

Exercises week 6 Friday 13 October 2017

Goal of the exercises

The goal of this week's exercises is to make sure that you can achieve good performance and scalability of lock-based concurrent software, using finer-grained locks, lock striping, the Java class library's atomically updatable numbers, immutability and the visibility effects of volatiles and atomics.

Do this first

Get and unpack this week's example code in zip file `pcpp-week06.zip` on the course homepage.

File `TestStripedMap.java` contains implementations of several thread-safe hash map classes:

- (A) A complete implementation of `SynchronizedMap<K,V>` which follows the Java monitor pattern: all mutable fields are private, all public methods are synchronized, and no internal data structures escape.
- (B) A partial implementation of `StripedMap<K,V>` which does not follow the Java monitor pattern, but divides the buckets table into stripes, locking each stripe both on read and write accesses. This is the subject of Exercise 6.1.
- (C) A partial implementation of `StripedWriteMap<K,V>` which also divides the buckets table into stripes, but locks each stripe only on write accesses. Read accesses do not take locks at all, but their correctness is assured — we hope — by (1) working on immutable item nodes, and (2) ensuring visibility of writes by careful use of atomics and volatiles. This is the subject of Exercise 6.2.
- (D) A simple wrapper `WrapConcurrentHashMap<K,V>` around Java's `ConcurrentHashMap<K,V>`, for comparison.

Exercise 6.1 The `SynchronizedMap<K,V>` implementation scales (and therefore performs) poorly on a multicore computer because of all the locking: only one thread at a time can read or write the hash map.

The lecture showed that scalability can be considerably improved by *lock striping*. Instead of locking on the entire table of buckets, one divides it into a number of stripes (here 32), and locks only the single stripe that is going to be read or updated.

This is the idea in the `StripedMap<K,V>` class, whose implementation in file `TestStripedMap.java` contain only methods `containsKey` and `put` and some auxiliary methods.

Your task below is to implement the remaining public methods, as described by interface `OurMap<K,V>`. They are very similar to the method implementations in class `SynchronizedMap<K,V>`, except that they do not lock the entire hash map, only the relevant stripe.

1. Implement method `V get(K k)` using lock striping. It is similar to `containsKey`, but returns the value associated with key `k` if it is in the map, otherwise `null`. It should use the `ItemNode.search` auxiliary method.
2. Implement method `int size()` using lock striping; it should return the total number of entries in the hash map. The size of stripe `s` is maintained in `sizes[s]`, so the `size()` method could simply compute the sum of these values, locking each stripe in turn before accessing its value.
Explain why it is important to lock stripe `s` when reading its size from `sizes[s]`.
3. Implement method `V putIfAbsent(K k, V v)` using lock striping. It must work as in Java's `ConcurrentHashMap`, that is, atomically do the following: check whether key `k` is already in the map; if it is, return the associated value; and if it is not, add `(k, v)` to the map and return `null`. The implementation can be similar to `putIfAbsent` in class `SynchronizedMap<K,V>` but should of course only lock on the stripe that will hold key `k`. It should use the `ItemNode.search` auxiliary method. Remember to increment the relevant `sizes[stripe]` count if any entry was added. Ignore `reallocateBuckets` for now.
4. Extend method `putIfAbsent` to call `reallocateBuckets` when the bucket lists grow too long. For simplicity, you can test whether the size of a stripe is greater than the number of buckets divided by the number of stripes, as in method `put`.

5. Implement method `V remove(K k)` using lock striping. Again very similar to `SynchronizedMap<K,V>`. Remember to decrement the relevant `sizes[stripe]` count if any entry was removed.
6. Implement method `void forEach(Consumer<K,V> consumer)`. Apparently, this may be implemented in two ways: either (1) iterate through the buckets as in the `SynchronizedMap<K,V>` implementation; or (2) iterate over the stripes, and for each stripe iterate over the buckets that belong to that stripe.

It seems that (1) requires locking on all stripes before iterating over the buckets. Otherwise an intervening `reallocateBuckets` call on another thread may replace the `buckets` array with one of a different size between observing the length of the array and accessing its elements. It does not work to just obtain a copy `bs` of the `buckets` field, because `reallocateBuckets` destructively redistributes item node lists from the old buckets array to the new one.

It seems that (2) can be implemented by locking only stripe-wise and then iterating over the stripe's buckets. While holding the lock on a stripe, no `reallocateBuckets` can happen.

Explain what version you have implemented and why.

7. You may use method `testMap(map)` for very basic single-threaded functional testing while making the above method implementations. See how to call it in method `testAllMaps`. To actually enable the assert statements, run with the `-ea` option:

```
java -ea TestStripedMap
```

8. Measure the performance of `SynchronizedMap<K,V>` and `StripedMap<K,V>` by timing calls to method `exerciseMap`. Report the results from your hardware and discuss whether they are as expected.
9. What advantages are there to using a small number (say 32 or 16) of stripes instead of simply having a stripe for each entry in the buckets table? Discuss.
10. Why can using 32 stripes improve performance even if one never runs more than, say, 16 threads? Discuss.
11. A comment in the example source code says that it is important for thread-safety of `StripedMap` and `StripedWriteMap` that the number of buckets is a multiple of the number of stripes.

Give a scenario that demonstrates the lack of thread-safety when the number of buckets is not a multiple of the number of stripes. For instance, use 3 buckets and 2 stripes, and consider two concurrent calls `put(k1, v1)` and `put(k2, v2)` where the hash code of `k1` is 5 and the hashcode of `k2` is 8.

Note that method `reallocateBuckets` has been provided for you. Its auxiliary method `lockAllAndThen` uses recursion to take all the stripe lock; this is the only way in Java to take a variable number of intrinsic locks.

Exercise 6.2 The striped hash map in class `StripedMap<K,V>` scales better with more threads than the `SynchronizedMap<K,V>` hash map. However, it can be further improved by locking a stripe only when writing, not when reading, so that many reads can proceed concurrently without locking. This idea is outlined in class `StripedWriteMap<K,V>`, which is a somewhat subtle undertaking, based on several ideas that are different from both `SynchronizedMap<K,V>` and `StripedMap<K,V>`.

First, the item nodes are made immutable, all fields of class `ItemNode<K,V>` are `final`. That means that as soon as a read access (`containsKey`, `get` or `forEach`) has obtained a reference to a list of item nodes in a bucket, it need not be concerned with atomicity or visibility: nothing it accesses can be affected by other threads.

Second, the slice sizes will now be represented by an `AtomicIntegerArray` so that no locking is needed when incrementing and decrementing the stripe sizes. It also ensures that a thread executing the `size()` method can see the increments and decrements made by threads that `put`, `putIfAbsent` and `remove` entries.

Third, the writes to and reads from the `sizes` array are (ab)used to ensure visibility of updates to the `buckets` array. After any write to an element of `buckets`, `sizes` is written also, and before any read of an element of `buckets`, `sizes` is read. This ensures that `containsKey`, `get`, `forEach` will see any writes performed by `put`, `putIfAbsent` and `remove`.

Fourth, making class `ItemNode<K,V>` immutable means that `put` and `remove` may need to copy part of the list of item nodes in a bucket, but those lists should in any case hold at most a few items (otherwise the hash map is slow), the allocation of a new item node is fast, the cost appears to be outweighed but the time saved on not locking, and parts of the code become much neater this way.

Some ideas in `StripedWriteMap<K,V>` are inspired by the implementation of Java's `ConcurrentHashMap`, which however uses many more sophisticated techniques.

1. Implement method `V get(K k)`. It is similar to `containsKey`, but returns the value associated with key `k` if it is in the map, otherwise `null`. It should use the `ItemNode.search` auxiliary method, and for visibility reasons, read the stripe's size before reading the buckets array entry.
2. Implement method `int size()`. This is very straightforward: simply compute the sum of the stripe sizes. Since these are represented in an `AtomicIntegerArray`, all writes are visible to this method's reads; no locking is needed.
3. Implement method `V putIfAbsent(K k, V v)`. You must lock on the relevant stripe. Use auxiliary method `ItemNode.search(bl, k, old)` to determine whether `k` is already in the hash map, where `bl` is the bucket list reference obtained from `buckets[hash]`. If yes, then do nothing; else create a new item node from `k, v` and `bl`, and update the `buckets` table with that. Remember to update the stripe size if an entry was added.

Why do you not need to write to the stripe size if nothing was added?

4. Implement method `V remove(K k)`. Lock on the relevant stripe. Use `ItemNode.delete(bl, k, old)` to delete the entry with key `k`, if any, from bucket list `bl`, and update the buckets table with the result. Remember to update the stripe size if an entry was removed.
5. Implement method `void forEach(Consumer<K,V> consumer)`. Unlike in Exercise 6.1.6 there seems to be no need to lock anything while you iterate over the buckets table and the item node lists, since the latter are immutable and assignment to the buckets array entries are atomic. But to protect against the redistribution effect of an intervening `reallocateBuckets`, you must either read the `buckets` reference into a local variable `bs` as the first thing and then refer to `bs` during the iteration over buckets or stripes; or iterate stripewise and read the `buckets` reference into a local variable `bs` before iterating over the stripe's buckets. Additionally, to ensure visibility of writes to the buckets array entries, you must read the stripe's size before iterating over its buckets.
6. Measure the performance of `SynchronizedMap<K,V>`, `StripedMap<K,V>`, `StripedWriteMap<K,V>` and `WrapConcurrentHashMap<K,V>` using method `exerciseAllMaps`. Report the results and discuss whether they are as expected.
7. (Optional, only really interesting if you have access to a computer with many cores) Measure the scalability of the four hash map implementations by running method `timeAllMaps`. Report the results, in tabular or graphical form, and discuss the results.

Exercise 6.3 File `TestLongAdders.java` contains several implementations of a long (64-bit) integer with thread-safe `add` and `get` operations: (a) Java's `AtomicLong`; (b) Java 8's `LongAdder`; (c) a simple `long` field with synchronized operations; (d) a number represented as the sum of multiple "stripes" densely allocated in an `AtomicLongArray`; and (e) a number represented as the sum of multiple "stripes" allocated as scattered `AtomicLong` objects.

1. Compile the file and run the code to measure, on your own hardware, the performance of the various atomic long implementations. Report the numbers and discuss whether they are plausible, eg. relative to the number of cores in your machine and the number of threads trying to access the thread-safe long integer.
2. Create a new class `NewLongAdderLessPadded` as a variant of the `NewLongAdderPadded` class, where you remove the strange `new Object()` creations in the constructor. Create a suitable version of the method `exerciseNewLongAdderPadded` where you measure the performance of this new class along with the others.

Do those `new Object()` allocations make any difference, positive or negative, on your version of the Java Virtual Machine, and your hardware?