \text{off}/\text{fst}_0$ and $\mathcal{D}_f = \text{off}/\text{snd}_0$ to the third case of the progress lemma, both of which will appeal to a canonical forms lemma.
- Add a case for $\mathcal{T} = t/\text{cont}_2(\mathcal{D}_f)$ to the proof of the progress lemma.

The essential structure of the safety proof is preserved under the extension of the language with parallel pairs. This is significant because, with the exception of the proof of the consumption lemma, the safety proof for \mathcal{L}_1 did not need to explicitly prepare for the possibility of parallel evaluation.

5.4 Exceptions

Exceptions and exception handling are another example of a feature that we can add to \mathcal{L}_1 or \mathcal{L}_2 in a modular fashion, though there is an important caveat. While we can add parallel evaluation *or* exceptions to \mathcal{L}_1 without reconsidering any aspects of the original specification, we cannot extend the base language with *both* exceptions and parallelism without considering their interaction. This should not be seen as a troublesome lack of modularity, however — rather, the SSOS specification in this section illustrates that exceptions and parallel evaluation are not truly orthogonal language features. Seen in this light, it is natural that we must clearly describe how the features interact.

The syntax of exceptions includes `error`, which raises an exception, and `trycatch E E'`, which includes a primary expression E and a secondary expression E' that is evaluated only if the evaluation of E returns an error. It would not be much more difficult to allow exceptions to carry a value, but we do not do so here.

```

error : exp.
trycatch : exp → exp → exp.
of/error : of error T.
of/trycatch : of E T → of E' T → of (trycatch E E') T.

```

5.4.1 Dynamic semantics

To give the dynamic semantics of exceptions and exception handling, we introduce two new ordered atomic propositions. The first, `handle E'`, represents an exception-handling stack frame. The second, `raise`, represents an uncaught exception.

```

raise : typeo.
handle : exp → typeo.
e/raise : eval error → raise.

```

$$\begin{aligned}
e/\text{trycatch} &: \text{eval}(\text{trycatch } E E') \rightarrow \text{handle } E' \cdot \text{eval } E. \\
e/\text{vhandle} &: \text{handle } E' \cdot \text{retn } V \rightarrow \text{retn } V. \\
e/\text{xhandle} &: \text{handle } E' \cdot \text{raise} \rightarrow \text{eval } E'. \\
e/\text{xcont} &: \text{cont } F \cdot \text{raise} \rightarrow \text{raise}.
\end{aligned}$$

As mentioned before, we must deal explicitly with the interaction of parallel evaluation and exception handling by explaining how exceptions interact with cont_2 frames. The following is one possibility in which both sub-computations must terminate before an error is returned.

$$\begin{aligned}
e/\text{xvcont} &: \text{cont}_2 F \cdot \text{raise} \cdot \text{retn } V \rightarrow \text{raise}. \\
e/\text{vxcont} &: \text{cont}_2 F \cdot \text{retn } V \cdot \text{raise} \rightarrow \text{raise}. \\
e/\text{xxcont} &: \text{cont}_2 F \cdot \text{raise} \cdot \text{raise} \rightarrow \text{raise}.
\end{aligned}$$

Another obvious candidate for the interaction between parallel evaluation of pairs and exceptions would be for a pair to immediately raise an exception if either of its components raises an exception. However, it is not obvious how to gracefully implement this in the ordered SSOS style we have presented so far. This is a symptom of a broader problem, namely that ordered SSOS specifications don't, in general, handle non-local transfer of control particularly well. This is a limitation of the specification style and not the framework; a different style of specification known as *linear destination passing* could easily handle the more general specification [24].

5.4.2 Static semantics

The static semantics of exception handling are extremely simple; there are just two rules. An uncaught exception has any type, and an exception handler must have the same type as its sub-computation.

$$\begin{aligned}
t/\text{raise} &: \text{raise} \rightarrow \text{abs } T. \\
t/\text{handle} &: \text{handle } E \cdot \text{abs } T \cdot \text{!of } E T \rightarrow \text{abs } T.
\end{aligned}$$

5.4.3 Safety


The existing structure of the safety proof also extends straightforwardly to handle the addition of exceptions to the language; the most significant change is to the definition of safety itself, as a safe state either steps, is a returned value $\text{retn } V$, or is an unhandled exception raise .

5.5 Another view: context generators

In the discussion of safety in Section 5.2, we needed a special “context invariant” lemma saying that the dynamic semantics preserves the property that a substructural context contains no $(\text{abs } T)$ propositions. Another way of specifying this context invariant is in terms of a transition system that *generates* the context. The degenerate “rewriting rules” that capture the context invariant for parallel evaluation are as follows:

$$\begin{aligned}
\text{gen}/\text{eval} &: \forall E. \text{eval } E. \\
\text{gen}/\text{retn} &: \forall E. \text{retn } E. \\
\text{gen}/\text{cont} &: \forall F. \text{cont } F. \\
\text{gen}/\text{cont}_2 &: \forall F. \text{cont}_2 F.
\end{aligned}$$

This changes the picture of safety presented in Section 5.2 to the following picture, which shows the two invariants given by the context generating rules and the static semantics. The context generating rules construct the context, and the static semantics then analyze it.




This account of context generating rules bears a strong resemblance to the *regular worlds* specifications mentioned in Section 3 and implemented in the LF/Twelf metalogical framework [27]; I am interested in exploring the use of transition rules as a logical justification for regular worlds specifications. In this setting, however, there is no reason why we need to limit ourselves to regular worlds-like context generators. Another possibility is the use of context generators that capture the tree-like structure of the ordered context:

$$\begin{aligned}
 \text{gen/eval} &: \forall E.\text{gen} \rightarrow \text{eval } E. \\
 \text{gen/retn} &: \forall E.\text{gen} \rightarrow \text{retn } E. \\
 \text{gen/cont} &: \forall F.\text{gen} \rightarrow \text{cont } F \cdot \text{gen}. \\
 \text{gen/cont}_2 &: \forall F.\text{gen} \rightarrow \text{cont}_2 F \cdot \text{gen} \cdot \text{gen}.
 \end{aligned}$$

But what if we don't stop there? We could also have the gen atomic proposition carry a type, and then only allow the generation of well-typed tree structures (we return to leaving quantification implicit):

$$\begin{aligned}
 \text{gen/eval} &: \text{gen } T \cdot \text{!of } E T \rightarrow \text{eval } E. \\
 \text{gen/retn} &: \text{gen } T \cdot \text{!value } E \cdot \text{!of } E T \rightarrow \text{retn } E. \\
 \text{gen/cont} &: \text{gen } T \cdot \text{!off } F T' T \rightarrow \text{cont } F \cdot \text{gen } T'. \\
 \text{gen/cont}_2 &: \text{gen } T \cdot \text{!off}_2 F T_1 T_2 T \rightarrow \text{cont}_2 F \cdot \text{gen } T_1 \cdot \text{gen } T_2.
 \end{aligned}$$

The static semantics of the language now have been captured entirely within the generation rules, seemingly removing any need for a second invariant that analyzes the state. The statement of the safety theorem would remain as it was in Theorem 1, but its proof would look like this:



The perspective suggested by this picture is quite different from our usual intuitions about static semantics. Usually, static semantics are seen as a way of *analyzing* or *abstractly evaluating* a state in order to generate an approximation (the type) that is invariant under evaluation and sufficient to ensure safety; here, the static semantics is more of a template for generating safe states. We have not explored this style of analysis in depth, but it does seem to address certain significant complications that have been encountered in the process of proving safety for the linear-destination-passing-style specifications [24].

6 Conclusion

In this report, I have presented the outlines of a logical framework based on a state-transition-based view of ordered linear logic and presented preliminary results indicating that the framework allows for both modular specification of programming languages and formal reasoning about their safety properties. The presentation of type safety for a SSOS specified language, the primary novel aspect of this report, is interesting for a number of reasons. In particular, the proofs retain much of the character of existing safety proofs — the general pattern of “Safety = Progress + Preservation” adapts well to the ordered SSOS specifications, and existing typing rules can be used without revision. The notion of what it means for an *expression* to be well typed, in other words, remains the same, and the notion of what it means for a *state* to be well typed is expressible in terms of an ordered logical specification. There is much future work in this area, but as that discussion is extensively explored in my thesis proposal [35], I will not discuss it further here. Instead, I will conclude by briefly discussing two strands of related work: work in similar logical frameworks and work in rewriting logic specifications of programming languages.

From the logical framework perspective, the most closely related work is the LF/LLF/OLF/CLF/HLF family of logical frameworks, though the focus that the framework presented here gives to state transitions exists elsewhere only in the CLF framework, and even there state transitions are second-class citizens. The only published work on proving properties of SSOS specifications in CLF is Schack-Nielsen’s proof of the equivalence of a big-step operational semantics and an original-style SSOS specification [33].

The notion that ideas at the intersection of abstract machine specifications and linear logic can capture stateful and concurrent programming language features in a modular way is one with a fairly long history [4, 1, 2, 24, 20, 28]. However, while the power of these frameworks for specification has long been recognized, formal reasoning about properties of these specifications has typically concentrated on properties like approximation [19, 37] and equivalence [33] rather than on familiar properties like progress and preservation that I demonstrate in the transition-based framework. I am aware of two exceptions to this pattern, though in both cases only preservation lemmas were formalized, not progress lemmas. Cervesato’s LLF specification of *MLR* can be seen as a transition-based specification that has been flipped around and turned into a backward-chaining specification (or, equivalently, as a small-step specification that has been forced to present itself as a big-step specification), and Reed’s HLF is able to mechanically verify Cervesato’s preservation theorem [30]. Felty and Momigliano used a similar style of specification in their work on Hybrid to encode abstract machines in ordered logic and reason about subject reduction in either Isabelle/HOL or Coq [8].

While the theoretical basis of this proposal is found in the study of logical frameworks, the most similar project in terms of goals and strategies is the rewriting logic semantics project [17], and in particular the K framework for language specifications [34, 32]. Based on the Maude rewriting framework, these two projects have proven successful in specifying, model-checking, and efficiently executing operational semantics of stateful and concurrent programming languages, including a number of large formalizations of object oriented [13] and functional [16] programming languages.

Many specifications in the K framework bear a strong resemblance to SSOS specifications, and the two approaches to language formalization seem to share a great deal of fundamental structure, even if they differ substantially in emphasis. I am only aware of one discussion of safety via progress and preservation for a K specification [7]; this approach was discussed further in Ellison’s masters thesis [6]. The primary limitation of K relative to our approach is a lack of LF-like canonical forms and higher-order abstract syntax, both of which were critical to our specifications and to reasoning about their type safety; in the K specifications, the existing “off the shelf” typing rules for expressions could not be used in the way we could.

On the other hand, the primary limitation of our approach relative to K is the absence of any mechanism for capturing arbitrary sets of propositions at the rule level — in K there is no distinction between terms and propositions, so it is as if we could write this rule in a specification of first-class continuations:

$$\text{CONT} \cdot \text{eval}(\text{callcc}(\lambda x. E x)) \rightarrow \text{CONT} \cdot \text{eval}(E(\text{continue CONT}))$$

where CONT captures *all* the propositions $\text{cont } F$ to the left of $\text{eval}(\text{callcc}(\lambda x. E x))$. This makes K 's approach to the modular specification of continuations infeasible in our framework. This does not mean that SSOS specifications cannot express first-class continuations, however! Linear-destination-passing style SSOS specifications can easily and naturally express first-class continuations [24]; previous work has furthermore indicated that there may be a formal connection between ordered SSOS specifications and linear-destination-passing SSOS specifications [38].

References

- [1] Iliano Cervesato and Frank Pfenning. A linear logical framework. *Information and Computation*, 179(1):19–75, 2002.
- [2] Iliano Cervesato, Frank Pfenning, David Walker, and Kevin Watkins. A concurrent logical framework II: Examples and applications. Technical Report CMU-CS-2002-002, Department of Computer Science, Carnegie Mellon University, March 2002. Revised May 2003.
- [3] Iliano Cervesato and Andre Scedrov. Relating state-based and process-based concurrency through linear logic. *Information and Computation*, 207:1044–1077, 2009.
- [4] Jawahar Lal Chirimar. *Proof Theoretic Approach To Specification Languages*. PhD thesis, University of Pennsylvania, 1995.
- [5] Henry DeYoung and Frank Pfenning. Reasoning about the consequences of authorization policies in a linear epistemic logic. In *Workshop on Foundations of Computer Security*, pages 9–23. Informal Proceedings, 2009.
- [6] Charles M. Ellison. A rewriting logic approach to defining type systems. Master's thesis, University of Illinois at Urbana-Champaign, 2008.
- [7] Charles M. Ellison, Traian Florin Șerbănuță, and Grigore Roșu. A rewriting logic approach to type inference. Technical Report UIUCDCS-R-2008-2934, University of Illinois at Urbana-Champaign, 2008.
- [8] Amy Felty and Alberto Momigliano. Hybrid: A definitional two level approach to reasoning with higher-order abstract syntax. Submitted for publication, available online: <http://www.site.uottawa.ca/~afelty/bib.html>, September 2008.
- [9] Matthew Fluet, Mike Rainey, John Reppy, and Adam Shaw. Implicitly-threaded parallelism in Mantecore. In *Proceedings of the International Conference on Functional Programming*, pages 119–130. ACM, 2008.
- [10] Jean-Yves Girard. Locus Solum: From the rules of logic to the logic of rules. *Mathematical Structures in Computer Science*, 11(3):301–506, 2001.

- [11] Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
- [12] Robert Harper and Daniel R. Licata. Mechanizing metatheory in a logical framework. *Journal of Functional Programming*, 17(4-5):613–673, 2007.
- [13] Mark Hills and Grigore Roşu. KOOL: An application of rewriting logic to language prototyping and analysis. In *Rewriting Techniques and Applications*, pages 246–256. Springer LNCS 4533, 2007.
- [14] Joshua S. Hodas and Dale Miller. Logic programming in a fragment of intuitionistic linear logic. *Information and Computation*, 110(2):327–365, 1994.
- [15] Per Martin-Löf. On the meanings of the logical constants and the justifications of the logical laws. *Nordic Journal of Philosophical Logic*, 1(1):11–60, 1996.
- [16] Patrick Meredith, Mark Hills, and Grigore Roşu. A K definition of Scheme. Technical Report UIUCDCS-R-2007-2907, University of Illinois at Urbana-Champaign, 2007.
- [17] José Meseguer and Grigore Roşu. The rewriting logic semantics project. *Theoretical Computer Science*, 373:213–237, 2007.
- [18] Dale Miller. A multiple-conclusion meta-logic. *Theoretical Computer Science*, 165(1):201–232, 1996.
- [19] Dale Miller. A proof-theoretic approach to the static analysis of logic programs. In Christoph Benzmueller, Chad E. Brown, Jörg Siekmann, and Richard Statman, editors, *Reasoning in Simple Type Theory: Festschrift in honour of Peter B. Andrews on his 70th birthday*. *Studies in Logic*, volume 17. College Publications, 2008.
- [20] Dale Miller. Formalizing operational semantic specifications in logic. *Electronic Notes in Theoretical Computer Science*, 246:147–165, 2009. Proceedings of WFLP 2008.
- [21] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML: Revised*. MIT Press, Cambridge, Massachusetts, 1997.
- [22] Peter D. Mosses. Modular structural operational semantics. *Journal of Logic and Algebraic Programming*, 60-61:195–228, 2004.
- [23] Peter D. Mosses and Mark J. New. Implicit propagation in structural operational semantics. *Electronic Notes in Theoretical Computer Science*, 229(4):49–66, 2009. Proceedings of SOS 2008.
- [24] Frank Pfenning. Substructural operational semantics and linear destination-passing style. In *Programming Languages and Systems*, page 196. Springer LNCS 3302, 2004. Abstract of invited talk.
- [25] Frank Pfenning. On linear inference. Unpublished note, available online: <http://www.cs.cmu.edu/~fp/papers/lininf08.pdf>, February 2008.
- [26] Frank Pfenning and Rowan Davies. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science*, 11(4):511–540, 2001. Notes to an invited talk at the *Workshop on Intuitionistic Modal Logics and Applications (IMLA'99)*, Trento, Italy, July 1999.

- [27] Frank Pfenning and Carsten Schürmann. System description: Twelf — a meta-logical framework for deductive systems. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, pages 202–206. Springer LNAI 1632, 1999.
- [28] Frank Pfenning and Robert J. Simmons. Substructural operational semantics as ordered logic programming. In *Proceedings of the 24th Annual Symposium on Logic in Computer Science (LICS'09)*, pages 101–110, Los Angeles, California, 2009.
- [29] Jeff Polakow. *Ordered Linear Logic and Applications*. PhD thesis, Carnegie Mellon University, 2001.
- [30] Jason Reed. *A Hybrid Logical Framework*. PhD thesis, Carnegie Mellon University, 2009.
- [31] Jason Reed. A judgmental deconstruction of modal logic. Submitted for publication, available online: <http://www.cs.cmu.edu/~jcreed/papers/jdml.pdf>, May 2009.
- [32] Grigore Roşu and Traian Florin Şerbănuţă. An overview of the K semantic framework. *Journal of Logic and Algebraic Programming*, 2010. Submitted for publication, available online: http://fsl.cs.uiuc.edu/index.php/An_Overview_of_the_K_Semantic_Framework.
- [33] Anders Schack-Nielsen. Induction on concurrent terms. *Electronic Notes in Theoretical Computer Science*, 196:37–51, 2008. Proceedings of LFMTTP 2007.
- [34] Traian Florin Şerbănuţă and Grigore Roşu. K-Maude: A rewriting based tool for semantics of programming languages. In *Rewriting Logic and its Applications*. Springer LNCS, 2010. To appear.
- [35] Robert J. Simmons. Logical frameworks for specifying and reasoning about stateful and concurrent languages. Thesis proposal at Carnegie Mellon University, July 2010.
- [36] Robert J. Simmons and Frank Pfenning. Linear logical algorithms. In *Automata, Languages and Programming*, pages 336–347. Springer LNCS 5126, 2008.
- [37] Robert J. Simmons and Frank Pfenning. Linear logical approximations. In *Partial Evaluation and Program Manipulation*, pages 9–20. ACM, 2009.
- [38] Robert J. Simmons and Frank Pfenning. Logical approximation for program analysis. Submitted for publication, available online: <http://www.cs.cmu.edu/~fp/papers/lapa10.pdf>, 2010.
- [39] Roberto Virga. *Higher-Order Rewriting with Dependent Types*. PhD thesis, Carnegie Mellon University, 1999.
- [40] Kevin Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. A concurrent logical framework I: Judgments and properties. Technical Report CMU-CS-2002-101, Department of Computer Science, Carnegie Mellon University, March 2002. Revised May 2003.
- [41] Noam Zeilberger. *The Logical Basis of Evaluation Order and Pattern-Matching*. PhD thesis, Carnegie Mellon University, 2009.