Filters



This is the advertised talk title.

Michael A. Bender



Time To Change Your Filter



This title is more to the point. (more explanation later).

Michael A. Bender

Now ... what's a filter.



Dictionary Data Structure

A dictionary maintains a set **S** from universe **U**.





A dictionary supports membership queries on S.

Filter Data Structure

A filter is an *approximate* dictionary.



A filter supports approximate membership queries on S.

A Filter Guarantees a False-Positive Rate ε



one-sided errors

False-positive rate enables filters to be compact

space = $\Omega(n \log |U|)$

space $\geq n \log(1/\varepsilon)$







Talk So Far

- Filter data structure
- Next: the Bloom filter [Bloom '70]



















Bloom filter: a bit array + k hash functions. (Here k=2.)



Bloom filters don't support delete.

Issue: on a delete, which 1s get decremented?

Bloom Filter Space Usage [Bloom '70]

Bloom filter space with false-positive rate ε: ≈1.44 lg (1/ε) bits/element.

Example: For $\varepsilon = 2\%$, bits/element ≈ 8 .

Common rule of thumb: Bloom filters take about 1 byte/element.



Bloom filters are ubiquitous

≥4300 citations

Computational biology



Databases



Networking





Storage systems



Streaming applications

Talk So Far

- Filter data structure
- The Bloom filter [Bloom '70]
- How filters are used



Most Common Filter Use

Filter out queries to a large remote dictionary.

Only an ε -fraction of negative queries don't get filtered out.



Filter

local, e.g., in RAM

Dictionary

remote, e.g., on disk

Speedup from Filter Use

Workload has *P* positive and *N* negative queries.

Dictionaries w/o	Dictionaries w/
Filters	Filters
P+N	$P+\varepsilon N$

Remote Accesses of Dictionary

Example: Filters Help Queries in LSM Trees

member(b)?

6

6

Log-structured merge tree (LSM)

• An LSM tree supports fast inserts by partitioning into independent dictionaries. [O'Neil, Cheng, Gawlick, O'Neil '96]

Point queries are slow without filters.



Talk So Far

- Filter data structure
- The Bloom filter [Bloom '70]
- How filters are used
- Time to change your filter (the talk title)



Limitations	Work-arounds
No deletes	Rebuild
No resizes	Guess N, and rebuild if wrong
No filter merging	???
nor enumeration of	
elements	
No values associated	Combine with other data structure
with keys	

Bloom filter limitations increase system complexity, waste space, and slow down application performance.

Bloom filters also have suboptimal guarantees

	Bloom filter	Optimal
Space	≈ 1.44 <i>n</i> lg(1/ε)	≈ <i>n</i> lg(1/ε) + O(1)
CPU cost	Ω(lg(1/ε))	O(1)
Data locality	$\Omega(lg(1/\epsilon))$ probes	O(1) probes

Bloom filter limitations increase system complexity, waste space, and slow down application performance.

Deletes + counting

[Bonomi, Mitzenmacher, Panigrahy, Singh, Varghese 06], [Yuan, Miao, Jia, Wang 08], [Pandey, Bender, Patro, Johnson SIGMOD 17],

Keys

[Chazelle, Kilian, Rubinfeld, Tal 04]

Filters on SSD

[Canim, Mihaila, Bhattacharhee, Lang, Ross 10], [Debnath, Sengupta, Lilja, Du 11], [Lu, Debnath, Du. 11], [Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12], [Pandey, Singh, Bender, Berry, Farach-Colton, Johnson, Kroeger, Phillips 20]

Optimizing asymptotics

[Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10], [Lovett & Porat 10], [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Engineering

[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12],
[Fan, Andersen, Kaminsky, Mitzenmacher 14],
[Pandey, Bender, Patro, Johnson SIGMOD 17],
[Pandey, Bender, Johnson, Patro 17],
[Breslow, Jayasena 18],
[Pandey, Conway, Durie, Bender, Martin, Johnson 21]

Adaptivity

[Mitzenmacher, Pontarelli, Reviriego 18] [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18] [Bender, Das, Farach-Colton, Mo, Wang 21]

When too much stuff is growing on your filter, it's time to change the filter.

Talk So Far

- Filter data structure
- The Bloom filter [Bloom '70]
- How filters are used
- Time to change your filter (the talk title)



"I see that you change the filter about as often as you change your mind."

I couldn't find the right talk title.

Then I received this email.





I couldn't find the right talk title.

Then I received this email.



🔦 Reply

Forward

M	Gmail	Q time to change your filter	×	-	?	(3)	* * * * * * * * *
	Any time 💌 🖙	Has attachment > To me Sunread	Advanc	ed seai	rch		
	- C :			1-31	of 31	<	>
	no-reply	Product Reminder - com" width="240" b	order="0">	∙ It's Tir	m	Sep	17
	FiltersFast.com	At home? You'll love the Home Filter Cl	ub - From	your #1	1	Sep	14
	FiltersFast.com	At home? You'll love the Home Filter Cl	ub - From	your #1	1	Aug	23
	Perform Better	Don't Miss This Week's FREE Webinars!	- If you ha	ive trou	I	Aug	17
	FiltersFast.com	At home? You'll love the Home Filter Cl	ub - From	your #1	1	Aug	g 1
	FiltersFast.com	At home? You'll love the Home Filter Cl	ub - From	your #1	1	Jul	11
	FiltersFast.com	At home? You'll love the Home Filter Cl	ub - From	your #1	1	May	29
	FiltersFast.com	At home? You'll love the Home Filter Cl	ub - From	your #1	1	May	y 7
	FiltersFast.com	Have you been forgetting to change you	ur filters?	- The H	0	Apr	15
	FiltersFast.com	Have you been forgetting to change you	ur filters?	- The H	0	Mar	24
	☆ no-reply	Product Reminder - com" width="240" b	order="0">	∙ It's Tir	n	Mar	17
	☆ FiltersFast.com	Have you been forgetting to change you	ur filters?	- The H	0	Ма	r 2
	FiltersFast.com	Have you been forgetting to change you	ur filters?	- The H	0	Fel	b 9
	FiltersFast.com	Have you been forgetting to change you	ur filters?	- The H	0	Jan	18
	FiltersFast.com	Have you been forgetting to change you	ur filters?	- The H	0	11/13/	19
	FiltersFast.com	Have you been forgetting to change you	ur filters?	- The H	0	10/23/	19
	FiltersFast.com	Have you been forgetting to change you	ur filters?	The H	0	10/1/	19
	FiltersFast.com	It's Time to Order Your Filters - Discoun	nt for Toda	y Only!	! - H	9/17/	19
	FiltersFast.com	Have you been forgetting to change you	ur filters?	+ The H	0	9/9/	19
	FiltersFast.com	Have you been forgetting to change you	ur filters?	- The H	0	8/19/	19
	FiltersFast.com	Have you been forgetting to change you	ur filters?	The H	0	7/28/	19
	FiltersFast.com	Have you been forgetting to change you	ur filters?	- The H	0	7/6/	19

And these emails.



M Gmail	Q time to change your filter	×	•	?	()	000 000 000
🖻 Any time 💌 🤇	₽ Has attachment ► To me 🔽 Is unread	Advance	ed searc	h		
□ - C :			1-31 of	f 31	<	>
🗌 📩 no-reply	Product Reminder - com" width="240"	border="0">	> It's Tim.		Sep	17
🗌 🕁 FiltersFast.con	n At home? You'll love the Home Filter C	Club - From	your #1 .		Sep	14
🗌 🗙 FiltersFast.con	At home? You'll love the Home Filter C	Club - From	your #1 .		Aug	23
🔲 🖈 Perform Better	Don't Miss This Week's FREE Webinars	s! - If you ha	ave trou		Aug	17
🗌 🕁 FiltersFast.con	n At home? You'll love the Home Filter C	Club - From	your #1 .		Aug	; 1
🗌 📩 FiltersFast.con	n At home? You'll love the Home Filter C	Club - From	your #1 .		Jul	11
☐ ☆ FiltersFast.con	n At home? You'll love the Home Filter C	Club - From	your #1		May	29
☐ ☆ FiltersFast.con	At home? You'll love the Home Filter C	Club - From	your #1 .		May	,7
☐ ☆ FiltersFast.con	Have you been forgetting to change yo	our filters?	- The Ho.		Apr	15
🔲 ☆ FiltersFast.con	Have you been forgetting to change yo	our filters?	- The Ho.		Mar	24
🔲 ☆ no-reply	Product Reminder - com" width="240"	border="0">	It's Tim.		Mar	17
🗌 🕁 FiltersFast.con	Have you been forgetting to change yo	our filters?	- The Ho.		Ma	r 2
☐ ☆ FiltersFast.con	Have you been forgetting to change yo	our filters?	- The Ho.		Feb	9
🗌 ☆ FiltersFast.con	Have you been forgetting to change yo	our filters?	- The Ho.		Jan	18
🗌 🕁 FiltersFast.con	Have you been forgetting to change yo	our filters?	- The Ho.		11/13/	19
🗌 🕁 FiltersFast.con	Have you been forgetting to change yo	our filters?	- The Ho.		10/23/	19
🗌 🕁 FiltersFast.con	Have you been forgetting to change yo	our filters?	The Ho.		10/1/	19
🗌 ☆ FiltersFast.con	n It's Time to Order Your Filters - Discou	Int for Toda	y Only! -	H	9/17/	19
🗌 🛣 FiltersFast.con	Have you been forgetting to change yo	our filters?	The Ho.		9/9/	19
🗌 ☆ FiltersFast.con	Have you been forgetting to change yo	our filters?	- The Ho.		8/19/	19
☐ ☆ FiltersFast.con	Have you been forgetting to change yo	our filters?	The Ho.		7/28/	19
🔲 ☆ FiltersFast.con	n Have you been forgetting to change yo	our filters?	- The Ho.		7/6/	19

And these emails.



This Talk: It's Time to Change Your Filter

- Tutorial-like introduction to filters.
- General techniques for solving filter problems.







fingerprinting

quotienting

collision resolution

This Talk: It's Time to Change Your Filter

- Tutorial-like introduction to filters.
- General techniques for solving filter problems.







fingerprinting

quotienting

This Talk

- Tutorial-like introduction to filters.
- General techniques for solving filter problems.











collision resolution

[Cleary 84] [Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10], [Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12], [Fan, Andersen, Kaminsky, Mitzenmacher 14], [Pandey, Bender, Johnson, Patro 17], [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Filter
$$F$$
 of set S : $\{h(x) | x \in S\}$





$$F = \{ h(x_1), h(x_2), \dots, h(x_n) \}$$

 $S = \{x_1, x_2, ..., x_n\}$

F can be stored more compactly than a Bloom filter.



[Cleary 84] [Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10], [Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12], [Fan, Andersen, Kaminsky, Mitzenmacher 14], [Pandey, Bender, Johnson, Patro 17], [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Filter
$$F$$
 of set S : $\{h(x) | x \in S\}$





$$F = \{ h(x_1), h(x_2), \dots, h(x_n) \}$$

 $S = \{x_1, x_2, ..., x_n\}$

F can be stored more compactly than a Bloom filter.

F is also just dictionary—just a more compact one.



[Cleary 84] [Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10], [Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12], [Fan, Andersen, Kaminsky, Mitzenmacher 14], [Pandey, Bender, Johnson, Patro 17], [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Filter
$$F$$
 of set S : $\{h(x) | x \in S\}$



 $F = \{ h(x_1), h(x_2), \dots, h(x_n) \}$



 $S = \{x_1, x_2, \dots, x_n\}$

$$\mathbf{member}(x) = \begin{cases} \checkmark & \text{if } h(x) \in \mathbf{F} \\ \times & \text{if } h(x) \notin \mathbf{F} \end{cases}$$



[Cleary 84] [Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10], [Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12], [Fan, Andersen, Kaminsky, Mitzenmacher 14], [Pandey, Bender, Johnson, Patro 17], [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]

Filter
$$F$$
 of set S : $\{h(x) | x \in S\}$



 $F = \{ h(x_1), h(x_2), \dots, h(x_n) \}$



$$\boldsymbol{S} = \{x_1, x_2, \dots, x_n\}$$

$$\mathbf{member}(x) = \begin{cases} \checkmark & \text{if } h(x) \in \mathbf{F} \\ \times & \text{if } h(x) \notin \mathbf{F} \end{cases}$$

F can be stored more compactly than a Bloom filter.

fingerprinting





 $F = \{ h(x_1), h(x_2), \dots, h(x_n) \}$

$$S = \{x_1, x_2, \dots, x_n\}$$

member(y)= \checkmark if $h(y) \in F$.

y is a false positive if $\exists x_i h(y)=h(x_i)$, but $y \notin S$.

Fingerprint collisions: only source of false positives.









 $\Pr[y \notin S \text{ is a false positive}] \leq \varepsilon$





 $\Pr[y \notin S \text{ is a false positive}] \leq \varepsilon$

n fingerprints can be stored compactly:

 $log(1/\varepsilon) + O(1)$ bits/element





 $\Pr[y \notin S \text{ is a false positive}] \leq \varepsilon$

n fingerprints can be stored compactly:

 $log(1/\varepsilon) + O(1)$ bits/element



Compact Storage Using Quotienting [Knuth]

Space: O(lg $(1/\epsilon)$) bits per element



fingerprinting

quotienting

Compact Storage Using Quotienting [Knuth]

Space: O(lg $(1/\epsilon)$) bits per element

n





q(x)=location in hash table r(x)=data stored in hash table

How to deal with collisions in the hash table?

Isn't this a solved problem?

What's wrong with out-ofthe-box linear probing?



Hash Collisions in Quotienting

Ex: 6 bit hash. 3 bits for address, 3 for data.



The hash is stored implicitly based on location. So how can we change its location?



Hash Collisions in Quotienting

Ex: 6 bit hash. 3 bits for address, 3 for data.



The hash is stored implicitly based on location. So how can we change its location?



Talk Structure

• Filters + how filters are used + Bloom limitations



Quotient filters variants	Cuckoo filter variants	Miscellaneous
[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok 12], [Pandey, Bender, Johnson, Patro 17], [Pandey, Bender, Conway, Farach-Colton, Johnson 21]	[Fan, Andersen, Kaminsky, Mitzenmacher 14], [Breslow, Jayasena 18]	[Pagh, Pagh, Rao 05], [Arbitman, Naor, Segev 10], [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18]
uses linear proving and	Uses Cuckoo hashing	Balls and bins + more advanced
[Celis, Larson, Munro 85]		nasining

Quotient Filters

[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok VLDB 12] [Pandey, Bender, Patro, Johnson SIGMOD 17]

2 metadata bits per slot let us recover original location.





1= "I'm hashed to the same place as the element before me"



Quotient Filters

[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok VLDB 12] [Pandey, Bender, Patro, Johnson SIGMOD 17]

2 metadata bits per slot let us recover original location.



Recall: no element is stored before its target position.

This is an empty slot.

Remainder is in correct slot.

Remainder is mapped to same location as previous remainder.

Remainder is mapped to same location as previous remainder.

Remainder is mapped to next array position.

Remainder would map to the next position. But there's none, so it's empty.

Remainder is in correct slot.

Remainder is mapped to same location as previous remainder.

1= "something is hashed here"



1= "I'm hashed to the same place as the element before me"



Quotient Filter Capabilities

[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok VLDB 12] [Pandey, Bender, Patro, Johnson SIGMOD 17]





Bloom limitations

No deletes or counting

No resizes

No element enumeration or merging of filters

Keys have no values

QF Copabilities counting and deletes resizing

element enumeration + filter merging values allowed



Quotient Filter Capabilities

[Bender, Farach-Colton, Johnson, Kraner, Kuszmaul, Medjedovic, Montes, Shetty, Spillane, Zadok VLDB 12] [Pandey, Bender, Patro, Johnson SIGMOD 17]





Reminder of General Approach

• **Fingerprinting**: key $k \in S \rightarrow h(k) \in F$.

False-positives only come from fingerprint collisions.

- Quotienting: Store fingerprints in a hash table implicitly.
 - Ig n bits of fingerprint depend on hash location.

Collision resolution:

- E.g., using linear probing/Robinhood hashing.
- Use metadata bits to recover each h(k).

- **Designs alternatives:** cuckoo, Morton, broom
 - Use a different hash table but the same general approach



Cuckoo Hashing → Cuckoo Filters

[Pagh, Rodler 01]

1] [Fan, Andersen, Kaminsky, Mitzenmacher 14] [Breslow, Jayasena 18]



Cuckoo hash table has 2 hash functions h_1 and h_2 .

Each hash bucket has 4 slots.



Cuckoo Hashing → Cuckoo Filters

[Pagh, Rodler 01]

[Fan, Andersen, Kaminsky, Mitzenmacher 14] [Breslow, Jayasena 18]



Cuckoo hash table has 2 hash functions h_1 and h_2 .

Each hash bucket has 4 slots.

If there's no space in any of the 8 slots: *kick out* an element, and move it to the alternative location (which may cause other kicks).



Cuckoo Hashing → Cuckoo Filters

[Pagh, Rodler 01]

[Fan, Andersen, Kaminsky, Mitzenmacher 14] [Breslow, Jayasena 18]



Cuckoo hash table has 2 hash functions h_1 and h_2 .

Each hash bucket has 4 slots.

If there's no space in any of the 8 slots: *kick out* an element, and move it to the alternative location (which may cause other kicks).



Obstacles for Cuckoo Filters

[Pagh, Rodler 01]

[Fan, Andersen, Kaminsky, Mitzenmacher 14] [Breslow, Jayasena 18]



Q: If *f*(*x*) is kicked, how to find an alternative location when we don't store *x*?

A: We give up on independent hash functions. The alternate location only depends on r(x).

We also give up on asymptotic correctness.

Amazingly, it works for (practical) *n* not too large.

Cuckoo hashing seemingly doesn't have metadata bits.

But because there are 4 slots per cell, 2 more fingerprint bits are stored explicitly.



Quotient Filter and Cuckoo Filter Comparison

Quotient filter	Cuckoo filter		
good space	good space		
very good locality	ok locality		
some degradation at high load factors	degradation at high load factors		
good searches	very good searches		
fast inserts	fast inserts		
supports counting, multisets, values, deletes			
complicated code	code is easier		
good for all <i>n</i> .	pre-asymptotic guarantees. fails w.h.p. large enough <i>n</i>		

Optimal Theoretical Guarantees

Theorem: There is an optimal filter with

- Space: $(1 + o(1)) n \log(1/\epsilon) + O(n)$
- **Error rate:** $\leq \varepsilon$
- Operations: O(1) insert, delete, query.

- [Pagh, Pagh, Rao 05]
- [Arbitman, Naor, and Segev 10]
- [Lovett & Porat 10]
- [Bender, Farach-Colton, Goswami, Johnson, McCauley, Singh 18] adaptive & worst case

Empirical Performance of a Vector Quotient Filter



(a) Insertion (Higher is better.)

Quotient filter plus minimum of 2 choice retains good performance even when almost full.

Empirical Performance of a Vector Quotient Filter



(c) Successful lookup (Higher is better.)

Cuckoo still kicks butt on queries.

Summery Slide

Summery Slide

It's time to change your filter.

Fingerprinting + quotienting + collision resolution is unifying theme in theory and practice.

Applications should demand a richer set of operations from their filter.