



**Jfokus**

**Kodewerk**  
Java™ Performance Services



# Concurrency and High Performance Reloaded

# Me

🚀 Work as independent (a.k.a. freelancer)

- performance tuning services

- benchmarking

- Java performance tuning course

🚀 [www.javaperformancetuning.com](http://www.javaperformancetuning.com)

🚀 [www.theserverside.com](http://www.theserverside.com)

🚀 Nominated Sun Java Champion

🚀 Other stuff



single-threaded, single-core



**Kodewerk**  
Java<sup>™</sup> Performance Services





how did we get better performance?



**Kodewerk**  
Java<sup>™</sup> Performance Services



concurrent programming is the norm

**Kodewerk**  
Java<sup>™</sup> Performance Services



sharing adds latency

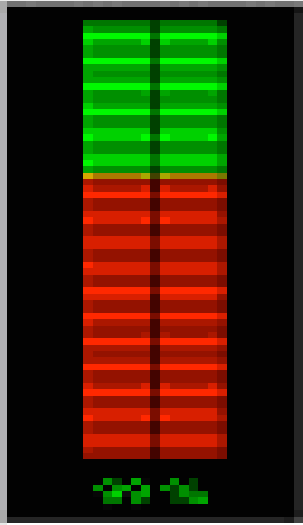


Kodewerk  
Java<sup>™</sup> Performance Services

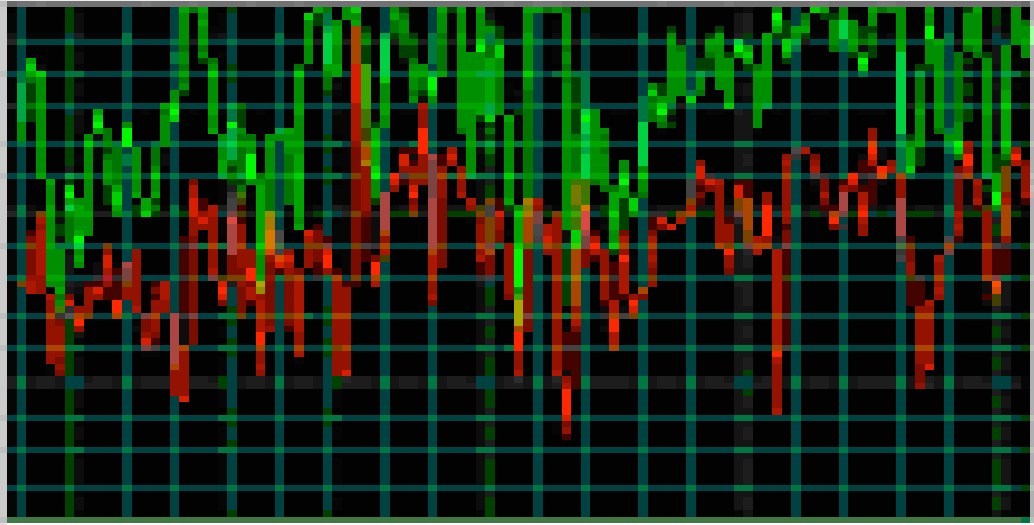




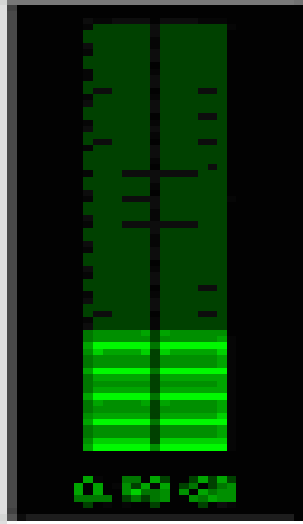
CPU Usage



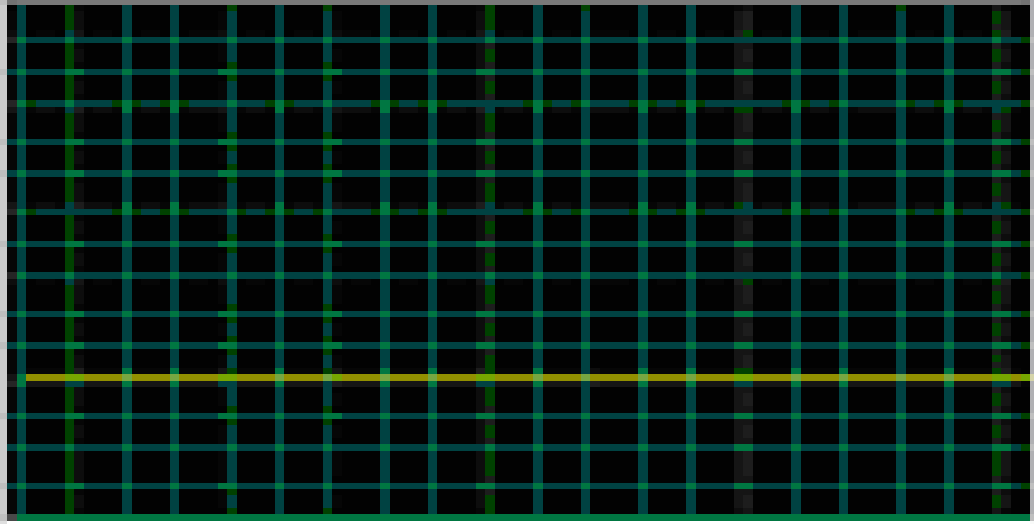
CPU Usage History



FP Usage



Page File Usage History



A white outline map of Sweden is centered on a black background. The text "multi-core is a fact of life!" is overlaid on the map.

multi-core is a fact of life!

A blue square logo is positioned to the left of the company name.

**Kodewerk**  
Java<sup>™</sup> Performance Services



we need to “deliver twice as much  
concurrency every 18 months”



hardware components are notsharable



access to shared data must be serialized



databases offer access to shared data



serialization limits scalability

🦋 Maths to explain relationship between serialized execution and processor utilization

$$\frac{1}{F + \frac{(1 - F)}{N}}$$

- F -> 0 number of utilized CPU -> N
- F -> 1 number of utilized CPU -> 1

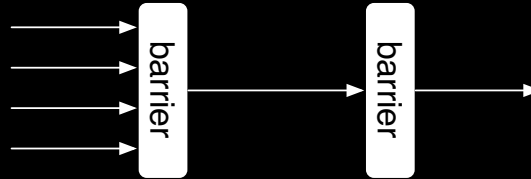
**Amdahl's Law**





serialization limits throughput

🦋 Maths explaining the relationship between locking and throughput



$$\lambda = 1 / \mu$$

$$\mu = 10\text{ms}, \lambda = 100 \text{ tps}$$

$$\mu = 100\text{ms}, \lambda = 10 \text{ tps}$$

**Little's Law**



locking is pessimistic



Kodewerk  
Java<sup>™</sup> Performance Services



getting better concurrency in the JVM

**Kodewerk**  
Java<sup>™</sup> Performance Services



# Java and system level locks

```
StringBuffer sb = new StringBuffer();
```

```
sb.append("a");
```

```
sb.append("b");
```

```
sb.append("c");
```

```
...
```

# Lock Coarsening

**Kodewerk**  
Java™ Performance Services

```
StringBuffer sb = new StringBuffer();
```

```
sb.append("a");
```

```
sb.append("b");
```

```
sb.append("c");
```

```
...
```

# Lock Coarsening

**Kodewerk**  
Java™ Performance Services

```
{
```

```
StringBuffer sb = new StringBuffer();
```

```
sb.append("a");
```

```
sb.append("b");
```

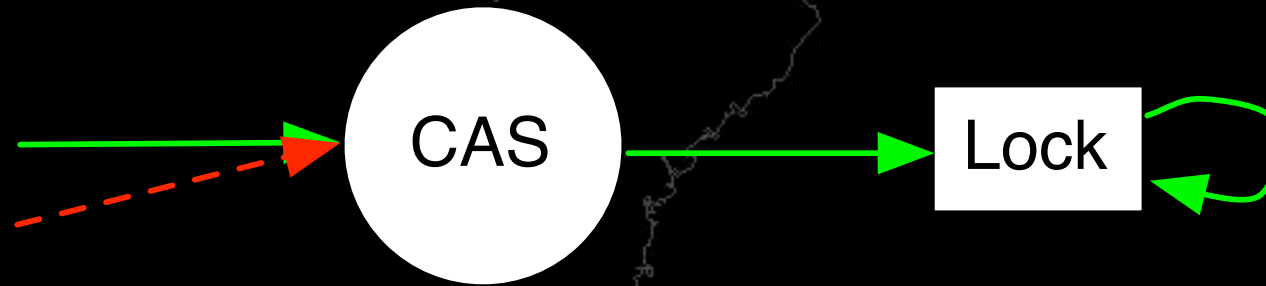
```
sb.append("c");
```

```
}
```

# Lock Elision

**Kodewerk**  
Java™ Performance Services





# Biased Locking



do these optimizations work?

**Kodewerk**  
Java<sup>™</sup> Performance Services



- A EliminateLocks
- B UseBiasedLocking (working)
- C UseBiasedLocking (not working)
- D EliminateLocks with UseBiasedLocking
- E DoEscapeAnalysis
- F EliminateLocks with UseBiasedLocking and DoEscapeAnalysis

■ StringBuffer   ■ StringBuilder





techniques we can use



**Kodewerk**  
Java<sup>™</sup> Performance Services



# Atomics to reduce lock contention

```
private int counter = 0;

Runnable mutator = new Runnable() {
    public void run() {
        long localCount = 0;
        while ( running ) {
            counter++;
            counter--;
            localCount++;
        }
        addToTotalCount( localCount);
    }
};
```

Baseline



Volatile



Synchronized



Lock



Atomic





# Lock striping



Thread ▶

lock

lock

lock



# ConcurrentHashMap

HashMap (no sync)



HashMap (sync)



HashTable



ConcurrentHashMap



**Blackboard**



teaching threads to steal

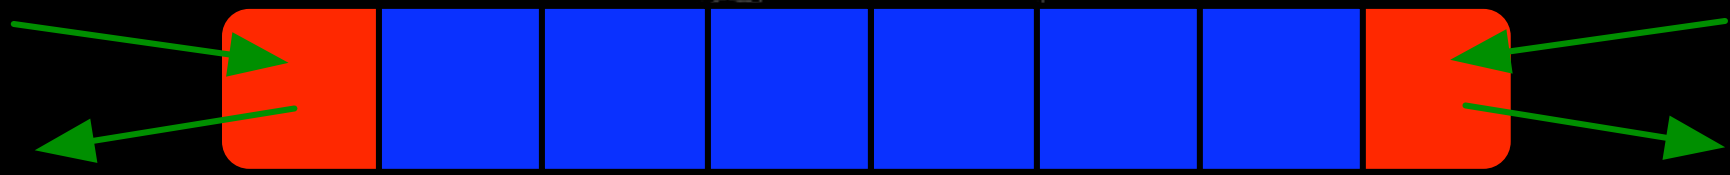
**Kodewerk**  
Java<sup>™</sup> Performance Services



Fork-Join

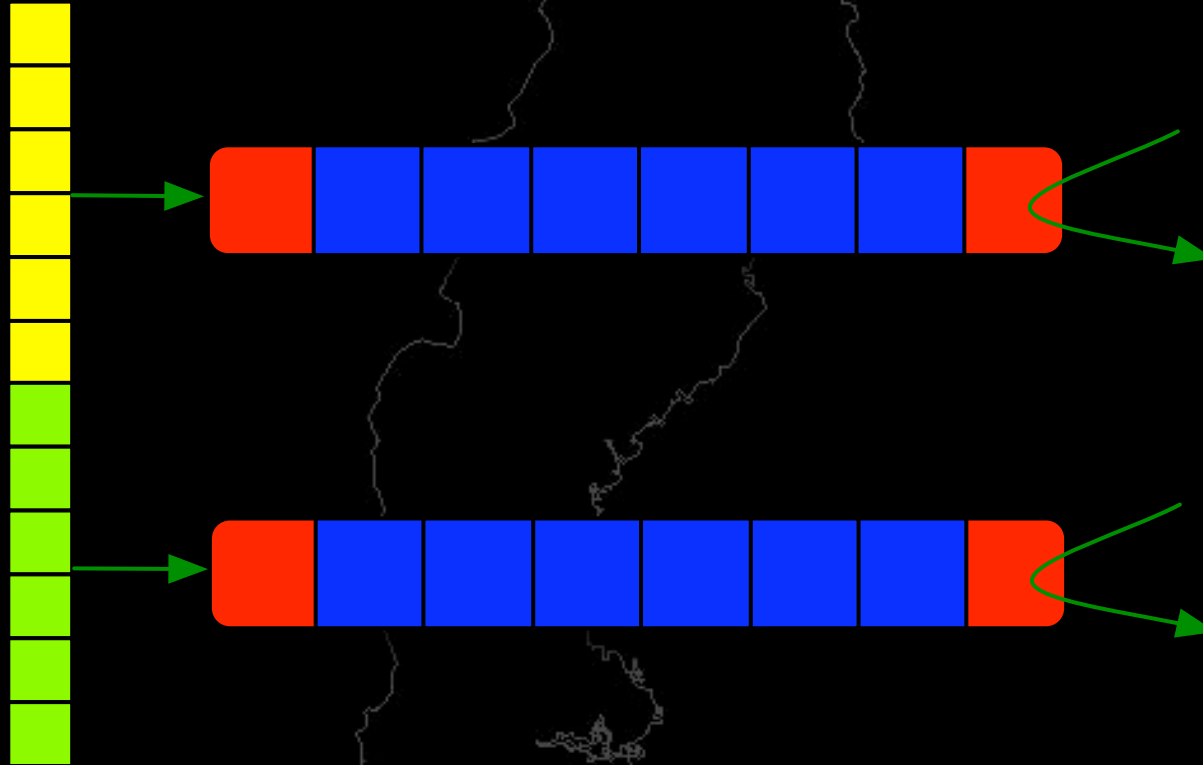


Kodewerk  
Java™ Performance Services

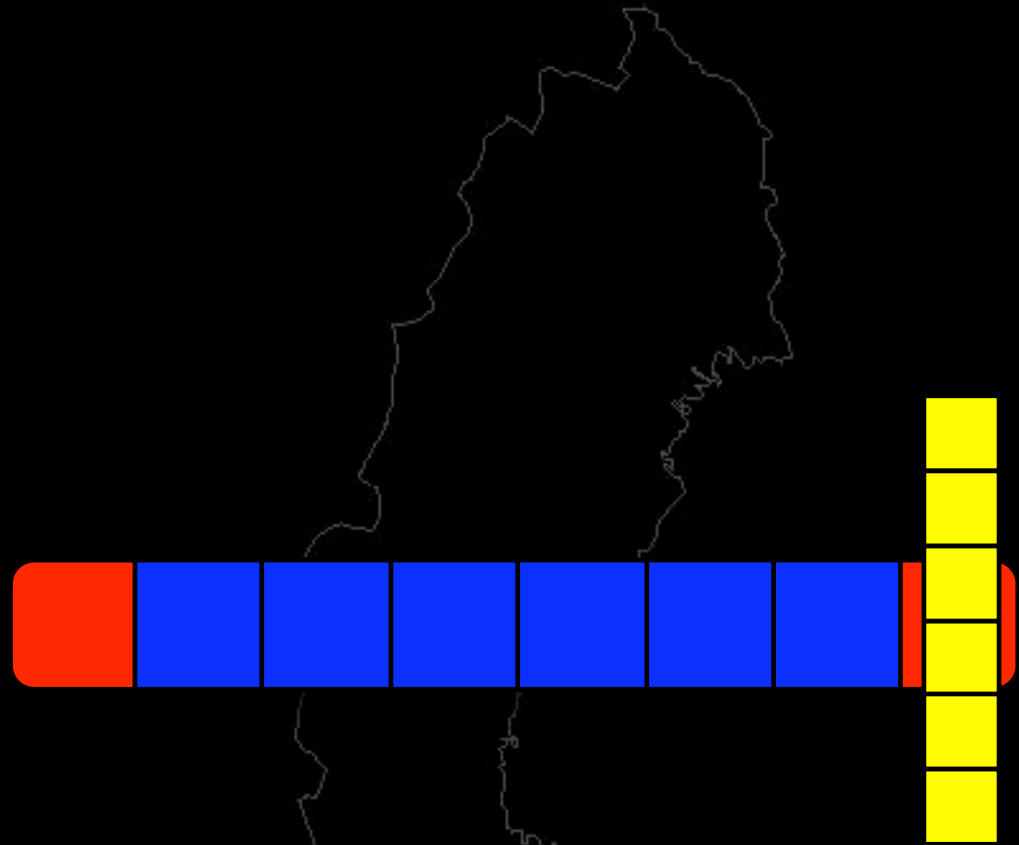


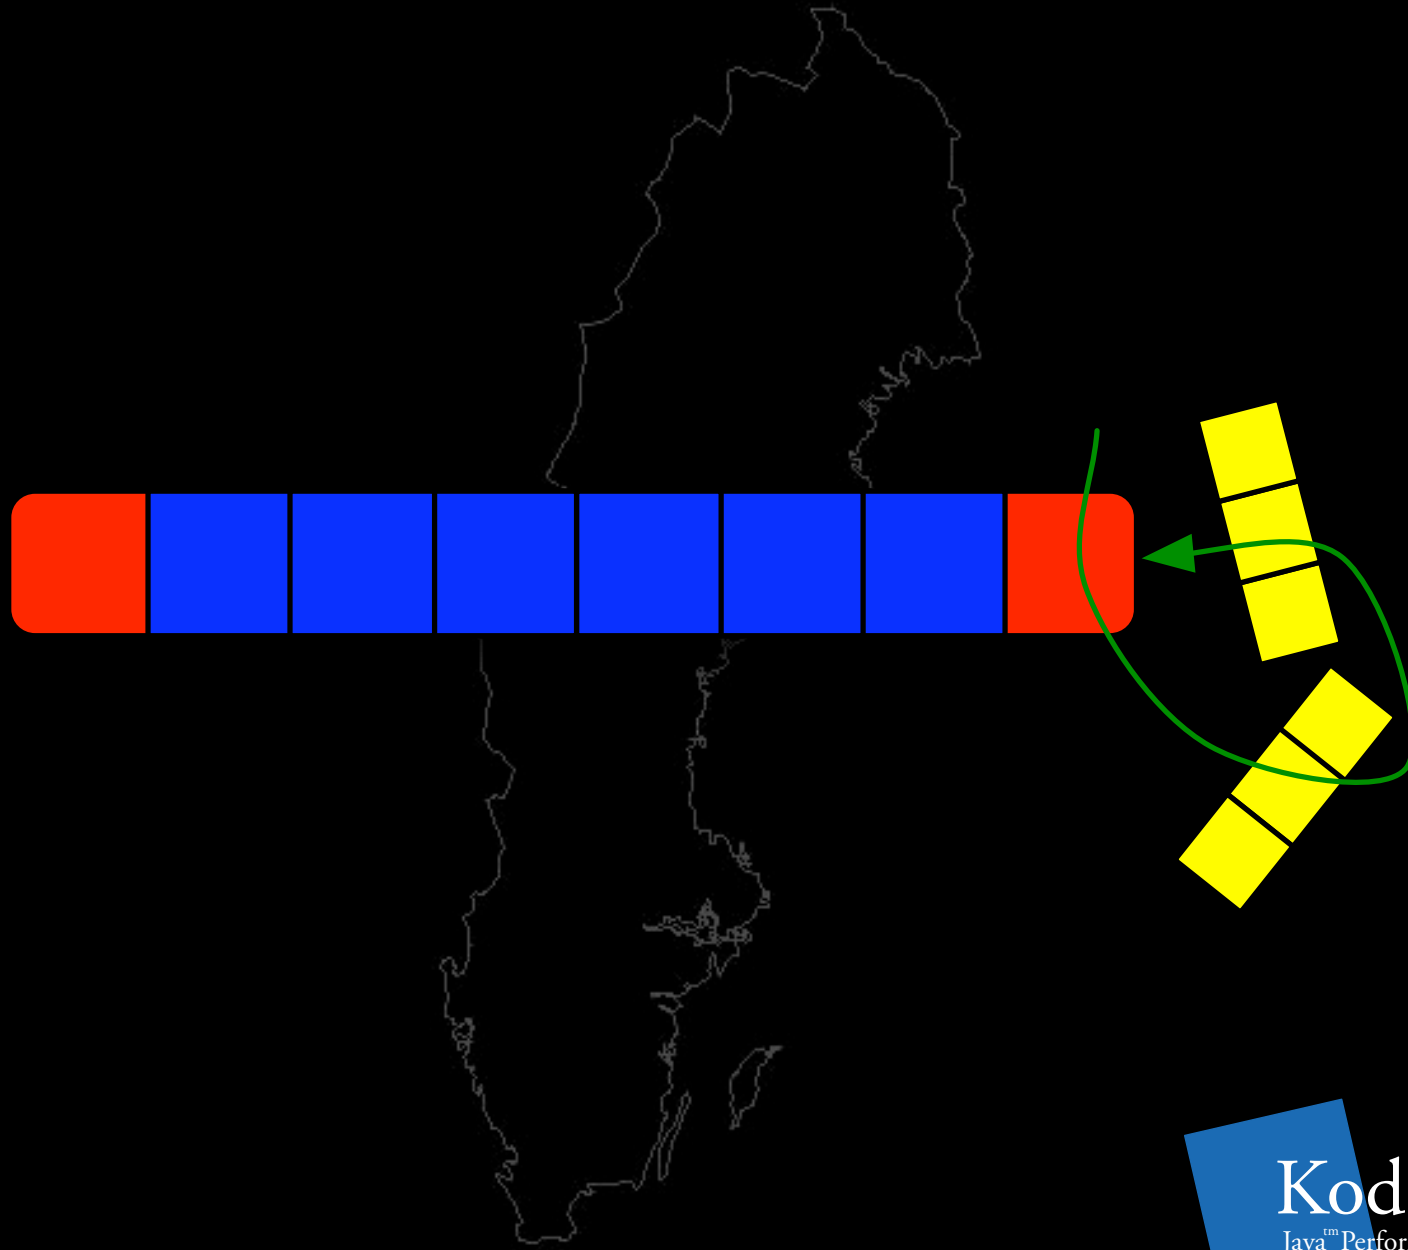
# Work Stealing Queue

Units of work

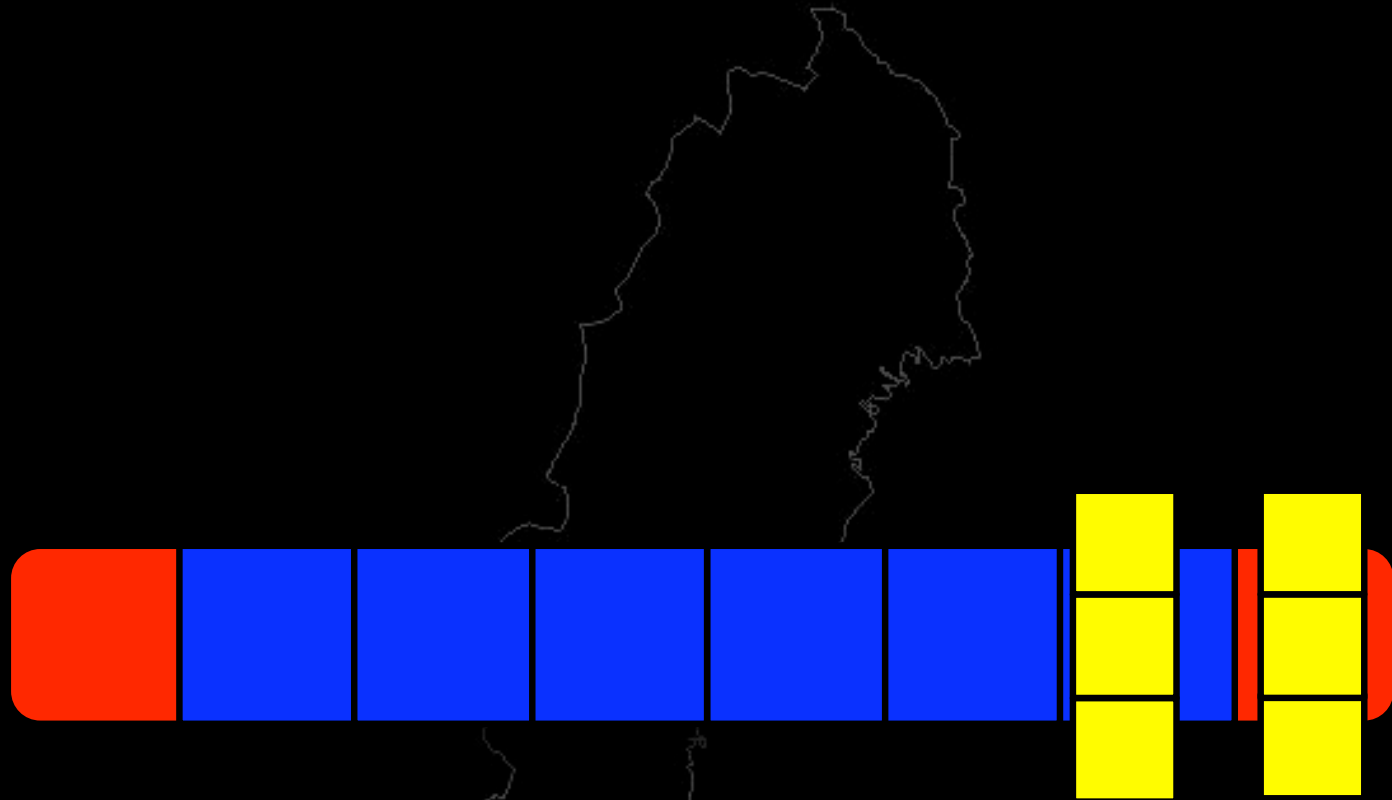


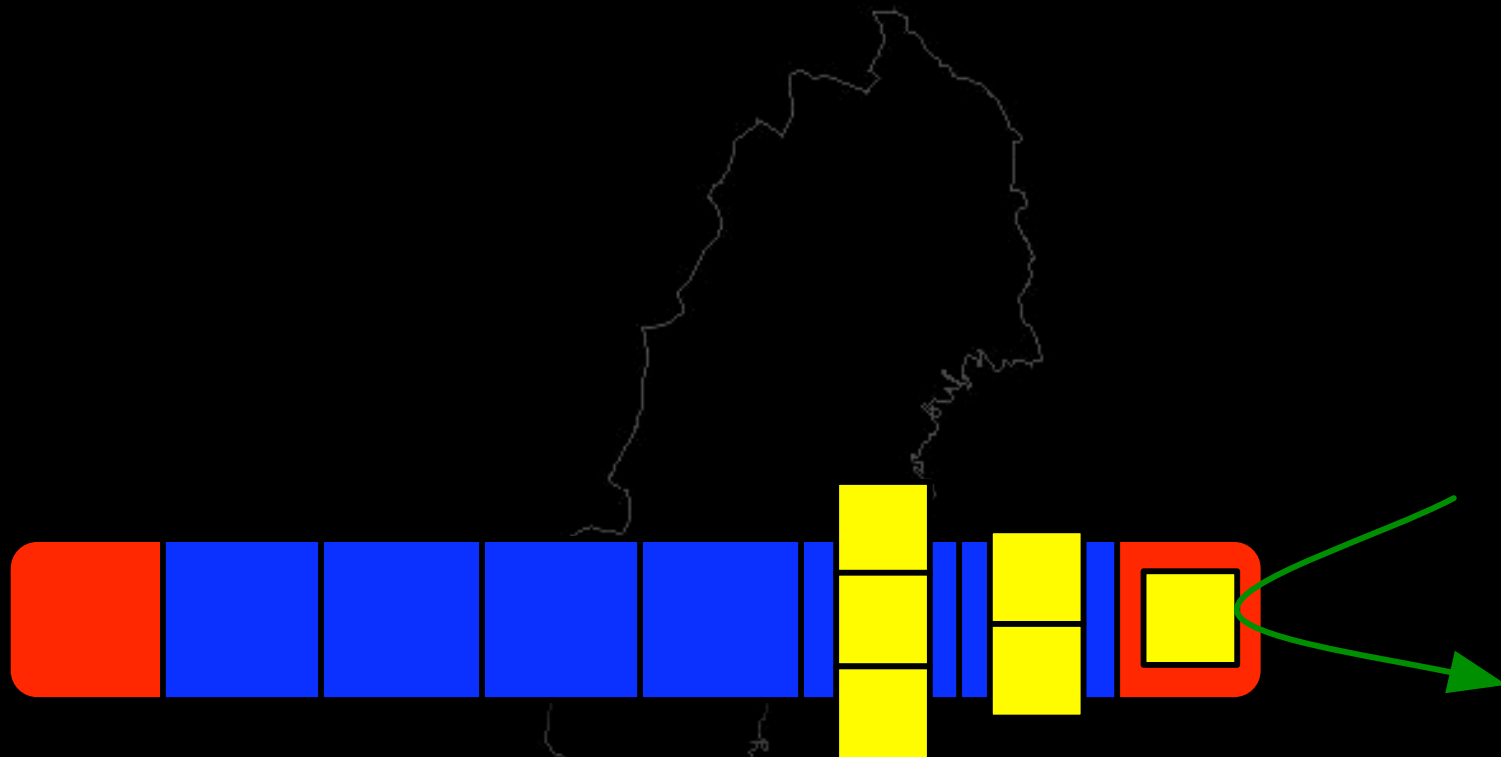
# Work splitting

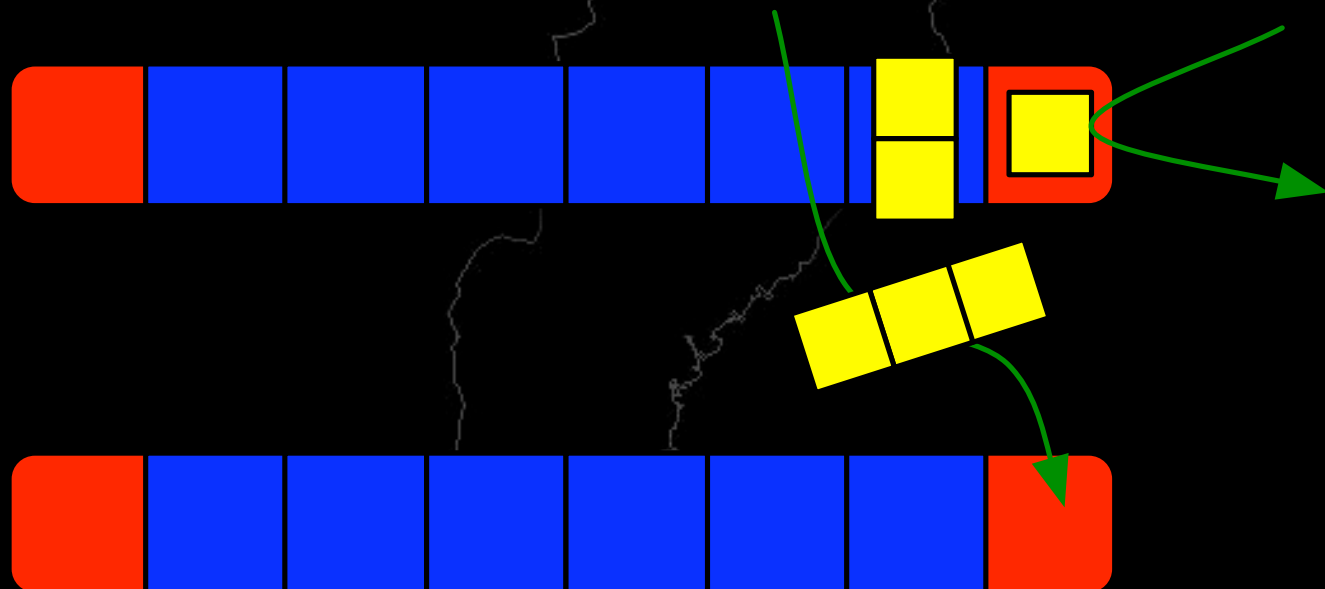














# Degrees of Scalability

**Kodewerk**  
Java™ Performance Services

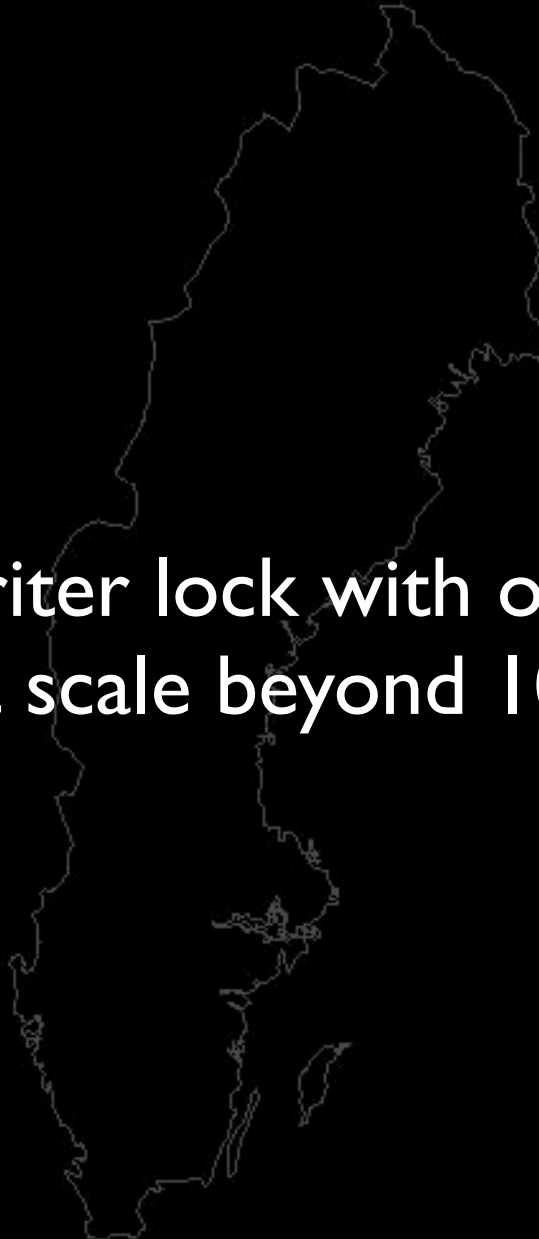


Lock free concurrency

**Kodewerk**  
Java<sup>™</sup> Performance Services



Parallel reads, serialized writes



Reader/Writer lock with only readers  
will not scale beyond 100 cpus



large arrays for concurrent





arrays to hold all data



resize cannot block



fully concurrent lock-less hashmap

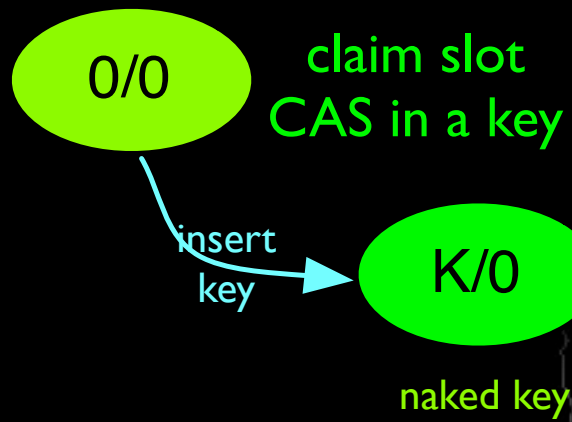
# Things we need

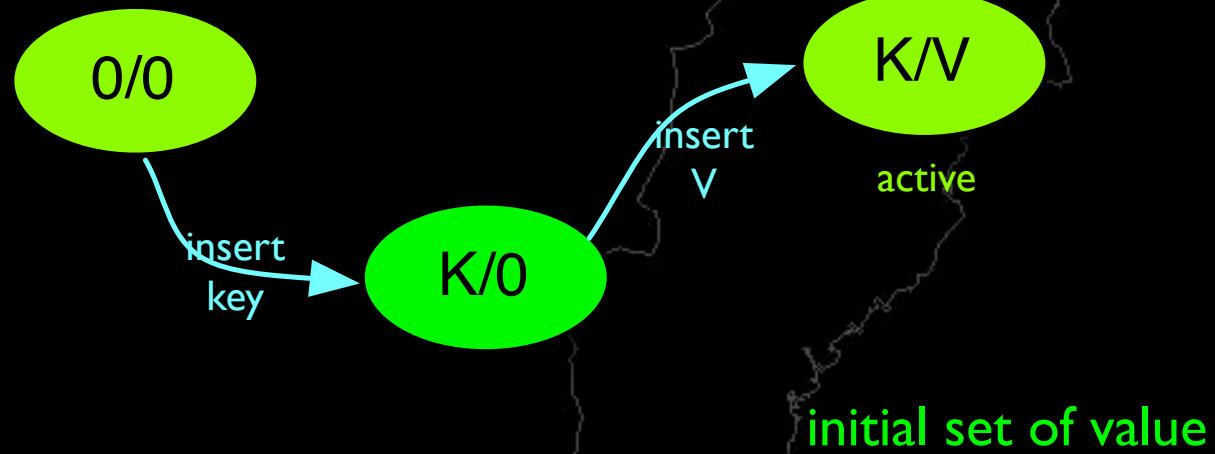
- 🚀 Large array to hold all data
  - alternating array of key value pairs
- 🚀 state machine for pair of words
  - CAS to manage state transition
- 🚀 Tombstone to mark deleted words
- 🚀 Use a box to mark values during resize
  - allows read access but prevents update
- 🚀 No single point of contention

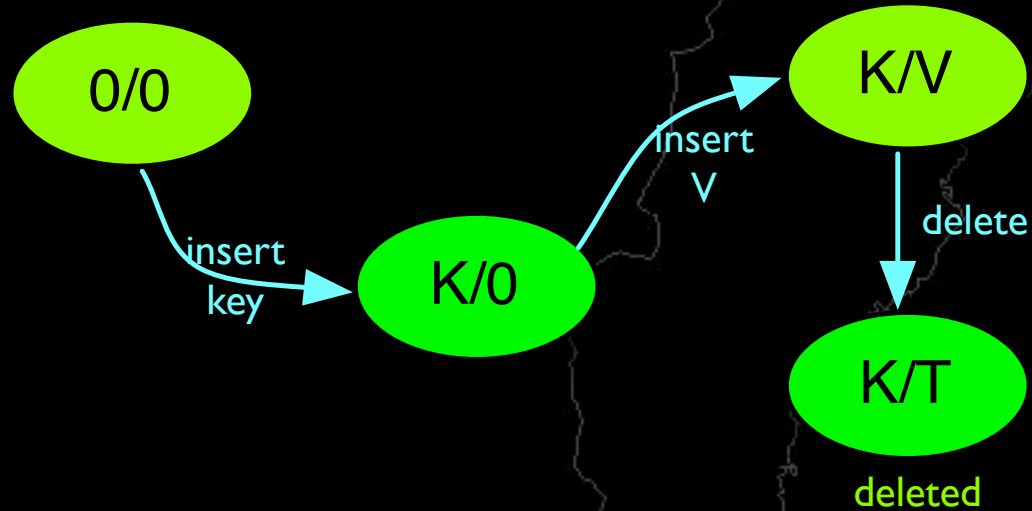
0/0

Initial

Inserting K/V pair  
Already probed table, missed  
Found proper empty K/V slot  
Ready to claim slot for this Key

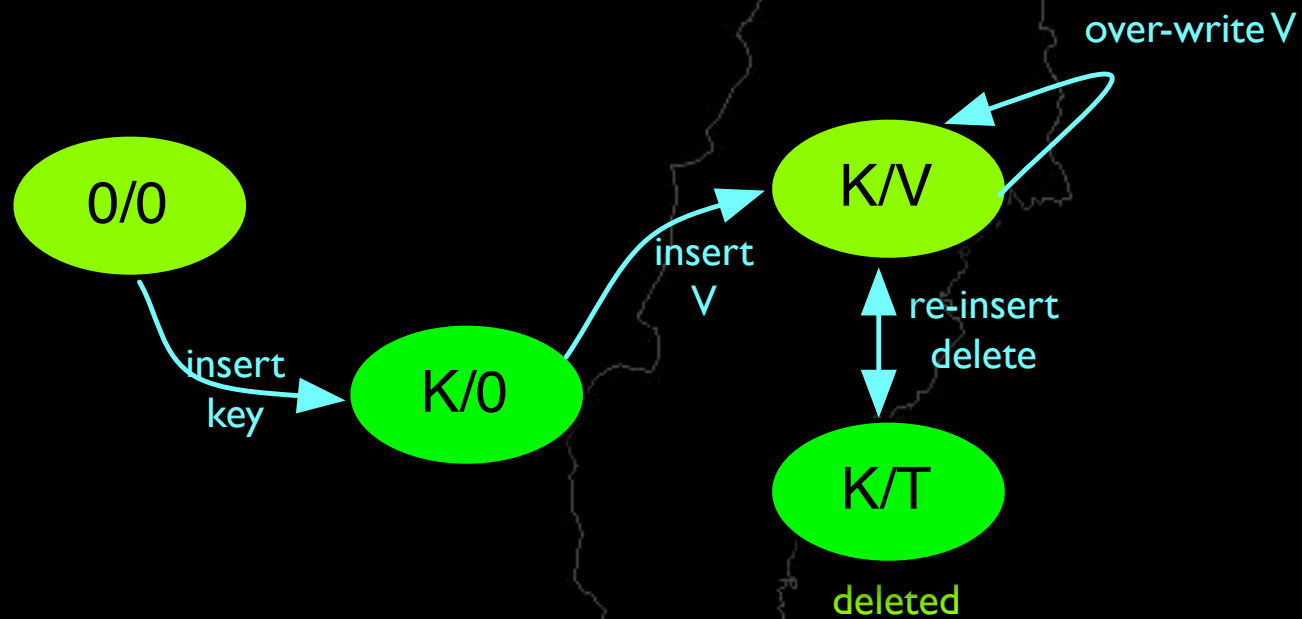




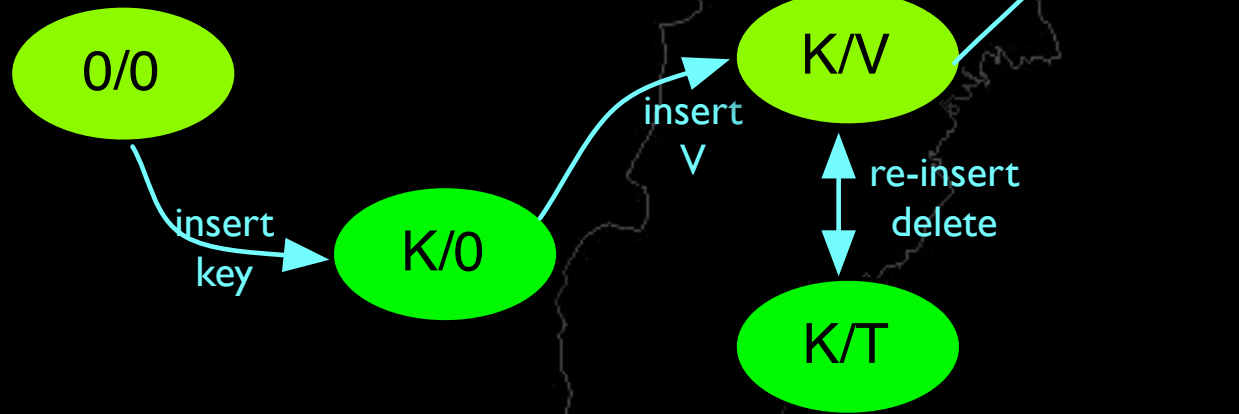


Tombstone marks delete, key remains



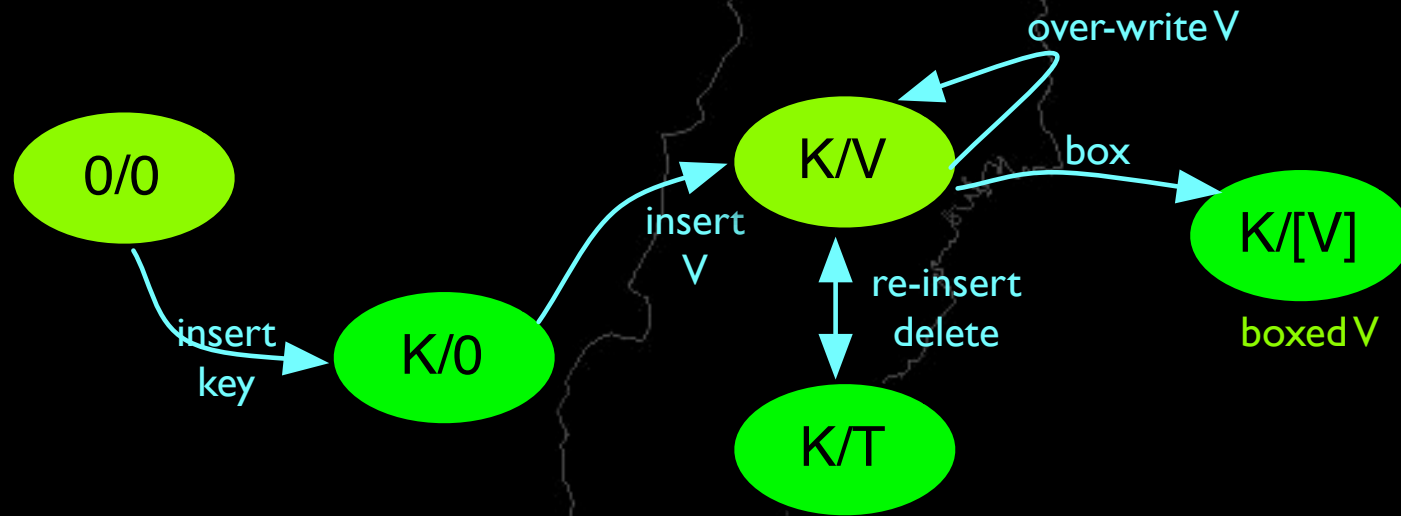


operations use same key slot  
all operations operate normally

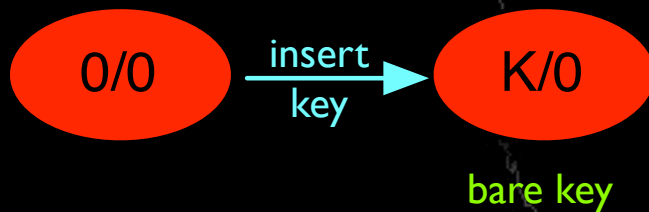
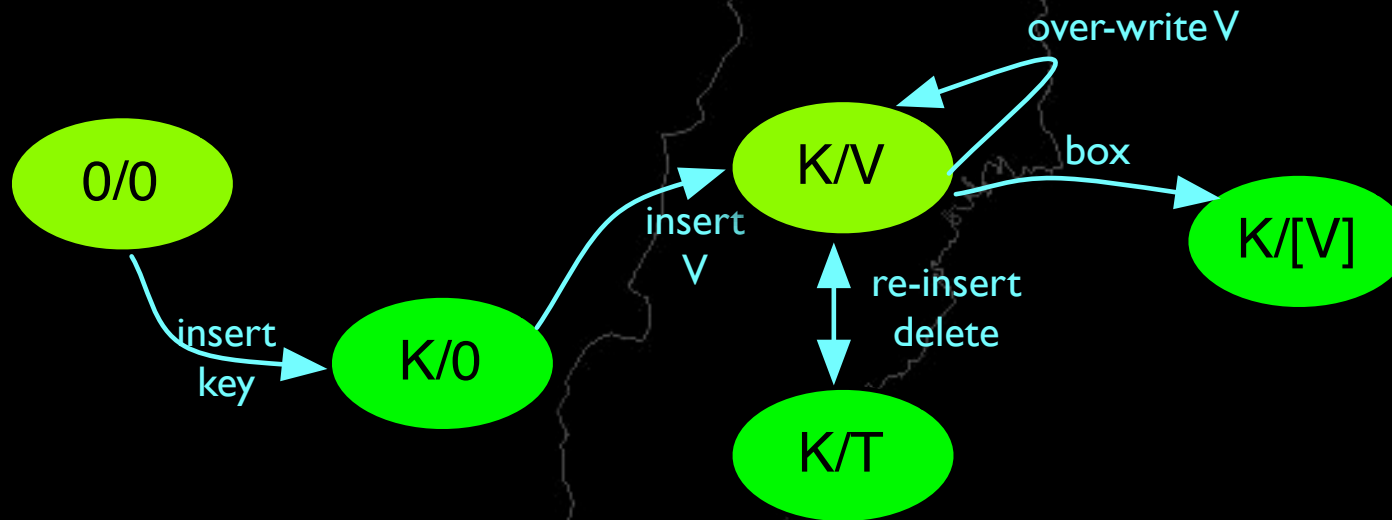


0/0

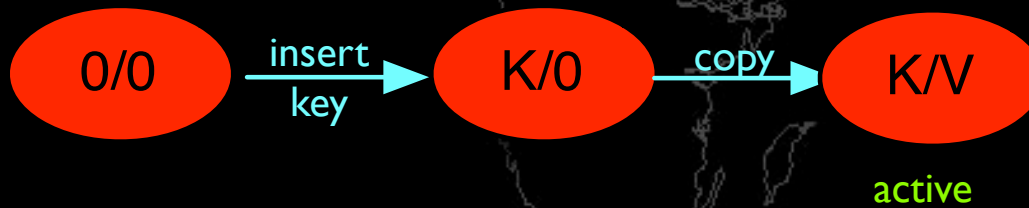
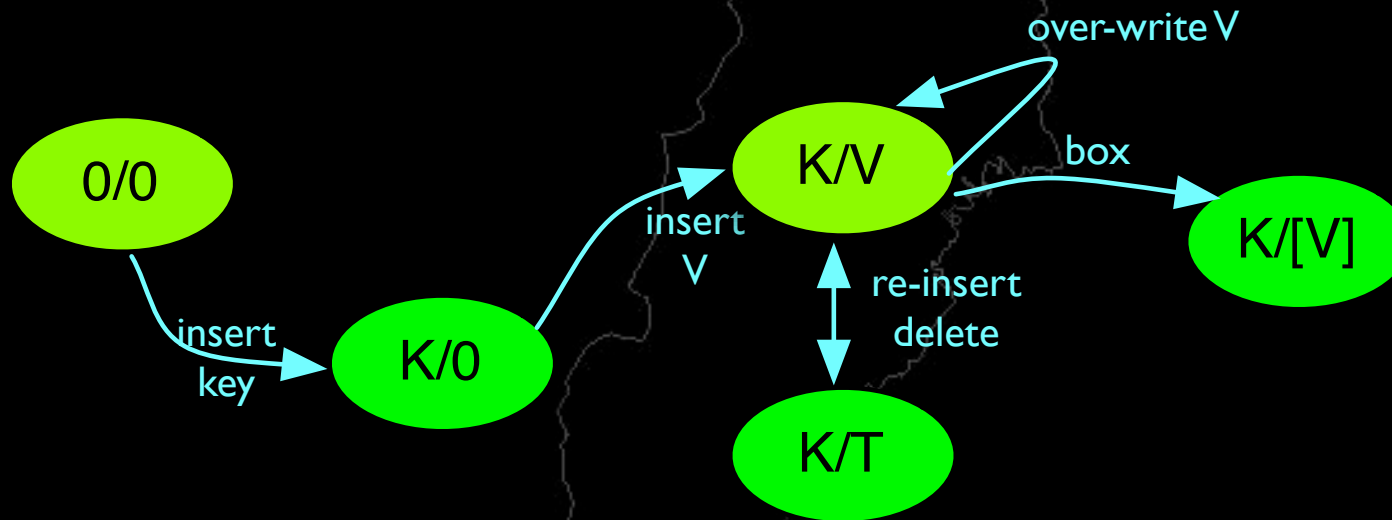
resize triggered  
new array created  
helper threads CAS a promise counter



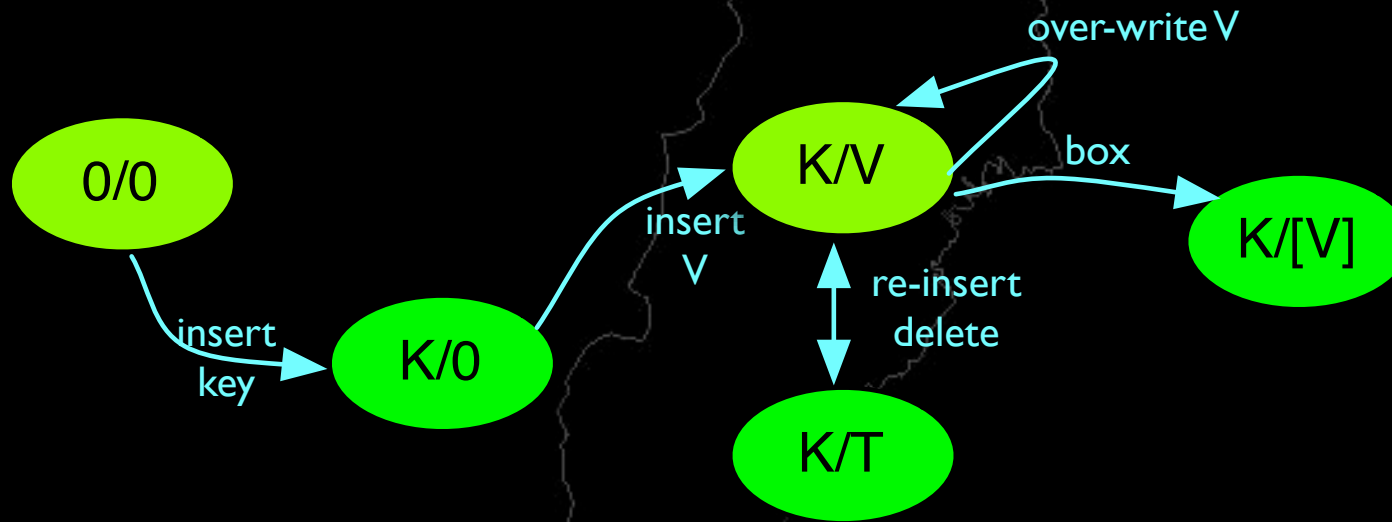
boxing prevents further changes



claim key slot in new table



copy V without box



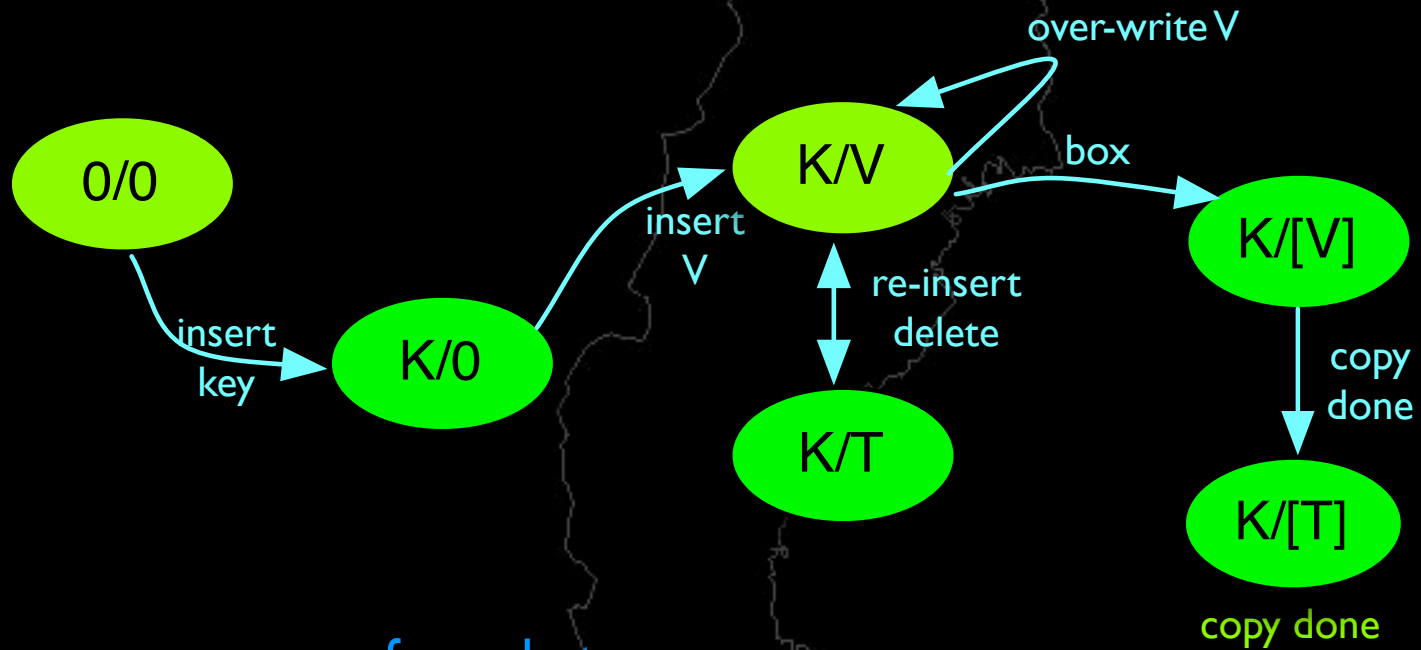
memory fence between arrays

---



active

fence after writing new array and before copy done

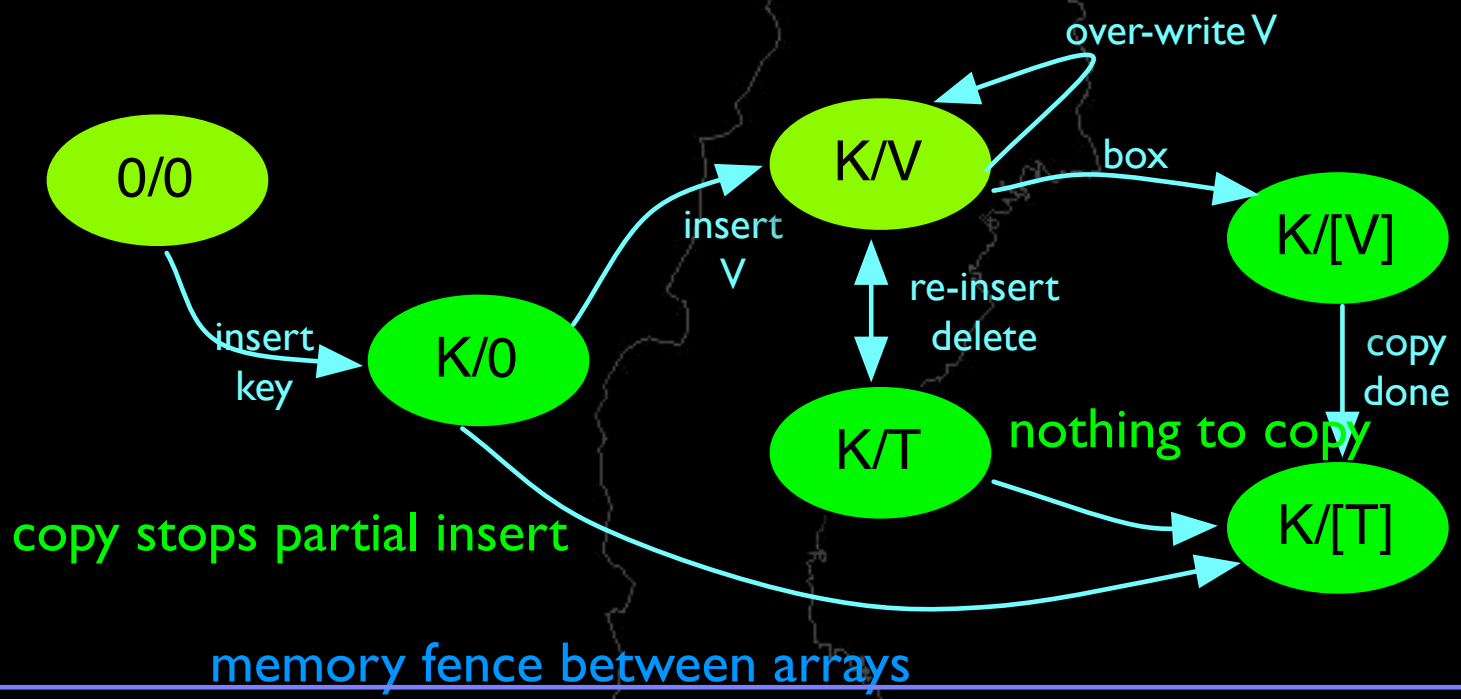


memory fence between arrays

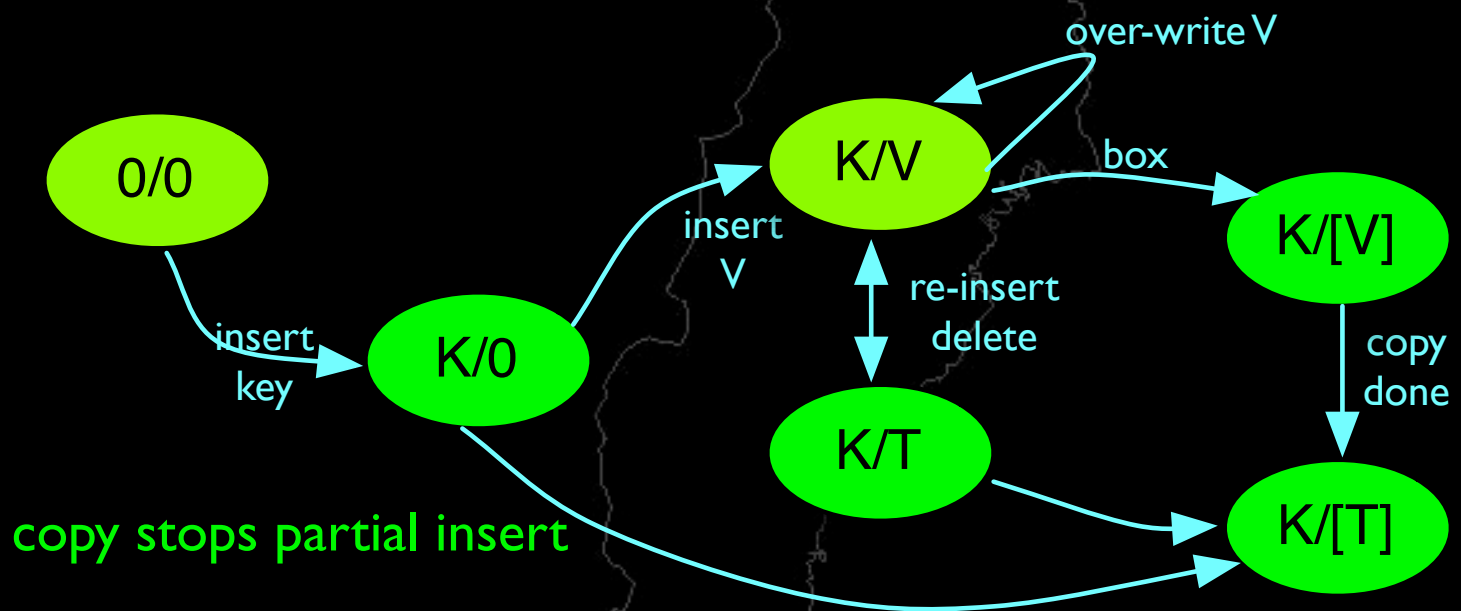
---



fence after writing new array and before copy done

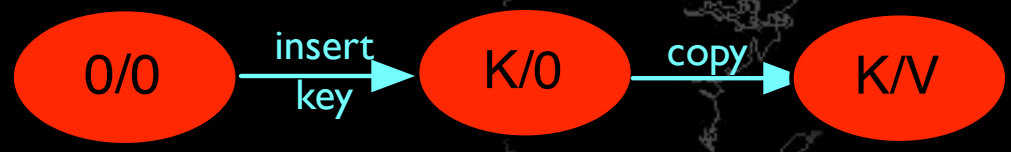






memory fence between arrays

---



HashMap (no sync)



HashMap (sync)



Hashtable



ConcurrentHashMap



NonBlockingHashMap



# Blackboard Reloaded

HashMap (no sync)



HashMap (sync)



Hashtable



ConcurrentHashMap



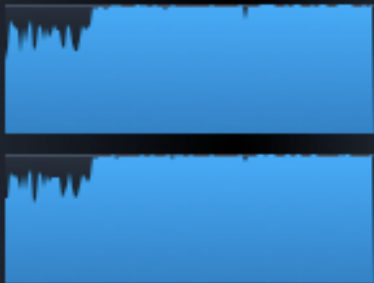
NonBlockingHashMap



# Blackboard Reloaded

**CPU**

User:	97%
System:	2%
Nice:	0%
Idle:	1%




**java** 191.1%  
firefox-bin 2.5%  
SystemUIServer 1.8%  
kernel\_task 1.1%  
Numbers 0.9%

Load Average  
7.18, 4.13, 2.71

Processes  
62 tasks, 301 threads

Uptime  
7d, 16h, 21m, 11s

Actual Running Time  
3d, 2h, 53m, 45s

 Activity Monitor



scales linearly up to 1000 CPUs



**Kodewerk**  
Java<sup>™</sup> Performance Services



Fully concurrent lock-less FIFO?



Stripe on queues and randomly pick one



stripe ad-absurdum



Kodewerk  
Java™ Performance Services





insert searches for null CAS down value



read searches for value and CAS down null



too large read spin, too small inserts spin



resize is earlier, promote when entire  
array is tombstoned



Questions?



**Kodewerk**  
Java<sup>™</sup> Performance Services