

Examples of Specifications and Implementations

This handout is a supplement for the first two lectures. It contains several example specifications and implementations, all written using Spec.

Section 1 contains a specification for sorting a sequence. Section 2 contains two specifications and one implementation for searching for an element in a sequence. Section 3 contains specifications for a read/write memory. Sections 4 and 5 contain implementations for a read/write memory based on caching and hashing, respectively. Finally, Section 6 contains an implementation based on replicated copies.

Sorting

The following specification describes the behavior required of a program that sorts integers.

```
APROC Sort(a: SET Int) -> SEQ Int = <<  
  VAR b: SEQ Int | (ALL i: Int | a.count(i) = b.count(i)) /\ Sorted(b) => RET b >>
```

This specification uses the auxiliary function `Sorted`, defined as follows.

```
FUNC Sorted(a: SEQ Int) -> Bool = RET (ALL i :IN a.dom - {0} | a(i-1) <= a(i))
```

If we made `Sort` a `FUNC` rather than a `PROC`, what would be wrong?

We could have written this more concisely as

```
APROC Sort(a: SET Int) -> SEQ Int = << VAR b :IN a.perms | Sorted(b) => RET b >>
```

using the `perms` method for sets that returns a set of sequences that contains all the possible permutations of the set.

Searching

Search specification

We begin with a specification for a procedure to search an array for a given element. This is an `APROC` rather than a `FUNC` because there can be several allowable results for the same inputs.

```
APROC Search(a: SEQ Int, x: Int) -> Int RAISES {NotFound} =  
  << IF VAR i: Int | (0 <= i /\ i < a.size /\ a(i) = x) => RET i
```

```

    [*] RAISE NotFound
  FI >>

```

Or, equivalently but slightly more concisely:

```

APROC Search(a: SEQ Int, x: Int) -> Int RAISES {NotFound} =
  << IF VAR i :IN a.dom | a(i) = x => RET i [*] RAISE NotFound FI >>

```

Sequential search implementation

Here is an implementation of the `Search` specification given above. It uses sequential search, starting at the first element of the input sequence.

```

APROC SeqSearch(a: SEQ Int, x: Int) -> Int RAISES {NotFound} = << VAR i := 0 |
  DO i < a.size => IF a(i) = x => RET i [*] i := i + 1 FI OD; RAISE NotFound >>

```

Alternative search specification

Some searching algorithms, for example, binary search, assume that the input argument sequence is sorted. Such algorithms require a different specification, one that expresses this requirement.

```

APROC Search1(a: SEQ Int, x: Int) -> Int RAISES {NotFound} = <<
  IF ~Sorted(a) => HAVOC
  [*] VAR i :IN a.dom | a(i) = x => RET i
  [*] RAISE NotFound
  FI >>

```

You might consider writing the specification to raise an exception when the array is not sorted:

```

APROC Search2(a: SEQ Int, x: Int) -> Int RAISES {NotFound, NotSorted} = <<
  IF ~Sorted(a) => RAISE NotSorted
  ...

```

This is not a good idea. The whole point of binary search is to obtain $O(\log n)$ time performance (for a sorted input sequence). But any implementation of the `Search2` specification requires an $O(n)$ check, even for a sorted input sequence, in order to verify that the input sequence is in fact sorted.

This is a simple but instructive example of the difference between defensive programming and efficiency. If `Search` were part of an operating system interface, it would be intolerable to have `HAVOC` as a possible transition, because the operating system is not supposed to go off the deep end no matter how it is called (though it might be OK to return the wrong answer if the input isn't sorted). On the other hand, the efficiency of a program often depends on assumptions that one part of it makes about another, and it's appropriate to express such an assumption in a spec by saying that you get `HAVOC` if it is violated. We don't care to be more specific about what happens because we intend to ensure that it doesn't happen. Obviously a program written in this style will be more prone to undetected or obscure errors than one that checks the assumptions.

Read/write memory

The simplest form of read/write memory is a single read/write register, say of type `D` (for data), with arbitrary initial value. This is described by the following Spec module:

```

MODULE Register [D] EXPORT Read, Write =

VAR x: D                                     % arbitrary initial value

APROC Read() -> D = << RET x >>
APROC Write(d) = << x := d >>

END Register

```

Now we give a specification for a simple addressable memory with elements of type `D`. This is like a collection of read/write registers, one for each address in a set `A`. For variety, we include new `Reset` and `Swap` operations in addition to `Read` and `Write`.

```

MODULE SimpleMemory [A, D] EXPORT Read, Write, Reset, Swap =

TYPE M = A -> D
VAR m := Init()

APROC Init() -> M = << VAR m' | (ALL a | m'!a) => RET m' >>
% Choose an arbitrary function that is defined everywhere.

FUNC Read(a) -> D = << RET m(a) >>
APROC Write(a, d) = << m(a) := d >>

APROC Reset(d) = << m := M{* -> d} >>
% Set all memory locations to d.

APROC Swap(a, d) -> D = << VAR d' := m(a) | m(a) := d; RET d' >>
% Set location a to the input value and return the previous value.

END SimpleMemory

```

The next three sections describe implementations of `SimpleMemory`.

Write-back cache implementation

Our first implementation is based on two memory mappings, a main memory `m` and a *write-back cache* `c`. The implementation maintains the invariant that the number of addresses at which `c` is defined is constant. A real cache would probably maintain a weaker invariant, perhaps bounding the number of addresses at which `c` is defined.

```

MODULE WBCache [A,D] EXPORT Read, Write, Reset, Swap =
% implements SimpleMemory

TYPE M          = A -> D
   C            = A -> D

VAR Csize      : Int := ...                % cache size; constant

VAR m          := InitM()
   c           := InitC()

APROC InitM() -> M = << VAR m' | (ALL a | m'!a) => m := m' >>
% Initializes all entries in m with arbitrary values.

APROC InitC() -> C = << VAR c' | c'.dom.size = CSize => c := c' >>
% Initializes exactly CSize entries in c with arbitrary values.

```

```

APROC Read(a) -> D = << Load(a); RET c(a) >>

APROC Write(a, d) = << IF ~c!a => FlushOne() [*] SKIP FI; c(a) := d >>
% Makes room in the cache if necessary, then writes to the cache.

APROC Reset(d) = <<...>> % exercise for the reader

APROC Swap(a, d) -> D = << VAR d' | Load(a); d' := c(a); c(a) := d; RET d' >>

APROC Load(a) = << IF ~c!a => FlushOne(); c(a) := m(a) [*] SKIP FI >>
% Ensures that address a appears in the cache.

APROC FlushOne() =
% Removes one (arbitrary) address from the cache, writing the data
% value back to main memory if necessary.
  << VAR a | c!a => IF Dirty(a) => m(a) := c(a) [*] SKIP FI; c := c{a -> } >>

FUNC Dirty(a) -> Bool = RET c!a /\ c(a) # m(a)
% Returns true if the cache is more up-to-date than the main memory.

END WBCache

```

The following Spec function is an abstraction function mapping a state of the WBCache module to a state of the SimpleMemory module:

```

FUNC AF(m, c, CSize: Int) -> M = RET (\ a | c!a => c(a) [*] m(a) )

```

Hash table implementation

Our second implementation of SimpleMemory uses a hash table for the representation.

```

MODULE HashMemory [A WITH {hf: A->Int}, D] EXPORT Read, Write, Reset, Swap =
% implements SimpleMemory

% The module expects that the hash function A.hf is total and that
% its range is a finite subset of the positive integers.

TYPE HashT      = SEQ Bucket
   Bucket      = SEQ Pair
   Pair        = [a, d]

VAR nb          := NumBuckets() % number of buckets
   m            := HashT.fill(Bucket{}, nb) % fills m with empty buckets
   default     : D % arbitrary default value

APROC Read(a) -> D = << VAR bucket := m(a.hf), i: Int |
   i := FindEntry(a, bucket) EXCEPT NotFound => RET default ; RET bucket(i).d >>

APROC Write(a, d) = << VAR bucket := DeleteEntry(a, m(a.hf)) |
   m(a.hf) := bucket + {Pair{a, d}} >>

APROC Reset(d) = << m := HashT.fill(Bucket{}, nb); default := d >>

APROC Swap(a, d) -> D = << VAR d' | d' := Read(a); Write(a, d); RET d' >>

FUNC NumBuckets() -> Int RAISES {BadHF} =
% Returns the number of buckets needed by the hash function;
% havoc if the hash function is not as expected.
  IF A.hf.rng.min >= 1 => RET A.hf.rng.max [*] HAVOC FI

```

```

APROC FindEntry(a, bucket) -> Int RAISES (NotFound) =
% If a appears in a pair in bucket, returns the index of some pair
% containing a; otherwise raises NotFound.
  << VAR i :IN bucket.dom | bucket(i).a = a => RET i [*] RAISE NotFound >>

APROC DeleteEntry(a, bucket) -> Bucket << VAR i: Int |
% Removes some pair with address a from bucket, if any exists.
  i := FindEntry(a, bucket) EXCEPT NotFound => RET bucket ;
  RET bucket.sub(1, i-1) + bucket.sub(i+1, bucket.size) >>

END HashMemory

```

Note that `FindEntry` and `DeleteEntry` are APROCS because they are not deterministic when given arbitrary `bucket` arguments.

The following is a key invariant that holds between invocations of the operations of `HashMemory`:

```

FUNC Inv(nb: Int, m: HashT, default: D) -> Bool = RET
(  nb > 0
/\ m.size = nb
/\ (ALL a | a.hf IN m.dom)
/\ (ALL i :IN m.dom, p :IN m(i).rng | p.a.hf = i)
/\ (ALL a | { j :IN m(a.hf) | m(a.hf)(j).a = a }.size <= 1) )

```

This says that the number of buckets is positive, that the hash function maps all addresses to actual buckets, that a pair containing address `a` appears only in the bucket at index `a.hf` in `m`, and that at most one pair for an address appears in the bucket for that address. Note that these conditions imply that in any reachable state of `HashMemory`, each address appears in at most one pair in the entire memory.

The following `Spec` function is an abstraction function between states of the `HashMemory` module and states of the `SimpleMemory` module:

```

FUNC AF(nb: Int, m: HashT, default: D) -> M = RET
(LAMBDA(a) -> D =
  IF  VAR i :IN m.dom, p :IN m(i).rng | p.a = a => RET p.d
  [*] RET default
FI)

```

That is, the data value for address `a` is any value associated with address `a` in the hash table; if there is none, the data value is the default value. `Spec` says that a function is undefined at an argument if its body can yield more than one result value. The invariants given above ensure that the `LAMBDA` is actually single-valued for all the reachable states of `HashMemory`.

Replicated copies

Our final implementation is based on some number $k \geq 1$ of copies of each memory location. Initially, all copies have the same default value. A `Write` operation only modifies an arbitrary *majority* of the copies. A `Read` reads an arbitrary majority, and selects and returns the most recent of the values it sees. In order to allow the `Read` to determine which value is the most recent, each `Write` records not only its value, but also a sequence number.

For simplicity, we just show the module for a single read/write register. It is parameterized by the constant `k`, which names the copies.

```

MODULE MajorityRegister [D] = % implements Register

TYPE N          = Int SUCHTHAT (\i: Int | i >= 0 ) % nonnegative ints
  KInt          = IN 1 .. k % ints in [1, k]
  Maj           = SET KInt SUCHTHAT (\m: Maj | m.size>k/2) % subsets of KInt
% of size > k/2

TYPE P          = [D, seqno: N]
  M             = KInt -> P
  S             = SET P

VAR default    : D
  m            := M{* -> P{d := default, seqno := 0}}

APROC Read() -> D = << RET ReadPair().d >>

APROC Write(d) = << VAR i: Int, maj |
% Determines the highest sequence number i, then writes d
% paired with i+1 to some majority maj of the copies.
  i := := ReadPair().seqno;
  DO VAR j :IN maj | m(j).seqno # i+1 => m(j) := P{d := d, seqno := i+1} OD >>

APROC ReadPair() -> P = << VAR s := ReadMajority() |
% Returns a pair with the largest sequence number from some majority of the copies.
  VAR p :IN s | p.seqno = {p' :IN s | | p'.seqno}.max => RET p >>

APROC ReadMajority() -> S = << VAR maj | RET { i :IN maj | | m(i) } >>
% Returns the set of pairs belonging to some majority of the copies.

END MajorityRegister

```

The following is a key invariant for MajorityRegister.

```

FUNC Inv(m: M) -> Bool = RET
  (ALL p :IN m.rng, p' :IN m.rng | p.seqno = p'.seqno ==> p.d = p'.d)
  /\ (EXISTS maj | (ALL i :IN maj, p :IN m.rng | m(i).seqno >= p.seqno))

```

The first conjunct says that any two pairs having the same sequence number also have the same data. The second conjunct says that the highest sequence number appears in some majority of the copies.

The following Spec function is an abstraction function between states of the MajorityRegister module and states of the Register module.

```

FUNC AF(m: M) -> D =
  VAR p :IN m.rng | p.seqno = {p' :IN m.rng | | p'.seqno}.max => RET p.d

```

That is, the abstract register data value is the data component of a copy with the highest sequence number. Again, because of the invariants, there is only one $p.d$ that will be returned.