

HOARE TYPE THEORY: DEPENDENT TYPES FOR STATE

LECTURE AND NOTES BY JAMES T. PERCONTI
MATERIAL TAKEN FROM WORK BY NANEVSKI ET AL.

In this lecture, we'll discuss Hoare Type Theory, a powerful system that lets us write down and enforce precise specifications for stateful, higher-order programs. It integrates ideas from Hoare Logic into a dependent type system. The system for Hoare Type Theory is fairly complex, so let's begin by reviewing separately some of the ideas from dependent types and Hoare logic.

1. DEPENDENT TYPES

Dependent types let us write down precise specifications for purely functional programs. For example, suppose we want to write a function on natural numbers that always returns a result greater than its argument. In a dependent type system, we can write a type that specifies exactly this:

$$\Pi m : \mathbf{nat}. \{n : \mathbf{nat}. n > m\}$$

$\Pi x : A.B$ is the dependent function type: it describes a function that takes an argument x of type A , and returns a value of type B . x can appear in B : for instance, in a refinement such as $\{y : A.y = x\}$. A refinement type, or subset type, $\{x : A.P\}$ is inhabited by terms of type A that also satisfy the proposition P .

We also have a dependent pair type $\Sigma x : A.B$. The type of the second component of the pair can depend on the value of the first component. That is, the first component has type A , and if it is a value v , then the second component will have type $B[v/x]$. For example, consider this type:

$$\Sigma x : \mathbf{nat}. \{y : \mathbf{nat}. y = x + 5\}.$$

This type is inhabited by values like $\langle 3, 8 \rangle$ and $\langle 1, 6 \rangle$, where the second component is 5 more than the first.

We abbreviate $\Pi x : A.B$ and $\Sigma x : A.B$ using the more usual notations $A \rightarrow B$ and $A \times B$, respectively, when the dependency is not being used ($x \notin \text{fv } B$).

Finally, a dependent type system can treat propositions and types as values and perform computation with them! In the system we will consider today, we will have a simple two-level stratification of types: types that appear in terms are called “monotypes” or “small types” and can't quantify over other types. The “large types” that are not used in computation can quantify over monotypes, giving them type `mono`.

We can do pretty crazy things in this system. For instance, we can have pairs containing a type and a value of that type:

$$\langle \mathbf{nat}, 4 \rangle : \Sigma \alpha : \mathbf{mono}. \alpha$$

Here is the syntax of the dependent type system we'll work with:

A, B, C	$::= K \mid \mathbf{nat} \mid \mathbf{bool} \mid \mathbf{prop} \mid 1 \mid \mathbf{mono} \mid \Pi x : A.B \mid \Sigma x : A.B \mid \{x : A.P\}$	<i>Types</i>
K, L	$::= x \mid K N \mid \mathbf{fst } K \mid \mathbf{snd } K \mid \mathbf{out } K \mid M : A$	<i>Elim Terms</i>
$M, N, P, Q, R, \tau, \sigma$	$::= K \mid n \mid \mathbf{true} \mid \mathbf{false} \mid () \mid \lambda x.M \mid \langle M, N \rangle \mid \mathbf{in } M$	<i>Intro Terms</i>
	$\mid M + N \mid M \times N \mid \mathbf{eq}_{\mathbf{nat}}(M, N)$	
	$\mid \mathbf{xid}_{A,B}(M, N) \mid \top \mid \perp \mid P \wedge Q \mid P \vee Q \mid P \supset Q \mid \neg P \mid \forall x : A.P \mid \exists x : A.P$	
	$\mid \mathbf{nat} \mid \mathbf{bool} \mid \mathbf{prop} \mid () \mid \Pi x : \tau.\sigma \mid \Sigma x : \tau.\sigma \mid \{x : \tau.P\}$	
Δ	$::= \cdot \mid \Delta, x : A, \Delta, P$	<i>Contexts</i>

We've touched on some of the constructs here already, but there is still plenty to talk about. First, note that large types are notated with the metavariables A, B, C , while small types are typically written as τ or σ . In the syntax of terms, we have so many metavariables in order to suggest whether we have a small type, a proposition (P, Q, R), or the more usual λ -calculus sort of term (M, N), but there is actually no distinction between them as far as the syntax is concerned. Any of these three kinds of terms can appear as the others. (Obviously, the type system will distinguish them.)

This brings us to the next observation: terms are explicitly split into introduction forms and elimination forms. This allows *bidirectional typechecking*, a technique that makes it explicit when we can and cannot synthesize types. Specifically, we can always generate the type of elim terms, so their typing judgment reflects this:

$$\Delta \vdash K \Rightarrow A$$

Intro terms, on the other hand, don't contain enough information in them to construct the correct type, so we only check them against a given type, as indicated by the shape of their typing judgment:

$$\Delta \vdash M \Leftarrow A$$

To make sure we're all comfortable with this and some of the other concepts of dependent types we've already touched on, let's write down some of the typing rules. Here are the rules for introduction and elimination of dependent functions:

$$\frac{\Delta, x:A \vdash M \Leftarrow B}{\Delta \vdash \lambda x.M \Leftarrow \Pi x:A.B} \qquad \frac{\Delta \vdash K \Rightarrow \Pi x:A.B \quad \Delta \vdash M \Leftarrow A}{\Delta \vdash K M \Rightarrow B[M/x]}$$

Note that the premise of the $\lambda x.M$ rule puts x into the environment for both the function body M and its type B . The application rule handles the dependency by directly substituting the argument supplied into the type of the expression.

Here are the rules for introduction and elimination of dependent pairs:

$$\frac{\Delta \vdash M \Leftarrow A \quad \Delta \vdash N \Leftarrow B[M/x]}{\Delta \vdash \langle M, N \rangle \Leftarrow \Sigma x:A.B} \qquad \frac{\Delta \vdash K \Rightarrow \Sigma x:A.B}{\Delta \vdash \text{fst } K \Rightarrow A} \qquad \frac{\Delta \vdash K \Rightarrow \Sigma x:A.B}{\Delta \vdash \text{snd } K \Rightarrow B[\text{fst } K/x]}$$

The dependency appears in both the premise of the intro rule and the result type of the `snd` rule.

Next, let's look at the rules for refinements. We use the intro form `in` M and the elim form `out` M as syntactic markers for refinements.

$$\frac{\Delta \vdash M \Leftarrow A \quad \Delta \vDash P[M/x]}{\Delta \vdash \text{in } M \Leftarrow \{x:A.P\}} \qquad \frac{\Delta \vdash K \Rightarrow \{x:A.P\}}{\Delta \vdash \text{out } K \Rightarrow A} \qquad \frac{\Delta \vdash K \Rightarrow \{x:A.P\}}{\Delta \vDash P[\text{out } K/x]}$$

The introduction rule here requires us to prove that the proposition P actually applies to the term M . Whenever we have a proof obligation, or *verification condition*, like this, we can pass it to an arbitrary theorem prover (or even ignore it altogether!). If we don't prove all the verification conditions, the type system is decidable, and still gives us a type safety guarantee similar to what a simple type system would give us. To ensure the precise specifications of our dependent type system are met, we do need to prove them (which is undecidable in general).

This separation between the decidable and undecidable components of our type system is also important in other places. For instance, we need to be able to reason about equality of terms. Our language of propositions includes a heterogeneous equality predicate $\text{xid}_{A,B}(M, N)$, which is true when A and B are the same type and M and N are the same term and have that type. This is our propositional, undecidable notion of equality.

We also need a decidable notion of equality (we directly need it to compare types, as we'll see momentarily, but terms can appear in types, so we need to compare terms as well). To get this, we reduce terms to normal form and then compare them syntactically. This works because our language is terminating, but setting it up correctly requires a lot of machinery, threaded through the entire

type system, that we will be leaving out. To give an idea of what is involved, here is the real type rule for applications:

$$\frac{\Delta \vdash K \Rightarrow \Pi x: A.B [N'] \quad \Delta \vdash M \Leftarrow A [M']}{\Delta \vdash K M \Rightarrow B[M'/x]_A^a [\text{apply}_A(N', M')]}$$

The bracketed term at the end of each judgment is the normalized form of the term that was input. The `apply` metafunction normalizes the application itself. The substitution performed by `apply` and the substitution in the synthesized type $B[M'/x]_A^a$ are *hereditary substitutions*. A hereditary substitution must renormalize any redexes generated by the substitution — which could in turn require more hereditary substitutions. It's a lot of work just to write down the definition of this operation, let alone to prove that it terminates! We won't worry about the details of normalization from here on, but it's important to note that it is going on behind the scenes.

The decidable equality we've been discussing, known as *definitional equality*, is used when we want to change the direction of typechecking. When an elim term appears in an intro position, we need to compare the type it synthesizes to the type we are checking. On the other hand, since intro terms can only appear in elim positions with explicit type annotations, we don't have to worry about equality in that rule:

$$\frac{\Delta \vdash K \Rightarrow A \quad A = B}{\Delta \vdash K \Leftarrow B} \quad \frac{\Delta \vdash A \text{ type} \quad \Delta \vdash M \Leftarrow A}{\Delta \vdash M: A \Rightarrow A}$$

OK, that's enough about the details of our dependent type system. By now at least one reason that adding effects to this language is problematic should be clear: effects destroy our ability to decide equality by normalizing terms, most obviously by giving us terms that diverge. To solve this, we'll encapsulate our effects in a monad based on another model for program specifications. That model is Hoare logic, a reasoning system designed for imperative programs.

2. HOARE LOGIC

Our goal is to examine the core idea of Hoare logic and integrate it into our dependent type system, so we won't go into much detail about Hoare logic as an independent entity. Essentially, we specify the behavior of a computation, or effectful program fragment, by writing a *Hoare triple*

$$\{P\} E \{Q\}.$$

Here E is the computation for which we are giving the specification. P is a precondition, an assertion giving requirements for the state of the program before E is run. Q is a postcondition, an assertion that describes the state of the program after E completes, assuming P was satisfied before E began to be computed.

For a simple example, let's write a program fragment that increments the value at some location ℓ (assumed to be in the heap already):

$$u = !\ell; \ell := u + 1$$

We can give it the following specification:

$$n: \text{nat}. \{\ell \mapsto n\} u = !\ell; \ell := u + 1 \{\ell \mapsto n + 1\}$$

The precondition says that ℓ maps to some natural number n in the heap, and in the postcondition ℓ will map to $n + 1$. The variable n that appears in both the precondition and postcondition, but not in the computation itself, is called a *logic variable* or *ghost variable*. Individual Hoare triples are often easier to read using logic variables, but it is easier to work with the metatheory without them. Fortunately, we will see in a moment that the binding structure implied by logic variables can be encoded in a system where the precondition and postcondition are self-contained.

We're making our way toward integrating Hoare triples into a type system, so let's separate the term from its specification and replace it with the type of value returned by the computation. Thus

we get a *Hoare type*, which has the form

$$\{P\} x : A \{Q\}.$$

To talk about the state of the program before the computation is run, the precondition needs to be given the initial heap. Since it must relate the state before the computation to the state after, the postcondition needs the initial and final heaps. Thus P and Q should be propositions of type $\text{heap} \rightarrow \text{prop}$ and $\text{heap} \rightarrow \text{heap} \rightarrow \text{prop}$ respectively. The postcondition can also depend on the return value x of the computation.

The Hoare type for our example program is

$$u = !\ell; \ell := u + 1 : n : \mathbf{nat}. \{ \ell \mapsto_{\mathbf{nat}} n \} r : 1 \{ \ell \mapsto_{\mathbf{nat}} n + 1 \}.$$

Now let's define our encoding of ghost variables in Hoare types:

$$x : A. \{P\} y : B \{Q\} \stackrel{\text{def}}{=} \{ \lambda i. \exists x : A. P \ i \} y : B \{ \lambda i. \lambda m. \forall x : A. (P \ i \supset Q \ m) \}$$

In the precondition, x is simply quantified as an existential, while in the postcondition, any value of type A that satisfied P under the initial heap should satisfy Q under the final heap. Thus the expanded type for our example is

$$\{ \lambda i. \exists n : \mathbf{nat}. (\ell \mapsto_{\mathbf{nat}} n) \ i \} r : 1 \{ \lambda i. \lambda m. \forall n : \mathbf{nat}. (\ell \mapsto_{\mathbf{nat}} n) \ i \supset (\ell \mapsto_{\mathbf{nat}} n + 1) \ m \}.$$

Note that the final precondition and postcondition have the correct types, assuming that the operation $x \mapsto_{\tau} v$ (which we will define later) has type $\text{heap} \rightarrow \text{prop}$.

3. HOARE TYPE THEORY

We will now add effectful computations to our dependently typed language. A computation E is encapsulated by the **do** construct, which suspends its evaluation until another computation forces it to run. Computations can dereference and assign to memory locations, run other computations, branch, return a pure value, and run recursive (possibly divergent) functions. To classify these new terms, we also add the Hoare type, as developed in the previous section. A computation of type $\{P\} x : A \{Q\}$ can be run in a heap i such that $P \ i$ holds, and will either diverge or reach a heap m and return a value $v : A$ such that $Q[v/x] \ i \ m$ holds.

Here is the full set of additions to our language. Note that the constructs **return** M and $x = c$; E correspond to monadic unit and bind operations, respectively.

A, B, C	$::= \dots \mid \{P\} x : A \{Q\}$	<i>Types</i>
$M, N, P, Q, R, \tau, \sigma$	$::= \dots \mid \mathbf{do} \ E \mid \{P\} \ \sigma : \tau \ \{Q\}$	<i>Intro Terms</i>
c	$::= !_{\tau} \ M \mid M :=_{\tau} \ N \mid \mathbf{if}_A \ M \ \mathbf{then} \ E_1 \ \mathbf{else} \ E_2$	<i>Commands</i>
	$\mid \mathbf{fix} \ x(y : A) : B = \mathbf{do} \ E \ \mathbf{in} \ \mathbf{eval} \ x \ M$	
E, F	$::= \mathbf{return} \ M \mid \mathbf{run} \ x = K \ \mathbf{in} \ E \mid x = c; \ E$	<i>Computations</i>

The type of our example program is already correctly formed, but for completeness, note that the program itself should be adjusted slightly to match this syntax:

$$\mathbf{do}(u = !_{\mathbf{nat}} \ell; z = (\ell :=_{\mathbf{nat}} u + 1); \mathbf{return} \ z).$$

With these additions to the language, we now have a complete system for Hoare type theory. It combines the strengths of dependent types and Hoare logic. Everything we need to write many useful programs and specifications can be encoded from the small set of primitives we have. For instance, let's develop an encoding for the type **heap** itself.

We model a heap as a set of locations paired with the values they map to. We can encode a set of terms of type A as an $A \rightarrow \text{prop}$ that returns a true proposition when its argument is in the set. To know what type of value we have at a location, we will just add the type itself as a component of the tuple. Finally, we use a refinement to require that the heap is finite and only maps each location at most once. Thus the full definition of the type of a heap is

$$\text{heap} = \{h : (\mathbf{nat} \times \Sigma \alpha : \text{mono}.\alpha) \rightarrow \text{prop}.\text{Finite}(h) \wedge \text{Functional}(h)\},$$

Properties like `Finite` and `Functional` are straightforward to define, though they in turn require some more definitions. The definitions of \leq , \in , `Injective`, and `Surjective` are left as exercises to the reader (or you can go look in the paper).

$$\begin{aligned} \text{Finite}_A &: (A \rightarrow \text{prop}) \rightarrow \text{prop} \\ &= \lambda x. \exists n : \text{nat}. \exists f : \{y : \text{nat}. y \leq n\} \rightarrow \{z : A. x z\}. \text{Injective}(f) \wedge \text{Surjective}(f) \\ \text{Functional}_{A,B} &: (A \times B \rightarrow \text{prop}) \rightarrow \text{prop} \\ &= \lambda R. \forall x : A. \forall y_1, y_2 : B. (\langle x, y_1 \rangle \in R \wedge \langle x, y_2 \rangle \in R) \supset \text{xid}_{B,B}(y_1, y_2) \end{aligned}$$

Now that we have our representation for heaps, we'll want some operations in order to work with them. Here are just a couple:

$$\begin{aligned} \text{empty} &: \text{heap} \\ &= \text{in}(\lambda h. \perp) \\ \text{upd} &: \Pi \alpha : \text{mono}. \text{heap} \rightarrow \text{nat} \rightarrow \alpha \rightarrow \text{heap} \\ &= \lambda \alpha. \lambda h. \lambda n. \lambda x. \text{in}(\lambda t. (n = \text{fst } t \supset \text{snd } t = \langle \alpha, x \rangle) \wedge (n \neq \text{fst } t \supset (\text{out } h) t)) \\ \text{seleq} &: \Pi \alpha : \text{mono}. \text{heap} \rightarrow \text{nat} \rightarrow \alpha \rightarrow \text{prop} \\ &= \lambda \alpha. \lambda h. \lambda n. \lambda x. (\text{out } h) \langle n, \langle \alpha, x \rangle \rangle \\ \mapsto &: \Pi \alpha : \text{mono}. \text{nat} \rightarrow \alpha \rightarrow \text{heap} \rightarrow \text{prop} \\ &= \lambda \alpha. \lambda n. \lambda x. \lambda h. \text{xid}_{\text{heap}, \text{heap}}(h, \text{upd } \alpha \text{ empty } n x) \end{aligned}$$

`empty` is of course the empty heap, and `upd` is used to extend or update a heap with a new value. `seleq` takes a heap, a location, a value, and that value's type, and returns the proposition that that value is in the heap at that location. \mapsto is similar in that it takes a heap entry (location, value, and type) and a heap and returns a proposition about that entry's presence in the heap, but it matches the given heap to one containing only the supplied entry.

Does the definition of \mapsto match our intuition from its use in the example? What if we want to run that program in a heap that contains entries besides ℓ ? Actually, we don't want to have to worry about other things in the heap, since our program doesn't touch them or care about them. We don't have to worry about them because the precondition and postcondition of the Hoare type are implicitly *small-footprint* specifications. This means they only need to describe the fragment of the heap that is actually affected by the computation being described. This fragment is assumed to be combined with any *frame*, or "rest of the heap," that is disjoint from it. Thus, a precondition $\{x \mapsto v\}$ accepts heaps that also contain locations other than x , but only as long as none of them is aliased to x , since the frame must be disjoint from the part of the heap described by the specification.

4. LOCAL STATE

We will now develop another example to show how Hoare type theory can hide the internal state invariants of a program. Let's start with a fairly simple computation that builds a counter:

$$\text{make_counter} = \text{do}(\text{run } x = (\text{alloc nat } 0) \text{ in do}(z = !_{\text{nat}} x; x :=_{\text{nat}} z + 1; \text{return } z + 1)).$$

Here, `alloc` is not something we've defined yet, but we can encode an allocator without too much trouble. Assume we have an allocator that relies on some invariant $I : \text{heap} \rightarrow \text{prop}$, and that the allocation function has the following type:

$$\text{alloc} : \Pi \alpha : \text{mono}. \Pi x : \alpha. \{I\} r : \text{nat} \{ \lambda i. (I * r \mapsto_{\alpha} x) \}$$

In the postcondition, the $*$ operation is a disjoint union of heap fragments.

So: `make_counter` allocates a new reference and returns a computation that will increment that reference and return its new value whenever it is called. Let's try to write down the type of this

program:

$$\begin{aligned} \text{make_counter} : \{I\} \\ y : (v : \text{nat} . \{x \mapsto_{\text{nat}} v\} r : \text{nat} \{ \lambda m . (x \mapsto_{\text{nat}} v + 1) m \wedge r = v + 1 \}) \\ \{ \lambda i . (I * x \mapsto_{\text{nat}} 0) \} \end{aligned}$$

This type isn't quite right, because we've left x unbound. But ignoring that for the moment, note that the outer computation just calls `alloc` and returns a value, so its precondition and postcondition essentially match `alloc`'s, while the value returned is a computation that increments and returns the value of x .

How should x be bound in this type? Ideally, we don't want it to show up at all, since we only have access to the allocated location inside the counter. As long as they are correct, different implementations of a counter should be indistinguishable. For example, what if we replace `make_counter` with

$$\begin{aligned} \text{do} (\text{run } x = \text{alloc nat } 0 \text{ in} \\ \text{run } y = \text{alloc nat } 0 \text{ in} \\ \text{do} (z =!_{\text{nat}} x; w =!_{\text{nat}} y; x :=_{\text{nat}} z + 1; y :=_{\text{nat}} w + 1; \text{return}(z + w)/2 + 1)) \end{aligned}$$

This version of a counter is silly, but clearly equivalent except in the way it locally uses state. We should be able to give different implementations the same type, even if they use local state differently.

To make our local state invariants abstract, we can modify our counter implementations to return the invariant itself. Note that there is no way to inspect a proposition at runtime, so passing the invariant cannot change the computational behavior of the program, but only allows the type to abstract over it.

$$\begin{aligned} \text{counter}' &= \text{do} (\text{run } x = (\text{alloc nat } 0) \text{ in} \\ &\quad \langle \langle \lambda v . x \mapsto_{\text{nat}} v \rangle, \text{do} (z =!_{\text{nat}} x; x :=_{\text{nat}} z + 1; \text{return } z + 1) \rangle \rangle) \\ \text{counter2}' &= \text{do} (\text{run } x = \text{alloc nat } 0 \text{ in} \\ &\quad \text{run } y = \text{alloc nat } 0 \text{ in} \\ &\quad \langle \lambda v . (x \mapsto_{\text{nat}} v) * (y \mapsto_{\text{nat}} v), \\ &\quad \text{do} (z =!_{\text{nat}} x; w =!_{\text{nat}} y; x :=_{\text{nat}} z + 1; y :=_{\text{nat}} w + 1; \text{return}(z + w)/2 + 1) \rangle \rangle) \end{aligned}$$

With this change, x is now well-scoped, as the invariant will get instantiated with the correct location once the counter is initialized. We can write the shared type of the two counters without referring to x (or y), as follows:

$$\begin{aligned} \text{counter}', \text{counter2}' : \{I\} \\ t : \Sigma \text{inv} : \text{nat} \rightarrow \text{heap} \rightarrow \text{prop}. \\ v : \text{nat} . \{ \text{inv } v \} r : \text{nat} \{ \lambda h . \text{inv } (v + 1) h \wedge r = v + 1 \} \\ \{ \lambda i . (I * \text{fst } t) \} \end{aligned}$$

Though our example of a counter is contrived, we can apply the same idea to get abstract handling of local state in many other situations.

5. TYPING RULES FOR HTT

We've gotten a taste of the kinds of specifications Hoare type theory lets us encode, so now let's see how the system manages to do all this. We add a pair of judgments for computations to the

dependent type system we had earlier:

$$\begin{aligned} \Delta; P \vdash E \Leftarrow x: A.Q \\ \Delta; P \vdash E \Rightarrow x: A.Q \end{aligned}$$

The first judgment checks if a computation E , under environment Δ and precondition P , will return a value of type A and leave the heap in a state that satisfies the postcondition Q . The second judgment still takes Δ , P , E , and A as inputs, but it synthesizes the postcondition. In particular, it synthesizes the strongest postcondition Q that could result from running E in a heap that satisfies P . The checking judgment has one inference rule, so let's look at it:

$$\frac{\Delta; P \vdash E \Rightarrow x: A.S \quad \Delta, x: A, i: \text{heap}, m: \text{heap}, (S \ i \ m) \models (Q \ i \ m)}{\Delta; P \vdash E \Leftarrow x: A.Q}$$

To check a computation against a given postcondition Q , we synthesize its strongest postcondition S using the \Rightarrow judgment, and then we have a proof obligation that S implies Q under an appropriate environment.

To see how the other judgment works, let's start with the simplest case: the computation that just returns a pure value.

$$\frac{\Delta \vdash M \Leftarrow A}{\Delta; P \vdash \text{return } M \Rightarrow x: A.(\lambda i. \lambda m. (P \ i \ m) \wedge \text{xid}_{A,A}(x, M))}$$

The postcondition we generate is the precondition, but with an additional requirement that specifies the value returned. The only premise needed is that the value we return typechecks.

One odd thing you might notice in the **return** rule is that P is given two arguments. Actually, preconditions in these judgments, unlike those in the Hoare type for a complete computation must always take two heaps, since as we will see in the remaining rules, the postcondition of one piece of a computation can become the precondition of the next piece. Think of these arguments as the heap at the beginning of the whole program and the heap right before the current piece of the computation. At the beginning of a computation, these will be the same. Let's look at the rule that gets us started:

$$\frac{\Delta; \lambda i. \lambda m. i = m \wedge (R * \lambda m'. \top) \ m \vdash E \Leftarrow x: A.R \gg Q}{\Delta \vdash \text{do } E \Leftarrow \{R\} \ x: A \ \{Q\}}$$

To typecheck a complete, **do**-encapsulated computation against its Hoare type, we build a two-argument precondition out of the given precondition R , by requiring the two arguments to be the same and then applying R to one of them. We also frame the fragment of the heap that R applies to with an arbitrary heap fragment (one that satisfies $\lambda m'. \top$). This is the interface between the small-footprint specifications for Hoare types we discussed earlier and the specifications in the judgments, which must describe the whole heap. The postcondition we check for E in the premise uses an operator \gg , whose formal definition we will omit here, but which essentially builds a large-footprint postcondition out of Q by requiring that if the initial heap i consists of a fragment described by R and a frame, then the final heap must leave the latter unchanged, while the former is described by $(Q \ i)$.

We'll omit the typing rules for conditionals and recursive functions, but let's look at the rules for dereferencing, assignment, and running a computation:

Here is the rule for dereference:

$$\frac{\Delta \vdash \tau \Leftarrow \text{mono} \quad \Delta \vdash M \Leftarrow \text{nat} \quad \Delta, i: \text{heap}, m: \text{heap}, (P \ i \ m) \models (M \hookrightarrow_{\tau} -) \ m \quad \Delta, x: \tau; \lambda i. \lambda m. (P \ i \ m) \wedge (M \hookrightarrow_{\tau} x) \ m \vdash E \Rightarrow y: B.Q}{\Delta; P \vdash x = !_{\tau} M; E \Rightarrow y: B. \lambda i. \lambda m. \exists x: \tau. (Q \ i \ m)}$$

A computation beginning with a dereference must be dereferencing a well-typed expression of type **nat** (a location), and then we check the rest of the computation in a context that adds x , the variable we bound, under a precondition that extends the earlier precondition by requiring that the location

we dereferenced maps to x in the heap (the \hookrightarrow is just a shorthand for `seleq` with the arguments rearranged; unlike \mapsto , it doesn't specify a singleton heap, which we need here because we can't do a disjoint union: P will already have something to say about location M). We also add a verification condition that in a heap satisfying P , M must point to something of the right type τ .

Next, we have the rule for assignment:

$$\frac{\Delta \vdash \tau \Leftarrow \text{mono} \quad \Delta \vdash M \Leftarrow \text{nat} \quad \Delta \vdash N \Leftarrow \tau \quad \Delta, i : \text{heap}, m : \text{heap}, (P \ i \ m) \models (M \hookrightarrow -) \ m \quad \Delta; P \circ ((M \mapsto -) \multimap (M \mapsto_{\tau} N)) \vdash E \Rightarrow y : B.Q}{\Delta; P \vdash M :=_{\tau} N; E \Rightarrow y : B.Q}$$

First, we typecheck the pure terms that make up the location to assign to and the value to assign. As in the dereference rule, we have a proof obligation that the precondition implies that the location M is mapped in the heap, and a premise that the rest of the computation typechecks under an appropriate precondition. In this case, the condition says to take the part of the original precondition P that describes M (that is, the part that looks like $M \mapsto -$) and replace it with $M \mapsto_{\tau} N$. (Formal definitions of \circ , \multimap , and \gg are in the appendix, but the details are not worth giving here.)

Finally, here is the rule for running a computation:

$$\frac{\Delta \vdash K \Rightarrow \{R\} x : A \{S\} \quad \Delta, i : \text{heap}, m : \text{heap}, (P \ i \ m) \models (R * \lambda m'. \top) \ m \quad \Delta, x : A; P \circ (R \gg S) \vdash E \Rightarrow y : B.Q}{\Delta; P \vdash \text{run } x = K \text{ in } E \Rightarrow y : B.\lambda i.\lambda m.\exists x : A.(Q \ i \ m)}$$

First, we must synthesize the type of the computation we want to run. Recall that an intro term such as `do E` appearing in an elim position must be annotated with a type, so we will often actually be checking K against a Hoare type that was given. Next, we check the rest of the current computation E , under a precondition that composes P with the large-footprint version of the postcondition of K , that is, $R \gg S$ (this is the same thing we did in the `do` rule).

None of the rules in the computation judgments required us to perform decidable checks for equality of specifications. Thus we don't need to normalize computations, and we can soundly handle their effects in our dependent type system.

To recap, we've seen how Hoare type theory merges Hoare logic-style specifications for effectful computations into a dependent type system, and taken a bit of a look at what kinds of precise specifications the system lets us encode. We've also had a look at the inner workings of the type system.

6. BIBLIOGRAPHY

Polymorphism and Separation in Hoare Type Theory; Aleksandar Nanevski, Greg Morrisett, Lars Birkedal; ICFP 2006

This earlier version of HTT is similar to the version we developed here in many ways, but it is only first-order; that is, we can't abstract over propositions in the same way as we saw, and thus we can't hide local state invariants.

Abstract Predicates and Mutable ADTs in Hoare Type Theory; Aleksandar Nanevski, Amal Ahmed, Greg Morrisett, Lars Birkedal; ESOP 2007.

This paper gives the version of Hoare Type Theory that was discussed in these notes. If you are looking for more detail, the technical report for this paper is a good place to go. It explicitly lays out pretty much everything going on in the system, including all the details needed for normalization that we alluded to, and all the other definitions we elided.

Ynot: Dependent Types for Imperative Programs; Aleksandar Nanevski, Greg Morrisett, Avi Shin-nar, Paul Govereau, Lars Birkedal; ICFP 2008

In this paper, the authors show how HTT can be implemented as an axiomatic extension to Coq, and discuss a bunch of the features of their implementation.

APPENDIX: COMPLETE TYPING RULES

These rules are still simplified: they don't perform normalization. For the *really* complete typing rules, see the tech report version of Nanevski et al. 2007.

Checking contexts.

 $\boxed{\vdash \Delta \text{ ctx}}$

$$\frac{}{\vdash \cdot \text{ ctx}} \quad \frac{\vdash \Delta \text{ ctx} \quad \Delta \vdash A \text{ type}}{\vdash \Delta, x:A \text{ ctx}} \quad \frac{\vdash \Delta \text{ ctx} \quad \Delta \vdash P \Leftarrow \text{prop}}{\vdash \Delta, P \text{ ctx}}$$

Checking types.

 $\boxed{\Delta \vdash A \text{ type}}$

$$\frac{\Delta \vdash K \Rightarrow \text{mono}}{\Delta \vdash K \text{ type}} \quad \frac{}{\Delta \vdash \text{nat type}} \quad \frac{}{\Delta \vdash \text{bool type}} \quad \frac{}{\Delta \vdash \text{prop type}} \quad \frac{}{\Delta \vdash \text{mono type}}$$

$$\frac{}{\Delta \vdash 1 \text{ type}} \quad \frac{\Delta \vdash A \text{ type} \quad \Delta, x:A \vdash B \text{ type}}{\Delta \vdash \Pi x:A. B \text{ type}} \quad \frac{\Delta \vdash A \text{ type} \quad \Delta, x:A \vdash B \text{ type}}{\Delta \vdash \Sigma x:A. B \text{ type}}$$

$$\frac{\Delta \vdash A \text{ type} \quad \Delta, x:A \vdash P \Leftarrow \text{prop}}{\Delta \vdash \{x:A. P\} \text{ type}}$$

$$\frac{\Delta \vdash P \Leftarrow \text{heap} \rightarrow \text{prop} \quad \Delta \vdash A \text{ type} \quad \Delta, x:A \vdash Q \Leftarrow \text{heap} \rightarrow \text{heap} \rightarrow \text{prop}}{\Delta \vdash \{P\} x:A \{Q\} \text{ type}}$$

Checking intro terms.

 $\boxed{\Delta \vdash M \Leftarrow A}$

$$\frac{\Delta \vdash K \Rightarrow A \quad A = B}{\Delta \vdash K \Leftarrow B} \quad \frac{}{\Delta \vdash \text{true} \Leftarrow \text{bool}} \quad \frac{}{\Delta \vdash \text{false} \Leftarrow \text{bool}} \quad \frac{}{\Delta \vdash n \Leftarrow \text{nat}}$$

$$\frac{}{\Delta \vdash () \Leftarrow 1} \quad \frac{\Delta \vdash M \Leftarrow \text{nat} \quad \Delta \vdash N \Leftarrow \text{nat}}{\Delta \vdash M + N \Leftarrow \text{nat}} \quad \frac{\Delta \vdash M \Leftarrow \text{nat} \quad \Delta \vdash N \Leftarrow \text{nat}}{\Delta \vdash M \times N \Leftarrow \text{nat}}$$

$$\frac{\Delta \vdash M \Leftarrow \text{nat} \quad \Delta \vdash N \Leftarrow \text{nat}}{\Delta \vdash \text{eq}_{\text{nat}}(M, N) \Leftarrow \text{nat}} \quad \frac{\Delta, x:A \vdash M \Leftarrow B}{\Delta \vdash \lambda x.M \Leftarrow \Pi x:A. B} \quad \frac{\Delta \vdash M \Leftarrow A \quad \Delta \vdash N \Leftarrow B[M/x]}{\Delta \vdash \langle M, N \rangle \Leftarrow \Sigma x:A. B}$$

$$\frac{\Delta \vdash M \Leftarrow A \quad \Delta \vdash P[M/x]}{\Delta \vdash \text{in } M \Leftarrow \{x:A. P\}} \quad \frac{\Delta; \lambda i. \lambda m. i = m \wedge (R * \lambda m'. \top) \quad m \vdash E \Leftarrow x:A. R \gg Q}{\Delta \vdash \text{do } E \Leftarrow \{R\} x:A \{Q\}}$$

$$\frac{\Delta \vdash A \text{ type} \quad \Delta \vdash B \text{ type} \quad \Delta \vdash M \Leftarrow A \quad \Delta \vdash M \Leftarrow B}{\Delta \vdash \text{xid}_{A,B}(M, N) \Leftarrow \text{prop}} \quad \frac{}{\Delta \vdash \top \Leftarrow \text{prop}}$$

$$\frac{}{\Delta \vdash \perp \Leftarrow \text{prop}} \quad \frac{\Delta \vdash M \Leftarrow \text{prop} \quad \Delta \vdash N \Leftarrow \text{prop}}{\Delta \vdash M \wedge N \Leftarrow \text{prop}} \quad \frac{\Delta \vdash M \Leftarrow \text{prop} \quad \Delta \vdash N \Leftarrow \text{prop}}{\Delta \vdash M \vee N \Leftarrow \text{prop}}$$

$$\frac{\Delta \vdash M \Leftarrow \text{prop} \quad \Delta \vdash N \Leftarrow \text{prop}}{\Delta \vdash M \supset N \Leftarrow \text{prop}} \quad \frac{\Delta \vdash M \Leftarrow \text{prop}}{\Delta \vdash \neg M \Leftarrow \text{prop}}$$

$$\frac{\Delta \vdash A \text{ type} \quad \Delta, x:A \vdash M \Leftarrow \text{prop}}{\Delta \vdash \forall x:A. M \Leftarrow \text{prop}} \quad \frac{\Delta \vdash A \text{ type} \quad \Delta, x:A \vdash M \Leftarrow \text{prop}}{\Delta \vdash \exists x:A. M \Leftarrow \text{prop}}$$

$$\begin{array}{c}
\overline{\Delta \vdash \text{nat} \Leftarrow \text{mono}} \quad \overline{\Delta \vdash \text{bool} \Leftarrow \text{mono}} \quad \overline{\Delta \vdash \text{prop} \Leftarrow \text{mono}} \quad \overline{\Delta \vdash 1 \Leftarrow \text{mono}} \\
\frac{\Delta \vdash \tau \Leftarrow \text{mono} \quad \Delta, x: \tau \vdash \sigma \Leftarrow \text{mono}}{\Delta \vdash \Pi x: \tau. \sigma \Leftarrow \text{mono}} \quad \frac{\Delta \vdash \tau \Leftarrow \text{mono} \quad \Delta, x: \tau \vdash \sigma \Leftarrow \text{mono}}{\Delta \vdash \Sigma x: \tau. \sigma \Leftarrow \text{mono}} \\
\frac{\Delta \vdash \tau \Leftarrow \text{mono} \quad \Delta, x: \tau \vdash P \Leftarrow \text{prop}}{\Delta \vdash \{x: \tau. P\} \Leftarrow \text{mono}} \\
\frac{\Delta \vdash P \Leftarrow \text{heap} \rightarrow \text{prop} \quad \Delta \vdash \tau \Leftarrow \text{mono} \quad \Delta, x: \tau \vdash Q \Leftarrow \text{heap} \rightarrow \text{heap} \rightarrow \text{prop}}{\Delta \vdash \{P\} x: \tau \{Q\} \Leftarrow \text{mono}}
\end{array}$$

Checking elim terms.

$$\boxed{\Delta \vdash K \Rightarrow A}$$

$$\begin{array}{c}
\overline{\Delta, x: A, \Delta' \vdash x \Rightarrow A} \quad \frac{\Delta \vdash K \Rightarrow \Pi x: A. B \quad \Delta \vdash M \Leftarrow A}{\Delta \vdash K M \Rightarrow B[M/x]} \quad \frac{\Delta \vdash K \Rightarrow \Sigma x: A. B}{\Delta \vdash \text{fst } K \Rightarrow A} \\
\frac{\Delta \vdash K \Rightarrow \Sigma x: A. B}{\Delta \vdash \text{snd } K \Rightarrow B[\text{fst } K/x]} \quad \frac{\Delta \vdash K \Rightarrow \{x: A. P\}}{\Delta \vdash \text{out } K \Rightarrow A} \quad \frac{\Delta \vdash A \text{ type} \quad \Delta \vdash M \Leftarrow A}{\Delta \vdash M: A \Rightarrow A}
\end{array}$$

Inference rules.

$$\boxed{\Delta \vDash P}$$

$$\begin{array}{c}
\overline{\Delta, P, \Delta' \vDash P} \quad \frac{\Delta, P \vDash Q}{\Delta \vDash P \supset Q} \quad \frac{\Delta \vDash P \supset Q \quad \Delta \vDash P}{\Delta \vDash Q} \quad \frac{\Delta, x: A \vDash P}{\Delta \vDash \forall x: A. P} \\
\frac{\Delta \vDash \forall x: A. P \quad \Delta \vdash M \Leftarrow A}{\Delta \vDash P[M/x]} \quad \frac{\Delta \vdash K \Rightarrow \{x: A. P\}}{\Delta \vDash P[K/x]}
\end{array}$$

There are also a lot of axioms giving the properties of primitives.

Checking computations.

$$\boxed{\Delta; P \vdash E \Rightarrow x: A. Q}$$

$$\boxed{\Delta; P \vdash E \Leftarrow x: A. Q}$$

$$\begin{array}{l}
P \circ Q = \lambda i. \lambda m. \exists h: \text{heap}. (P \ i \ h) \wedge (Q \ h \ m) \\
R_1 \multimap R_2 = \lambda i. \lambda m. \forall h: \text{heap}. (R_1 * \lambda h'. h' = h) \ i \supset (R_2 * \lambda h'. h' = h) \ m \\
R \gg Q = \lambda i. \lambda m. \forall h: \text{heap}. (\lambda h'. R(h') \wedge h = h') \multimap Q \ h
\end{array}$$

$$\begin{array}{c}
\frac{\Delta; P \vdash E \Rightarrow x: A.S \quad \Delta, x: A, i: \text{heap}, m: \text{heap}, (S \ i \ m) \vDash (Q \ i \ m)}{\Delta; P \vdash E \Leftarrow x: A.Q} \\
\\
\frac{\Delta \vdash M \Leftarrow A}{\Delta; P \vdash \text{return } M \Rightarrow x: A.(\lambda i. \lambda m. (P \ i \ m) \wedge \text{xid}_{A,A}(x, M))} \\
\\
\frac{\Delta \vdash K \Rightarrow \{R\} \ x: A \ \{S\} \quad \Delta, i: \text{heap}, m: \text{heap}, (P \ i \ m) \vDash (R * \lambda m'. \top) \ m \quad \Delta, x: A; P \circ (R \gg S) \vdash E \Rightarrow y: B.Q}{\Delta; P \vdash \text{run } x = K \text{ in } E \Rightarrow y: B. \lambda i. \lambda m. \exists x: A. (Q \ i \ m)} \\
\\
\frac{\Delta \vdash \tau \Leftarrow \text{mono} \quad \Delta \vdash M \Leftarrow \text{nat} \quad \Delta, i: \text{heap}, m: \text{heap}, (P \ i \ m) \vDash (M \hookrightarrow_{\tau} -) \ m \quad \Delta, x: \tau; \lambda i. \lambda m. (P \ i \ m) \wedge (M \hookrightarrow_{\tau} x) \ m \vdash E \Rightarrow y: B.Q}{\Delta; P \vdash x = !_{\tau} M; E \Rightarrow y: B. \lambda i. \lambda m. \exists x: \tau. (Q \ i \ m)} \\
\\
\frac{\Delta \vdash \tau \Leftarrow \text{mono} \quad \Delta \vdash M \Leftarrow \text{nat} \quad \Delta \vdash N \Leftarrow \tau \quad \Delta, i: \text{heap}, m: \text{heap}, (P \ i \ m) \vDash (M \hookrightarrow -) \ m \quad \Delta; P \circ ((M \mapsto -) \multimap (M \mapsto_{\tau} N)) \vdash E \Rightarrow y: B.Q}{\Delta; P \vdash M :=_{\tau} N; E \Rightarrow y: B.Q} \\
\\
\frac{\Delta \vdash M \Leftarrow \text{bool} \quad \Delta \vdash A \ \text{type} \quad \Delta; \lambda i. \lambda m. (P \ i \ m) \wedge \text{xid}_{\text{bool}, \text{bool}}(M, \text{true}) \vdash E_1 \Rightarrow x: A.P_1 \quad \Delta; \lambda i. \lambda m. (P \ i \ m) \wedge \text{xid}_{\text{bool}, \text{bool}}(M, \text{false}) \vdash E_2 \Rightarrow x: A.P_2 \quad \Delta, x: A; \lambda i. \lambda m. (P_1 \ i \ m) \vee (P_2 \ i \ m) \vdash E \Rightarrow y: B.Q}{\Delta; P \vdash x = \text{if}_A M \text{ then } E_1 \text{ else } E_2; E \Rightarrow y: B. \lambda i. \lambda m. \exists x: A. Q \ i \ m} \\
\\
\frac{T = \{R\} \ y: B \ \{S\} \quad \Delta \vdash M \Leftarrow A \quad \Delta, x: A \vdash T \ \text{type} \quad \Delta, i: \text{heap}, m: \text{heap}, (P \ i \ m) \vDash (R[M/x] * \lambda h. \top) \ m \quad \Delta, x: A, f: \Pi x: A. T; \lambda i. \lambda m. i = m \wedge (R * \lambda h. \top) \ m \vdash E \Leftarrow y: B. R \gg S \quad \Delta, y: B[M/x]; P \circ (R \gg S)[M/x] \vdash F \Rightarrow z: C.Q}{\Delta; P \vdash y = \text{fix } f(x: A) : T = \text{do } E \text{ in eval } f \ M; F \Rightarrow z: C. (\lambda i. \lambda m. \exists y: B[M/x]. Q \ i \ m)}
\end{array}$$