

Policy Exploration for JITDs (Java)

By

Team Datum


PAST : Studying Existing Implementation

- Understanding different types of Cogs: ConcatCog, ArrayCog, BTreeCog, SortedArrayCog
- Current implementation modes: Database Cracking, Adaptive Merge
- Hybrid modes : Swap, Transition


PAST: Replication

- Replicated the results from the paper by running the current implementation on many data sets of uniform distribution.
- Written Python scripts to generate graphs from the generated output file.
- Achieved near to accurate results as compared to the graphs presented in the paper for the modes:
 - Cracking
 - Adaptive Merge
 - Swap
 - Transition

PAST: Graphs – 1 (Cracking)



(a) Primitive Policy: Cracking



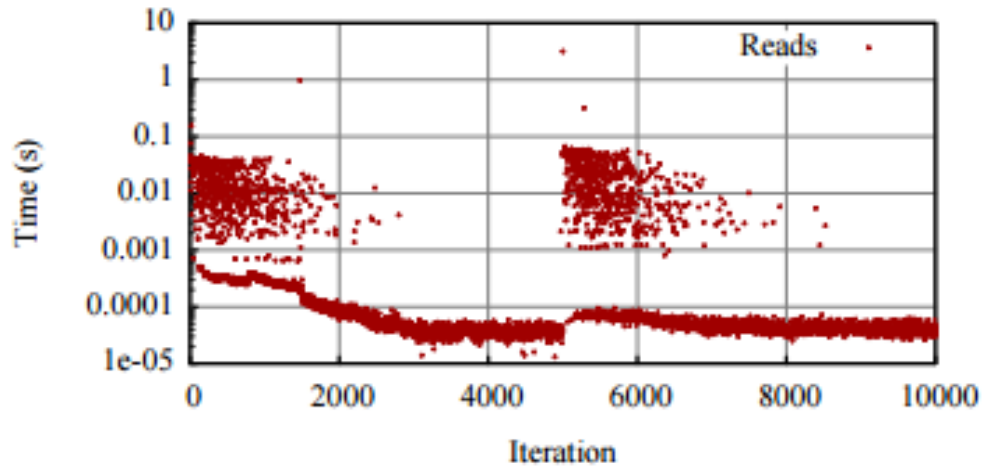
System Specifications:

Dual core 2.40 GHz Intel i5
8 GB of RAM
Ubuntu 14.10 and JDK 1.7
JVM heap size was set to 5 GB
100MB of stack space.


Tested with :

mode cracker
init 100000000
seqread 5000
write 10000000
seqread 5000

PAST: Graphs – 2 (Adaptive Merge)



(b) Primitive Policy: Adaptive Merge



System Specifications:


Dual core 2.40 GHz Intel i5
8 GB of RAM
Ubuntu 14.10 and JDK 1.7
JVM heap size was set to 5 GB
100MB of stack space.

Tested with :

mode simplemerge
init 10000000
seqread 5000
write 10000000
seqread 5000



PAST: Challenges Faced

- Performance breakdown due to Garbage Collection
- Performance issue with Push Down Adaptive Merge raising the runtimes to exponentially increasing curve.
- Tuned the java garbage collector parameters to allocate high heap sizes and avoid GC invocation at runtime. Obtained similar performance.



PAST: YCSB's Zipfian Workload

- Operated the current implementation on other workloads: Zipfian
- Challenges Faced: Irregular Cracking behavior on Zipfian workload with / without splaying.



*Figure : Cracking operation with Zipfian Distribution,
Tested as:
KeyRange1000000
Load 1000000
Reads 1000*

CURRENT : Profiling

- Studying the low level implementation along with Profiling to check the space and time complexities at various blocks of code.
- Used the Profilers :
 - Visual VM
 - JProfiler

CURRENT : Profiling

jtd.ScriptDriver (pid 7087)

Monitor

CPU Memory Classes Threads

Uptime: 9 min 58 sec


Perform GC Heap Dump

CPU

Heap Metaspace

CPU usage: 17.8%


GC activity: 0.0%



CPU usage GC activity

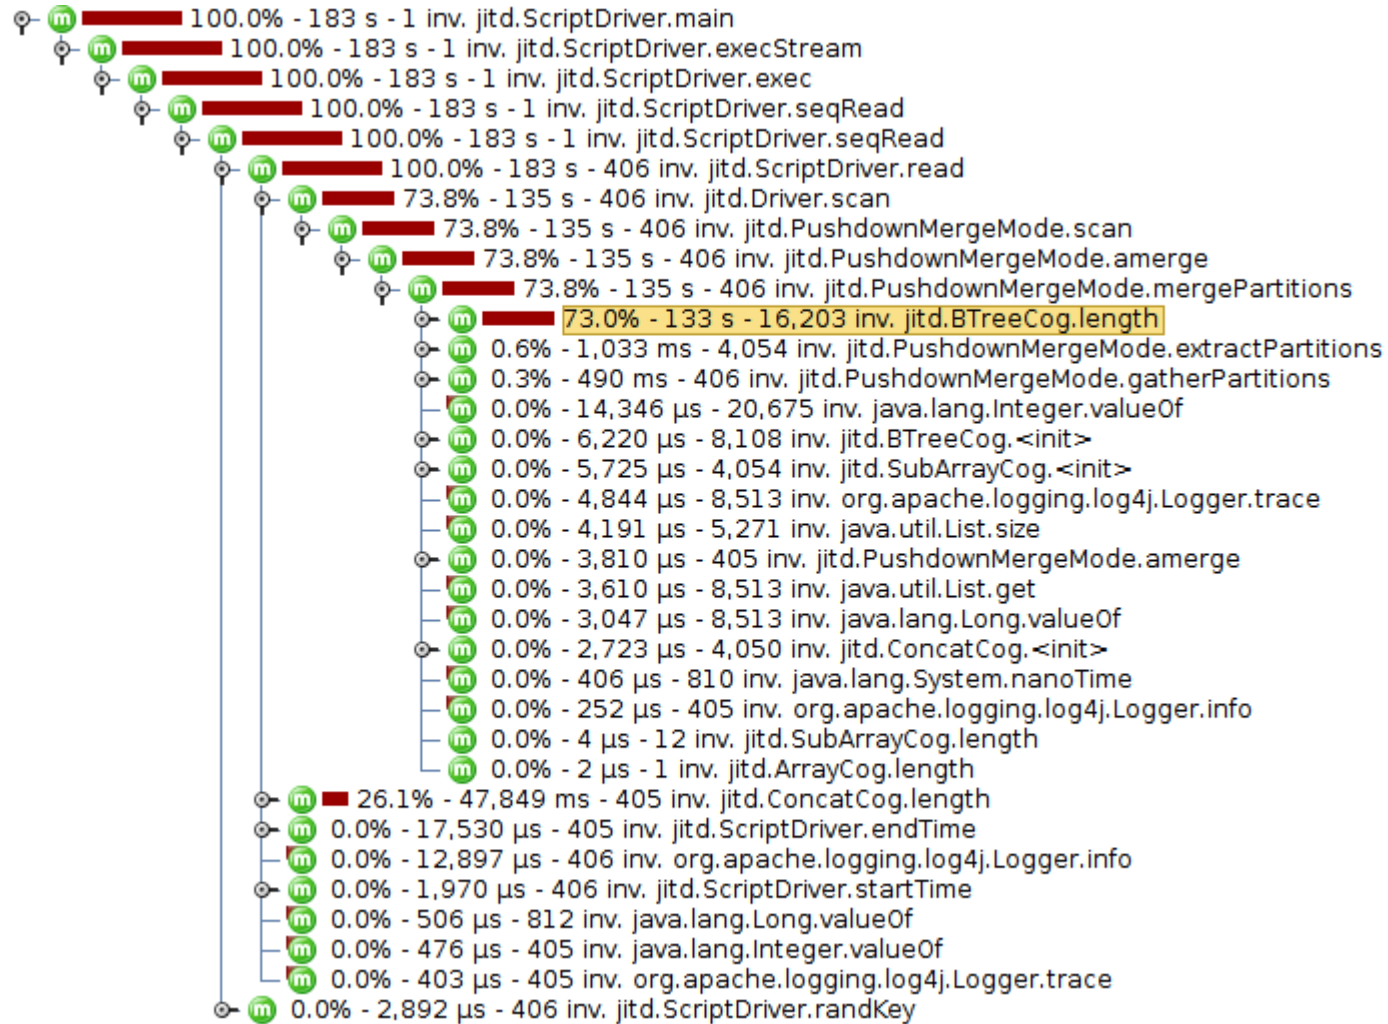
Size: 7,516,192,768 B
Max: 7,516,192,768 B

Used: 1,877,581,336 B



Heap size Used heap

CURRENT : Profiling



Splaying

- Added splay operations to the BTree Cog to make the root cog nearly balanced.
- It follows the traditional approaches of Zig and Zag operations of a Splay Tree.
- Currently, splay policy has been integrated into the Cracking mode which splays the root cog with the given lower bound of the range query.

Splaying Algorithm

Algorithm Splaying algorithm

```
1: procedure SPLAY(key, cog)
2:   if cog is BTree then
3:     if cog.separator > key then
4:       if cog.left is BTree then
5:         if cog.left.separator > key then
6:           cog.left.left ← Splay(key, cog.left.left)
7:           cog ← RotateRight(cog)
8:         else if cog.left.separator < key then
9:           cog.left.right ← Splay(key, cog.left.right)
10:          cog.left ← RotateLeft(cog.left)
11:        else
12:          cog ← RotateRight(cog)
13:        end if
14:      end if
```

```
15:     else if cog.separator < key then
16:       if cog.right is BTree then
17:         if cog.right.separator > key then
18:           cog.right.left ← Splay(key, cog.right.left)
19:           cog.right ← RotateRight(cog.right)
20:         else if cog.right.separator < key then
21:           cog.right.right ← Splay(key, cog.right.right)
22:           cog ← RotateLeft(cog)
23:         else
24:           cog ← RotateLeft(cog)
25:         end if
26:       end if
27:     else
28:       return cog
29:     end if
30:   end if
31: end procedure
```

Work in Progress..

- Debugging splaying along with cracking.
- Analyzing the issue for null pointers during merge phase, which happens for the initial load of 100m.
- Understanding the behavior while running the implementation over YCSB's Zipfian workloads with and without splaying.

Further works to do..

- Debugging splaying along with adaptive merge.
- Splaying with Adaptive Merge on different read / write ratios given by YCSB (Zipfian).
- Comparing the above results with uniform workloads to determine final policies.

Thank you..!!

Questions?