

One more thing we need: in pure linear logic the only thing we pass is channels; we encode data by having processes that represent data. But for practical programming, we'd like to be able to pass data: arbitrary data in some sense. So what is a logical construct that corresponds to being able to pass arbitrary data?

The right way to do this is quantification.

*revisable types*

$$\begin{aligned} \tau &::= \tau_1 \multimap \tau_2 \mid \dots \mid \{A_1, \dots, A_n \vdash A\} \\ A &::= A_1 \otimes A_2 \mid \mathbb{1} \mid A_1 \multimap B \mid \&\{l_i : A_i\} \mid \oplus \{l_i : A_i\} \\ M &::= \lambda x. M \mid M_1 M_2 \mid \dots \mid \{c \leftarrow P \leftarrow c_1, \dots, c_n\} \end{aligned}$$

*a contextual monad*

*generalized to finitely many w/ labels*

As usual, we'll explicate the behavior of quantifiers using their left and right rules.

*persistent*      *ephemeral*

$$\frac{\Psi, n : \tau ; \Delta \vdash P_n^x :: x : A}{\Psi ; \Delta \vdash (n \leftarrow \text{recv } x ; P_n^x) :: x : \forall n : \tau. A} \quad \forall R^n$$

*non-dependent versions*

*syntax for forming monad expressions*

*n shows up in P*

*ordinary type theory*

$$\frac{\Psi \vdash N : \tau \quad \Psi ; \Delta, x :: [N/n] A \vdash Q_x^z :: z : C}{\Psi ; \Delta, x :: \forall n : \tau. A \vdash (\text{send } x N ; Q_x^z) :: z : C} \quad \forall L^n$$

Intuitively, forall "inputs a value of type t", and exists "outputs a value of type t."

There is a general pattern to left and right rules in sequent calculus, where both rules move towards the identity proof. In natural deduction, the elimination rule moves in the other direction; thus elimination and left rule are "inverses" in some sense.

This would be clearer with an example.

$$\begin{array}{ccc} \text{ND} & \begin{array}{c} E \downarrow \\ \hline \end{array} & \text{id} \\ \text{SEQ} & \begin{array}{c} \hline \\ \uparrow \\ \text{L} \\ = E^{-1} \end{array} & \text{id} \\ & \begin{array}{c} \uparrow \\ \text{R} \end{array} & \end{array}$$

At this point, we can write processes which compute functional values, but we cannot write functions with compute processes. So the idea is to use a monad to let us embed processes into the functional language. In fact, we will need a contextual monad which indicates which processes it uses. Let us state the language for processes, recast in monadic language:

$$\begin{array}{l}
 P ::= c \leftarrow a \\
 | x \leftarrow M \leftarrow c_1 \dots c_n ; Q_x \\
 | \dots
 \end{array}
 \begin{array}{l}
 (\text{id}) \\
 (\text{comp}) \text{ also known as bind}
 \end{array}$$

This is classical monads, except that the variables inside the monad have some extra binding structure.

Now for some examples:

$$\begin{aligned}
 \text{natstream} &= \exists n : \text{nat}. \text{natstream} \\
 &= \text{nat} \wedge \text{natstream} \\
 &= \mu \alpha. \text{nat} \wedge \alpha
 \end{aligned}$$

$$\text{nats} : \text{nat} \rightarrow \{ \vdash \text{natstream} \}$$

$$c \leftarrow \text{nats } n =$$

$$\{ \text{send } c \ n ;$$

$$c' \leftarrow \text{nats}(n+1) ;$$

$$c \leftarrow c' \}$$

send the current number over the stream


make a recursive call

forward c' on c

these two lines are common, so they can be abbreviated to:

$$c \leftarrow \text{nats}(n+1)$$

Our goal is to generate a stream of prime numbers, using the Sieve of Eratosthenes.



We have a process per prime number we've outputted so far, and it checks numbers streaming to see if they are divisible (dropping them if they are). If a number makes it through, it's prime, so we generate a new process for it.

First, we need to build a filter transducer:

$\text{filter} : (\text{nat} \rightarrow \text{bool}) \rightarrow \{ \text{natstream} \vdash \text{natstream} \}$

$c \leftarrow \text{filter } p \leftarrow d =$

$\{ n \leftarrow \text{recv } d;$   
 $\text{if } (p \ n)$

$\text{then send } c \ n;$

abbreviation!  $c \leftarrow \text{filter } p \leftarrow d$

$\text{else } c \leftarrow \text{filter } p \leftarrow d \}$

Q: Is the send and recursive call concurrent?

A: No, it's a synchronous language.

Q: Can I pull out the filter  $p$  recursive call?

A: No, there is a dependence. **Interesting!**

Q: Where does the if come from?

A: It is in the term language. This is a standard monadic programming trick.

Note that if you try to actually implement this in Coq, you won't be able to do this. The problem is that filter is not productive. Now, in the case of a true prime sieve, it will be productive, since there are infinitely many primes, but that's not obvious! There's a paper on how to do this.

Now let's do the sieve. The sieve takes a stream, knowing that the first element is a prime number (by some invariant), and then starts building processes.

$\text{sieve} : \{ \text{natstream} \vdash \text{natstream} \}$

$c \leftarrow \text{sieve} \leftarrow d =$

$\{ k \leftarrow \text{recv } d; \quad \text{we know } k \text{ is prime}$

$\text{send } c \ k;$

$x \leftarrow \text{filter } (\lambda n. k \nmid n) \leftarrow d;$

$c \leftarrow \text{sieve} \leftarrow x \quad \leftarrow \text{not divisible by}$   
 $\}$


Finally, we generate the primes stream:

```
primes : { ⊢ natstream }
c ← primes =
{ x ← nats 2 ;
  c ← sieve ← x
}
```

Q: Could you do sieve using cut instead of monads?

A: Well, cut doesn't let you get the result of a function. We've just gotten rid of cut with monadic composition, since there is not really any reason to have it any more.

Our next example is a binary number counter.



operations: inc  
val  
dealloc

```
bin = & { inc : bin
        val : nat ∧ bin
        dealloc : 1
      }
```

Q: Why isn't it tensor?

A: I'm outputting a value, not a new channel, and there is no dependence (so we didn't use an existential).


```
val : (nat ∧ 1) ⊗ bin
```

This version sends a channel, and doesn't block! Which is potentially interesting.

```
ε : { ⊢ bin }
```

```
c ← ε =
{ case recv c of
  | inc ⇒ x ← ε ;
  | val ⇒ send 0 c ;
  | dealloc ⇒ close c
}
```

this is a bit confusing, because it's in reverse



```
bit : bool → { bit ⊢ bin }
```

```
c ← bit b ← d =
{ case recv of
  | inc ⇒ if b
    then send d inc ;
    c ← bit 0 ← d
  else c ← bit 1 ← d
  | val ⇒ send d val ;
  n ← recv d ;
  send c (b + 2 * n) ;
  c ← bit b ← d
  | dealloc ⇒ send d dealloc ;
  wait d ;
  close c
}
```

if we omit the wait it's not linearly typed!

Note that this is concurrent; you can send multiple increments, and they will all be in flight simultaneously!

Homework: implement stacks.

$$\text{stack} = \&\{ \text{push} : \text{nat} \supset \text{stack}, \\ \text{pop} : \oplus\{ \text{empty} : \text{stack}, \\ \text{some} : \text{nat} \wedge \text{stack} \}, \\ \text{dealloc} : 1 \}$$

Remark: when you say 1, it really does mean that EVERYTHING is deallocated, because there must not be anything in the context. So if we made a mistake in implementing dealloc, type error!

$$\text{new} : \{ \vdash \text{stack} \}$$

One last note: what is the meaning of persistent assumptions and persistent channels in this context? (I.e. bang-A) Interestingly, it doesn't seem to come up very often in this kind of program.

We have an implementation of this, but it's not quite stable enough. Note that Concurrent ML is almost expressive enough to write all these programs, but it is completely undisciplined. Session typing with these systems guarantees absence of deadlock, etc; and with restrictions on recursion, also guarantee termination.

$$\underbrace{u_1 : B_1, \dots, u_n : B_n}_{\text{shared}}; \underbrace{x_1 : A_1, \dots, x_n : A_n}_{\text{linear}} \vdash P :: x : A$$

can show up many times

Recall that for ordinary channels, their type changes as the program evaluates. But for a shared channel, you can't allow the type to change because other references may get mixed up and type preservation would not hold. Linearity prevents this from happening, but for things that can be used multiple times, you need some restrictions. So instead, you spawn a copy of the service in question, and use that!

$$\frac{\Gamma, u : A; \Delta, x : A \vdash Q_{u,x} :: z : C}{\Gamma, u : A; \Delta \vdash x \leftarrow \text{copy } u; Q_{u,x} :: z : C} \text{ copy}$$

Q: So, is the reason this doesn't show up frequently because we can simulate it using the monadic syntax?

A: Well, sometimes you want the use of the shared channel to show up in the type. Consider a file storage system:

$u: \text{file} \supset (\text{file} \wedge 1)$

$x \leftarrow \text{copy } u;$   
 $\text{send } x \text{ } F;$   
 $G \leftarrow \text{recv } x$

If it's correctly implemented,  $g$  and  $g'$  should be equal. With dependent types, you can actually express that these should be equal; this was a paper last year.

$u: \text{file} \supset !(\text{file} \wedge 1)$

$x \leftarrow \text{copy } u;$   
 $\text{send } x \text{ } F;$   
 $w \leftarrow \text{recv } x;$   
 $y \leftarrow \text{copy } w$   
 $g \leftarrow \text{recv } y$   
 $y' \leftarrow \text{copy } w$   
 $g' \leftarrow \text{recv } y'$

} as many  
as you  
want

$u: \forall x: \text{file}. \exists y: \text{file}. \exists p: x \approx y. 1$

↖ e.g. the index is similar

This is a dependent specification; notice that without the "input a file, output a file" there are a lot of implementations, but with the equality now we know what it should be.

$u: \forall x: \text{file}. !\exists y: \text{file}. \exists p: x = y. 1$

Note that in a distributed setting, you will need to do some runtime typechecking to ensure that all proof obligations are satisfied, if you don't trust the nodes that you are interacting with. So the proofs will have some runtime content! (On a single machine, nothing can go wrong).