




Continuous Monitoring


Verena Bitto
Philipp Lengauer

2016-03-13


Trace Growth




Trace Size vs Disk Limit




Compression




Similar Problem: Objects vs Heap Size



Rethink Trace Size vs Disk Limit



Split trace into n files, overwrite oldest file first.



Every trace file may be eventually be the oldest.


What is the state of the heap at the beginning of the oldest file?



Need-to-know

Synchronization Points


Use GCs as synchronization points



What if no GC occurs at the right point?

Trace Size Deviation

Trigger “**Emergency GCs**” after max deviation is reached.



MaxSize=16GB Deviation=25%


file count = 100% / Deviation

target file size = MaxSize * Deviation


max file size = file size + file size * Deviation

Synchronization GC


Optimized move event




Optimized move sync event




Replace all *move* events with *move sync* events.

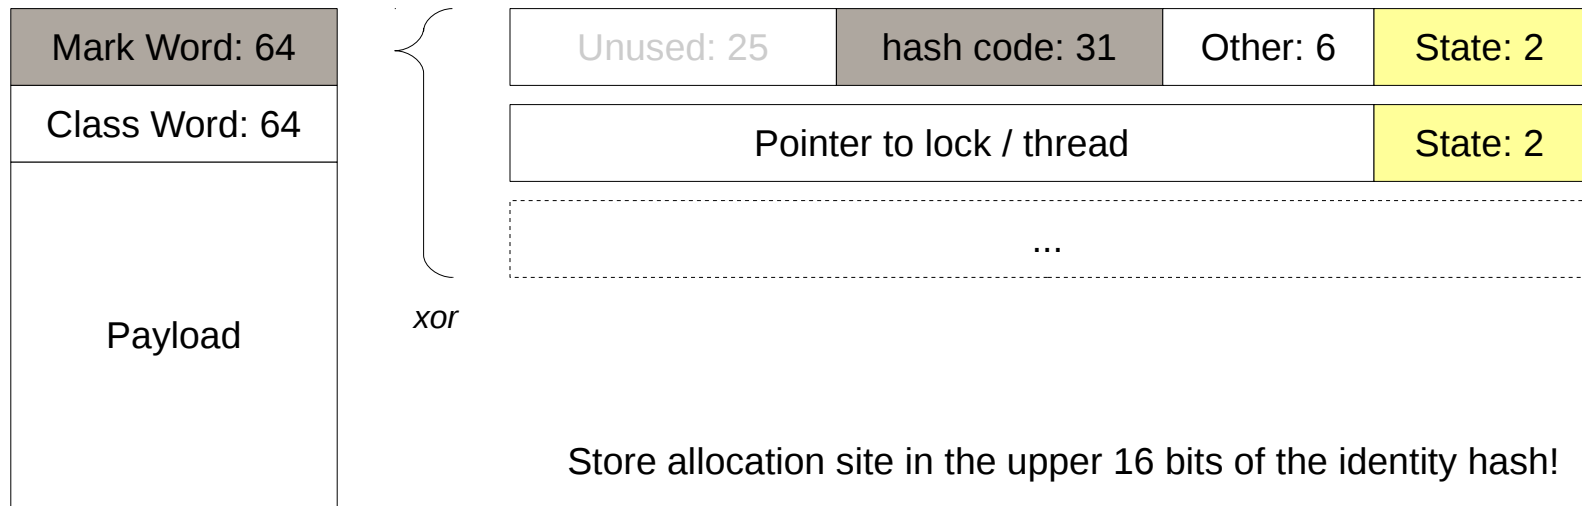


Overhead Rotation





Restoring Allocation Sites



- Must generate hash code eagerly for every (!) object.
- Reduces entropy of the hash to **0.0015%**.



Artificial Worst Case

```


void main() {
    Set<Object> set = new HashSet<>();
    for(int i = 0; i < 1_000_000_000; i++) {
        set.put(create());
    }
    for(int i = 0; i < 1_000_000_000; i++) {
        set.contains(create());
    }
}

Object create() {
    // all objects have same allocation site
    return new Object();
}

```

→ run time +2191%

Overhead with Saving Allocation Sites



Reducing Hash Code Generation

Assumption: classes overwriting `hashCode()` will *most likely* never use the identity hash!

