



**TECHNION**  
Israel Institute  
of Technology

# Adaptive Software Cache Management

---

Gil Einziger<sup>1</sup>, **Ohad Eytan**<sup>2</sup>, Roy Friedman<sup>2</sup> and Ben Manes<sup>3</sup>

June 3, 2019

Middleware '18

---

<sup>1</sup>Ben Gurion University of the Negev & Nokia Bell Labs

<sup>2</sup>Technion - Israel Institute of Technology

<sup>3</sup>Independent

# The Essence of Caching

---

# The Essence of Caching



# The Essence of Caching



- A **fast** but relatively **small** memory
- Can temporarily store some items of the "real storage"
- Improves **performance** if **hit-ratio** is high



# LRU [5]

## Least Recently Used

- Idea: recently requested items probably will be requested again
- Policy: evict the oldest item from the cache
- Simple & efficient
- Easily polluted



# LFU [5]

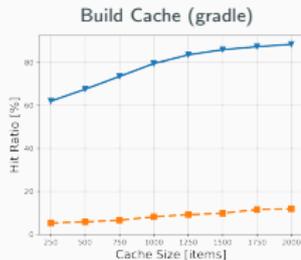
## Least Frequently Used

- Idea: most popular items probably will be requested again
- Policy: evict the item with the lowest access count from cache
- Complex to implement efficiently
- No freshness mechanism



# Problem

- Different workloads have different access patterns:
  - ★ Some are recency biased
  - ★ Some are frequency biased.
  - ★ In fact, most are mixed.



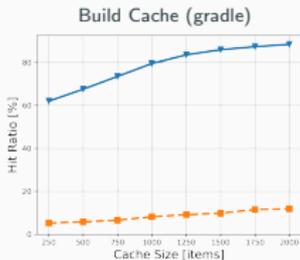
➡ LRU



➡ LFU

# Problem

- Different workloads have different access patterns:
  - ★ Some are recency biased
  - ★ Some are frequency biased.
  - ★ In fact, most are mixed.



➡ LRU

➡ LFU

- Can we develop a silver bullet policy?



# Modern Cache Management Policies

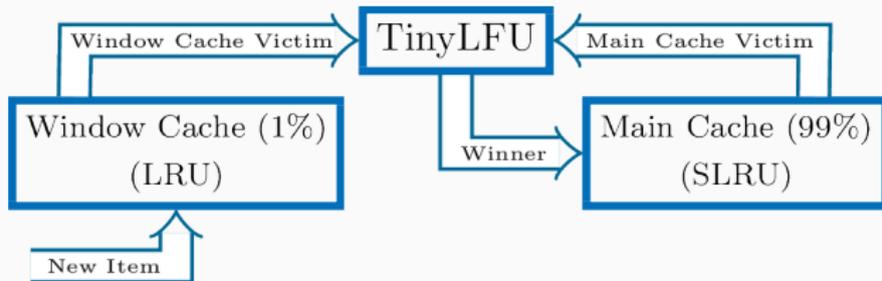
---





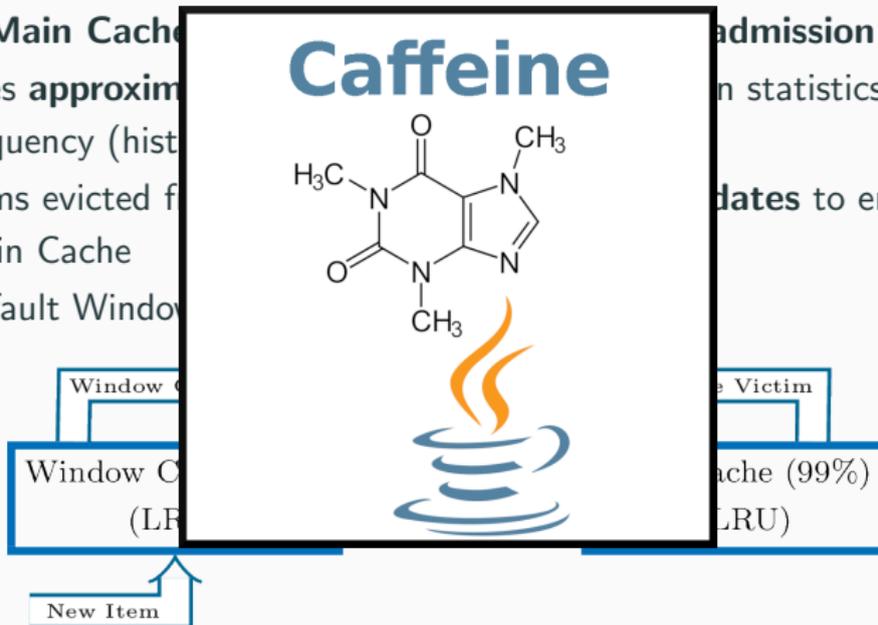
## W-TinyLFU [2]

- The cache consists of two areas:
  - ★ **Window Cache** which is a simple LRU cache
  - ★ **Main Cache** which is a SLRU cache with an **admission policy**
- Uses **approximate counting scheme** to maintain statistics of items frequency (histogram) with **periodic aging**
- Items evicted from the Window Cache are **candidates** to enter the Main Cache
- Default Window Cache is **1%** of the cache



# W-TinyLFU [2]

- The cache consists of two areas:
  - ★ **Window Cache** which is a simple LRU cache
  - ★ **Main Cache**
- Uses **approximate** frequency (histogram)
- Items evicted from Window Cache go to Main Cache
- Default Window Cache size is 1000 items



admission policy

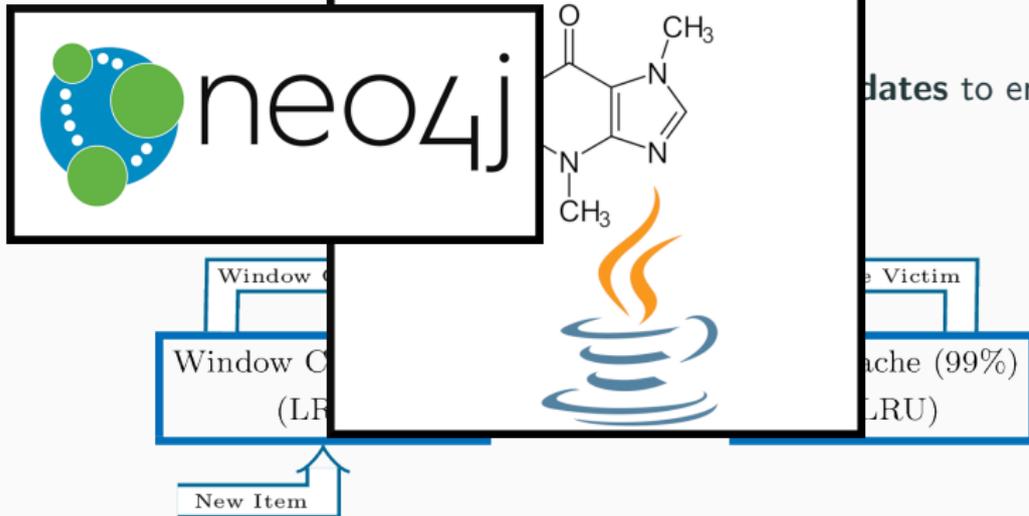
in statistics of items

dates to enter the

# W-TinyLFU [2]

- The cache consists of two areas:
  - ★ **Window Cache** which is a simple LRU cache
  - ★ **Main Cache**
- Uses **approxim**

admission policy  
n statistics of items  
dates to enter the

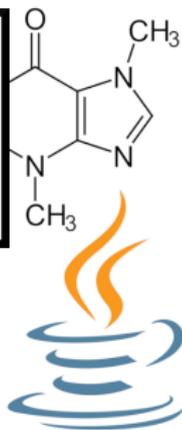


# W-TinyLFU [2]

- The cache consists of two areas:
  - ★ **Window Cache** which is a simple LRU cache
  - ★ **Main Cache** which uses an admission policy and maintains statistics of items
- Uses approx



Caffeine



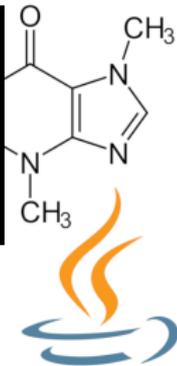
# W-TinyLFU [2]

- The cache consists of two areas:
  - ★ **Window Cache** which is a simple LRU cache
  - ★ **Main Cache**
- Uses **approxim**

admission policy

n statistics of items

## Caffeine



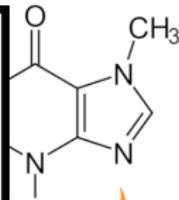
# W-TinyLFU [2]

- The cache consists of two areas:
  - ★ **Window Cache** which is a simple LRU cache
  - ★ **Main Cache**
- Uses approxi

admission policy

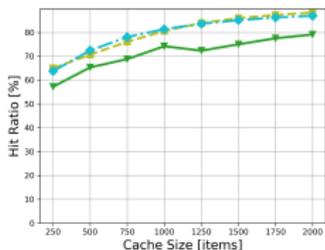
n statistics of items

## Caffeine

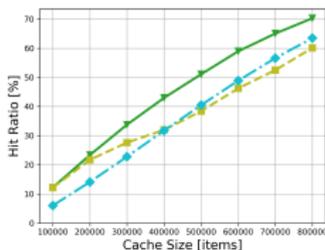


# Unsatisfying

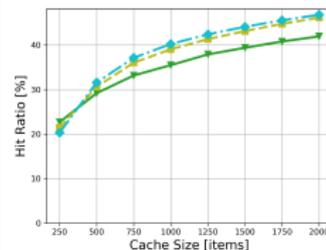
## Build cache (gradle)



## Search engine (S3)



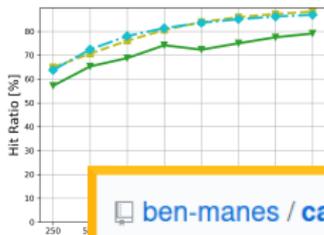
## Financial (OLTP)



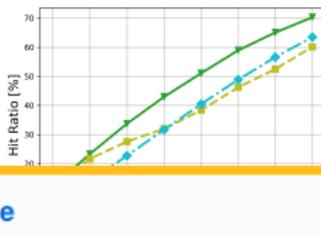
▼ W-TinyLFU (1%)    ■ ARC    ◆ FRD

# Unsatisfying

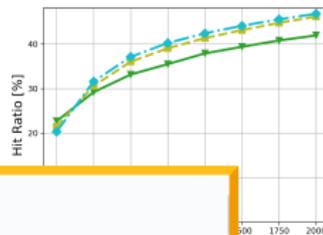
## Build cache (gradle)



## Search engine (S3)



## Financial (OLTP)



 [ben-manes / caffeine](#)

 Code

 Issues **14**

 Pull requests **3**

 Projects **0**

 Wiki

## cache performing worse than LRU #106

 Open

phrakte opened this issue on Aug 7, 2016 · 38 comments

# Our Adaptive Caching

---

- Dynamically adjust a selected tuning parameter

- Dynamically adjust a selected tuning parameter
- Suggested tuning **parameters**:
  - ★ For W-TinyLFU: change the ratio between the cache areas
  - ★ For W-TinyLFU: change the sketch increment parameter
- Suggested **adaptation approaches**:
  - ★ Hill climbing: try and see what happens
  - ★ Indicator: track statistics and decide directly
- We end up with **4** suggested policies

# Parameters: Areas Ratio

---

# Parameters: Areas Ratio

- The partition between the cache areas implies a trade-off between recency and frequency:

★ Frequency biased configuration:



★ Recency biased configuration:



# Parameters: Areas Ratio

- The partition between the cache areas implies a trade-off between recency and frequency:

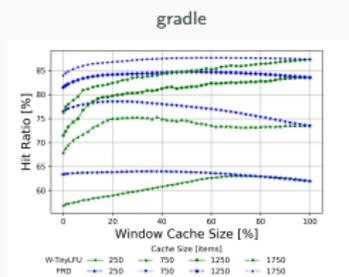
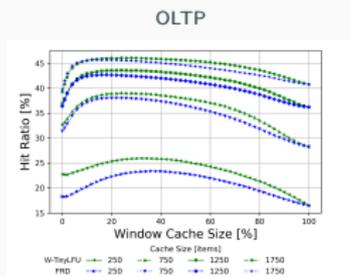
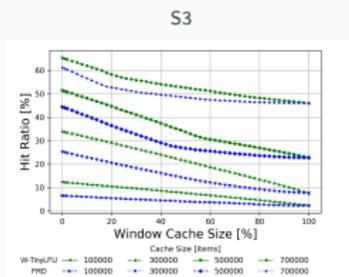
★ Frequency biased configuration:



★ Recency biased configuration:

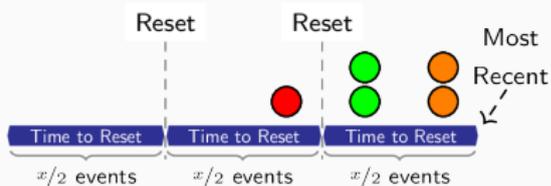


- Very effective:

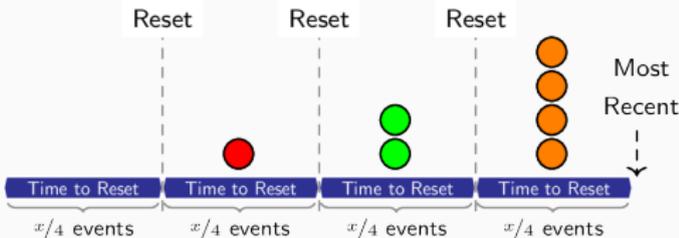


# Parameters: Sketch

- TinyLFU sketch:
  - ★ **Aging** mechanism divides all the counters by 2 each  $S$  steps.
  - ★ The counters are **bounded** by 16.
- Enlarging the **counters increment** on each item's access from 1 to a larger value **favors recency**:
  - ★ Increment of 2:

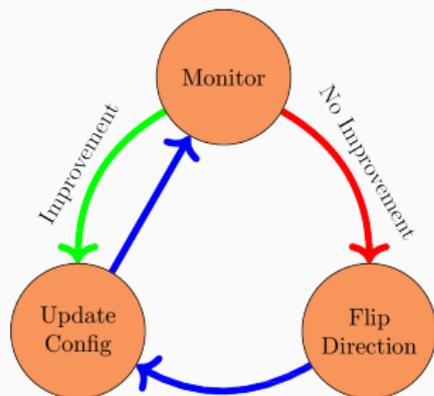


- ★ Increment of 4:



# Adaptation Techniques: Hill Climbing

- Well known optimization technique:



- Step size: 5% or 1.
- Almost no overhead.
- Constantly changes.

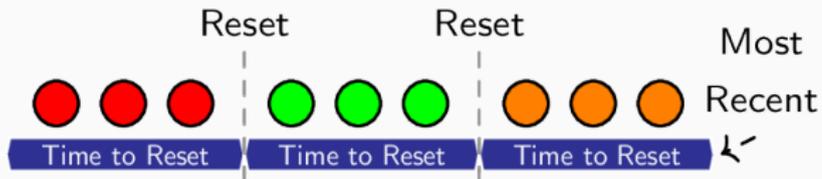
# Adaptation Techniques: Indicator

- Composed from two ingredients:
  - ★ **Hint** - the average of the sketch estimation for all of the accesses.



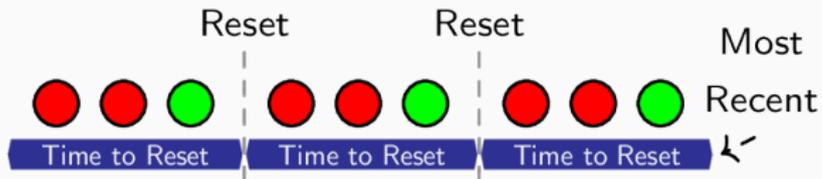
# Adaptation Techniques: Indicator

- Composed from two ingredients:
  - ★ **Hint** - the average of the sketch estimation for all of the accesses.



# Adaptation Techniques: Indicator

- Composed from two ingredients:
  - ★ **Hint** - the average of the sketch estimation for all of the accesses.



# Adaptation Techniques: Indicator

- Composed from two ingredients:
  - ★ **Hint** - the average of the sketch estimation for all of the accesses.



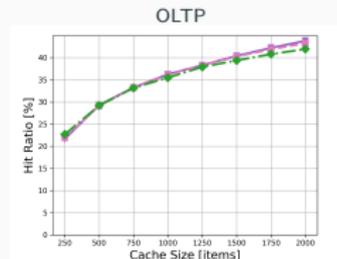
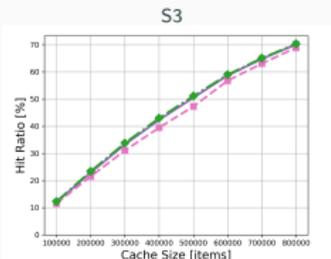
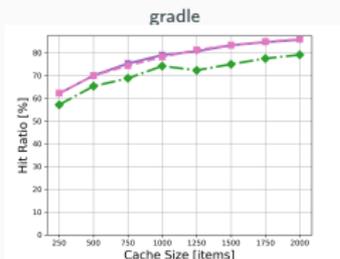
- ★ **Skew** - an estimation of the skewness of the items.
- We define:

$$indicator \triangleq \frac{hint \cdot (1 - \min \{1, skew^3\})}{maxFreq}$$

Which gives us a value in  $[0, 1]$ .



## Adaptive W-TinyLFU Sketch:

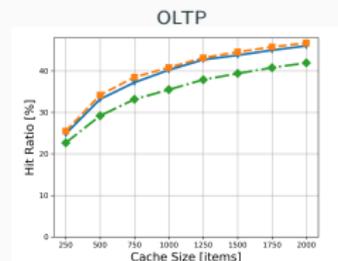
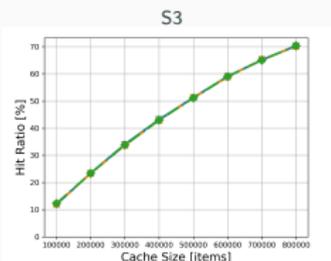
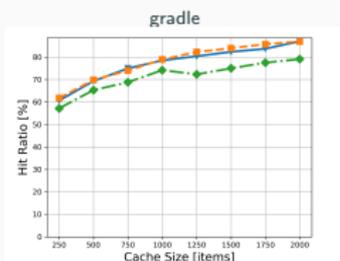


SC-W-TinyLFU

SI-W-TinyLFU

W-TinyLFU (1%)

## Adaptive W-TinyLFU Window:



WC-W-TinyLFU

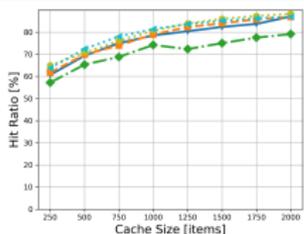
WI-W-TinyLFU

W-TinyLFU (1%)

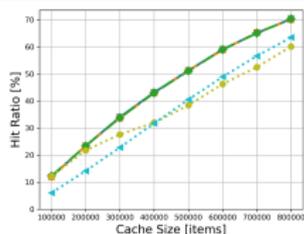
# Results

Competitive for all tested workloads: ✓

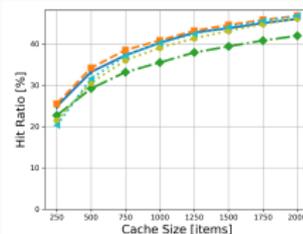
gradle



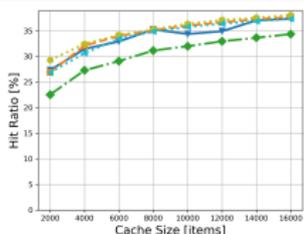
S3



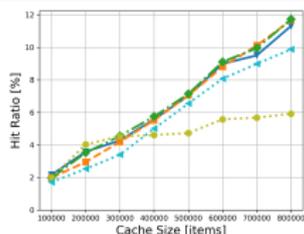
OLTP



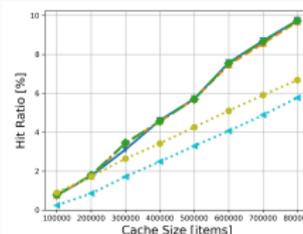
F1



DS1



WS1

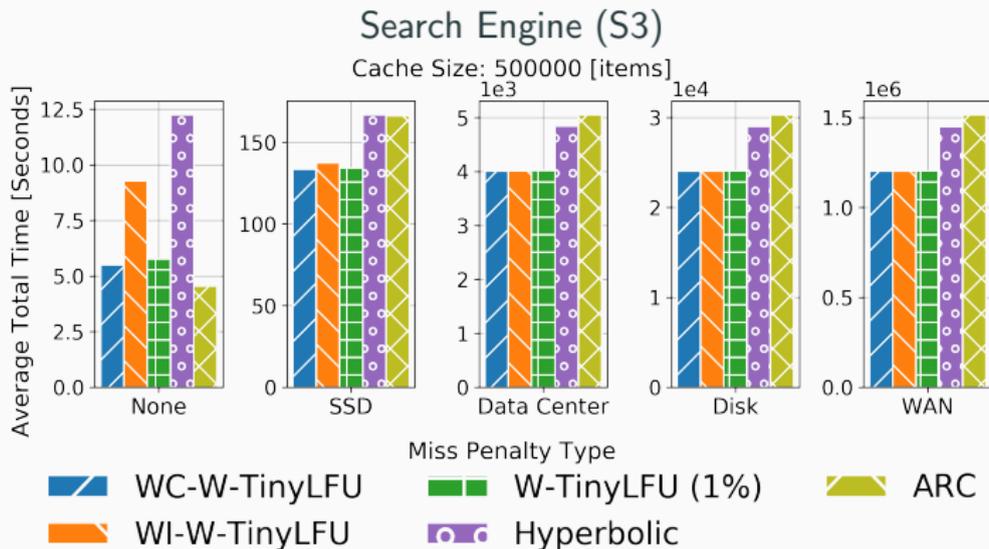


WC-W-TinyLFU  
WI-W-TinyLFU

W-TinyLFU (1%)  
ARC

FRD

# Results: Completion Time



# Conclusions

- Adaptation works 😊
- Window adaptation better than sketch adaptation
- Indicator adapt quicker

## **But**

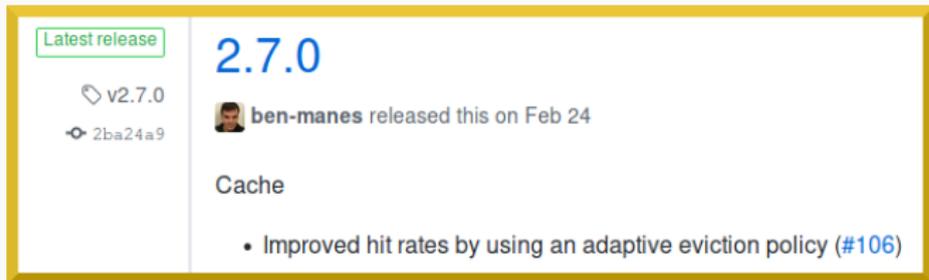
Hill climber is simpler to implement and requires no extra space

# Conclusions

- Adaptation works 😊
- Window adaptation better than sketch adaptation
- Indicator adapt quicker

## But

Hill climber is simpler to implement and requires no extra space



The screenshot shows a GitHub release page for version 2.7.0. On the left, there is a 'Latest release' badge, the version number 'v2.7.0', and a commit hash '2ba24a9'. The main content area features the version number '2.7.0' in large blue text, followed by the release information: 'ben-manes released this on Feb 24'. Below this is a 'Cache' section with a bullet point: 'Improved hit rates by using an adaptive eviction policy (#106)'.

**Thank You**  
**Questions/Ideas?**

**P.S. If you could share a trace with variable item sizes for further research, please contact me at [ohadey@cs.technion.ac.il](mailto:ohadey@cs.technion.ac.il)**

# References i



L. A. Belady.

**A study of replacement algorithms for a virtual-storage computer.**

*IBM Systems Journal*, 5(2):78–101, 1966.



G. Einziger, R. Friedman, and B. Manes.

**Tinylfu: A highly efficient cache admission policy.**

*CoRR*, abs/1512.00727, 2015.



N. Megiddo and D. S. Modha.

**ARC: A self-tuning, low overhead replacement cache.**

In J. Chase, editor, *Proceedings of the FAST '03 Conference on File and Storage Technologies, March 31 - April 2, 2003, Cathedral Hill Hotel, San Francisco, California, USA*. USENIX, 2003.



S. Park and C. Park.

**Frd: A filtering based buffer cache algorithm that considers both frequency and reuse distance.**

2017.



A. Silberschatz, P. B. Galvin, and G. Gagne.

**Operating system concepts, 7th Edition.**

Wiley, 2005.