
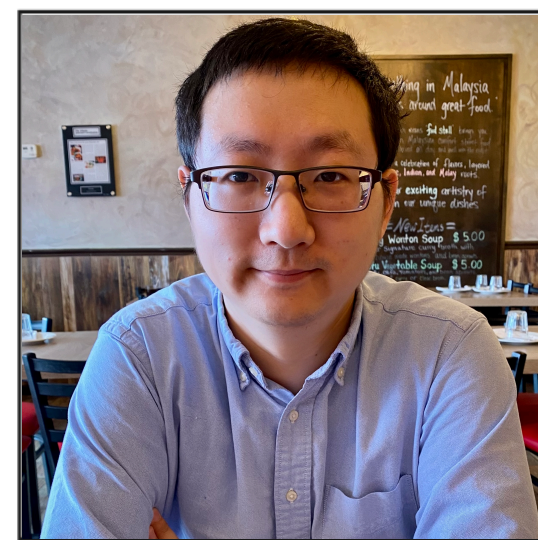


# SIEVE is Simpler than LRU:

An Efficient Turn-Key Eviction Algorithm for Web Caches





Yazhuo Zhang, Juncheng Yang, Yao Yue, Ymir Vigfusson, K.V. Rashmi

*Emory University, Carnegie Mellon University, Pelikan Foundation*

# Caching System is Important

## Page Cache



Cachelib



Pelican



## Web Caches

Limited Space!


Core:  
Eviction Algorithm

# Cache Metrics


Efficiency

Scalability

*Cache Miss Ratio*

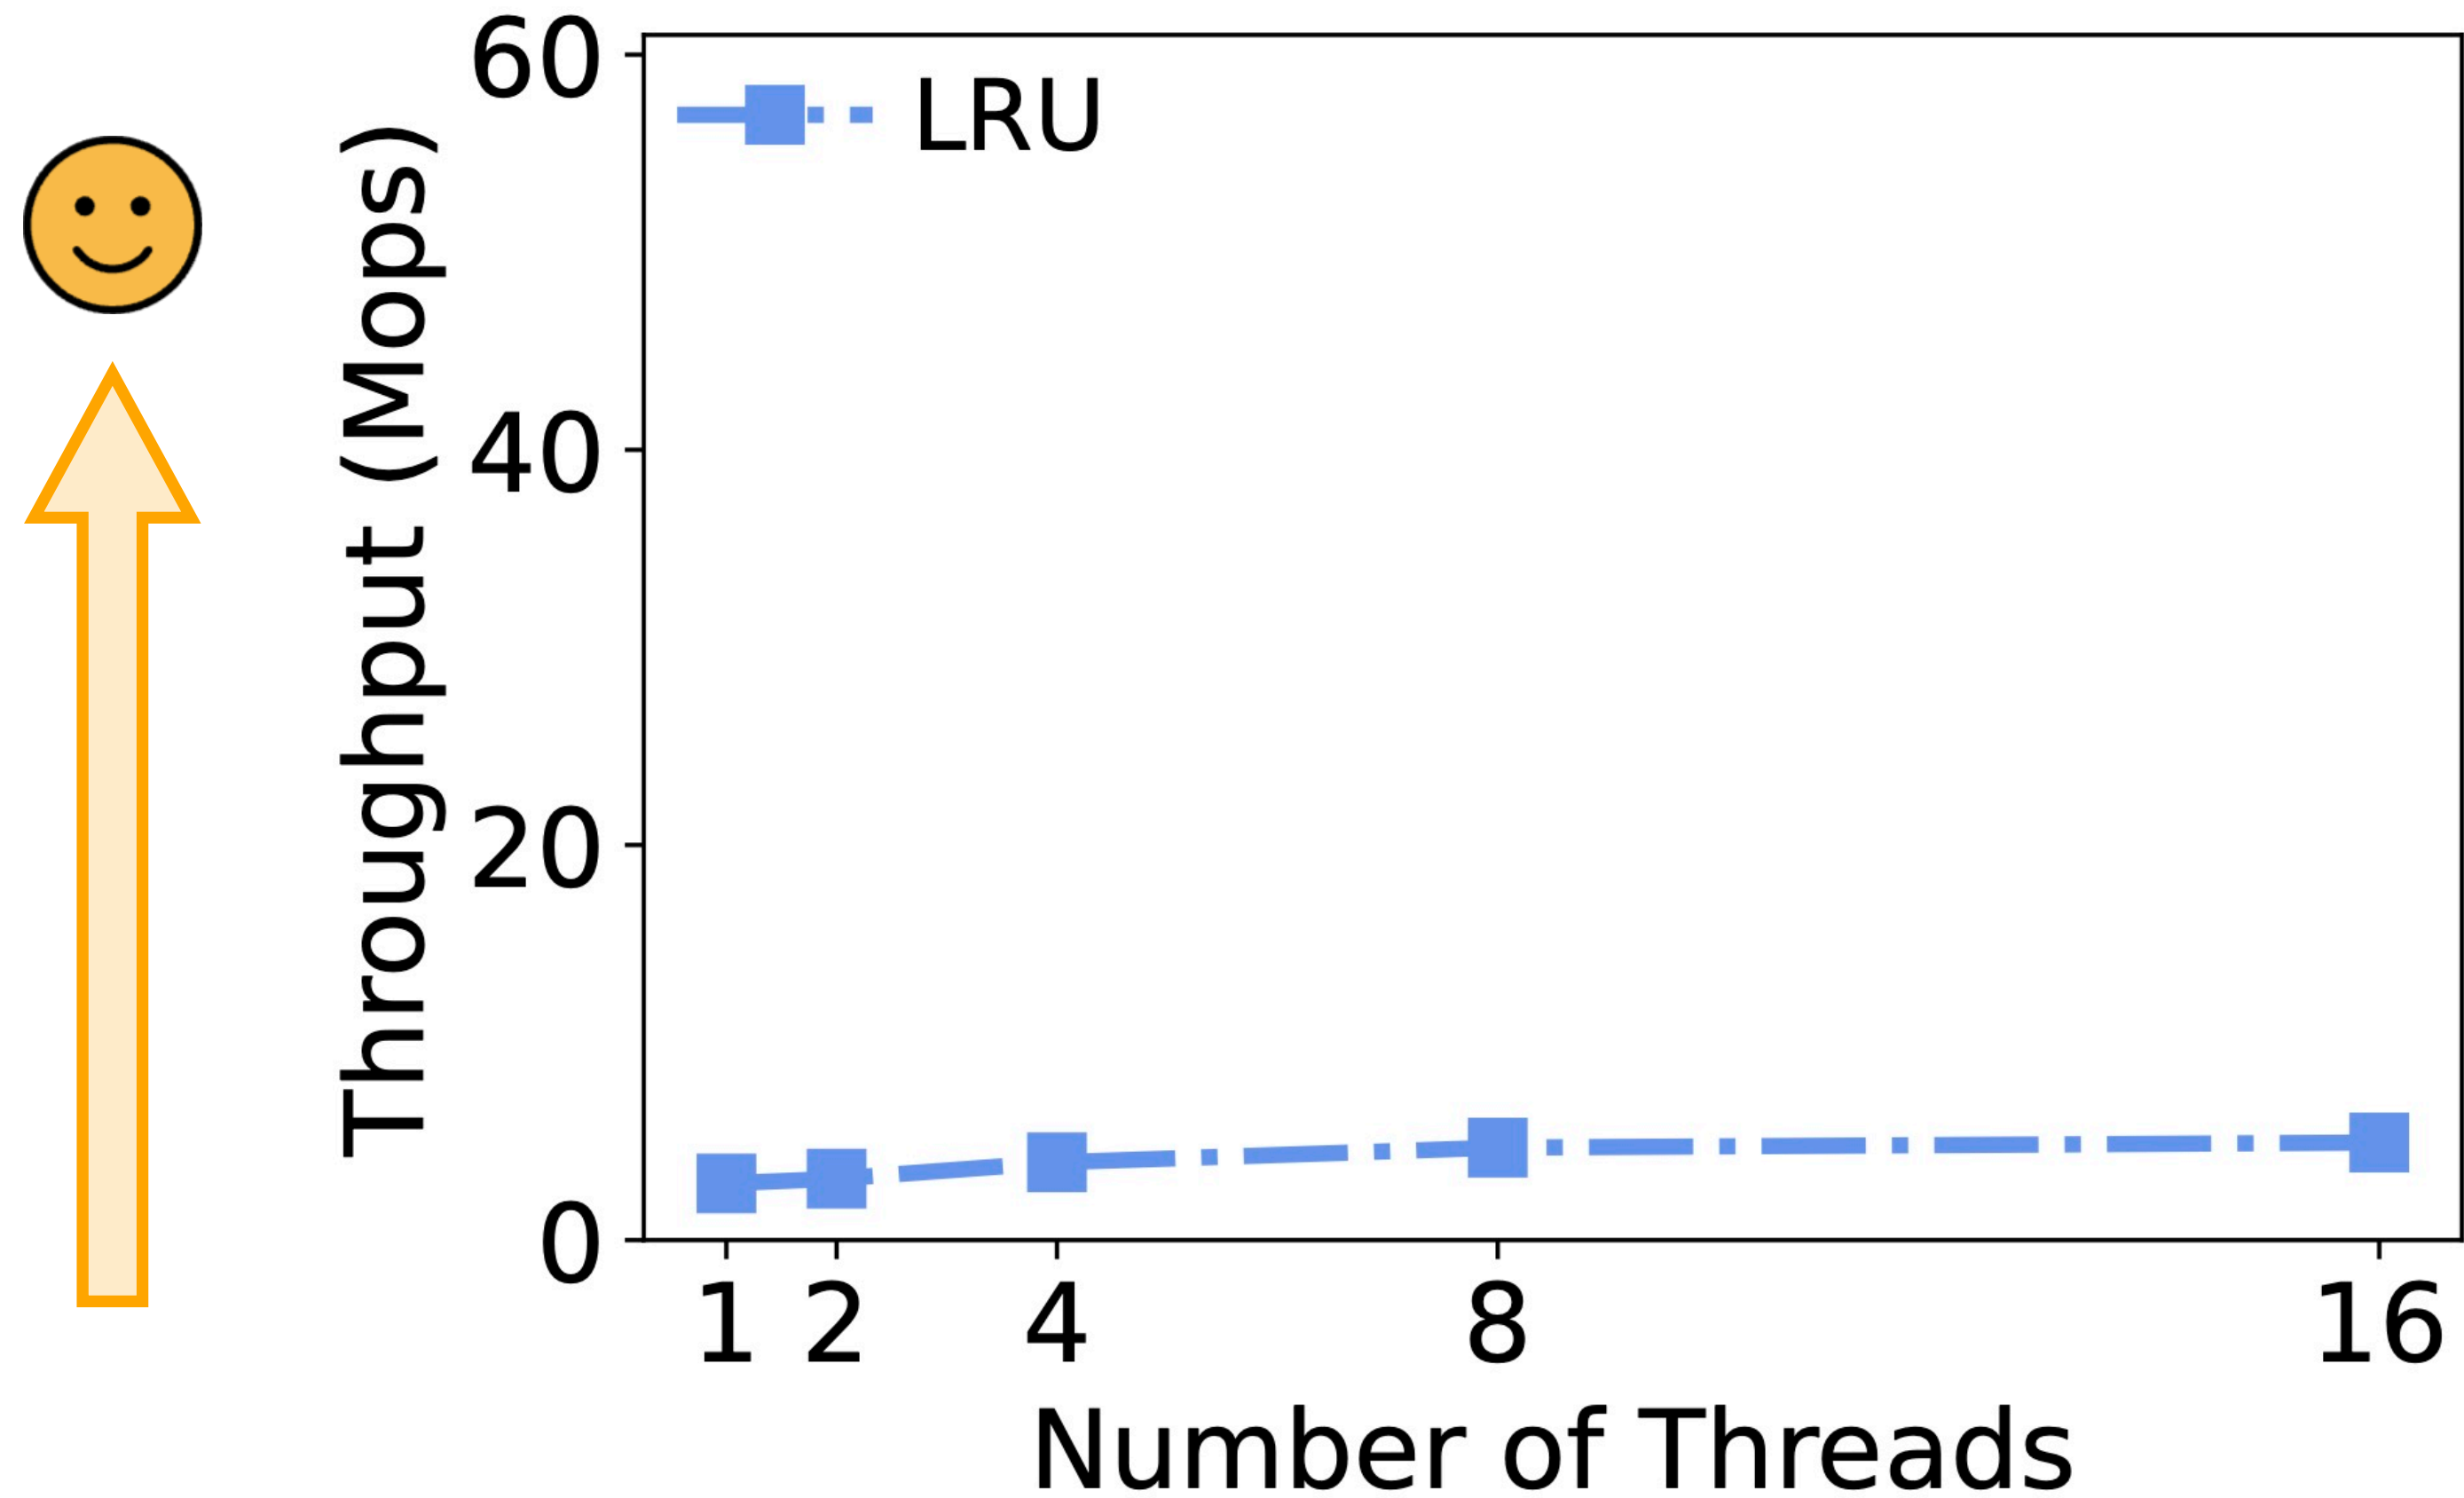


*Reqs/Second*




# Throughput Measured in Cachelib

*Twitter workload*



# Throughput Measured in Cachelib

*Twitter workload*




# Cache Metrics

Efficiency


Scalability

Simplicity


*Cache Miss Ratio*



*Reqs/Second*



# A Rich Literature of Eviction Algorithms



# The Trouble with Complexity


- Difficult to debug and maintain
- Difficult to tune the parameters

*“Predicting which pages will be accessed in the near future is tricky, and the kernel has evolved many mechanisms to improve its chances of guessing right. But the kernel **not only often gets it wrong**, but also spends a lot of CPU time to make the incorrect choice.”*

*-- Linux kernel developer*



# SIEVE: a Simple and Efficient Cache Eviction Algorithm



\* Measured by lines of code

† Measured by average object miss ratio reduction from FIFO

# SIEVE Design

# The Secret to Designing Efficient Eviction Algorithms



Lazy promotion



Quick demotion


[HotOS'23] FIFO queues can be better than LRU: the Power of Lazy Promotion and Quick Demotion

[SOSP'23] FIFO Queues are all You Need for Cache Eviction

# The Secret to Designing Efficient Eviction Algorithms



## Lazy promotion




Retain popular objects with minimal effort

- Improve throughput due to less computation
- Improve efficiency due to more information at eviction

[HotOS'23] FIFO queues can be better than LRU: the Power of Lazy Promotion and Quick Demotion

[SOSP'23] FIFO Queues are all You Need for Cache Eviction

# The Secret to Designing Efficient Eviction Algorithms




Quickly remove most new objects,  
such as one-hit-wonders (no request after insertion)

[HotOS'23] FIFO queues can be better than LRU: the Power of Lazy Promotion and Quick Demotion

[SOSP'23] FIFO Queues are all You Need for Cache Eviction

# The Secret to Designing Efficient Eviction Algorithms



## Lazy promotion

Retain popular objects with minimal effort



## Quick demotion

Remove unpopular objects fast, such as one-hit-wonders

# Efficiency

# Scalability



- Eager promotion
- No quick demotion



- No promotion
- No quick demotion



# Efficiency


# Scalability



- Eager promotion
- No quick demotion



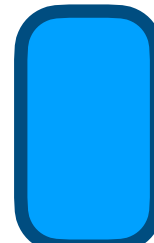
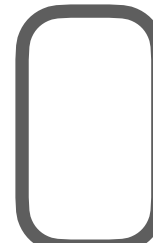
- No promotion
- No quick demotion



- Lazy promotion
- No quick demotion

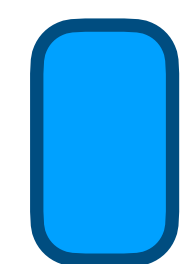
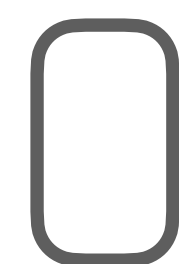




 visited     not visited

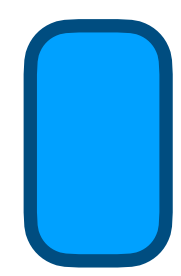
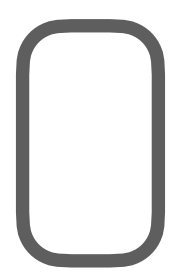
cache hit on D



 visited     not visited

# cache miss



 visited     not visited

# Efficiency


# Scalability



Eager promotion  
No quick demotion



No promotion  
No quick demotion



Lazy promotion  
No quick demotion



# Efficiency


# Scalability



Eager promotion  
No quick demotion



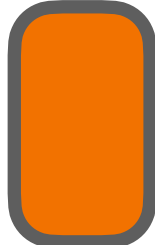
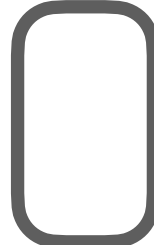
No promotion  
No quick demotion



Lazy promotion  
No quick demotion


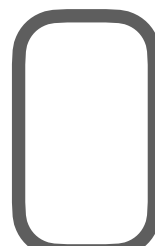




 visited     not visited

cache hit on D



 visited     not visited

# cache miss

reset visited bit & move hand

1



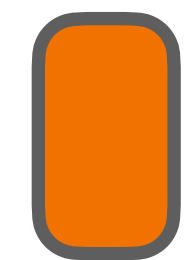
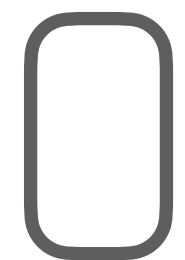
2



3

insert



 visited  not visited



# Efficiency


# Scalability



Eager promotion  
No quick demotion



No promotion  
No quick demotion



Lazy promotion  
No quick demotion



Lazy promotion






# Quickly remove new objects

insert





# Separate new and old objects



 "survived" object

 newly inserted object

# Efficiency


# Scalability



Eager promotion  
No quick demotion



No promotion  
No quick demotion



Lazy promotion  
No quick demotion



Lazy promotion  
Quick demotion




# SIEVE Evaluation

# Web Cache Workloads


1559 traces | 247,017 million requests | 14, 852 million objects

- Simulator: libCacheSim
- Prototype: Cachelib
- Testbed: Cloudfab




trace collection	collection time	#traces	cache type	# request (million)	# object (million)
CDN1	2021	1273	object	37,460	2,652
CDN2	2018	219	object	3,728	298
Tencent Photo	2018	2	object	5,650	1,038
Wiki CDN	2019	3	object	2,863	56
Twitter KV	2020	54	KV	195,441	10,560
Meta KV	2022	5	KV	1,644	82
Meta CDN	2023	3	object	231	76

# SIEVE: Efficiency




SIEVE reduces FIFO's miss ratio by more than 42% on 10% of the traces (top whisker) with a mean of 21%




CDN1, 1273 traces (37,460 million requests)

# SIEVE: Efficiency





**SIEVE achieves the best efficiency on the well-studied Zipfian workloads**



CDN1, 1273 traces (37,460 million requests)



# SIEVE: Throughput



# SIEVE: Simplicity

Cache library	Language	Lines of change
groupcache	Golang	21
mnemonist	Javascript	12
lru-rs	Rust	16
lru-dict	Python + C	21

# SIEVE: Simplicity

Cache library	Language	Lines of change
groupcache	Golang	21
mnemonist	Javascript	12
lru-rs	Rust	16
lru-dict	Python + C	21

## Adoption

Large systems:  Pelikan  Nyrkiö  SkiftOS  DragonFly

 DNSCrypt-proxy  encrypted-dns-resolver

Cache libraries:  golang-fifo  js-sieve  rust-sieve-cache  go-sieve

 sieve\_cache (Ruby)  zig-sieve (Zig)  sieve (Swift)

 sieve (JavaScript)  sieve (Elixir)  sieve (Nim)

 sieve-cache (Java)  sieve (Python)  sieve-cache-in-rust


 sieve-cache (JavaScript)  gosieve,  sieve (typescript)

# SIEVE: Primitive

LeCaR: LRU + LFU + ML

TwoQ: LRU + FIFO

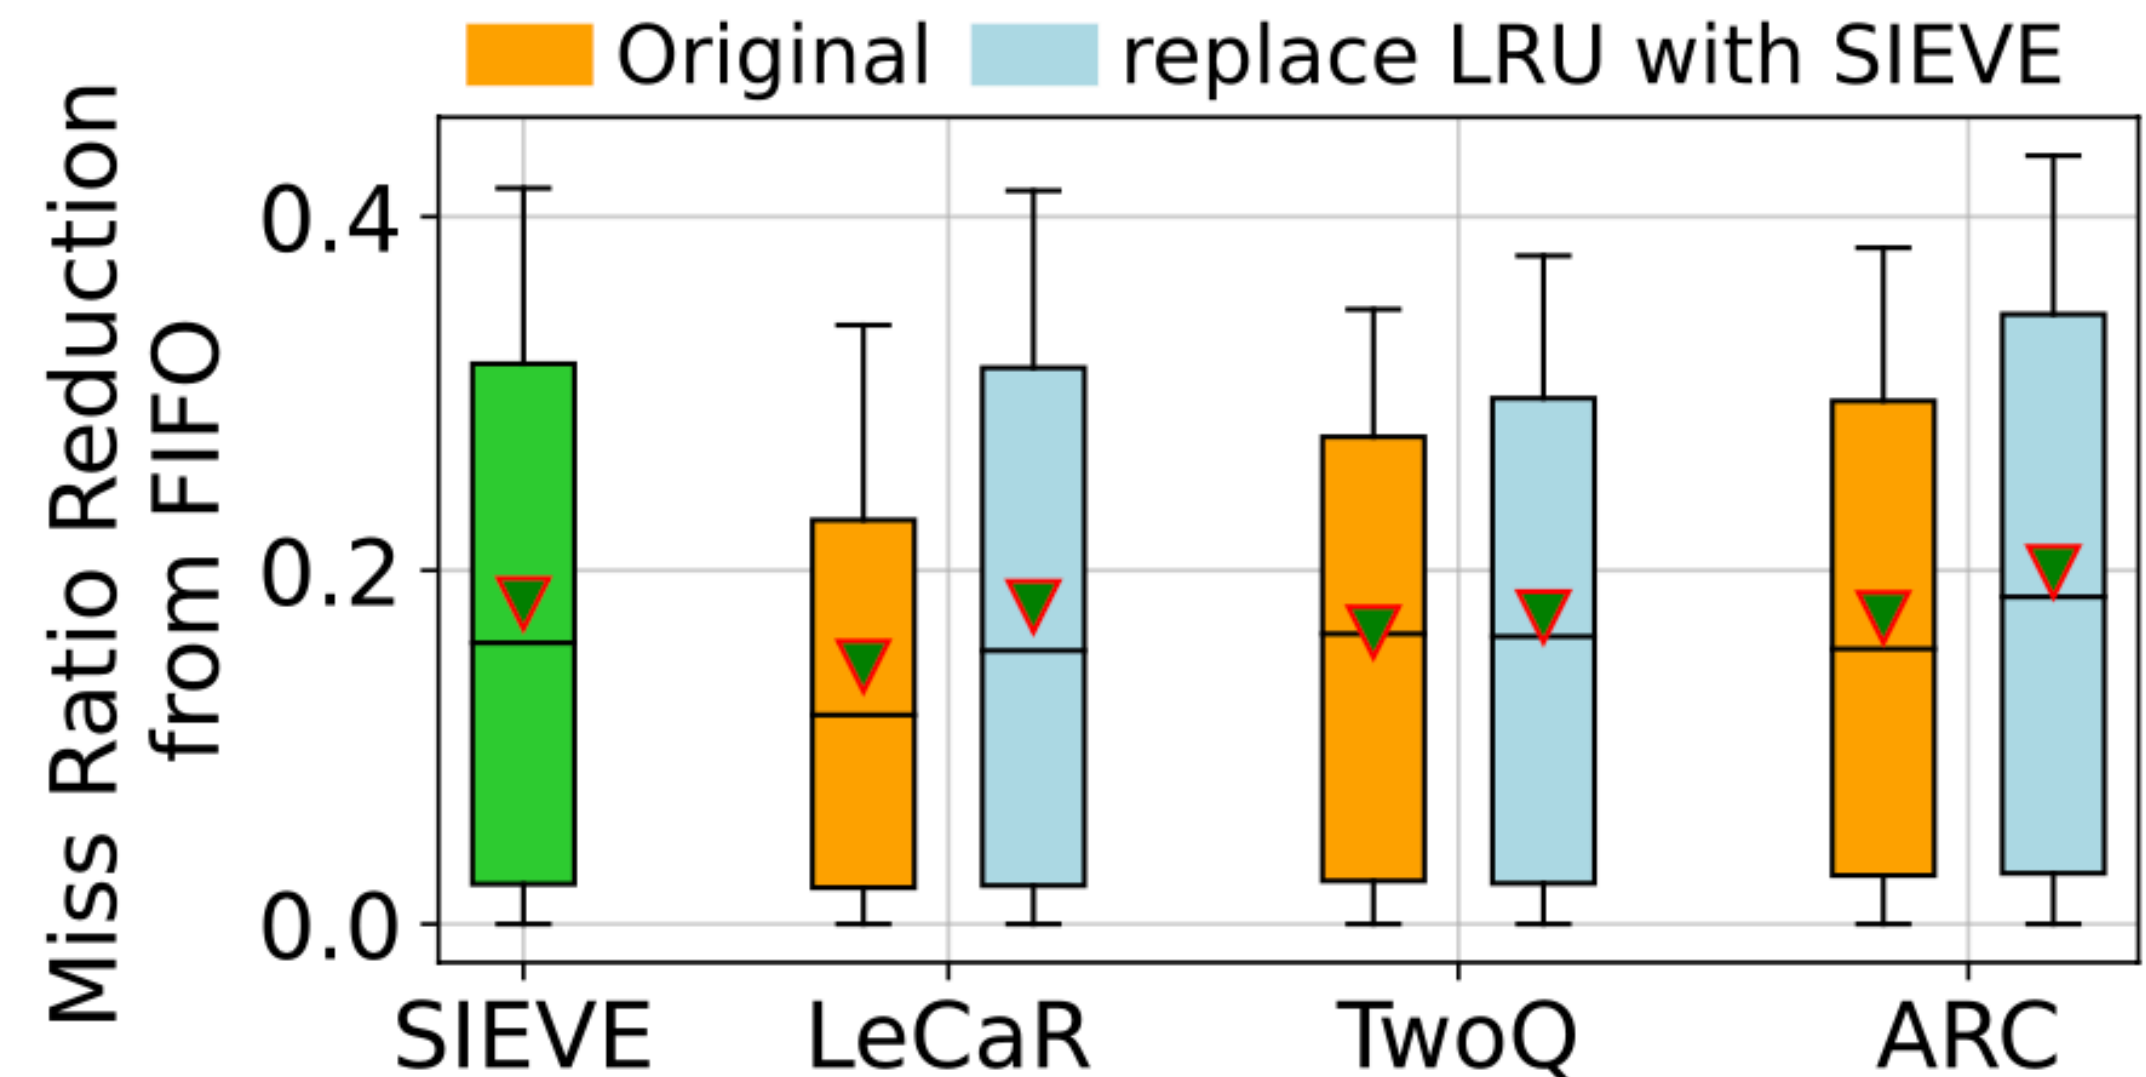
ARC: LRU + LRU + 2 ghost queues



# SIEVE: Primitive

LeCaR: LRU + LFU + ML  
TwoQ: LRU + FIFO  
ARC: LRU + LRU + 2 ghost queues

Replace LRU with SIEVE




# More in the paper

- Why SIEVE is effective
- Byte miss ratio
- When SIEVE is not effective
- Comparison to ML algorithms

# SIEVE Adoption

- SIEVE is available in over **20** cache libraries with **10+** programming languages
- Production systems start integrating SIEVE: Pelican, SkiftOS, DragonFly, and etc




**Richard Artoul** ✓  
@richardartoul

Turns out Ristretto cache is *\*async\**... I switched WarpStream's footer cache from Ristretto to golang-fifo (Sieve algo) and got a 33x reduction in cache misses and 16% CPU savings...

## Cache Loads

Richard Artoul | Updated 4 minutes ago

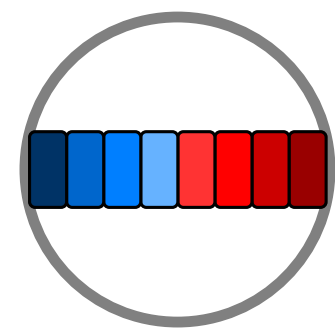
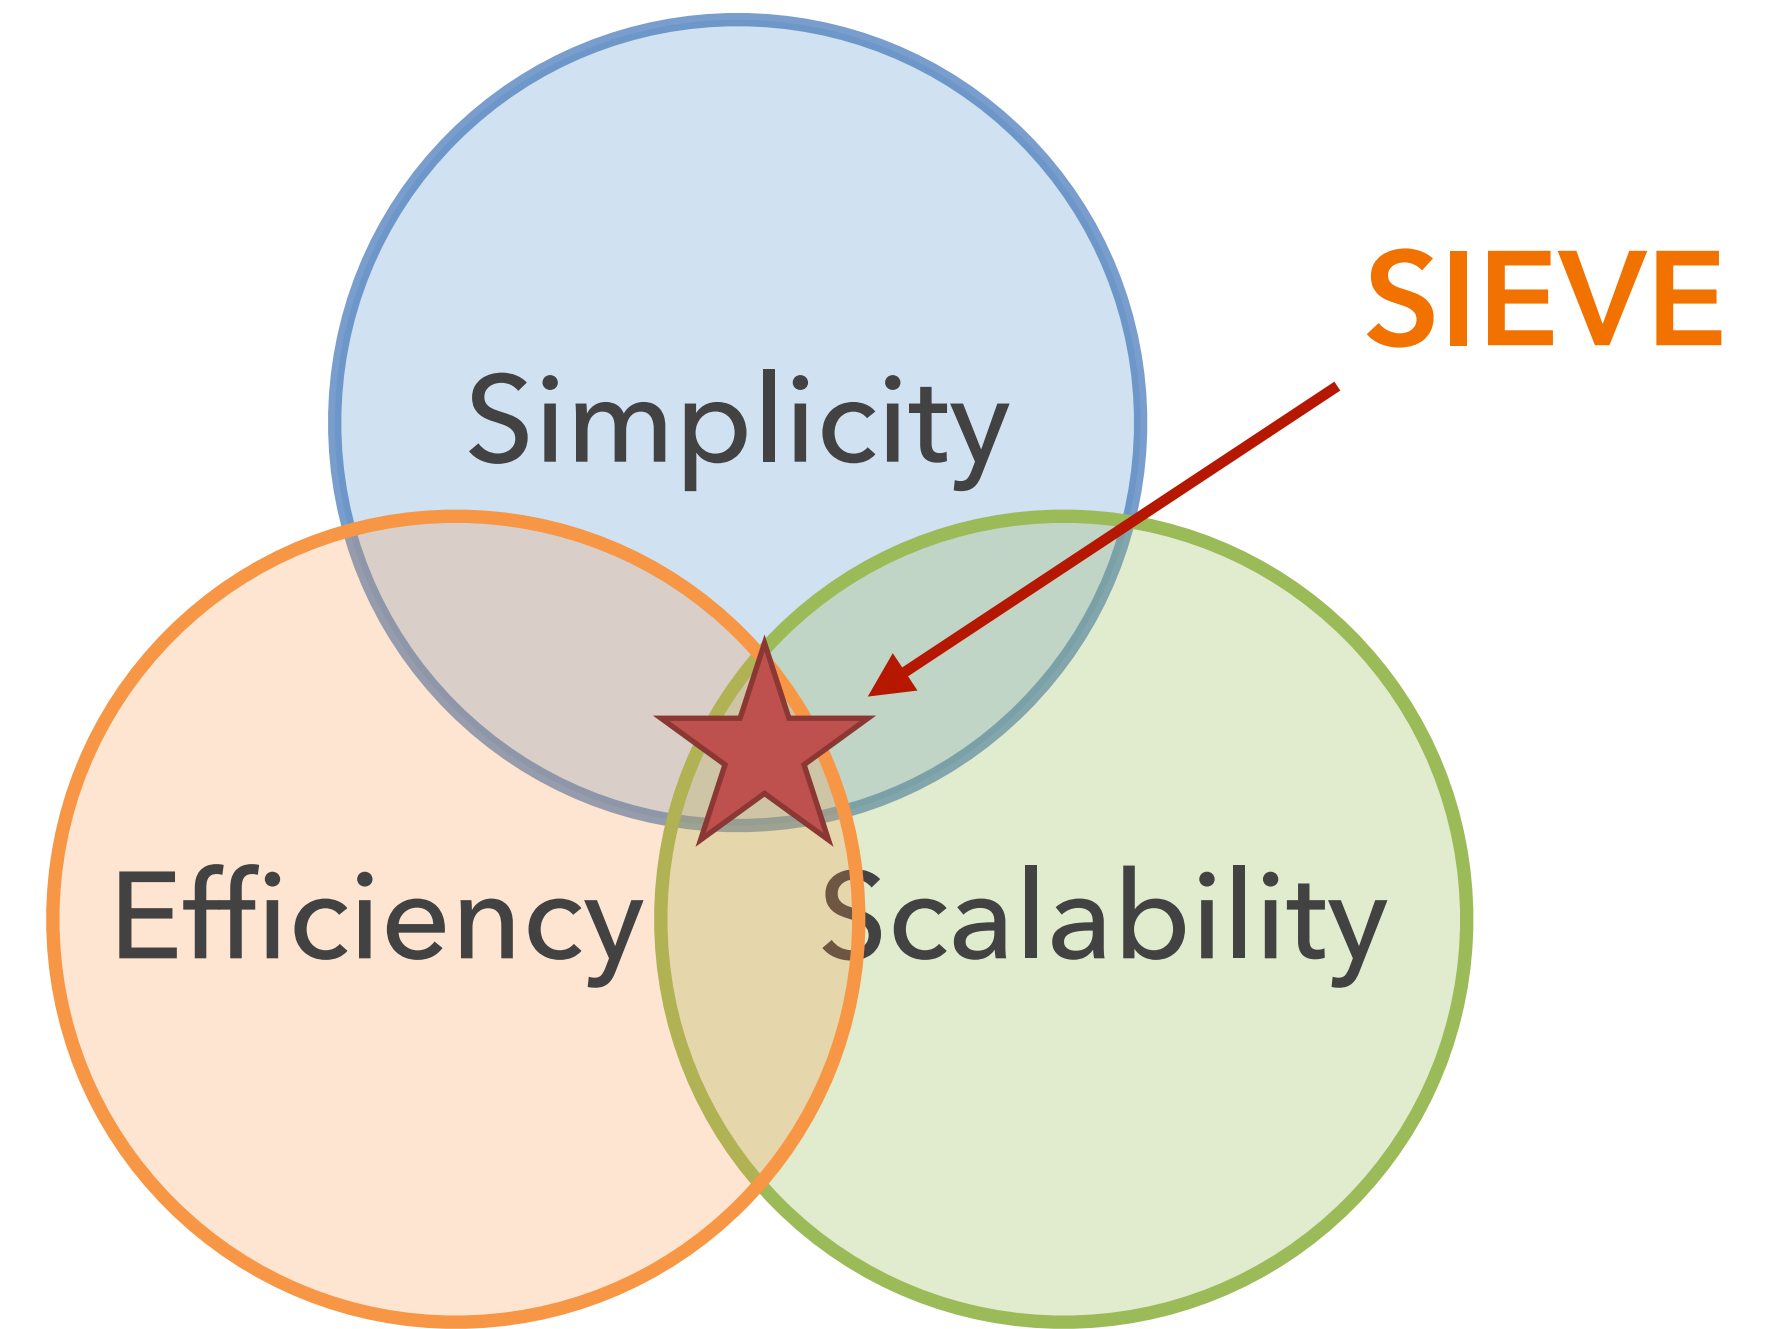


Tags in sum:warpstream_loading_cache_loader_outcome(service:warp-agent,env...	Avg	Min	Max	Sum	Value
cache_name:acls,host:i-0cc491918ccbaf6e7	6e-3 /s	0 /s	0.10 /s	0.20 /s	—
cache_name:acls,host:i-0ddd25eb52bf97b6f	0.01 /s	0 /s	0.30 /s	0.50 /s	—
cache_name:acls,host:i-0ee0b6128679455ed	0.013 /s	0 /s	0.40 /s	0.40 /s	—
cache_name:cluster,host:i-0ddd25eb52bf97b6f	4e-3 /s	0 /s	0.10 /s	0.10 /s	—

9:35 PM · Jan 20, 2024 · 17.3K Views

# Takeaway

- Lazy promotion and quick demotion are key to efficient eviction algorithm
- SIEVE uses a moving hand to 1) retain popular objects in place, and 2) remove unpopular objects quickly
- The simplest algorithm with state-of-the-art efficiency and scalability



<https://sievecache.com>