

# Bounded Buffer

Shankar

September 18, 2014

# Overview

- Bounded-buffer service
  - input functions: `void put(x)`, `Val get()`, `void end()`
  - no output functions
- Implementations using **standard** synchronization constructs
  - locks, condition variables
  - semaphores
- Implementations using **await** synchronization constructs
  - more powerful, convenient
  - mechanical transformation to standard synch constructs
- Reduce blocking (increase parallelism) in implementations
- Cancelling blocked calls to remote systems
  - allows a caller no longer interested in call to retrieve itself

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT

- Fifo bounded buffer of size N
  - input functions: put(x), get(), end()
  - no output functions
- Main
  - buff: sequence of items in buffer
  - ending: true iff end() has been called
  - putBusy: true iff put call ongoing
  - getBusy: true iff get call ongoing
- void mysid.put(x)
  - ic {not ending and no ongoing put call}
  - oc {buff has space}
  - append x to buff; return

- `Val mysid.get()`
  - `ic {not ending and no ongoing get call}`
  - output `rval`
    - `oc {buff has item, rval is buff.head}`
    - `behead buff; return rval`
- `void mysid.end()`
  - `ic {not ending}`
    - `set ending`
  - `oc {true} return`
- Progress assumption
  - `put call returns if buff has space` // uses `putBusy`
  - `get call returns if buff has item` // uses `getBusy`
  - `end call returns` // uses “thread in `mysid.end`”

```
// main
ic {N ≥ 1}
buff ← [];
ending ← false;
putBusy ← false;
getBusy ← false;
return mysid;

input void mysid.put(Val x)
  ic {not ending and
      not putBusy}
  putBusy ← true;
  oc {buff.size < N}
  buff.append(x);
  putBusy ← false;
  return;
```

```
input Val mysid.get()
  ic {not ending and
      not getBusy}
  getBusy ← true;
  output(Val rval)
    oc {buff.size > 0
        and rval = buff[0]}
  buff.remove();
  getBusy ← false;
  return rval;

input void mysid.end()
  ic {not ending}
  ending ← true;
  oc {true}
  return;
```

atomicity assumption {input and output parts}

progress assumption {

// thread in put returns if buffer has space  
(putBusy and buff.size < N) *leads-to* not putBusy;

// thread in get returns if buffer has an item  
(getBusy and buff.size > 0) *leads-to* not getBusy;

// thread in end returns  
(thread u in mysid.end) *leads-to* (not u in mysid.end);

}

}

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT



- `BoundedBufferInverse(N, bb)` // bb: sid of implementation
  - main: buff, ending, putBusy, getBusy
  - output functions: `doPut(x)`, `doGet()`, `doEnd()`
- `doPut(x)`
  - oc {not ending and no ongoing put call}  
bb.put(x)
  - ic {buff has space}  
append x to buff
- `doGet()`
  - oc {not ending and no ongoing get call}  
rval ← bb.get()
  - ic {buff has item, rval is buff.head}  
behead buff

- doEnd()

- oc {not ending}  
set ending; bb.end()
- ic {true} return

- progress condition

- put call returns if buff has space
- get call returns if buff has item
- end call returns

// uses putBusy

// uses getBusy

// uses “thread in **bb**.end”

```
// main
ic {N ≥ 1}
buff ← [];
putBusy ← false;
getBusy ← false;
return mysid;

output doPut(Val x) {
  oc {not ending and
      not putBusy}
  putBusy ← true;
  bb.put(x);
  ic {buff.size < N}
  buff.append(x);
  putBusy ← false;
```

```
output doGet() {
  oc {not ending and not getBusy}
  getBusy ← true;
  Val x ← bb.get(x);
  ic {buff.size > 0 and x = buff[0]}
  buff.remove();
  getBusy ← false;

output doEnd() {
  oc {not ending}
  ending ← true;
  lck.end();
  ic {true}

progress condition {... mysid bb...}
```

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT

- Await: powerful synchronization construct not provided by PLs
  - convenient for writing programs
  - implementable by standard synchronization constructs
- `await (B) S`
  - B is predicate, S is non-blocking code
  - atomically execute S only if B holds, otherwise wait
  - weak or strong fairness
- `await (B) S`: more general than `oc {B} S`
  - S can make (non-blocking) output calls, use return values
- `atomic S`: short for `await (true) S`

- **Await-structured** program
  - awaits are the only synchronization construct
  - code outside awaits does not conflict with code executed by other threads
- Await-structured program
  - easier to understand than equivalent program with standard synchronization constructs
  - can be mechanically transformed to program that uses standard synchronization constructs

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT

```
program BBufAwait(int N) {  
  ia {N ≥ 1}  
  Seq buff ← seq();  
  return mysid;
```

```
input void mysid.put(Val x)  
  • await (buff.size < N)  
    buff.append(x);  
  return;
```

```
input Val mysid.get()  
  • await (buff.size > 0) {  
    Val x ← buff[0];  
    buff.remove();  
    return x;
```

```
input void mysid.end()  
  endSystem();  
  return;
```

atomicity assumption  
awaits

progress assumption  
weak fairness  
for threads



```
■ program Z(int N) {  
    ...  
    bb ← startSystem( BBufferAwait(N) );  
    si ← startSystem( BoundedBufferInverse(N,bb) );  
    ...  
  
    atomicity assumption {}  
    progress assumption {weak fairness}  
}
```

- To not violate `si.doPut.ic`, want  $Inv\ C_0$  to hold

$$C_0 : ((\text{thread in bb.put}) \text{ and } \text{bb.buff.size} < N) \Rightarrow \text{si.buff.size} < N$$

- To not violate `si.doGet.ic`, want  $Inv\ C_1$  to hold

$$C_1 : ((\text{thread at bb.get}) \text{ and } \text{bb.buff.size} > 0) \Rightarrow (\text{si.buff.size} > 0 \text{ and } \text{bb.buff}[0] = \text{si.buff}[0])$$

- Hold because  $Inv\ C_2$  holds (via invariance rule)

$$C_2 : \text{bb.buff} = \text{si.buff}$$

- Want  $B_2$ – $B_4$  to hold

$B_2$  : (putBusy and si.buff.size < N) *leads-to* not putBusy

$B_3$  : (getBusy and si.buff.size > 0) *leads-to* not getBusy

$B_4$  : (thread u in bb.end) *leads-to* (not u in bb.end)

- $B_2$  holds via weak fairness and  $Inv\ C_2$

- only a thread in bb.put can falsify  $B_2.lhs$

- only one such thread at any time

- so it eventually executes, establishing  $B_2.rhs$

// doPut.oc

// wfair

- $B_3$  holds via weak fairness and  $Inv\ C_2$

- $B_3$  holds via weak fairness

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT

- A lock is either **acquired** by a thread or **free**
- **Lock lck:**       // initially free
- **lck.acq():**       // acquire
  - caller must not hold lock
  - atomically acquire lck only if free, o/w wait
- **lck.rel():**       // release
  - caller must hold lock
  - atomically free lck
- **Progress:** a thread at lck.acq() eventually gets past if
  - lock is free continuously                       // wfair; **weak lock**
  - lock is free continuously or repeatedly       // sfair **strong lock**

- `Condition(lck) cv:` // cond var cv associated with lock lck
- `cv.wait():` // always blocks
  - caller must hold lck
  - atomically release lck and wait on cv;  
when awakened: acquire lck; return
- `cv.signal():`
  - caller must hold lck
  - atomically awaken a thread (if any) waiting on cv; return
- Progress: a thread at `cv.wait()` eventually gets past if
  - cv is signalled, and no other process is waiting on cv // weak
  - cv is repeatedly signalled // strong

- Await-structured program with distinct await guards  $B_1, \dots, B_N$ 
  - works even if guards are not distinct
- Introduce `lck` and associated  $cv_1, \dots, cv_N$
- Replace `await ( $B_i$ ) S` by

```
lck.acq()
while (not  $B_i$ )
     $cv_i$ .wait()
S
for k in 1, ..., N
    if ( $B_k$ )
         $cv_k$ .signal()
lck.rel()
```
- For more parallelism
  - partition awaits into “non-conflicting” groups
  - use separate lock for each group

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT



```
program BBufLockCv(int N) {  
  ia {N ≥ 1}  
  Seq buff ← [];  
  Lock lck;                                // protects buffer  
  Condition(lck) cvItem;                  // signaled when buffer not empty  
  Condition(lck) cvSpace;                 // signaled when buffer not full  
  return mysid;  
  
input void mysid.put(Val x)  
  lck.acq();                               // Note: no '•'  
  while (buff.size = N)  
    • cvSpace.wait();  
  buff.append(x);  
  cvItem.signal();  
  lck.rel();  
  return;
```

```
input Val mysid.get()
  lck.acq();                      // Note: no '•'
  while (buff.size = 0)
    • cvItem.wait();
  Val x ← buff[0];
  buff.remove();
  cvSpace.signal();
  lck.rel();
  return x;
```

```
input void mysid.end()
  endSystem();
  return;
```

```
atomicity assumption {lck, cvItem, cvSpace}
progress assumption {weak fairness for threads}
```

```
}
```

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT

- Combines mutual exclusion + conditional wait
- Counting semaphores
  - Semaphore(N) sem: sem initialized to  $N \geq 0$
  - sem.P(): atomically  $\text{sem} \leftarrow \text{sem} - 1$  only if  $\text{sem} > 0$ , o/w wait
  - sem.V(): atomically  $\text{sem} \leftarrow \text{sem} + 1$
- Binary semaphores
  - Semaphore(N) sem: sem initialized to  $N$  in  $0..1$
  - sem.P(): atomically  $\text{sem} \leftarrow 0$  only if  $\text{sem} = 1$ , o/w wait
  - sem.V(): atomically  $\text{sem} \leftarrow 1$
- Progress: condition in which a thread at P() eventually gets past
  - $\text{sem} > 0$  holds continuously // wfair; weak sem
  - $\text{sem} > 0$  continuously or intermittently // sfair; strong sem

- Program with locks and condition variables
- For every lock `lck`
  - introduce binary semaphore, say `lckMutex`, initialized to 1
  - `lck.acq()`  $\longrightarrow$  `lckMutex.P()`
  - `lck.rel()`  $\longrightarrow$  `lckMutex.V()`
- For every condition variable `cv` associated with lock `lck`
  - introduce binary semaphore, say `cvGate`, initialized to 0
  - `cv.wait()`  $\longrightarrow$  `lckMutex.V()`; `cvGate.P()`; `lckMutex.P()`
  - `cv.signal()`  $\longrightarrow$  `cvGate.V()`
- To have waiting thread come before entering thread
  - skip `lckMutex.P()` after `cvGate.P()`
  - skip `lckMutex.V()` after `cvGate.V()`
  - ...

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT

- Consider an await-structured program
- Define two awaits to be **strongly** nonconflicting if they do not conflict even without atomicity of awaits
- Technique 1 to increasing parallelism
  - partition awaits into non-conflicting groups of awaits
  - use separate locks for the groups
  - to avoid deadlock, obtain locks in increasing order
- Technique 2 to increasing parallelism
  - modify code to increase extent of strongly-nonconflicting awaits
  - duplicate hot spots into separate memory areas
  - loosen coupling between duplicates



- Implement buff as a circular array

- buffA[N]

- in  $\leftarrow$  0

// next put call accesses buffA[in]

- out  $\leftarrow$  0

// next get call accesses buffA[out]

- cnt  $\leftarrow$  0

// # items in buffA

- input void mysid.put(x)

p1:  $\bullet$ await (cnt < N);

p2:  $\bullet$ buffA[in]  $\leftarrow$  x;

await (true)

cnt  $\leftarrow$  cnt+1;

in  $\leftarrow$  mod(in+1,N);

return

- input Val mysid.get(x)

g1:  $\bullet$ await (cnt > 0);

g2:  $\bullet$ x  $\leftarrow$  buffA[out];

await (true)

cnt  $\leftarrow$  cnt-1;

out  $\leftarrow$  mod(out+1,N);

return

- If statements p2 and g2 do not conflict, we can remove their  $\bullet$ 's

- Let  $X$  be  $\text{BBuffPar}$  with statements  $p2$  and  $g2$  replaced by  $\text{skip}$
- Can remove  $\bullet$ 's at  $p2$  and  $g2$  if  $X$  satisfies  $\text{Inv } D_0$   
 $D_0 : (\text{thread at } p2) \text{ and } (\text{thread at } g2) \Rightarrow \text{in} \neq \text{out}$
- $D_1$ – $D_5$  satisfies invariance rule and implies  $D_0$ , //  $\text{Inv } D_0$  holds  
 $D_1 : (\text{at most one thread in put})$   
 $D_2 : (\text{at most one thread in get})$   
 $D_3 : (\text{thread on } p2) \Rightarrow \text{cnt} < N$   
 $D_4 : (\text{thread at } g2) \Rightarrow \text{cnt} > 0$   
 $D_5 : \text{cnt} = \text{mod}(\text{in} - \text{out}, N)$
- So  $\bullet$ 's at  $p2$  and  $g2$  can be removed from  $\text{BBuffPar}$
- Now easy to show  $\text{BBuffPar}$  implements  $\text{BoundedBuffer}$

Bounded-Buffer Service

Bounded-Buffer Service Inverse

Awaits

Bounded-Buffer Implementation using Awaits

Locks and Condition Variables

Bounded-Buffer Implementation using Locks and Condition Variables

Semaphores

Bounded-Buffer Implementation using Semaphores // SEE TEXT

Increasing Parallelism

Canceling Blocked Calls // SEE TEXT