

Lecture Notes on Queues and Stacks

15-816: Substructural Logics
Frank Pfenning

Lecture 9
Tuesday, September 27, 2016

In the last lecture we introduced lists with arbitrary elements and wrote ordered programs for *nil* (the empty list), *cons* (adding an element to the head of a list) and *append* to append two lists. The representation was in the form of an *internal choice*

$$\text{list}_A = \oplus\{\text{cons} : A \bullet \text{list}_A, \text{nil} : \mathbf{1}\}$$

We might think of this as the usual functional data structure of lists, but we should keep in mind that it is really just an interface specification for processes. It does not imply any particular representation.

Today, we will look at a data structure in which we can insert and delete channels of arbitrary type. The interface is different because it is in the form of an *external choice*, more in the style of object-oriented programming or signatures in module systems for functional languages.

1 Storing Channels

Here is our simple interface to a storage service for channels:

$$\text{store}_A = \&\{\text{ins} : A \setminus \text{store}_A, \\ \text{del} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \bullet \text{store}_A\}\}$$

Using our operational interpretation, we can read this as follows:

A store for channels of type A offers a client a choice between insertion (label ins) and deletion (label del).

When inserting, the clients sends a channel of type A which is added

to the store.

When deleting, the store responds with the label `none` if there are no elements in the store and terminates, or with the label `some`, followed by an element.

When an element is actually inserted or deleted the provider of the storage service then waits for the next input (again, either an insertion or deletion).

In this reading we have focused on the operations, and intentionally ignored the restrictions order might place on the use of the storage service. Hopefully, this will emerge as we write the code and analyze what the restrictions might mean.

First, we have to be able to create an empty store. We will write the code in stages, because I believe it is much harder to understand the final program than it is to follow its construction.

$$\text{store}_A = \&\{\text{ins} : A \setminus \text{store}_A, \\ \text{del} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \bullet \text{store}_A\}\}$$

First, the header of the process definition.

$$\cdot \vdash \text{empty} :: (s : \text{store}_A) \\ s \leftarrow \text{empty} = \dots$$

Because a store_A is an external choice, we begin with a case construct, branching on the received label.

$$\cdot \vdash \text{empty} :: (s : \text{store}_A) \\ s \leftarrow \text{empty} = \text{case } s \text{ (ins } \Rightarrow \dots \qquad \% \cdot \vdash s : A \setminus \text{store}_A \\ \qquad \qquad \qquad | \text{del } \Rightarrow \dots \qquad \% \cdot \vdash s : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \bullet \text{store}_A\} \\ \qquad \qquad \qquad)$$

The case of deletion is actually easier: since this process represents an empty store, we send the label `none` and terminate.

$$\cdot \vdash \text{empty} :: (s : \text{store}_A) \\ s \leftarrow \text{empty} = \text{case } s \text{ (ins } \Rightarrow \dots \qquad \% \cdot \vdash s : A \setminus \text{store}_A \\ \qquad \qquad \qquad | \text{del } \Rightarrow s.\text{none} ; \text{close } s)$$

In the case of an insertion, the type dictates that we receive a channel of type A which we call x . It is added at the left end of the antecedents. Since they are actually none, both $A \setminus \text{store}_A$ and store_A / A would behave the same way here.

$$\begin{array}{l} \cdot \vdash \text{empty} :: (s : \text{store}_A) \\ s \leftarrow \text{empty} = \text{case } s \text{ (ins } \Rightarrow \\ \quad x \leftarrow \text{recv } s ; \quad \% \quad \cdot \vdash s : A \setminus \text{store}_A \\ \quad \dots \quad \% \quad x:A \vdash s : \text{store}_A \\ \quad | \text{del } \Rightarrow s.\text{none} ; \text{close } s) \end{array}$$

At this point it seems like we are stuck. We need to start a process implementing a store with *one* element, but so far we just writing the code for an empty store. We need to define a process *elem*

$$(x:A) (t:\text{store}_A) \vdash \text{elem} :: (s : \text{store}_A)$$

which holds an element $x:A$ and also another store $t:\text{store}_A$ with further elements. In the singleton case, t will then be the empty store. Therefore, we first make a recursive call to create another empty store, calling it n for *none*.

$$\begin{array}{l} \cdot \vdash \text{empty} :: (s : \text{store}_A) \\ s \leftarrow \text{empty} = \text{case } s \text{ (ins } \Rightarrow x \leftarrow \text{recv } s ; \quad \% \quad x:A \vdash s : \text{store}_A \\ \quad n \leftarrow \text{empty} ; \quad \% \quad (x:A) (n:\text{store}_A) \vdash s : \text{store}_A \\ \quad \dots \\ \quad | \text{del } \Rightarrow s.\text{none} ; \text{close } s) \end{array}$$

$$\begin{array}{l} (x:A) (t:\text{store}_A) \vdash \text{elem} :: (s : \text{store}_A) \\ s \leftarrow \text{elem} \leftarrow x t = \dots \end{array}$$

Postponing the definition of *elem* for now, we can invoke *elem* to create a singleton store with just x , calling the resulting channel e . This call will consume x and n , leaving e as the only antecedent.

$$\begin{array}{l} \cdot \vdash \text{empty} :: (s : \text{store}_A) \\ s \leftarrow \text{empty} = \text{case } s \text{ (ins } \Rightarrow x \leftarrow \text{recv } s ; \quad \% \quad x:A \vdash s : \text{store}_A \\ \quad n \leftarrow \text{empty} ; \quad \% \quad (x:A) (n:\text{store}_A) \vdash s : \text{store}_A \\ \quad e \leftarrow \text{elem} \leftarrow x n ; \quad \% \quad e:\text{store}_A \vdash s : \text{store}_A \\ \quad \dots \\ \quad | \text{del } \Rightarrow s.\text{none} ; \text{close } s) \end{array}$$

$$\begin{array}{l} (x:A) (t:\text{store}_A) \vdash \text{elem} :: (s : \text{store}_A) \\ s \leftarrow \text{elem} \leftarrow x t = \dots \end{array}$$

At this point we can implement s by e (the singleton store), which is just an application of the identity rule.

```

· ⊢ empty :: (s : storeA)
s ← empty = case s (ins ⇒ x ← recv s ;           % (x:A) ⊢ s : storeA
                  n ← empty ;                   % (x:A) (n:storeA) ⊢ s : storeA
                  e ← elem ← x n                 % e:storeA ⊢ s : storeA
                  s ← e
                  | del ⇒ s.none ; close s)
(x:A) (t:storeA) ⊢ elem :: (s : storeA)
s ← elem ← x t = ...

```

It remains to write the code for the process holding an element of the store. We suggest you reconstruct or at least read it line by line the way we developed the definition of *empty*, but we will not break it out explicitly into multiple steps. However, we will still give the types after each interaction. For easy reference, we repeat the type definition for store_A .

$$\text{store}_A = \&\{\text{ins} : A \setminus \text{store}_A, \\ \text{del} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \bullet \text{store}_A\}\}$$

$$(x:A) (t:\text{store}_A) \vdash \text{elem} :: (s : \text{store}_A)$$

```

1  s ← elem ← x t =
2      case s (ins ⇒ y ← recv s ;           % (y:A) (x:A) (t:store_A) ⊢ s : store_A
3          t.ins ;                          % (y:A) (x:A) (t:A \ store_A) ⊢ s : store_A
4          send t x ;                       % (y:A) (t:store_A) ⊢ s : store_A
5          r ← elem ← y t ;                 % r:store_A ⊢ s : store_A
6          s ← r
7      | del ⇒ s.some ;                     % (x:A) (t:store_A) ⊢ s : A • store_A
8          send s x ;                       % t:store_A ⊢ s : store_A
9          s ← t)
```

A few notes on this code. Look at the type at the end of the *previous* line to understand the next line.

- In line 2, we add $y:A$ at the left end of the context since $s : A \setminus \text{store}_A$.
- In line 4, we can only pass x to t but not y , due restrictions of $\setminus L^*$.
- In line 5, y and t are in the correct order to call *elem* recursively.
- In line 8, we can pass x along s since it is at the left end of the context.

How does this code behave? Assume we have a store s holding elements x_1 and x_2 it would look like

$$\text{proc}(s, s \leftarrow \text{elem} \leftarrow x_1 t_1) \quad \text{proc}(t_1, t_1 \leftarrow \text{elem} \leftarrow x_2 t_2) \quad \text{proc}(t_2, t_2 \leftarrow \text{empty})$$

where we have indicated the code executing in each process without unfolding the definition. If we insert an element along s (by sending *ins* and then a new y) then the process $s \leftarrow \text{elem} \leftarrow x_1 t_1$ will insert x_1 along t_1 and then, in two steps, become $s \leftarrow \text{elem} \leftarrow y t_1$. Now the next process will pass x_2 along t_2 and hold on to x_1 , and finally the process holding no element will spawn a new one (t_3) and itself hold on to x_2 .

$$\text{proc}(s, s \leftarrow \text{elem} \leftarrow y t_1) \quad \text{proc}(t_1, t_1 \leftarrow \text{elem} \leftarrow x_1 t_2) \\ \text{proc}(t_2, t_2 \leftarrow \text{elem} \leftarrow x_2 t_3) \quad \text{proc}(t_3, t_3 \leftarrow \text{empty})$$

If we next delete an element, we will get y back and the store will effectively revert to its original state, with some (internal) renaming.

$\text{proc}(s, s \leftarrow \text{elem} \leftarrow x_1 \ t_2) \quad \text{proc}(t_2, t_2 \leftarrow \text{elem} \leftarrow x_2 \ t_3) \quad \text{proc}(t_3, t_3 \leftarrow \text{empty})$

In essence, the store behaves like a *stack*: the most recent element we have inserted will be the first one deleted. If you carefully look through the intermediate types in the *elem* process, it seems that this behavior is forced. We conjecture that any implementation of the store interface we have given will behave like a stack or might at some point not respond to further messages. We do not yet have the means to carry out such a proof. Some related prior work might provide hints on how this might be proved using parametricity [Rey83, CPPT13].¹

¹If I or someone else in the class can prove or refute this conjecture, we may return to it in a future lecture.

2 Tail Calls

Let's look again at the two pieces of code we have written.

```
storeA = &{ ins : A \ storeA,
           del : ⊕{none : 1, some : A • storeA}}
```

```
· ⊢ empty :: (s : storeA)
```

```
1  s ← empty =
2  case s (ins ⇒ x ← recv s ;      % (x:A) ⊢ s : storeA
3           n ← empty ;           % (x:A) (n:storeA) ⊢ s : storeA
4           e ← elem ← x n      % e:storeA ⊢ s : storeA
5           s ← e
6           | del ⇒ s.none ; close s)
```

```
(x:A) (t:storeA) ⊢ elem :: (s : storeA)
```

```
7  s ← elem ← x t =
8  case s (ins ⇒ y ← recv s ;      % (y:A) (x:A) (t:storeA) ⊢ s : storeA
9           t.ins ;                % (y:A) (x:A) (t:A \ storeA) ⊢ s : storeA
10          send t x ;              % (y:A) (t:storeA) ⊢ s : storeA
11          r ← elem ← y t ;      % r:storeA ⊢ s : storeA
12          s ← r
13          | del ⇒ s.some ;       % (x:A) (t:storeA) ⊢ s : A • storeA
14          send s x ;            % t:storeA ⊢ s : storeA
15          s ← t)
```

empty starts two new processes, in lines 3 and 4 and then terminates in line 5 by forwarding. *elem* spawns only one new process, in line 11, and then terminates in line 12 by forwarding. Intuitively, spawning a new process and then immediately forwarding to this process is wasteful, especially if process creation is an expensive operation.

It would be nice if the process executing *empty* could effectively just continue by executing *elem*, and similarly, if *elem* could continue as the same process once *x* has been sent along *t*. This can be achieved if we treat *tail calls* specially. So instead of writing

```
4  e ← elem ← x n
5  s ← e
```

we write

```
4  s ← elem ← x n
```

and similarly in the definition of *elem*.

In general, we compress a cut in the form of a process invocation followed by an identity simply as a process invocation:

$$\begin{aligned} y &\leftarrow X \leftarrow y_1 \dots y_n \\ x &\leftarrow y \end{aligned}$$

becomes

$$x \leftarrow X \leftarrow y_1 \dots y_n$$

This is analogous to the so-called *tail-call optimization* in functional languages where instead of f calling a function g and immediately returning its value, f just continues as g . This is often represented as saving stack space since it can be implemented as a jump instead of a call. Here, too, recursively defined processes executing a sequence of interactions can simply continue without spawning a new process and then forwarding the result immediately, thereby saving process invocations.

From now on, we will often silently use the compressed form. Of course, its purely logical meaning can be recovered by expanding it into a cut followed by an identity.

3 Analyzing Parallel Complexity

We can analyze various complexity measures of our implementations. For example, we can count the number of processes that execute. Any call (except for a tail call) will spawn a new process, and any forward and close will terminate a process. Looking at the code below we can see that inserting an element into a store will spawn exactly one new process, namely when we eventually insert the last element into the empty store. Deleting an element will terminate exactly one process: either the empty one, or the one holding the element we are returning. Therefore in a store with n elements there will be *exactly* $n + 1$ processes.

```

storeA = &{ ins : A \ storeA,
           del : ⊕{ none : 1, some : A • storeA } }

· ⊢ empty :: (s : storeA)
1  s ← empty =
2    case s (ins ⇒ x ← recv s ;           % (x:A) ⊢ s : storeA
3              n ← empty ;               % (x:A) (n:storeA) ⊢ s : storeA
4              s ← elem ← x n
5              | del ⇒ s.none ; close s)

(x:A) (t:storeA) ⊢ elem :: (s : storeA)
6  s ← elem ← x t =
7    case s (ins ⇒ y ← recv s ;           % (y:A) (x:A) (t:storeA) ⊢ s : storeA
8              t.ins ;                   % (y:A) (x:A) (t:A \ storeA) ⊢ s : storeA
9              send t x ;                 % (y:A) (t:storeA) ⊢ s : storeA
10             s ← elem ← y t
11             | del ⇒ s.some ;           % (x:A) (t:storeA) ⊢ s : A • storeA
12             send s x ;                 % t:storeA ⊢ s : storeA
13             s ← t)

```

Another interesting measure is the *reaction time* which is analogous to the *span* complexity measure for parallel programs. If we try to carry out two consecutive operations, how many steps must elapse between them, assuming maximal parallelism? Here it is convenient to count every interaction as a step and no other costs.

Looking at the code for *elem* we see that there are only two interactions along channel t until the *elem* process can interact again along s after it has received *ins* and y . For *empty* there is only one spawn but no other

interactions. Moreover, there is no delay for a deletion, since the process will respond immediately along s .

In aggregate, when we store n elements consecutively, the constant reaction time means that there will be n elements building up the internal data structure simultaneously. No matter how many insertions and deletions we carry out, the reaction time (measured in total system interactions assuming maximal parallelism) is always constant.

On the other hand, if we count the total number of interactions of the system taking place (ignoring any question of parallelism) we see that for n insertions it will be $O(n^2)$, since each new element initiates a chain reaction that reaches to the end of the chain of elements. This is usually called the *work* performed by the algorithm.

4 Queues

As notes, our implementation so far ended up behaving like a stack, and we conjectured that the type of the interface itself forced this behavior. Can we modify the type to allow (and perhaps force) the behavior of the store as a queue, where the first element we store is the first one we receive back? I encourage you to try to work this out before reading on . . .

The key idea is to change the type

$$\text{store}_A = \&\{\text{ins} : A \setminus \text{store}_A, \\ \text{del} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \bullet \text{store}_A\}\}$$

to

$$\text{queue}_A = \&\{\text{ins} : \text{queue}_A / A, \\ \text{del} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \bullet \text{queue}_A\}\}$$

We will not go through this in detail, but reading the following code and the type after each interaction should give you a sense for what this change entails.

$\cdot \vdash \text{empty} :: (s : \text{queue}_A)$

```

1  s ← empty =
2    case s (ins ⇒ x ← recv s ;           % x:A ⊢ s : queue_A
3          n ← empty ;                   % (x:A) (n:queue_A) ⊢ s : queue_A
4          s ← elem ← x n
5          | del ⇒ s.none ; close s)

```

$(x:A) (t:\text{queue}_A) \vdash \text{elem} :: (s : \text{queue}_A)$

```

6  s ← elem ← x t =
7    case s (ins ⇒ y ← recv s ;           % (x:A) (t:queue_A) (y:A) ⊢ s : queue_A
8          t.ins ;                       % (x:A) (t:queue_A / A) (y:A) ⊢ s : queue_A
9          send t y ;                     % (x:A) (t:queue_A) ⊢ s : queue_A
10         s ← elem ← x t
11         | del ⇒ s.some ;               % (x:A) (t:queue_A) ⊢ s : A • queue_A
12         send s x ;                     % t:queue_A ⊢ s : queue_A
13         s ← t)

```

The critical changes are in line 7 (where y is added to the *right end* of the antecedents instead of the left) and line 9 (where consequently y instead of x must be sent along t).

The complexity of all the operations remains the same, since the only difference is whether the current x or the new y is sent along t , but the implementation now behaves like a queue rather than a stack.

Exercises

Exercise 1 In this exercise we explore an alternative implementation of stacks. First, consider type of stacks (renamed from store_A in this lecture)

$$\text{stack}_A = \&\{\text{ins} : A \setminus \text{stack}_A, \\ \text{del} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \bullet \text{stack}_A\}\}$$

1. Provide definitions for

$$\begin{aligned} \cdot &\vdash \text{stack_new} :: (s : \text{stack}_A) \\ l : \text{list}_A &\vdash \text{stack} :: (s : \text{stack}_A) \end{aligned}$$

which represents the elements of the stack in a list. If you need auxiliary process definitions for lists, please state them clearly, including their type.

2. Repeat the analysis of [Section 3](#):
 - (a) How many processes execute for a stack with n elements?
 - (b) What is the reaction time for an insertion or deletion given a stack with n elements?
 - (c) What is the total work for each insertion or deletion given a stack with n elements?

Exercise 2 In this exercise we explore an alternative implementation of queues. First, recall the type of queues from [Section 4](#).

$$\text{queue}_A = \&\{\text{ins} : \text{queue}_A / A, \\ \text{del} : \oplus\{\text{none} : \mathbf{1}, \text{some} : A \bullet \text{queue}_A\}\}$$

1. Provide definitions for

$$\begin{aligned} \cdot &\vdash \text{queue_new} :: (s : \text{queue}_A) \\ l : \text{list}_A &\vdash \text{queue} :: (s : \text{queue}_A) \end{aligned}$$

which represents the elements of the queue in a list. If you need auxiliary process definitions for lists, please state them clearly, including their type.

2. Repeat the analysis of [Section 3](#):

- (a) How many processes execute for a queue with n elements?
- (b) What is the reaction time for an insertion or deletion given a queue with n elements?
- (c) What is the total work for each insertion or deletion given a queue with n elements?

Exercise 3 In this exercise we will “turn around” Exercise 1. Write a process definition

$$s:\text{stack}_A \vdash \text{to_list} :: (l:\text{list}_A)$$

which converts a stack into a list. As far as you can tell, is the order of the elements that are sent along l fixed?

Exercise 4 Consider the standard functional programming technique of implementing a queue with two lists. Just briefly, we have an *input list* in to which we add elements when they are enqueued and an *output list* out from which we take elements when they are dequeued. When the output list becomes empty, we reverse the input list, adding each element in turn onto the output list. Initially, both lists are empty.

Explore if you can write such an implementation against the queue interface from Section 4. The implementation should have one of the two types

$$\begin{aligned} (in:\text{list}_A) (out:\text{list}_A) \vdash \text{queue2} &:: (s : \text{queue}_A) \\ (out:\text{list}_A) (in:\text{list}_A) \vdash \text{queue2} &:: (s : \text{queue}_A) \end{aligned}$$

References

- [CPPT13] Luís Caires, Jorge A. Pérez, Frank Pfenning, and Bernardo Toninho. Behavioral polymorphism and parametricity in session-based communication. In M.Felleisen and P.Gardner, editors, *Proceedings of the European Symposium on Programming (ESOP'13)*, pages 330–349, Rome, Italy, March 2013. Springer LNCS 7792.
- [Rey83] John C. Reynolds. Types, abstraction, and parametric polymorphism. In R.E.A. Mason, editor, *Information Processing 83*, pages 513–523. Elsevier, September 1983.