

Substructural Types

Justin Slepak

1 Three Structural Lemmas

Simply-typed λ -calculus follows three “structural” lemmas.

Lemma 1.1. Exchange:

If $\Gamma \vdash e : \tau$ and Γ' is a permutation of Γ , then $\Gamma' \vdash e : \tau$.

Lemma 1.2. Weakening:

If $\Gamma \vdash e : \tau$ and $x \notin \text{dom}(\Gamma)$, then $\Gamma, x : \tau' \vdash e : \tau$.

Lemma 1.3. Contraction:

If $(\Gamma_1, x_2 : \tau', x_3 : \tau', \Gamma_2) \vdash e : \tau$, then $(\Gamma_1, x_1 : \tau', \Gamma_2) \vdash e[x_1/x_2, x_1/x_3] : \tau$.

Exchange allows the environment to be rearranged without affecting typing judgments. Weakening means that the type environment can be extended with unused entries without affecting typing judgments. Contraction lets multiple environment entries with the same type be merged.

Substructural type systems do not uphold these properties. A *relevant* type system eliminates only weakening and requires that every variable be used at least once. An *affine* type system eliminates only contraction and prevents variables from being used more than once. A *linear* type system eliminates both weakening and contraction and mandates exactly one use of every variable. An *ordered* type system removes all three properties and requires variables to be used once each, in the order of their introduction.

2 λ^{URAL}

For practical purposes, it is more convenient to allow different variables to have different usage restrictions. ATTAPL [3] presents a language which allows both linear and unrestricted types. We will use λ^{URAL} [1], which allows unrestricted, relevant, affine, and linear types. We start with a simple core language based on simply-typed λ -calculus:

2.1 Syntax

$$\begin{aligned} e &::= q n \mid q \lambda x : \tau. e \mid \text{if0 } e e e \mid (e e) \mid (op e e) \mid x && \text{(expressions)} \\ v &::= q n \mid q \lambda x : \tau. e \mid x && \text{(values)} \\ op &::= + \mid - \mid * && \text{(base operations)} \\ \tau &::= q p && \text{(types)} \\ p &::= \text{Int} \mid \tau \rightarrow \tau && \text{(pretypes)} \\ q &::= \text{U} \mid \text{R} \mid \text{A} \mid \text{L} && \text{(qualifiers)} \\ E &::= E e \mid v E \mid (op E e) \mid (op v E) \mid (\text{if0 } E e e) && \text{(evaluation contexts)} \end{aligned}$$

2.2 Operational semantics

The operational semantics are essentially the same as for simply-typed λ -calculus. Certain terms, specifically λ -abstractions and numeric literals, must carry explicit type qualifiers. The only way λ -abstractions can be introduced (syntactically) necessitates specifying a qualifier, but $(op\ e\ e)$ introduces a new number without one, so we define it to have an unrestricted result:

$$(op\ (q_1\ n_1)\ (q_2\ n_2)) \mapsto (\mathbf{U}\ \llbracket n_1\ op\ n_2 \rrbracket)$$

2.3 Infrastructure for typing

A \preceq order relation is needed to express the typing rules. It is defined first for pairs of qualifiers, and then lifted to operate on (type, qualifier) and (environment, qualifier) pairs. The relation on qualifiers is the reflexive-transitive closure of:

$$\begin{aligned} \mathbf{U} &\preceq \mathbf{R} \\ \mathbf{U} &\preceq \mathbf{A} \\ \mathbf{R} &\preceq \mathbf{L} \\ \mathbf{A} &\preceq \mathbf{L} \end{aligned}$$

The relation is lifted for types by $(q\ p) \preceq q' \iff q \preceq q'$ and lifted for environments by $\Gamma \preceq q \iff \forall (x : q' p) \in \Gamma. q' \preceq q$, i.e. every entry in Γ must have a qualifier no more restrictive than q . The typing rules also rely on an “environment splitting” judgment, \boxplus .

$$\begin{array}{c} \overline{\emptyset \boxplus \emptyset = \emptyset} \\ \frac{\Gamma = \Gamma_1 \boxplus \Gamma_2}{\Gamma, x : \tau = (\Gamma_1, x : \tau) \boxplus \Gamma_2} \\ \frac{\Gamma = \Gamma_1 \boxplus \Gamma_2}{\Gamma, x : \tau = \Gamma_1 \boxplus (\Gamma_2, x : \tau)} \\ \frac{\Gamma = \Gamma_1 \boxplus \Gamma_2 \quad \tau \preceq \mathbf{R}}{\Gamma, x : \tau = (\Gamma_1, x : \tau) \boxplus (\Gamma_2, x : \tau)} \end{array}$$

Intuitively, an environment can be split by partitioning into two separate environments. This does not introduce any weakening property, as every binding must appear on at least one side of the split. Depending on how a type derivation is examined, it may be useful to think of them as a “merge” operation: if a pair of environments can typecheck e_1 and e_2 , their merge can typecheck $(e_1\ e_2)$. The merge of two environments is undefined if they share any linear or affine bindings, but relevant and unrestricted bindings may be shared by the two environments being merged. This introduces a limited form of contraction that only applies to relevant and unrestricted variables.

2.4 Typing

$$\frac{\Gamma_1 \vdash e : \tau \quad \Gamma_2 \preceq \mathbf{A}}{\Gamma_1 \boxplus \Gamma_2 \vdash e : \tau} (\text{T-WEAKEN})$$

$$\frac{}{\bullet, x : \tau \vdash x : \tau} (\text{T-VAR}) \quad \frac{}{\bullet \vdash q \ n : q \ \mathbf{Int}} (\text{T-INT})$$

Typing a numeric literal requires the empty environment, and typing a variable requires an environment containing only that variable. The explicit weakening rule allows the empty or singleton environment to be expanded by adding bindings to match the larger environment that actually describes the current scope, but we can only add variables which are neither relevant nor linear (it also gives us the exchange property in the same way as TAPL's T-VAR rule does). ATTAPL's presentation of the typing rules rolls the use of weakening into the variable and base-value literal rules.

$$\frac{\Gamma_1 \vdash e_{if} : q \ \mathbf{Int} \quad \Gamma_2 \vdash e_{then} : \tau \quad \Gamma_2 \vdash e_{else} : \tau}{\Gamma_1 \boxplus \Gamma_2 \vdash (\mathbf{if0} \ e_{if} \ e_{then} \ e_{else}) : \tau} (\text{T-IF0})$$

Any affine and linear variables used in e_{if} are unavailable for use in e_{then} and e_{else} . Relevant and linear variables not used in e_{if} must be used in e_{then} and e_{else} . The environment splitting handles this arrangement.

$$\frac{\Gamma \preceq q \quad \Gamma, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash q \lambda x : \tau. e : q(\tau_1 \rightarrow \tau_2)} (\text{T-ABST})$$

A closure that references already-bound variables must respect their linear, affine, or relevant restrictions, so the same restrictions must apply to the closure itself. T-WEAKEN allows variables not referenced by this closure to be excluded from the environment used to typecheck it.

$$\frac{\Gamma_1 \vdash e_1 : q(\tau \rightarrow \tau') \quad \Gamma_2 \vdash e_2 : \tau}{\Gamma_1 \boxplus \Gamma_2 \vdash (e_1 \ e_2) : \tau'} (\text{T-APP})$$

$$\frac{\Gamma_1 \vdash e_1 : q_1 \ \mathbf{Int} \quad \Gamma_2 \vdash e_2 : q_2 \ \mathbf{Int}}{\Gamma_1 \boxplus \Gamma_2 \vdash (op \ e_1 \ e_2) : \mathbf{U} \ \mathbf{Int}} (\text{T-ARITH})$$

The reasoning in T-APP and T-ARITH is similar to that for T-IF0. Variables consumed by e_1 are unavailable for e_2 . Arithmetic operators are defined to produce unrestricted values.

3 Language Extensions

We now extend the language with several new features, all of which have been seen in class as extensions for λ -calculus, but some complications are introduced by substructural typing.

3.1 Pairs and product types

$$\begin{aligned}
 e & ::= \dots \mid q \langle e, e \rangle \mid \underline{\quad} && \text{(expressions)} \\
 v & ::= \dots \mid q \langle v, v \rangle && \text{(values)} \\
 p & ::= \dots \mid \tau \times \tau && \text{(pretypes)} \\
 E & ::= \dots \mid q \langle E, e \rangle \mid q \langle v, E \rangle && \text{(evaluation contexts)}
 \end{aligned}$$

The pair introduction form is obvious (simply add a qualifier); the immediate problem is how to construct an elimination form for pairs. Before, we used **fst** and **snd** operators to extract individual elements, but this keeps us from being able to access both terms of a linear or affine pair. Instead of **fst** and **snd**, we introduce a new binding form:

$$\begin{aligned}
 e & ::= \dots \mid \mathbf{letpair} \ e \langle x, x \rangle e && \text{(expressions)} \\
 E & ::= \dots \mid \mathbf{letpair} \ E \langle x, x \rangle e && \text{(evaluation contexts)}
 \end{aligned}$$

This form simultaneously binds both elements of a pair in another expression, so both elements can be accessed with a single use of the pair. We extend the operational semantics:

$$(\mathbf{letpair} \ (q \langle v_1, v_2 \rangle) \langle x_1, x_2 \rangle e) \mapsto e[v_1/x_1, v_2/x_2]$$

Type checking is substantially similar to the structurally-typed λ -calculus, but we have a few restrictions on qualifiers. We should not allow, for example, an unrestricted pair of affine values, as the pair could be duplicated, bypassing the affine restrictions for its contents. So we require that a pair's qualifier be at least as restrictive as the qualifiers of its contents.

$$\frac{\Gamma_1 \vdash e_l : \tau_l \quad \Gamma_2 \vdash e_r : \tau_r \quad \tau_l \preceq q \quad \tau_r \preceq q}{\Gamma_1 \boxplus \Gamma_2 \vdash (q \langle e_l, e_r \rangle) : q (\tau_l \times \tau_r)} \text{(T-PAIR)}$$

Type checking a **letpair** looks similar to type checking a λ -abstraction, except two variables are introduced into the environment used to check the body.

$$\frac{\Gamma_1 \vdash e_{arg} : q (\tau_1 \times \tau_2) \quad \Gamma_2, x_1 : \tau_1, x_2 : \tau_2 \vdash e_{body} : \tau}{\Gamma_1 \boxplus \Gamma_2 \vdash (\mathbf{letpair} \ e_{arg} \langle x_1, x_2 \rangle e_{body}) : \tau} \text{(T-LETPAIR)}$$

3.2 Sum types

$$\begin{aligned}
e &::= \dots \mid q \text{ inL}_p e \mid q \text{ inR}_p e && (\text{expressions}) \\
&\quad \mid \text{case } e (\text{inL } x \Rightarrow e) (\text{inR } x \Rightarrow e) \\
v &::= \dots \mid q \text{ inL}_p v \mid q \text{ inR}_p v && (\text{values}) \\
p &::= \dots \mid \tau + \tau && (\text{pretypes}) \\
E &::= \dots \mid q \text{ inL}_p E \mid q \text{ inR}_p E && (\text{evaluation contexts})
\end{aligned}$$

The associated semantic extension is straightforward:

$$\begin{aligned}
(\text{case } (q \text{ inL}_p v) (\text{inL } x_l \Rightarrow e_l) (\text{inR } x_r \Rightarrow e_r)) &\mapsto e_l[v/x_l] \\
(\text{case } (q \text{ inR}_p v) (\text{inL } x_l \Rightarrow e_l) (\text{inR } x_r \Rightarrow e_r)) &\mapsto e_r[v/x_r]
\end{aligned}$$

Again, typing a sum requires preventing usage restrictions from being bypassed by wrapping a value in a less restrictive sum type (but consider what happens if we modify our typing rules to allow something like $\text{U } (\text{U } p1) + (\text{L } p2)$ or $\text{U } (\text{L } p1) + (\text{L } p2)$ to be a well-formed type).

$$\begin{aligned}
&\frac{\Gamma \vdash e : \tau_l \quad \tau_l \preceq q \quad \tau_r \preceq q}{\Gamma \vdash (q \text{ inL}_{\tau_l + \tau_r} e) : q (\tau_l + \tau_r)} (\text{T-INL}) \\
&\frac{\Gamma \vdash e : \tau_r \quad \tau_l \preceq q \quad \tau_r \preceq q}{\Gamma \vdash (q \text{ inR}_{\tau_l + \tau_r} e) : q (\tau_l + \tau_r)} (\text{T-INR}) \\
&\frac{\Gamma_1 \vdash e : q (\tau_l + \tau_r) \quad \Gamma_2, x_l : \tau_l \vdash e_l : \tau \quad \Gamma_2, x_r : \tau_r \vdash e_r : \tau}{\Gamma_1 \boxplus \Gamma_2 \vdash (\text{case } e (\text{inL } x_l \Rightarrow e_l) (\text{inR } x_r \Rightarrow e_r)) : \tau} (\text{T-CASE})
\end{aligned}$$

3.3 Recursive types

$$\begin{aligned}
e &::= \dots \mid \text{fold}_p e \mid \text{unfold } e && (\text{expressions}) \\
v &::= \dots \mid \text{fold}_p v && (\text{values}) \\
p &::= \dots \mid \alpha \mid \mu\alpha.\tau && (\text{pretypes}) \\
E &::= \dots \mid \text{fold}_p E \mid \text{unfold } E && (\text{evaluation contexts})
\end{aligned}$$

$$(\text{unfold } (\text{fold}_p v)) \mapsto v$$

Again, the concern is to prevent wrapping one type in a less restrictive type. We will let a `folded` expression keep the same qualifier as the `pre-fold` expression, i.e. the qualifier outside a $\mu\alpha.\tau$ pretype is ignored with use restrictions coming from τ 's qualifier. Because of the addition of type variables, we must introduce a type variable environment, Δ , to check that types are well-formed.

$$\frac{\Gamma \vdash e : \tau[\mu\alpha.\tau/\alpha] \quad \bullet \vdash \tau[\mu\alpha.\tau/\alpha] \preceq q}{\Gamma \vdash (\mathbf{fold}_{\mu\alpha.\tau} e) : q \mu\alpha.\tau} \text{(T-FOLD)}$$

$$\frac{\Gamma \vdash e : (q \mu\alpha.\tau)}{\Gamma \vdash (\mathbf{unfold} e) : \tau[\mu\alpha.\tau/\alpha]} \text{(T-UNFOLD)}$$

The typing rules for `fold` and `unfold` are straightforward and similar to the rules without substructural types. We cannot have type variables appear anywhere but a $\mu\alpha.\tau$ pretype, so we require that the empty environment be enough to prove our types to be well-formed. In proving this, it is necessary to alternate (mutually recursively) between proving well-formedness of a type and well-formedness of a pretype. We introduce this as a judgment that p is a well-formed pretype, $p : P$, or that τ is a well-formed type, $\tau : T$, but Δ contains only pretype variables. The rules for well-formedness are fairly intuitive, and there is only one way for variables to be introduced into the environment:

$$\frac{\Delta \vdash p : P}{\Delta \vdash (q p) : T} \quad \frac{\Delta, \alpha \vdash \tau : T}{\Delta \vdash \mu\alpha.\tau : T} \quad \frac{\Delta \vdash \tau : T \quad \Delta \vdash \tau' : T}{\Delta \vdash q (\tau \rightarrow \tau') : T}$$

$$\frac{\Delta \vdash \tau : T \quad \Delta \vdash \tau' : T}{\Delta \vdash q (\tau + \tau') : T} \quad \frac{\Delta \vdash \tau : T \quad \Delta \vdash \tau' : T}{\Delta \vdash q (\tau \times \tau') : T}$$

We also bring Δ into our judgments for \preceq . At this point, we needn't reference Δ when comparing qualifiers, but we use it in lifting \preceq for types and environments.

3.4 Polymorphism

$$\begin{aligned} e ::= \dots & \mid q\Lambda\alpha.e \mid e\langle p \rangle \mid q\Lambda\xi.e \mid e\langle q \rangle && \text{(expressions)} \\ v ::= \dots & \mid q\Lambda\alpha.e \mid q\Lambda\xi.e && \text{(values)} \\ p ::= \dots & \mid \forall\alpha.\tau \mid \forall\xi.\tau && \text{(pretypes)} \\ q ::= \dots & \mid \xi && \text{(qualifiers)} \\ E ::= \dots & \mid E\langle p \rangle \mid E\langle q \rangle && \text{(evaluation contexts)} \end{aligned}$$

We quantify separately over pretypes and qualifiers. This allows a quantified type to express that a function “works on `Ints` of any restriction level” (e.g. a purely arithmetic function) or “works for any unrestricted α ” (e.g. a type signature for Church numerals) or “works on any types” (e.g. `curry` and `uncurry`). This syntax extension keeps α as the generic pretype variable and introduces ξ as the generic qualifier variable. The corresponding semantic extension follows:

$$\begin{aligned} (q\Lambda\alpha.e)\langle p \rangle &\mapsto e[p/\alpha] \\ (q_0\Lambda\xi.e)\langle q_1 \rangle &\mapsto e[q_1/\xi] \end{aligned}$$

Rather than requiring the programmer to explicitly specify a qualifier and pretype which will only be used with each other, we could even allow quantification over full types, introducing type variables and type application.

$$\begin{aligned}
e & ::= \dots \mid q\Lambda\alpha.e \mid e\langle p \rangle && \text{(expressions)} \\
v & ::= \dots \mid q\Lambda\alpha.e \mid q\Lambda\xi.e && \text{(values)} \\
\tau & ::= \dots \mid \sigma && \text{(types)} \\
E & ::= \dots \mid E\langle \tau \rangle && \text{(evaluation contexts)}
\end{aligned}$$

Checking pretype-polymorphic terms is similar to checking them in System F. Permitting qualifier polymorphism in addition to pretype polymorphism requires that our Δ also track qualifier variables (if pretype, qualifier, and type variables are to be drawn from the same set of identifiers, Δ must track $\alpha : \mathbf{P}$, $\xi : \mathbf{Q}$, and $\sigma : \mathbf{T}$ bindings, indicating whether a type-level variable represents a pretype, qualifier, or full type). At this point, we must also convert our old typing rules to use Δ in addition to Γ , as type-level variables may appear in expressions. We thread Δ through the old typing rules to ensure that types are well-formed, but none of the old rules are able to modify Δ (T-FOLD must be changed to use Δ in place of \bullet). The introduction of qualifier variables means that we now need Δ to determine whether $\Gamma = \Gamma_1 \boxplus \Gamma_2$; our old typing rules must replace the $\Gamma_1 \boxplus \Gamma_2 \vdash e : \tau$ presentation by adding a premise $\Delta \vdash \Gamma = \Gamma_1 \boxplus \Gamma_2$ and rewriting the conclusion as $\Gamma \vdash e : \tau$. In order to handle qualifier variables in both new and old typing rules, we must also extend our \preceq relation. We have \mathbf{U} as the least restrictive qualifier and \mathbf{L} as the most restrictive, so we can add $\mathbf{U} \preceq \xi$ and $\xi \preceq \mathbf{L}$ to the relation. However, knowing nothing more about ξ means we cannot extend the order any further.

$$\begin{aligned}
& \frac{\Gamma \preceq q \quad \Delta, \alpha : \mathbf{P}; \Gamma \vdash e : \tau}{\Delta; \Gamma \vdash q\Lambda\alpha.e : q\forall\alpha.\tau} \text{(T-PREABST)} \\
& \frac{\Delta; \Gamma \vdash e : q\forall\alpha.\tau \quad \Delta \vdash p : \mathbf{P}}{\Delta; \Gamma \vdash e\langle p \rangle : \tau[p/\alpha]} \text{(T-PREAPP)} \\
& \frac{\Gamma \preceq q \quad \Delta, \xi : \mathbf{Q}; \Gamma \vdash e : \tau}{\Delta; \Gamma \vdash q\Lambda\xi.e : q\forall\xi.\tau} \text{(T-QUALABST)} \\
& \frac{\Delta; \Gamma \vdash e : q'\forall\xi.\tau \quad \Delta \vdash q : \mathbf{Q}}{\Delta; \Gamma \vdash e\langle q \rangle : \tau[q/\xi]} \text{(T-QUALAPP)} \\
& \frac{\Gamma \preceq q \quad \Delta, \sigma : \mathbf{T}; \Gamma \vdash e : \tau}{\Delta; \Gamma \vdash q\Lambda\sigma.e : q\forall\sigma.\tau} \text{(T-TYPEABST)} \\
& \frac{\Delta; \Gamma \vdash e : q\forall\alpha.\tau' \quad \Delta \vdash \tau' : \mathbf{T}}{\Delta; \Gamma \vdash e\langle \tau \rangle : \tau'[\tau/\alpha]} \text{(T-TYPEAPP)}
\end{aligned}$$

3.5 State

$$\begin{aligned}
e &::= \dots \mid q \text{ new } e \mid \text{free } e \mid \text{rd } e \mid \text{wr } e e \mid \text{sw } e e && \text{(expressions)} \\
v &::= \dots \mid l \in \text{LOCATIONS} && \text{(values)} \\
p &::= \dots \mid \text{ref } \tau && \text{(pretypes)} \\
E &::= \dots \mid \text{new } q E \mid \text{free } E \mid \text{rd } E \mid \text{wr } E e \mid \text{wr } v E && \text{(evaluation contexts)} \\
&\quad \mid \text{sw } E e \mid \text{sw } v E
\end{aligned}$$

This extension places “locations,” i.e. names of reference cells, which can be used alongside the store’s variable mappings or in a store machine which simply operates on the control string and does not introduce variable mappings into the store.

The machine-oriented operational semantics for state maps a location to both its usage restrictions and its contents. The reason for having **sw**, a swap operation, in addition to **rd** and **wr** arises from typing restrictions and will be explained with the typing rules.

$$\begin{aligned}
(S; q \text{ new } v) &\mapsto (S, l \Rightarrow (q v); l) \\
((S_1, l \Rightarrow (q v), S_2); \text{free } l) &\mapsto ((S_1, S_2); v) \\
((S_1, l \Rightarrow (q v), S_2); \text{rd } l) &\mapsto ((S_1, l \Rightarrow (q v), S_2); L < l, v >) \\
((S_1, l \Rightarrow (q v), S_2); \text{wr } l v') &\mapsto ((S_1, l \Rightarrow (q v'), S_2); l) \\
((S_1, l \Rightarrow (q v), S_2); \text{sw } l v') &\mapsto ((S_1, l \Rightarrow (q v'), S_2); L < l, v >)
\end{aligned}$$

A reference cell that can only be accessed once is not useful, so accesses to a reference cell include a new instance of the reference in their return value. In the case of **rd** and **sw**, this necessitates constructing a pair so that the contents of the cell can also be returned. The pair is linear to avoid illicit duplication both of the reference and of the returned cell contents (recall that a pair type must be at least as restrictive as either of its component types). If either of these is legal to duplicate, that duplication can be performed after splitting the pair. It is important to remember that the q in a $q(\text{ref } \tau)$ refers not to the reference cell or the data it contains but to the reference itself. The τ behaves according to its own qualifier, and the reference cell remains in the heap until freed (explicitly or by a garbage collector). This brings up the question of what operations are compatible with what combinations of reference and data qualifiers.

$$\frac{\Delta \vdash q \preceq A \quad \Delta \vdash q' \preceq A \quad \Delta; \Gamma \vdash e : q' p}{\Delta; \Gamma \vdash \text{new } q e : q (\text{ref } q' p)} \text{(T-NEWUA)}$$

$$\frac{\Delta \vdash R \preceq q \quad \Delta; \Gamma \vdash e : \tau}{\Delta; \Gamma \vdash q \text{ new } e : q (\text{ref } \tau)} \text{(T-NEWRL)}$$

Separate typing rules are needed for constructing unrestricted/affine references and relevant/linear references. An unrestricted or affine reference cell may never be accessed, so it cannot contain relevant or linear data. On the other hand, a relevant or linear reference cell is guaranteed to be accessed, so it can contain anything (though the set of allowable operations will still depend on the qualifier of the cell contents). Notice that we are not confined according to the \preceq lattice: it is legal for a relevant reference cell to contain affine or linear data or for an unrestricted cell to contain affine data.

$$\frac{\Delta; \Gamma \vdash e : q(\mathbf{ref} \tau) \quad \Delta \vdash \mathbf{A} \preceq q}{\Delta; \Gamma \vdash \mathbf{free} e : \tau} (\text{T-FREE})$$

A relevant or unrestricted reference might still be used again later in the program, so freeing it is unsafe. Because this is a use of the reference, we can be sure that an affine or linear reference will not be used after it is freed.

$$\frac{\Delta; \Gamma \vdash e : q(\mathbf{ref} \tau) \quad \Delta \vdash \tau \preceq \mathbf{R}}{\Delta; \Gamma \vdash \mathbf{rd} e : \mathbf{L}((q \mathbf{ref} \tau) \times \tau)} (\text{T-READ})$$

Multiple read references (possible due to the non-destructive nature of \mathbf{rd} , as opposed to \mathbf{free}) makes it possible to duplicate the contents of a reference cell. This is not permissible if the cell contains affine or linear data.

$$\frac{\Delta \vdash \Gamma = \Gamma_1 \boxplus \Gamma_2 \quad \Delta; \Gamma_1 \vdash e_1 : q(\mathbf{ref} \tau) \quad \Delta; \Gamma_2 \vdash e_2 : \tau \quad \Delta \vdash \tau \preceq \mathbf{A}}{\Delta; \Gamma \vdash \mathbf{wr} e_1 e_2 : q \mathbf{ref} \tau} (\text{T-WEAKWRITE})$$

The typing rule for weak (i.e. type-preserving) writes must ensure compatibility between the reference and contents expressions, i.e. that the program is writing a τ into a $\mathbf{ref} \tau$. It must also ensure that the program is not overwriting a relevant or linear value because that would allow such values to be destroyed without ever being used.

$$\frac{\Delta \vdash \Gamma = \Gamma_1 \boxplus \Gamma_2 \quad \Delta; \Gamma_1 \vdash e_1 : q(\mathbf{ref} \tau) \quad \Delta; \Gamma_2 \vdash e_2 : \tau}{\Delta; \Gamma \vdash \mathbf{sw} e_1 e_2 : \mathbf{L}((q \mathbf{ref} \tau) \times \tau)} (\text{T-WEAKSWAP})$$

What can we do with reference cells that contain linear data? Reading requires the contents to be no more restricted than \mathbf{R} . Writing requires the contents to be no more restricted than \mathbf{A} . Tracking whether a cell has been referenced already in order to decide whether to allow reading/writing is too complicated. Instead, we require that the cell contents always keep a “not-yet-used” status. Then we can take a linear (or less restricted) value out of a cell by exchanging it with another value of the same type. This allows us to write the cell contents without the $\tau \preceq \mathbf{A}$ requirement and to read the cell contents without the $\tau \preceq \mathbf{R}$ requirement.

$$\frac{\begin{array}{c} \Delta \vdash \Gamma = \Gamma_1 \boxplus \Gamma_2 \quad \Delta; \Gamma_1 \vdash e_1 : q \text{ (ref } \tau) \\ \Delta; \Gamma_2 \vdash e_2 : \tau \quad \Delta \vdash \tau \preceq \mathbf{A} \\ \Delta \vdash \mathbf{A} \preceq q \quad \Delta \vdash \tau' \preceq q \end{array}}{\Delta; \Gamma \vdash \mathbf{wr} \ e_1 \ e_2 : q \text{ ref } \tau} \text{(T-STRONGWRITE)}$$

In order to make a strong (i.e. type-altering) write, we must be certain that the reference in question is unique (i.e. affine or linear) so that typechecking other parts of the program does not need to consider whether this reference has changed its type. We also require that the new content type be permissive enough to be allowed in this reference (e.g. a relevant reference cannot be updated to contain affine data, but a linear reference can).

$$\frac{\begin{array}{c} \Delta \vdash \Gamma = \Gamma_1 \boxplus \Gamma_2 \quad \Delta; \Gamma_1 \vdash e_1 : q \text{ (ref } \tau) \quad \Delta; \Gamma_2 \vdash e_2 : \tau \\ \Delta \vdash \mathbf{A} \preceq q \quad \Delta \vdash \tau' \preceq q \end{array}}{\Delta; \Gamma \vdash \mathbf{sw} \ e_1 \ e_2 : \mathbf{L} \ ((q \text{ ref } \tau) \times \tau)} \text{(T-STRONGSWAP)}$$

The same restrictions added for strong writes also apply to strong swaps.

4 Progress and preservation for λ^{UAL}

Our type system is meant to ensure certain values are not used multiple times. In order to state progress and preservation for this system in such terms, we must express the operational semantics in terms of a machine with a variable store. To show how the store is updated by accesses to affine and linear variables, we must define a \sim_q operation:

$$\begin{aligned} S \sim_U x &= S \\ (S_1, x \mapsto v, S_2) \sim_A x &= (S_1, S_2) \\ (S_1, x \mapsto v, S_2) \sim_L x &= (S_1, S_2) \end{aligned}$$

This removes affine/linear entries from the store upon use. The machine's strategy will be to select an evaluation context and then evaluate the expression in its hole.

$$\frac{(S; e) \mapsto_\beta (S'; e')}{(S; E[e]) \mapsto (S'; E[e'])}$$

The machine-oriented semantics reduces expressions to variables, which refer to data in the store. Not all of the rules are given here, as they are straightforward transformations from the expression-only semantics. The general theme is that steps which compute new values create new store entries with fresh names.

$$\begin{aligned}
& (S ; q n) \mapsto_{\beta} (S, x \mapsto q n ; x), \\
& \quad \text{with fresh } x \\
& ((S_1, x \mapsto q 0, S_2); (\text{if} 0 x e_1 e_2)) \mapsto_{\beta} ((S_1, x \mapsto q 0, S_2) \sim_q x ; e_1) \\
& ((S_1, x \mapsto q n, S_2); (\text{if} 0 x e_1 e_2)) \mapsto_{\beta} ((S_1, x \mapsto q 0, S_2) \sim_q x ; e_2), \\
& \quad \text{where } n \neq 0 \\
& (S ; q \lambda x_1 : \tau. e) \mapsto_{\beta} (S, x_0 \mapsto q \lambda x_1 : \tau. e ; x_0), \\
& \quad \text{with fresh } x_0 \\
& ((S_1, x_0 \mapsto q \lambda x_1 : \tau. e, S_2); (x_0 x_2)) \mapsto_{\beta} ((S_1, x_0 \mapsto q \lambda x_1 : \tau. e, S_2) \sim_q x_0 ; \\
& \quad e[x_2/x_1]) \\
& (S ; (q < x_l, x_r >)) \mapsto_{\beta} ((S, x_0 \mapsto (q < x_l, x_r >)) ; x_0), \\
& \quad \text{with fresh } x_0 \\
& \textit{etc.}
\end{aligned}$$

We also require a notion of store typing. In essence, we map a store to an environment, which should then be capable of typechecking the expression the machine has at the same time as it has that store. Thus, already-used affine and linear variables are excluded from the environment.

$$\frac{}{\vdash \emptyset : \emptyset} \quad \frac{\vdash S : \Gamma \boxplus \Gamma' \quad \Gamma' \vdash x : \tau}{\vdash S, x \mapsto v : \Gamma, x : \tau}$$

Finally, we combine our rules in order to judge well-typedness for machine states, i.e. store/expression pairs:

$$\frac{\vdash S : \Gamma \quad \Gamma \vdash e : \tau}{\vdash (S ; e)}$$

This gives us the machinery we need to state progress and preservation.

Theorem 4.1. Progress:

If $\vdash (S ; e)$, then $\exists S', e'. (S ; e) \mapsto (S' ; e')$ or e is a value.

Theorem 4.2. Preservation:

If $\vdash (S ; e)$ and $(S ; e) \mapsto (S' ; e')$, then $\vdash (S' ; e')$.

But what about relevant types? Our machine knows exactly when to eliminate linear and affine variables, but removing a relevant variable is unsafe as it may be used again in the future (i.e. progress would be lost). We also cannot simply leave it there because the machine will eventually reach a state where the corresponding Γ will fail to typecheck its control expression (i.e. preservation would be lost). We can introduce a “used” flag into each store entry. On creation, the flag is clear, and it is set when the variable is used by the control string redex (this flag only needs to be managed for relevant variables). The environment we get from store typing would also track whether relevant entries are already-used, and they can be considered as subject to T-WEAKEN.

5 Algorithmic type checking for λ^{UAL}

In order to derive $\Gamma \vdash e : \tau$, the type checker must often choose how to split Γ into Γ_1 and Γ_2 . If a mechanical type checker fails to prove $\Gamma \vdash e : \tau$, then it must backtrack and try again with different Γ_1 and Γ_2 until it knows that no choice of Γ_1 and Γ_2 will work. To avoid nondeterminism in type checking, we introduce the “output environment” associated with a typing judgment: $\Gamma \vdash e : \tau // \Gamma'$. Γ' contains all entries from Γ not “used up” in deriving $e : \tau$. We then replace \boxplus , the “environment splitting” operation, with \div , a deterministic “environment difference” operation:

$$\begin{array}{c} \overline{\Gamma \div \emptyset = \Gamma} \\ \frac{\Gamma_1 \div \Gamma_2 = \Gamma_3 \quad (x : \tau) \notin \Gamma_3 \quad \mathbf{A} \preceq \tau}{\Gamma_1 \div (\Gamma_2, x : \tau) = \Gamma_3} \\ \frac{\Gamma_1 \div \Gamma_2 = \Gamma_3 \quad \Gamma_3 = \Gamma_4, x : \tau, \Gamma_5 \quad \tau \preceq \mathbf{U}}{\Gamma_1 \div (\Gamma_2, x : \tau) = \Gamma_4, \Gamma_5} \end{array}$$

We use this operation to remove variables at the end of their scope, and we ensure at removal that the variables are not affine or linear variables left untouched during their scope. We accept a term as well-typed if the output environment contains no linear variables.

$$\begin{array}{c} \overline{\Gamma \vdash q \ n : q \ \mathbf{Int} // \Gamma} \text{(A-INT)} \\ \frac{\tau \preceq \mathbf{U}}{\Gamma_1, x : \tau, \Gamma_2 \vdash x : \tau // \Gamma_1, x : \tau, \Gamma_2} \text{(A-UVAR)} \\ \frac{\mathbf{A} \preceq \tau}{\Gamma_1, x : \tau, \Gamma_2 \vdash x : \tau // \Gamma_1, x : \tau, \Gamma_2} \text{(A-ALVAR)} \end{array}$$

Using an integer or unrestricted variable does not alter the environment, whereas using a linear variable requires that it be removed from the output environment.

$$\begin{array}{c} \frac{\Gamma_1 \vdash e_1 : \tau \rightarrow \tau' // \Gamma_2 \quad \Gamma_2 \vdash e_2 : \tau // \Gamma_3}{\Gamma_1 \vdash (e_1 \ e_2) : \tau' // \Gamma_3} \text{(A-APP)} \\ \frac{\Gamma_1 \vdash e_1 : q \ \mathbf{Int} // \Gamma_2 \quad \Gamma_2 \vdash e_2 : q' \ \mathbf{Int} // \Gamma_3}{\Gamma_1 \vdash (op \ e_1 \ e_2) : \mathbf{U} \ \mathbf{Int} // \Gamma_3} \text{(A-ARITH)} \end{array}$$

Checking multiple subexpressions chains the environment changes together. Variables consumed by e_1 are unavailable for e_2 . Variables consumed by either e_1 or e_2 are considered consumed by $(e_1 \ e_2)$ or $(op \ e_1 \ e_2)$.

$$\frac{\Gamma_1, x : \tau \vdash e : \tau' // \Gamma_2 \quad q \preceq \mathbf{U} \Rightarrow \Gamma_2 \div (x : \tau) = \Gamma_1}{\Gamma_1 \vdash q\lambda x : \tau.e : q(\tau \rightarrow \tau') // \Gamma_2 \div (x : \tau)} \text{(A-ABST)}$$

Giving $\Gamma_2 \div (x : \tau)$ as the output environment for typechecking a λ -abstraction removes the parameter variable from the environment at the end of its scope. In the case of an unrestricted closure, we use the $\Gamma_2 \div (x : \tau) = \Gamma_1$ requirement to ensure the closure does not use any linear variables from its environment. If it does, Γ_2 will have them removed, and $\Gamma_2 \div (x : \tau)$ will exclude those bindings.

$$\frac{\Gamma_1 \vdash e_{if} : q \mathbf{Int} // \Gamma_2 \quad \Gamma_2 \vdash e_{then} : \tau // \Gamma_3 \quad \Gamma_2 \vdash e_{else} : \tau // \Gamma_3}{\Gamma_1 \vdash (\mathbf{if0} \ e_{if} \ e_{then} \ e_{else}) : \tau // \Gamma_3} \text{(A-IF0)}$$

This rule is similar to T-APP and T-ARITH. By mandating the same output environment when checking e_{then} and e_{else} , we require that they consume the same linear variables.

References

- [1] Amal Ahmed. A step-indexed model of substructural state. In *In: Proc. International Conference on Functional Programming. (2005) 7891*, pages 78–91. ACM Press, 2005.

This paper describes λ^{URAL} . The implications of substructural reference types are the most important insight for this lecture. Included is a table describing what data may be stored in and operations may be performed on reference cells according to their qualifiers. Instead of progress and preservation, type soundness is proven via a step-indexed logical relation.

- [2] Greg Morrisett, Amal Ahmed, and Matthew Fluet. L 3 : A linear language with locations. In *In Seventh International Conference on Typed Lambda Calculi and Applications*, pages 293–307, 2005.

This paper describes a language which uses substructural state to allow strong updates of reference cells. Because a linear reference cell is only accessible from one part of the program at a time, the type checker is able to track type-changing updates to the cell.

- [3] Benjamin C. Pierce, editor. *Advanced Topics in Types and Programming Languages*. MIT Press, 2005.

Chapter 1 covers substructural types, with an emphasis on a type system which includes linear and unrestricted types (but not affine or relevant types). Topics not covered in this lecture include: ordered types, which eliminate all three substructural lemmas mentioned at the beginning; reference-counted types, which introduce a controlled way to duplicate affine or linear references; a language system which uses affine types to control the time a function may expend while running; and compiler optimizations which can take advantage of the information carried by type qualifiers.