



























```

1  wCQ * LHead = <empty wCQ>, LTail = LHead;
2  void Enqueue_Unbounded(void * p)
3  wCQ * ltail = hp.protectPtr(HPTail, Load(&LTail));
4  // Enqueue_Ptr() returns false if full or finalized
5  if (!ltail->next and ltail->Enqueue_Ptr(p, finalize=True))
6  | hp.clear();
7  | return;
8  wCQ * cq = alloc_wCQ(); // Allocate wCQ
9  cq->init_wCQ(p); // Initialize & put p
10 enqueueers[TID] = cq; // == Slow path (CRTurn) ==
11 for i = 0 .. NUM_THRDS-1 do
12 | if (enqueueers[TID] = null)
13 | | hp.clear();
14 | | return;
15 wCQ * ltail = hp.protectPtr(HPTail, Load(&LTail));
16 if (ltail ≠ Load(&LTail)) continue;
17 if (enqueueers[ltail->enqTid] = ltail)
18 | CAS(&enqueueers[ltail->enqTid], ltail, null);
19 for j = 1 .. NUM_THRDS do
20 | cq = enqueueers[(j + ltail->enqTid) mod NUM_THRDS];
21 | if (cq = null) continue;
22 | finalize_wCQ(ltail); // Duplicate finalize is OK since
23 | CAS(&ltail->next, null, cq); // cq or another node follows
24 | break;
25 wCQ * lnext = Load(&ltail->next);
26 if (lnext ≠ null)
27 | finalize_wCQ(ltail); // Duplicate finalize is OK since
28 | CAS(&LTail, ltail, lnext); // lnext or another node follows
29 enqueueers[TID] = null;
30 hp.clear();
31 bool dequeue_rollback(wCQ * prReq, wCQ * myReq)
32 deqself[TID] = prReq;
33 giveUp(myReq, TID);
34 if (deqhelp[TID] ≠ myReq)
35 | deqself[TID] = myReq;
36 | return False;
37 hp.clear();
38 | return True;
39 void finalize_wCQ(wCQ * cq)
40 | OR(&cq->Tail, { .Value=0, .Finalize=1 });
41 void * Dequeue_Unbounded()
42 wCQ * lhead = hp.protectPtr(HPHead, Load(&LHead));
43 // skip_last modifies the default behavior for Dequeue on
44 // the last element in aq, as described in the text.
45 void * p = lhead->Dequeue_Ptr(skip_last=True);
46 if (p ≠ last)
47 | if (p ≠ null or lhead->next = null)
48 | | hp.clear();
49 | | return p;
50 wCQ * prReq = deqself[TID]; // == Slow path (CRTurn) ==
51 wCQ * myReq = deqhelp[TID];
52 deqself[TID] = myReq;
53 for i = 0 .. NUM_THRDS-1 do
54 | if (deqhelp[TID] != myReq) break;
55 wCQ * lhead = hp.protectPtr(HPHead, Load(&LHead));
56 if (lhead ≠ Load(&LHead)) continue;
57 void * p = lhead->Dequeue_Ptr(skip_last=True);
58 if (p ≠ last)
59 | if (p ≠ null or lhead->next = null)
60 | | if (!dequeue_rollback(prReq, myReq)) break;
61 | | return p;
62 | Store(&lhead->aq.Threshold, 3n - 1);
63 | p = lhead->Dequeue_Ptr(skip_last=True);
64 | if (p ≠ last and p ≠ null)
65 | | if (!dequeue_rollback(prReq, myReq)) break;
66 | | return p;
67 wCQ * lnext = hp.protectPtr(HPNext, Load(&lhead->next));
68 if (lhead ≠ Load(&LHead)) continue;
69 if (searchNext(lhead, lnext) ≠ NOIDX) casDeqAndHead(lhead, lnext);
70 wCQ * myCQ = deqhelp[TID];
71 wCQ * lhead = hp.protectPtr(HPHead, Load(&LHead));
72 if (lhead = Load(&LHead) and myCQ = Load(&lhead->next))
73 | CAS(&LHead, lhead, myCQ);
74 hp.clear();
75 hp.retire(prReq);
76 return myCQ->Locate_Last_Ptr();

```

Figure 13: Adapting CRTurn to an unbounded wCQ-based queue design (high-level methods).

## A APPENDIX: UNBOUNDED QUEUE

LSCQ and LCRQ implement unbounded queues by using an outer layer of M&S lock-free queue which links ring buffers together. Since operations on the outer layer are very rare, the cost is dominated by ring buffer operations. wCQ can follow the same idea.

Although the outer layer does not have to be performant, it still must be wait-free with bounded memory usage. However, M&S queue is only lock-free. The (non-performant) CRTurn wait-free queue [37, 39] does satisfy the aforementioned requirements. Moreover, CRTurn already implements wait-free memory reclamation by using hazard pointers in a special way. wCQ and CRTurn combined together would yield a fast queue with bounded memory usage.

Because CRTurn’s design is non-trivial and is completely orthogonal to the wCQ presentation, its discussion is beyond the scope of this paper. We refer the reader to [37, 39] for more details about CRTurn. Below, we sketch expected changes to CRTurn (assuming prior knowledge of CRTurn) to link ring buffers rather than individual nodes. In Figure 13, we present pseudocode (assuming that entries are pointers) with corresponding changes to enqueue and dequeue operations in CRTurn. For convenience, we retain the same variable and function names as in [37] (e.g., *giveUp* that is not

shown here). Similar to [37], we assume memory reclamation API based on hazard pointers (the *hp* object).

The high-level idea is that we create a wait-free queue (list) of wait-free ring buffers, where *LHead* and *LTail* represent corresponding head and tail of the list. *Enqueue\_Unbounded* will first attempt to insert an entry to the last ring buffers as long as *LTail* is already pointing to the last ring buffer. Otherwise, it allocates a new ring buffer and inserts a new element. It then follows CRTurn’s procedure to insert the ring buffer to the list. The only difference is that when helping to insert the new ring buffer, threads will make sure that the previous ring buffer is finalized.

*Dequeue\_Unbounded* will first attempt to fetch an element from the first ring buffer. wCQ’s *Dequeue* for *aq* needs to be modified to detect the very last entry in a *finalized* ring buffer. (Note that it can only be done for finalized ring buffers, where no subsequent entries can be inserted.) Instead of returning the true entry, *Dequeue\_Ptr* returns a special *last* value. This approach helps to retain CRTurn’s wait-freedom properties as every single ring buffer contains at least one entry. Helper methods must also be modified accordingly.