# Java Concurrent Collections: Designing a Memoizer with ConcurrentHashMap



Douglas C. Schmidt <u>d.schmidt@vanderbilt.edu</u> www.dre.vanderbilt.edu/~schmidt

> Institute for Software Integrated Systems Vanderbilt University Nashville, Tennessee, USA



### Learning Objectives in this Lesson

apply(T)

- Understand the capabilities of Java's concurrent collections
- Recognize the capabilities of Java's ConcurrentHashMap & BlockingQueue
- Know how to apply the Java Concurrent HashMap class to design a "memoizer"





<<Java Interface>>

Memoizer caches function call results & returns cached results for same inputs

## **Overview of Memoizer**

• Memoization is optimization technique used to speed up programs



See en.wikipedia.org/wiki/Memoization

- Memoization is optimization technique used to speed up programs
  - It caches the results of expensive function calls
    - V computeIfAbsent(K key, Function func) {
      - 1. If key doesn't exist in cache perform a long-running function associated w/key & store the resulting value via the key
      - 2. Return value associated with key





- Memoization is optimization technique used to speed up programs
  - It caches the results of expensive function calls
    - V computeIfAbsent(K key, Function func) {
      - 1. If key doesn't exist in cache perform a long-running function associated w/key & store the resulting value via the key
      - 2. Return value associated with key





- Memoization is optimization technique used to speed up programs
  - It caches the results of expensive function calls
    - V computeIfAbsent(K key, Function func) {
      - 1. If key doesn't exist in cache perform a long-running function associated w/key & store the resulting value via the key
      - 2. Return value associated with key





- Memoization is optimization technique used to speed up programs
  - It caches the results of expensive function calls
  - When the same inputs occur again the cached results are simply returned
    - V computeIfAbsent(K key, Function func) {
      - 1. If key already exists in cache return cached value associated w/key



Memoizer



- Memoization is optimization technique used to speed up programs
  - It caches the results of expensive function calls
  - When the same inputs occur again the cached results are simply returned
    - V computeIfAbsent(K key, Function func) {
      - 1. If key already exists in cache return cached value associated w/key





 Memoizer defines a cache that returns a value produced by applying a (longrunning) function to a key

<<Java Interface>>

#### Function<T,R>

apply(T)

ocompose(Function<? super V,? extends T>):Function<V,R>

andThen(Function<? super R,? extends V>):Function<T,V>

Sidentity():Function<T,T>



See <a href="https://www.see.org/linewide-see">PrimeExecutorService/app/src/main/java/vandy/mooc/prime/utils/Memoizer.java</a>

 Memoizer defines a cache that returns a value produced by applying a (longrunning) function to a key





This class is based on "Java Concurrency in Practice" by Brian Goetz et al.



- Memoizer defines a cache that returns a value produced by applying a (longrunning) function to a key
  - A value that's been computed for a key is returned, rather than applying the function to recompute it



#### Function<T,R>

#### apply(T)

compose(Function<? super V,? extends T>):Function<V,R>
 andThen(Function<? super R,? extends V>):Function<T,V>
 <sup>S</sup>identity():Function<T,T>



- Memoizer defines a cache that returns a value produced by applying a (longrunning) function to a key
  - A value that's been computed for a key is returned, rather than applying the function to recompute it
  - A memoizer can be used whenever a Function is expected

Function<Long, Long> func =
 doMemoization

? new Memoizer<>

(PrimeCheckers::isPrime,

new ConcurrentHashMap());

: PrimeCheckers::isPrime;

<<Java Interface>>

Interface>>

#### apply(T)

compose(Function<? super V,? extends T>):Function<V,R>
 andThen(Function<? super R,? extends V>):Function<T,V>
 sidentity():Function<T,T>



new PrimeCallable(randomNumber, func));

See <a href="https://docs/api/java/util/function/Function.html">docs.oracle.com/javase/8/docs/api/java/util/function/Function.html</a>

- Memoizer defines a cache that returns a value produced by applying a (longrunning) function to a key
  - A value that's been computed for a key is returned, rather than applying the function to recompute it
  - A memoizer can be used whenever a Function is expected

Function<Long, Long> func =

? new Memoizer<>

doMemoization -



<<Java Interface>>

compose(Function<? super V,? extends T>):Function<V,R> andThen(Function<? super R,? extends V>):Function<T,V> Sidentity():Function<T,T>





(PrimeCheckers::isPrime,

PrimeCheckers::isPrime;

new PrimeCallable(randomNumber, func)); ...

Use memoizer

- Memoizer defines a cache that returns a value produced by applying a (longrunning) function to a key
  - A value that's been computed for a key is returned, rather than applying the function to recompute it
  - A memoizer can be used whenever a Function is expected

Function<Long, Long> func =
 doMemoization

? new Memoizer<>

(PrimeCheckers::isPrime,

new ConcurrentHashMap());

: PrimeCheckers::isPrime;

Don't use memoizer

<<Java Interface>>

#### Function<T,R>

#### apply(T)

compose(Function<? super V,? extends T>):Function<V,R>
 andThen(Function<? super R,? extends V>):Function<T,V>
 sidentity():Function<T,T>



new PrimeCallable(randomNumber, func)); ...

- Memoizer defines a cache that returns a value produced by applying a (longrunning) function to a key
  - A value that's been computed for a key is returned, rather than applying the function to recompute it
  - A memoizer can be used whenever a Function is expected

Function<Long, Long> func =

doMemoization

? new Memoizer<>

(PrimeCheckers::isPrime,

new ConcurrentHashMap());

: PrimeCheckers::isPrime;

<<Java Interface>>

#### Function<T,R>

#### apply(T)

compose(Function<? super V,? extends T>):Function<V,R>
 andThen(Function<? super R,? extends V>):Function<T,V>
 <sup>s</sup>identity():Function<T,T>



func is identical, regardless of which branch is chosen

new PrimeCallable(randomNumber, func)); ...

See upcoming lesson on "*Java ExecutorCompletion Service: Application to PrimeChecker App*"

 Memoizer uses a ConcurrentHashMap to minimize synchronization overhead

<<Java Interface>>

#### Function<T,R>

apply(T)

ocompose(Function<? super V,? extends T>):Function<V,R>

andThen(Function<? super R,? extends V>):Function<T,V>

identity():Function<T,T>



See docs.oracle.com/javase/8/docs/api/java/util/concurrent/ConcurrentHashMap.html

- Memoizer uses a ConcurrentHashMap to minimize synchronization overhead
  - A group of locks guard different subsets of the hash buckets



#### ConcurrentHashMap

<<Java Interface>>

apply(T)

compose(Function<? super V,? extends T>):Function<V,R>
 andThen(Function<? super R,? extends V>):Function<T,V>
 <sup>S</sup>identity():Function<T,T>



See <a href="https://www.ibm.com/developerworks/java/library/j-jtp08223">www.ibm.com/developerworks/java/library/j-jtp08223</a>

- Memoizer uses a ConcurrentHashMap to minimize synchronization overhead
  - A group of locks guard different subsets of the hash buckets



In contrast, a SynchronizedMap uses a single lock

#### SynchronizedMap

# Function<T,R> apply(T) compose(Function<? super V,? extends T>):Function<V,R> andThen(Function<? super R,? extends V>):Function<T,V> <sup>S</sup>identity():Function<T,T>

<<Java Interface>>



See <a href="codepumpkin.com/hashtable-vs-synchronizedmap-vs-concurrenthashmap">codepumpkin.com/hashtable-vs-synchronizedmap-vs-concurrenthashmap</a>

- Memoizer uses a ConcurrentHashMap to minimize synchronization overhead
  - A group of locks guard different subsets of the hash buckets
  - apply() uses computeIfAbsent() to ensure a function only runs when key/value pair is added to cache



<<Java Interface>>

apply(T)

ocompose(Function<? super V,? extends T>):Function<V,R>

- andThen(Function<? super R,? extends V>):Function<T,V>
- Sidentity():Function<T,T>



See <u>docs.oracle.com/javase/8/docs/api/java/util/concurrent/</u> <u>ConcurrentHashMap.html#computeIfAbsent</u>

- Memoizer uses a ConcurrentHashMap to minimize synchronization overhead
  - A group of locks guard different subsets of the hash buckets
  - apply() uses computeIfAbsent() to ensure a function only runs when key/value pair is added to cache, e.g.
    - This method implements "atomic check-then-act" semantics

```
return map.computeIfAbsent
  (key,
```

```
k -> mappingFunc(k));
```

<<Java Interface>>

#### apply(T)

compose(Function<? super V,? extends T>):Function<V,R>
 andThen(Function<? super R,? extends V>):Function<T,V>
 <sup>S</sup>identity():Function<T,T>



See <a href="mailto:dig.cs.illinois.edu/papers/checkThenAct.pdf">dig.cs.illinois.edu/papers/checkThenAct.pdf</a>

- Memoizer uses a ConcurrentHashMap to minimize synchronization overhead
  - A group of locks guard different subsets of the hash buckets
  - apply() uses computeIfAbsent() to ensure a function only runs when key/value pair is added to cache, e.g.
    - This method implements "atomic check-then-act" semantics
    - Here's the equivalent sequence of non-atomic/-optimized Java code

```
V value = map.get(key);
if (value == null) {
  value = mappingFunc.apply(key);
  if (value != null) map.put(key, value);
}
return value;
```

See dig.cs.illinois.edu/papers/checkThenAct.pdf

### <<Java Interface>>

#### apply(T)

- compose(Function<? super V,? extends T>):Function<V,R>
   andThen(Function<? super R,? extends V>):Function<T,V>
- Sidentity():Function<T,T>



- Memoizer uses a ConcurrentHashMap to minimize synchronization overhead
  - A group of locks guard different subsets of the hash buckets
  - apply() uses computeIfAbsent() to ensure a function only runs when key/value pair is added to cache



Only one computation per key is performed even if multiple threads simultaneously call computeIfAbsent() using the same key End of Java Concurrent Collections: Designing a Memoizer with ConcurrentHashMap