

Scaling Log-Structured KV-Stores

featuring

Monkey, Dostoevsky and Wacky

SIGMOD17/18/19



DASlab
@ Harvard SEAS

Niv Dayan



BigTable



Amazon
DynamoDB

 **riak**

WT

Log-Structured KV-Stores





Couchbase



levelDB

Log-Structured KV-Stores




Why Log-Structured KV-Stores?

Why Log-Structured KV-Stores?



fast writes

Why Log-Structured KV-Stores?





memory





storage


Why Log-Structured KV-Stores?




Why Log-Structured KV-Stores?




Why Log-Structured KV-Stores?




write data




write data



write data




In-Place Writes



write data




In-Place Writes




Log-Structured Writes


buffer writes




Log-Structured Writes



Log-Structured Writes



Log-Structured Writes



Log-Structured KV-Stores



fast writes

Log-Structured KV-Stores




fast writes





fast reads



massive data







Background




Background




The Log-Structured Merge-Tree

Background



LSM-tree


buffer




writes





buffer




buffer gets full








buffer




exponentially increasing capacities


level 1



level 2




level 3



find X

buffer




binary searching

X

where's
Waldo

buffer



pointers

one I/O per run

X

where's
Waldo

buffer

**Bloom
filters**

pointers

SSD

X

where's
Waldo

buffer

**Bloom
filters**

true
negative

pointers

SSD


X

where's
Waldo

buffer

**Bloom
filters**

pointers



true
negative

false
positive

X

where's
Waldo

buffer

**Bloom
filters**

pointers


SSD

true
negative

false
positive

true
positive

X



merging

writes

reads



writes



merging



reads

merging


Tiering
write-optimized




Leveling
read-optimized




Tiering
write-optimized




Leveling
read-optimized




Tiering
write-optimized




Leveling
read-optimized




Tiering
write-optimized




Leveling
read-optimized




Tiering
write-optimized




Leveling
read-optimized




Tiering
write-optimized




Leveling
read-optimized



Tiering
write-optimized



Leveling
read-optimized




Tiering
write-optimized



gather



Leveling
read-optimized




merge




Tiering
write-optimized

Leveling
read-optimized



Tiering
write-optimized


Leveling
read-optimized



size ratio

Tiering
write-optimized

Leveling
read-optimized



Tiering
write-optimized



Leveling
read-optimized



↪ size ratio R

Tiering
write-optimized



Leveling
read-optimized



size ratio $R \gg$

Tiering
write-optimized



Leveling
read-optimized



↪ size ratio R

Tiering
write-optimized

Leveling
read-optimized

$O(N)$ runs per level



1 run per level





log


**sorted
array**


size ratio $R \gg$























Monkey: **O**ptimal **N**avigable **K**ey-Value Store

SIGMOD17



Monkey: **O**ptimal **N**avigable **K**ey-Value Store

SIGMOD17




Niv Dayan
Manos Athanassoulis
Stratos Idreos

Monkey: **O**ptimal **N**avigable **K**ey-Value Store

SIGMOD17

data



**Bloom
filters**



data



Bloom
filters



bits/entry

x

x

x

data



Bloom
filters



bits/entry

x

x

x

data



Bloom
filters




**false
positive rate**

$$O(e^{-x})$$

$$O(e^{-x})$$

$$O(e^{-x})$$

Bloom
filters




false
positive rate

$$O(e^{-x})$$

$$O(e^{-x})$$

$$O(e^{-x})$$

$$\left. \begin{array}{l} O(e^{-x}) \\ O(e^{-x}) \\ O(e^{-x}) \end{array} \right\} = O(\mathbf{e^{-x}} \cdot \mathbf{\log_R(N)})$$



false
positive rate

$$O(e^{-x})$$

$$O(e^{-x})$$

$$O(e^{-x})$$

$$\left. \begin{array}{l} O(e^{-x}) \\ O(e^{-x}) \\ O(e^{-x}) \end{array} \right\} = O(\mathbf{e^{-x}} \cdot \mathbf{\log_R(N)})$$


A red thumbs down icon is positioned below the final equation, indicating a negative or undesirable result.

Bloom
filters



most memory

false
positive rate

$$O(e^{-x})$$

$$O(e^{-x})$$

$$O(e^{-x})$$

Bloom
filters



most memory




saves at most 1 I/O!

false
positive rate

$$O(e^{-x})$$

$$O(e^{-x})$$


$$O(e^{-x})$$



reallocate



same memory - fewer false positives



relax

false positive rates

$$0 < p_0 < 1$$

$$0 < p_1 < 1$$

$$0 < p_2 < 1$$

relax

false positive rates

$$0 < p_0 < 1$$

$$0 < p_1 < 1$$

$$0 < p_2 < 1$$

model

$$\text{read cost} = f(\mathbf{p}_0, \mathbf{p}_1 \dots)$$

$$\text{memory footprint} = f(\mathbf{p}_0, \mathbf{p}_1 \dots)$$

relax

false positive rates

$$0 < p_0 < 1$$

$$0 < p_1 < 1$$

$$0 < p_2 < 1$$

model

$$\text{read cost} = \sum_1^L p_i$$

$$\text{memory footprint} = - \sum_i^L \frac{N}{T^{L-i}} \cdot \frac{\ln(p_i)}{\ln(2)^2}$$

relax

false positive rates

$$0 < p_0 < 1$$

$$0 < p_1 < 1$$

$$0 < p_2 < 1$$

model


read cost

$$= \sum_1^L p_i$$

memory footprint

$$= - \sum_i^L \frac{N}{T^{L-i}} \cdot \frac{\ln(p_i)}{\ln(2)^2}$$

optimize



false
positive rate



$$p_0 \approx O(e^{-x}/\mathbf{R}^2)$$

$$p_1 \approx O(e^{-x}/\mathbf{R}^1)$$

$$p_2 \approx O(e^{-x}/\mathbf{R}^0)$$



false
positive rate

$$O(e^{-x}/R^2)$$

$$O(e^{-x}/R^1)$$

$$O(e^{-x}/R^0)$$





**geometric
progression**

$$= O(e^{-x})$$




$$O(\mathbf{e^{-x}} \cdot \log_R(\mathbf{N})) > O(\mathbf{e^{-x}})$$






$$O(\mathbf{e^{-x}} \cdot \mathbf{\log_R(N)})$$


$$O(\mathbf{e^{-x}})$$






$$O(\mathbf{e^{-x}} \cdot \mathbf{\log_R(N)})$$

$$O(\mathbf{e^{-x}})$$





I/O overheads with leveling

●
point

→
long range

→
short range

↶
writes



●
point

false positive rates

exponentially
decreasing



$$O(\mathbf{e^{-x}/R^2})$$

$$O(\mathbf{e^{-x}/R})$$

$$O(\mathbf{e^{-x}})$$



false positive rates

$$O(e^{-x}/R^2)$$

$$O(e^{-x}/R)$$

●
point

→ $O(e^{-x})$

largest level



point

largest level

$O(e^{-x})$


long range

 **short range**


writes





long range

target range

$$O(\textcolor{red}{s}/\textcolor{red}{R}^2)$$

$$O(\textcolor{red}{s}/\textcolor{red}{R})$$

$$O(\textcolor{red}{s})$$



target key range


target range

$$O(s/R^2)$$

$$O(s/R)$$

$$O(s)$$

long range →



largest level

target key range

●
point

largest level

$O(e^{-x})$

→
long range

largest level

$O(s)$


→
short range

↖
writes




short → range

target range



short range →

target range



1

1

1

all levels
 $O(\log_R(N))$

●
point

largest level

$O(e^{-x})$

→
long range

largest level

$O(s)$

→
short range

all levels

$O(\log_R(N))$

↖
writes




writes



**exponentially
more work**





writes




exponentially
more work




**exponentially
less frequent**


writes





more work
↓
less frequent



writes




all levels











write-amplification



$O(\textcolor{red}{R})$

$O(\textcolor{red}{R})$

$O(\textcolor{red}{R})$


writes



$$\left. \begin{array}{l} O(R) \\ O(R) \\ O(R) \end{array} \right\} O(R \cdot \log_R(N))$$

●
point

→
long range

→
short range

↶
writes

largest level
 $O(e^{-x})$

largest level
 $O(s)$

all levels
 $O(\log_R(N))$

all levels
 $O(R \cdot \log_R(N))$

=

=

=

=

$O(e^{-x}/R^2)$

$O(s/R^2)$

1

$O(R)$

+

+

+

+

$O(e^{-x}/R)$

$O(s/R)$

1

$O(R)$

+

+

+

+

$O(e^{-x})$

$O(s)$

1

$O(R)$



point

long range

writes

largest level

largest level

all levels

$O(e^{-x})$

$O(s)$



$O(R)$



$O(R)$



$O(R)$

point

long range

writes


largest level

largest level


all levels

$O(e^{-x})$


$O(s)$




●  for point lookups and long range lookups
merging at smaller levels is superfluous




worse as data grows!





poor performance



poor performance

lower device lifetime (on SSD)



Dostoevsky

SIGMOD18



Dostoevsky: **S**pace-**T**ime **O**ptimized **E**volvable **S**calable **K**ey-Value Store

very write-optimized



Dostoevsky: **S**pace-**T**ime **O**ptimized **E**volvable **S**calable **K**ey-Value Store

Tiering
write-optimized


Leveling
read-optimized

Tiering
write-optimized

Lazy Leveling
mixed-optimized

Leveling
read-optimized

Lazy Leveling



Lazy Leveling



●
point

→
long range

→
short range

↶
writes



●
point

false positive rates

$$O(\mathbf{e^{-x}/R^3})$$



$$O(\mathbf{e^{-x}/R^2})$$



$$O(\mathbf{e^{-x}})$$



●
point

false positive rates

exponentially
decreasing

↑ $O(\mathbf{e^{-x}/R^3})$

$O(\mathbf{e^{-x}/R^2})$

$O(\mathbf{e^{-x}})$



false positive rates

$$O(e^{-x}/R^3)$$

$$O(e^{-x}/R^2)$$

●
point

→ $O(e^{-x})$




largest level

●
point

$O(e^{-x})$

with uniform FPRs



●
point

$O(e^{-x})$

→
long range

→
short range

↶
writes



→
long range

target range


$O(\mathbf{s}/R^2)$



$O(\mathbf{s}/R)$



$O(\mathbf{s})$



target key range


→
long range

target range

$$O(s/R^2)$$

$$O(s/R)$$

$$O(s)$$



largest level

target key range

●
point


$O(e^{-x})$

→
long range


$O(s)$

→
short range

↖
writes



→
short range



●
point

$O(e^{-x})$


→
long range

$O(s)$

→
short range


$O(1 + R \cdot (\log_R(N) - 1))$

↖
writes






write-amplification





write-amplification



$$\left. \begin{array}{l} O(1) \\ O(1) \\ O(R) \end{array} \right\} O(R + \log_R(N))$$

●
point

$O(e^{-x})$

→
long range

$O(s)$

→
short range

$O(1 + R \cdot (\log_R(N) - 1))$

↖
writes

$O(R + \log_R(N))$






point



long range



short range



writes

Lazy Leveling

$$O(e^{-x})$$

$$O(s)$$

$$O(1 + R \cdot (\log_R(N) - 1))$$

$$O(R + \log_R(N))$$

=

=

$\vee$

$\wedge$

Leveling

$$O(e^{-x})$$

$$O(s)$$

$$O(\log_R(N))$$

$$O(R \cdot \log_R(N))$$

●
point

→
long range

→
short range

↖
writes

Tiering

$$O(R \cdot e^{-x})$$

$$O(R \cdot s)$$

$$O(R \cdot \log_R(N))$$

$$O(\log_R(N))$$

$\vee$

$\vee$

$\vee$

$\wedge$

Lazy Leveling

$$O(e^{-x})$$

$$O(s)$$

$$O(1 + R \cdot (\log_R(N) - 1))$$

$$O(R + \log_R(N))$$

Leveling


$$O(e^{-x})$$

$$O(s)$$

$$O(\log_R(N))$$


$$O(R \cdot \log_R(N))$$

●
point



↖
writes

short range →




Tiering

Lazy Leveling

Leveling

writes



Tiering

Lazy Leveling

Leveling

Tiering
writes



Lazy Leveling

Leveling

Tiering
writes 

Lazy Leveling

Leveling
short range 

Tiering
writes 

Lazy Leveling
writes & point 

Leveling
short range 


Lazy Leveling



Tiering


Leveling

Fluid LSM-Tree



Fluid LSM-Tree

Lazy Leveling




point

long range

short range

writes

Lazy Leveling




point

long range

optimize
short range

writes

Lazy Leveling




point

long range

optimize
short range

writes

Lazy Leveling




point

long range

optimize
short range

writes

Leveling




point

long range

short range

optimize
writes

Leveling




point

long range

short range

optimize
writes

Lazy Leveling




point

long range

short range

optimize
writes

Lazy Leveling




point

long range

short range

optimize
writes

Tiering




optimize
point

long range

short range

writes

Tiering




optimize
point

long range

short range

writes

Lazy Leveling




optimize
point


long range

short range

writes

 R size ratio

Lazy Leveling




optimize
point

long range

short range

writes



Lazy Leveling

Fluid LSM-Tree



R size ratio




K runs at smaller levels



Z runs at largest level




Lazy Leveling




Fluid LSM-Tree


Tiering

Leveling







normalized
throughput





normalized
throughput



normalized
throughput







New Work


Lazy Leveling


writes




$$O(R + \log_R(N))$$

Lazy Leveling




writes



$O(\textcolor{red}{1})$

$O(\textcolor{red}{1})$

$O(\textcolor{red}{1})$

$O(\textcolor{red}{1})$


$O(\textcolor{red}{R})$


grow





point reads


writes







memory



point reads



writes

$$O(X/R^4)$$

$$O(e^{-X}/R^5)$$

$$O(1)$$

$$O(X/R^3)$$

$$O(e^{-X}/R^4)$$

$$O(1)$$

$$O(X/R^2)$$

$$O(e^{-X}/R^3)$$

$$O(1)$$

$$O(X/R^1)$$

$$O(e^{-X}/R^2)$$


$$O(1)$$

$$O(X)$$

$$O(e^{-X})$$

$$O(R)$$

growing as $O(\log_R(N))$



memory



point reads



writes

**exponentially
diminishing
returns**

$O(X/R^4)$

$O(e^{-X}/R^5)$

$O(1)$

$O(X/R^3)$

$O(e^{-X}/R^4)$

$O(1)$

$O(X/R^2)$

$O(e^{-X}/R^3)$

$O(1)$

$O(X/R^1)$

$O(e^{-X}/R^2)$

$O(1)$

$O(X)$

$O(e^{-X})$

$O(R)$

growing as $O(\log_R(N))$

Wacky SIGMOD19



Wacky: **A**morphous **C**alculable **K**ey-Value Store



Wacky: **A**morphous **C**alculable **K**ey-Value Store




introducing the **Log-Structured Merge-Bush**


Wacky: **A**morphous **C**alculable **K**ey-Value Store





LSM-bush

runs












memory

$$O(X)$$


point reads

$$O(e^{-X})$$




memory


$O(X)$


point reads

$O(e^{-X})$


range

$O(N^{1/2})$




writes

$O(\log_2 \log_R(N))$

Tiering

Lazy
Leveling

LSM-bush

Leveling



Wacky generalizes

Tiering Lazy Leveling LSM-bush Leveling

Conclusion

Conclusion

Bloom
filters



LSM-tree



Conclusion



**optimizes
memory
allocation**

Bloom
filters



LSM-tree



Conclusion



optimizes
memory
allocation

Bloom
filters



LSM-tree



**lazy smaller
levels**



Conclusion



optimizes
memory
allocation

Bloom
filters



LSM-tree




lazy smaller
levels




**increasing
ratios**




Conclusion




Conclusion




Conclusion



Conclusion





Conclusion




🔧 little memory



Conclusion




 ample memory


Conclusion





  ample memory

Conclusion




  ample memory

Conclusion




Conclusion




 ample memory


Conclusion




Conclusion





Conclusion



Conclusion







Thanks!