Poor Man's URCU

Pedro Ramalhete
Cisco Systems
pramalhe@gmail.com

Andreia Correia
Concurrency Freaks
andreiacraveiroramalhete@gmail.com

Abstract

RCU is, among other things, a well known mechanism for memory reclamation that is meant to be used in languages without an automatic Garbage Collector, unfortunately, it requires operating system support, which is currently provided only in Linux. An alternative is to use Userspace RCU (URCU) which has two variants that can be deployed on other operating systems, named *Memory Barrier* and *Bullet Proof*.

We present a novel algorithm that implements the three core APIs of RCU: `rcu_read_lock()`, `rcu_read_unlock()`, and `synchronize_rcu()`. Our algorithm uses one mutual exclusion lock and two reader-writer locks with `trylock()` capabilities, which means it does not need a language with a memory model or atomics API, and as such, it can be easily implemented in almost any language, regardless of the underlying CPU architecture, or operating system.

Keywords RCU, locks

1. Introduction

When implementing lock-free and wait-free data structures in languages without an automatic Garbage Collector (GC), or just to manage object lifetime, some kind of memory reclamation technique is required. Several such techniques have been discovered, with RCU (McKenney et al. 2001) one of the most well known. RCU has two interesting properties for its non-reclaiming methods, i.e. `rcu_read_lock()` and `rcu_read_unlock()`, one being its wait-free progress condition, and the other its low impact on latency with high scalability. RCU is unfortunately not a generic technique because it requires operating system support.

Fortunately, Userspace RCU (URCU) (Desnoyers et al. 2012) has two variants that are generic enough to be implemented regardless of operating system support, named Memory Barrier and Bullet Proof. Even so, these two variants of URCU are currently implemented in C99 and although many architectures are supported, the support is not universal, and other languages may have difficulty linking with the library.

We present a new algorithm which we named Poor Man's URCU which can be easily implemented in any language, as long as the language provides a mutual exclusion lock and a reader-writer lock with `trylock()` functionality. This means our algorithm can be implemented across a large number of languages, without need for a memory model, or atomics API, and regardless of the operating system on which it runs. Our algorithm implements only the core functionality of RCU, namely `rcu_read_lock()`, `rcu_read_unlock()`, and `synchronize_rcu()`, and it does so providing only lock-free progress for the first two, while RCU and URCU typically provide wait-free progress.

2. Algorithm

Poor Man's URCU is composed of one mutual exclusion lock and two reader-writer lock instances. Threads that call `rcu_read_lock()` and `rcu_read_unlock()` (readers) will attempt to acquire the reader lock on the first of the reader-writer locks (`rwlock1`) with a `trylock()`, and if it fails, do a `trylock()` on the second reader-writer lock (`rwlock2`), and if that also fails, try again the first rwlock, and so on. This procedure can go on indefinitely, but it will fail only if there is another thread making progress, namely, a thread acquiring the write-lock in one of the reader-writer locks, which means the operation is lock-free.

A thread that calls `synchronize_rcu()` (updater) will start by acquiring the lock on the mutual exclusion lock to guarantee only one updater at a time is present. The updater will then acquire the write-lock on the second reader-writer lock and release it as soon as it has obtained it, followed by acquiring the write-lock on the first reader-writer lock and release it as soon as it has obtained it. This is enough to introduce a linearization point, and it is now safe for the updater to free/delete the memory of the object that is meant to be reclaimed, because it has the guarantee that any ongoing readers can no longer hold a pointer to such object.

Algorithm 1 shows the implementation in C99 of *Poor Man's URCU* using the Pthreads library API. Notice that

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PPoPP '17 February 04-08, 2017, Austin, TX, USA
© 2017 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-4493-7/17/02…$15.00
DOI: http://dx.doi.org/10.1145/3018743.3019021

Algorithm 1 Pthread implementation of Poor Man's URCU

```
1  typedef struct {
2    pthread_mutex_t  updaterMutex;
3    pthread_rwlock_t rwlock1;
4    pthread_rwlock_t rwlock2;
5  } poorman_urcu_t;
6
7  // The return value should be passed as arg 'whichone'
8  int rcu_read_lock(poorman_urcu_t* self) {
9    while (1) {
10     if (pthread_rwlock_tryrdlock(&self->rwlock1) == 0) {
11       return 1;
12     }
13     if (pthread_rwlock_tryrdlock(&self->rwlock2) == 0) {
14       return 2;
15     }
16   }
17 }
18
19 void rcu_read_unlock(poorman_urcu_t* self, int whichone) {
20   if (whichone == 1) {
21     pthread_rwlock_unlock(&self->rwlock1);
22   } else {    // whichone == 2
23     pthread_rwlock_unlock(&self->rwlock2);
24   }
25 }
26
27 void synchronize_rcu(poorman_urcu_t* self) {
28   pthread_mutex_lock(&self->updaterMutex);
29   pthread_rwlock_wrlock(&self->rwlock2);
30   pthread_rwlock_unlock(&self->rwlock2);
31   pthread_rwlock_wrlock(&self->rwlock1);
32   pthread_rwlock_unlock(&self->rwlock1);
33   pthread_mutex_unlock(&self->updaterMutex);
34 }
```

in this implementation the `rcu_read_lock()` API returns which reader-writer lock was acquired in the form of an integer, which is then passed to `rcu_read_unlock()` as an argument, so that it knows which reader-writer lock to unlock, but an alternative implementation can be done using a thread-local variable to pass this information.

In our implementation the order of acquiring the lock on `synchronize_rcu()` is first on `rwlock2` then on `rwlock1`, but it would be equally correct to do the opposite order.

3. Benchmarks and Discussion

Figure 1 shows the median for 5 runs of the number of operations per second, on a microbenchmark that iterates for 20 seconds per run. Each read-iteration calls `rcu_read_lock()`, reads all items in an array of 100 user-defined objects (read-critical section), then calls `rcu_read_unlock()`. Each update-iteration creates a new user-defined object, then it replaces a random object in the array of 100 using `compare_exchange_strong()`, calls `synchronize_rcu()`, and then it safely `frees` the memory of the previous instance. The benchmark was ran on a PowerPC Power8E with 8 cores, running Ubuntu Linux 14.04 64 bit, with GCC 4.9.2. The Memory Barrier implementation of URCU requires prior thread registration, but the Bullet Proof does not, therefore, in order to have a fair comparison with Poor Man's URCU which also requires no registration of threads, we compared only against Bullet Proof.

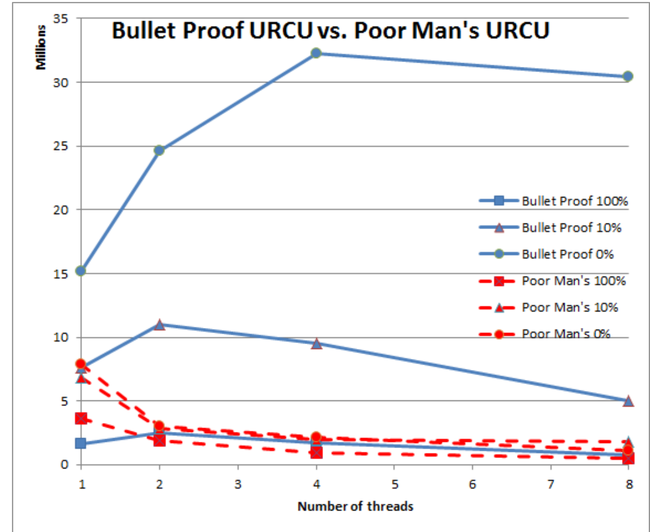

Figure 1. The plot above shows the total number of operations as a function of the number of threads, on a PowerPC machine with 8 cores. The 100% lines means that each thread would only call `synchronize_rcu()`. The 10% lines means that each thread would call `synchronize_rcu()` 10% of the times, and the remaining 90% of the times it would call `rcu_read_lock()` followed by `rcu_read_unlock()`.

Even though there is a performance advantage of the Bullet Proof algorithm, the Poor Man's implementation is not that far behind and it surpasses in the special case of 1 thread at 100%, when the thread is doing only calls to `synchronize_rcu()`. We should keep in mind that the performance of Poor Man's URCU is tightly coupled to the underlying implementation of the Reader-Writer Locks used, and recent advances in this area (Calciu et al. 2013) should provide significant improvement in throughput.

Although not particularly efficient, and with a limited API, our *Poor Man's URCU* algorithm is highly portable across Operating Systems and programming languages, does not require linking with an external library, does not require kernel/OS support, and has no licensing constraint. The simplicity and ease-of-use of our algorithm can foster the usage of RCU in domains where before it was impractical.

References

I. Calciu, D. Dice, Y. Lev, V. Luchangco, V. J. Marathe, and N. Shavit. NUMA-aware reader-writer locks. *PPoPP 2013*, 2013.

M. Desnoyers, P. E. McKenney, A. S. Stern, M. R. Dagenais, and J. Walpole. User-level implementations of read-copy update. *Parallel and Distributed Systems, IEEE Transactions on*, 23(2): 375–382, 2012.

P. E. McKenney, J. Appavoo, A. Kleen, O. Krieger, R. Russell, D. Sarma, and M. Soni. Read-copy update. In *AUUG Conference Proceedings*, page 175. AUUG, Inc., 2001.