

Concurrent Expandable AMQs on the Basis of Quotient Filters

Conference Presentation · June 17, 2020 Tobias Maier, Peter Sanders, Robert Williger

INSTITUTE OF THEORETICAL INFORMATICS · ALGORITHMICS GROUP



KIT – University of the State of Baden-Wuerttemberg and National Laboratory of the Helmholtz Association

www.kit.edu



















- insert
- contains



contains

false

constructed with capacity and false positive rate

insert













similar to hashing with *linear probing* **BUT sorted by** *fingerprint*





























































status bits of changed slots need to be updated





limited growing





Quotient Filter – Growing Karlsruhe Ins limited growing Fingerprint $f(\cdot)$ f(e) = 11010е quotient remainder q bits quotient r bits remainder q' = q + 1 r' = r - 1• ٠





limited growing





limited growing







limited growing







limited growing

unlimited growing (scalable quotient filter)



growing exponentially

therefore logarithmic number of tables

exponentially smaller fp rates

overall bounded fp rate





limited growing







limited growing







limited growing







limited growing







limited growing







limited growing







limited growing







limited growing





Quotient Filter – Trivial Locking





Quotient Filter – Trivial Locking























thrd-1 find *x x* must be in 2nd run







thrd-1 find *x x* must be in 2nd run

thrd-2 insert y







thrd-1 find x x must be in 2nd run

thrd-2 insert y



















without status bits assuming no remainder = 0

- no run reconstruction no sorted order
- linear probing like insertion

• more comparisons = more false positives? $E[fp \ rate] = \frac{1}{2} \left(1 + \frac{1}{1 - \frac{n}{m}}\right) \cdot \frac{1}{2^{r+3} - 1}$

counteract that with 3 more fingerprint bits (same slot size)







without status bits assuming no remainder = 0

- no run reconstruction
 no sorted order
- linear probing like insertion

• more comparisons = more false positives? $E[fp \ rate] = \frac{1}{2} \left(1 + \frac{1}{1 - \frac{n}{m}} \right) \cdot \frac{1}{2^{r+3} - 1}$

counteract that with 3 more fingerprint bits (same slot size)







counteract that with 3 more fingerprint bits (same slot size)







1 0 1 and 1 are impossible status-bit-combinations 0 0 i.e., new cluster but no new run















Institute of Theoretical Informatics

Algorithmics Group





1 0 1 and 1 are impossible status-bit-combinations 0 0 i.e., new cluster but no new run







1 0
 1 and 1 are impossible status-bit-combinations
 0 0 i.e., new cluster but no new run







- 3. execute operation
 - 1 \mathbf{O} and 1 are impossible status-bit-combinations 1 0 0 i.e., new cluster but no new run







1 0 1 and 1 are impossible status-bit-combinations 0 0 i.e., new cluster but no new run







4. release both locks

1 0
 1 and 1 are impossible status-bit-combinations
 0 0 i.e., new cluster but no new run









Experimental Results





executed with 80 threads 4-socket Intel Xeon Gold 6138 (20 cores per socket)



Experimental Results





executed with 80 threads 4-socket Intel Xeon Gold 6138 (20 cores per socket)



Experimental Results







Conclusion



- trivial (external) locking is not enough
- Iocal locking with inherent status bits previously unused combinations
- Iock-free linear probing quotient filter more bits counteract more comparisons
- growing implementations for status bit variants
- unlimited growing combines migration and multi table

