

CIS 500

Software Foundations

Fall 2003

27 October

Administrivia

- ◆ Changes to recitation schedule:
 - ◆ The advanced recitations on Wednesday 6:00 - 7:30 and Thursday 6:30 - 8:00 are cancelled. The Wednesday 3:30 - 5:00 advanced recitation continues as before.
 - ◆ Two new review recitations are now available at [...]
 - ◆ You are welcome to attend any recitation(s) you find convenient. (If we have capacity problems, we will figure something out later.)
The complete list of recitation times and meeting places is available on the course web page.

References (continued)

Evaluation

What is the **value** of the expression `ref 0`?

Evaluation

What is the **value** of the expression `ref 0`?

Crucial observation: evaluating `ref 0` must **do** something.

Otherwise,

```
r = ref 0
```

```
s = ref 0
```

and

```
r = ref 0
```

```
s = r
```

would behave the same.

Evaluation

What is the **value** of the expression `ref 0`?

Crucial observation: evaluating `ref 0` must **do** something.

Otherwise,

```
r = ref 0
```

```
s = ref 0
```

and

```
r = ref 0
```

```
s = r
```

would behave the same.

Specifically, evaluating `ref 0` should **allocate some storage** and yield a **reference** (or **pointer**) to that storage.

Evaluation

What is the **value** of the expression `ref 0`?

Crucial observation: evaluating `ref 0` must **do** something.

Otherwise,

```
r = ref 0
```

```
s = ref 0
```

and

```
r = ref 0
```

```
s = r
```

would behave the same.

Specifically, evaluating `ref 0` should **allocate some storage** and yield a **reference** (or **pointer**) to that storage.

So what is a reference?

The Store

A reference names a **location** in the **store** (also known as the **heap** or just the **memory**).

What is the store?

The Store

A reference names a **location** in the **store** (also known as the **heap** or just the **memory**).

What is the store?

- ◆ **Concretely:** An array of 8-bit bytes, indexed by 32-bit integers.

The Store

A reference names a **location** in the **store** (also known as the **heap** or just the **memory**).

What is the store?

- ◆ **Concretely:** An array of 8-bit bytes, indexed by 32-bit integers.
- ◆ **More abstractly:** an array of **values**

The Store

A reference names a **location** in the **store** (also known as the **heap** or just the **memory**).

What is the store?

- ◆ **Concretely:** An array of 8-bit bytes, indexed by 32-bit integers.
- ◆ **More abstractly:** an array of **values**
- ◆ **Even more abstractly:** a partial function from **locations** to **values**.

Locations

Syntax of values:

$v ::=$

unit

$\lambda x:T.t$

l

values

unit constant

abstraction value

store location

... and since all values are terms...

Syntax of Terms

$t ::=$

unit

x

$\lambda x:T.t$

$t t$

ref t

! t

$t := t$

l

terms

unit constant

variable

abstraction

application

reference creation

dereference

assignment

store location

Aside

Does this mean we are going to allow programmers to write explicit locations in their programs??

No: This is just a modeling trick. We are enriching the “source language” to include some run-time structures, so that we can continue to formalize evaluation as a relation between source terms.

Aside: If we formalize evaluation in the big-step style, then we can add locations to the set of values (results of evaluation) without adding them to the set of terms.

Evaluation

The result of evaluating a term now depends on the store in which it is evaluated. Moreover, the result of evaluating a term is not just a value — we must also keep track of the changes that get made to the store.

I.e., the evaluation relation should now map a term and a store to a reduced term and a new store.

$$t \mid \mu \longrightarrow t' \mid \mu'$$

We use the metavariable μ to range over stores.

Evaluation

An assignment $t_1 := t_2$ first evaluates t_1 and t_2 until they become values...

$$\frac{t_1 \mid \mu \longrightarrow t'_1 \mid \mu'}{t_1 := t_2 \mid \mu \longrightarrow t'_1 := t_2 \mid \mu'} \quad (\text{E-ASSIGN1})$$

$$\frac{t_2 \mid \mu \longrightarrow t'_2 \mid \mu'}{v_1 := t_2 \mid \mu \longrightarrow v_1 := t'_2 \mid \mu'} \quad (\text{E-ASSIGN2})$$

... and then returns `unit` and updates the store:

$$l := v_2 \mid \mu \longrightarrow \text{unit} \mid [l \mapsto v_2]\mu \quad (\text{E-ASSIGN})$$

A term of the form $\text{ref } t_1$ first evaluates inside t_1 until it becomes a value...

$$\frac{t_1 \mid \mu \longrightarrow t'_1 \mid \mu'}{\text{ref } t_1 \mid \mu \longrightarrow \text{ref } t'_1 \mid \mu'} \quad (\text{E-REF})$$

... and then chooses (allocates) a fresh location l , augments the store with a binding from l to v_1 , and returns l :

$$\frac{l \notin \text{dom}(\mu)}{\text{ref } v_1 \mid \mu \longrightarrow l \mid (\mu, l \mapsto v_1)} \quad (\text{E-REFV})$$

A term $!t_1$ first evaluates in t_1 until it becomes a value...

$$\frac{t_1 \mid \mu \longrightarrow t'_1 \mid \mu'}{!t_1 \mid \mu \longrightarrow !t'_1 \mid \mu'} \quad (\text{E-DEREF})$$

... and then looks up this value (which must be a location, if the original term was well typed) and returns its contents in the current store:

$$\frac{\mu(l) = v}{!l \mid \mu \longrightarrow v \mid \mu} \quad (\text{E-DEREFLOC})$$

Evaluation rules for function abstraction and application are augmented with stores, but don't do anything with them directly.

$$\frac{t_1 | \mu \longrightarrow t'_1 | \mu'}{t_1 \ t_2 | \mu \longrightarrow t'_1 \ t_2 | \mu'} \quad (\text{E-APP1})$$

$$\frac{t_2 | \mu \longrightarrow t'_2 | \mu'}{v_1 \ t_2 | \mu \longrightarrow v_1 \ t'_2 | \mu'} \quad (\text{E-APP2})$$

$$(\lambda x:T_{11}.t_{12}) \ v_2 | \mu \longrightarrow [x \mapsto v_2]t_{12} | \mu \quad (\text{E-APPABS})$$

Aside: garbage collection

Note that we are not modeling garbage collection — the store just grows without bound.

Aside: pointer arithmetic

We can't do any!

Store Typings

Typing Locations

Q: What is the **type** of a **location**?

Typing Locations

Q: What is the **type** of a **location**?

A: It depends on the store!

E.g., in the store $(l_1 \mapsto \text{unit}, l_2 \mapsto \text{unit})$, the term $!l_2$ has type `Unit`.

But in the store $(l_1 \mapsto \text{unit}, l_2 \mapsto \lambda x:\text{Unit}.x)$, the term $!l_2$ has type `Unit→Unit`.

Typing Locations — first try

Roughly:

$$\frac{\Gamma \vdash \mu(l) : T_1}{\Gamma \vdash l : \text{Ref } T_1}$$

Typing Locations — first try

Roughly:

$$\frac{\Gamma \vdash \mu(l) : T_1}{\Gamma \vdash l : \text{Ref } T_1}$$

More precisely:

$$\frac{\Gamma \mid \mu \vdash \mu(l) : T_1}{\Gamma \mid \mu \vdash l : \text{Ref } T_1}$$

I.e., typing is now a **four**-place relation (between contexts, **stores**, terms, and types).

Problem

However, this rule is not completely satisfactory. For one thing, it can make typing derivations very large!

E.g., if

$$\begin{aligned}(\mu = & \mathbf{l}_1 \mapsto \lambda x:\text{Nat}. \text{999}, \\ & \mathbf{l}_2 \mapsto \lambda x:\text{Nat}. \text{!l}_1 (\text{!l}_1 x), \\ & \mathbf{l}_3 \mapsto \lambda x:\text{Nat}. \text{!l}_2 (\text{!l}_2 x), \\ & \mathbf{l}_4 \mapsto \lambda x:\text{Nat}. \text{!l}_3 (\text{!l}_3 x), \\ & \mathbf{l}_5 \mapsto \lambda x:\text{Nat}. \text{!l}_4 (\text{!l}_4 x)),\end{aligned}$$

then how big is the typing derivation for !l_5 ?

Problem!

But wait... it gets worse. Suppose

$$(\mu = l_1 \mapsto \lambda x:\text{Nat}. !l_2 x, \\ l_2 \mapsto \lambda x:\text{Nat}. !l_1 x),$$

Now how big is the typing derivation for $!l_2$?

Store Typings

Observation: The typing rules we have chosen for references guarantee that a given location in the store is **always** used to hold values of the **same** type.

These intended types can be collected into a **store typing** — a partial function from locations to types.

E.g., for

$$\begin{aligned}\mu = (& l_1 \mapsto \lambda x:\text{Nat}. 999, \\ & l_2 \mapsto \lambda x:\text{Nat}. !l_1 (!l_1 x), \\ & l_3 \mapsto \lambda x:\text{Nat}. !l_2 (!l_2 x), \\ & l_4 \mapsto \lambda x:\text{Nat}. !l_3 (!l_3 x), \\ & l_5 \mapsto \lambda x:\text{Nat}. !l_4 (!l_4 x)),\end{aligned}$$

A reasonable store typing would be

$$\begin{aligned}\Sigma = (& l_1 \mapsto \text{Nat} \rightarrow \text{Nat}, \\ & l_2 \mapsto \text{Nat} \rightarrow \text{Nat}, \\ & l_3 \mapsto \text{Nat} \rightarrow \text{Nat}, \\ & l_4 \mapsto \text{Nat} \rightarrow \text{Nat}, \\ & l_5 \mapsto \text{Nat} \rightarrow \text{Nat})\end{aligned}$$

Now, suppose we are given a store typing Σ describing the store μ in which we intend to evaluate some term t . Then we can use Σ to look up the types of locations in t instead of calculating them from the values in μ .

$$\frac{\Sigma(l) = T_1}{\Gamma \mid \Sigma \vdash l : \text{Ref } T_1} \quad (\text{T-Loc})$$

I.e., typing is now a four-place relation between contexts, **store typings**, terms, and types.

Final typing rules

$$\frac{\Sigma(l) = T_1}{\Gamma \mid \Sigma \vdash l : \text{Ref } T_1} \quad (\text{T-LOC})$$

$$\frac{\Gamma \mid \Sigma \vdash t_1 : T_1}{\Gamma \mid \Sigma \vdash \text{ref } t_1 : \text{Ref } T_1} \quad (\text{T-REF})$$

$$\frac{\Gamma \mid \Sigma \vdash t_1 : \text{Ref } T_{11}}{\Gamma \mid \Sigma \vdash !t_1 : T_{11}} \quad (\text{T-DEREF})$$

$$\frac{\Gamma \mid \Sigma \vdash t_1 : \text{Ref } T_{11} \quad \Gamma \mid \Sigma \vdash t_2 : T_{11}}{\Gamma \mid \Sigma \vdash t_1 := t_2 : \text{Unit}} \quad (\text{T-ASSIGN})$$

Q: Where do these store typings come from?

Q: Where do these store typings come from?

A: When we first typecheck a program, there will be no explicit locations, so we can use an empty store typing.

So, when a new location is created during evaluation,

$$\frac{l \notin \text{dom}(\mu)}{\text{ref } v_1 \mid \mu \longrightarrow l \mid (\mu, l \mapsto v_1)} \quad (\text{E-REFV})$$

we can observe the type of v_1 and extend the “current store typing” appropriately.

Safety

[on board]