Tutorial on Circular Logs & Cuckoo Filters

Database System Technology

Niv Dayan

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. Comment on any trade-offs or downsides.

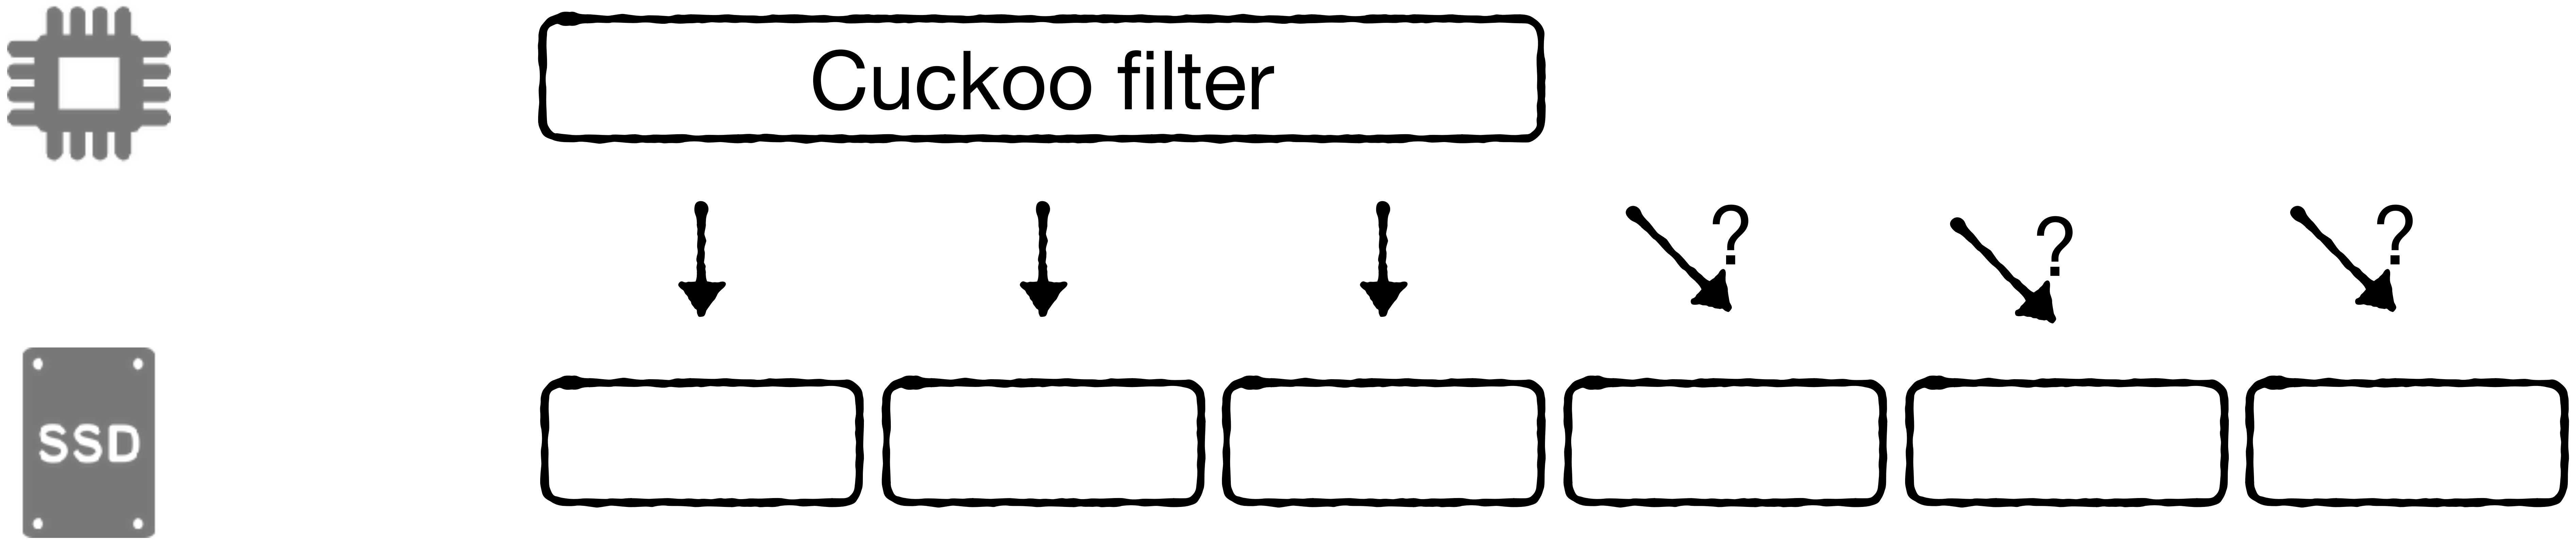

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. Comment on any trade-offs or downsides.

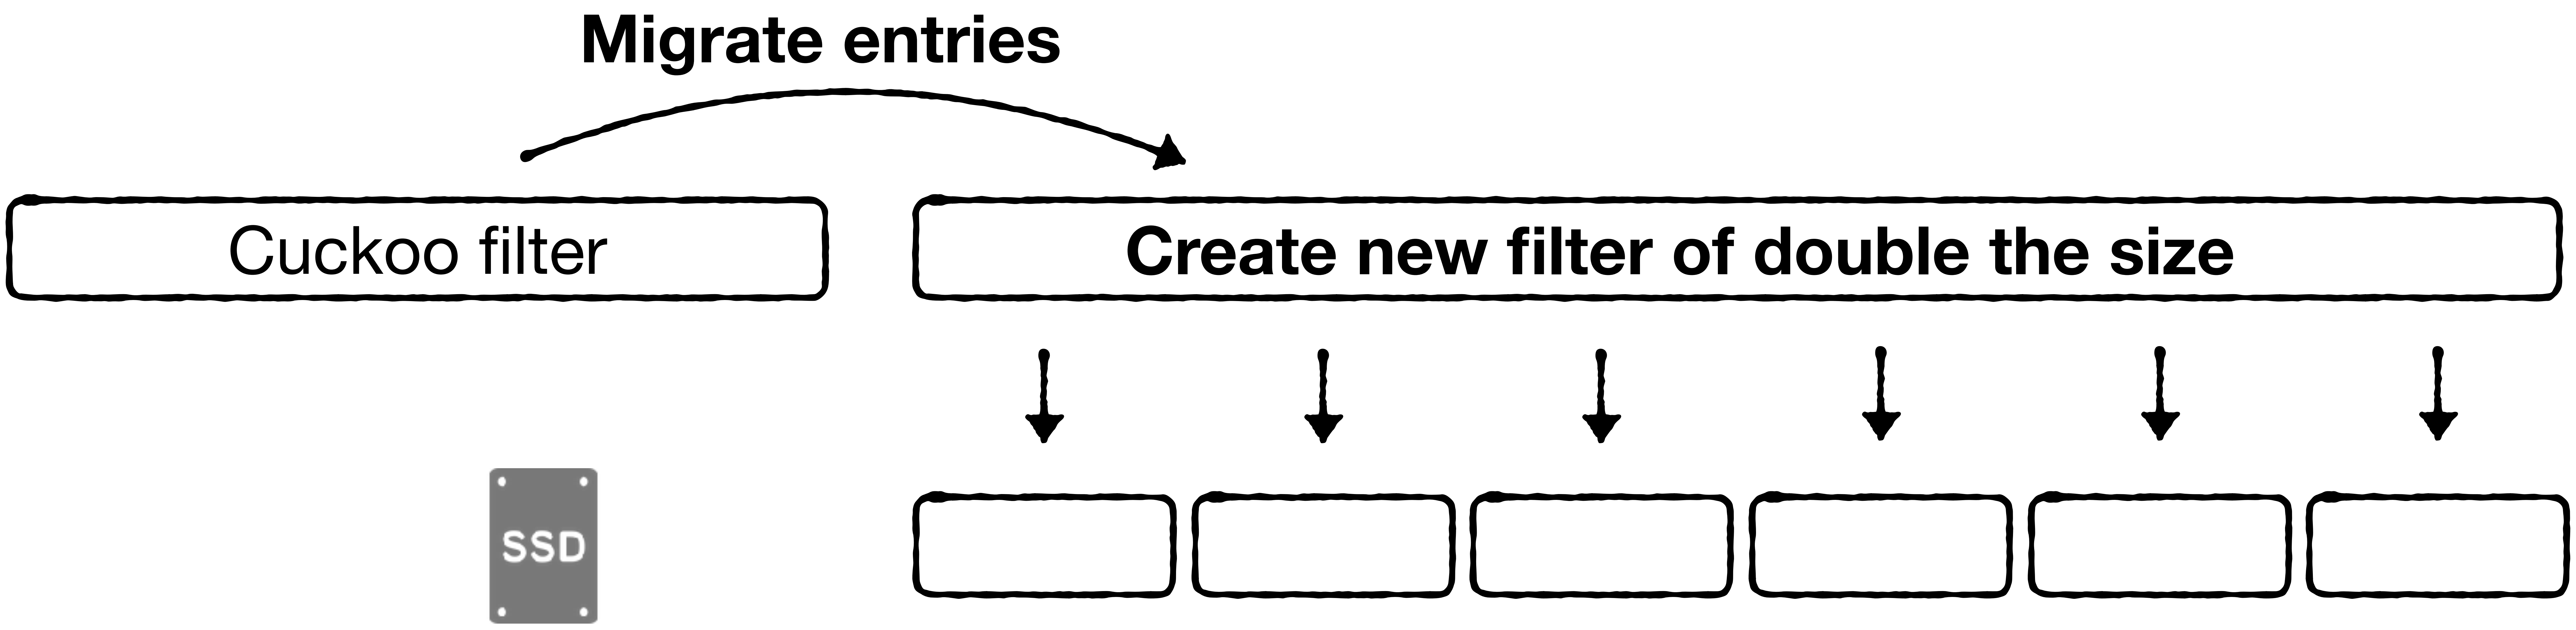

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. Comment on any trade-offs or downsides.

Challenge: we do not have the full keys to rehash. We only have fingerprints.

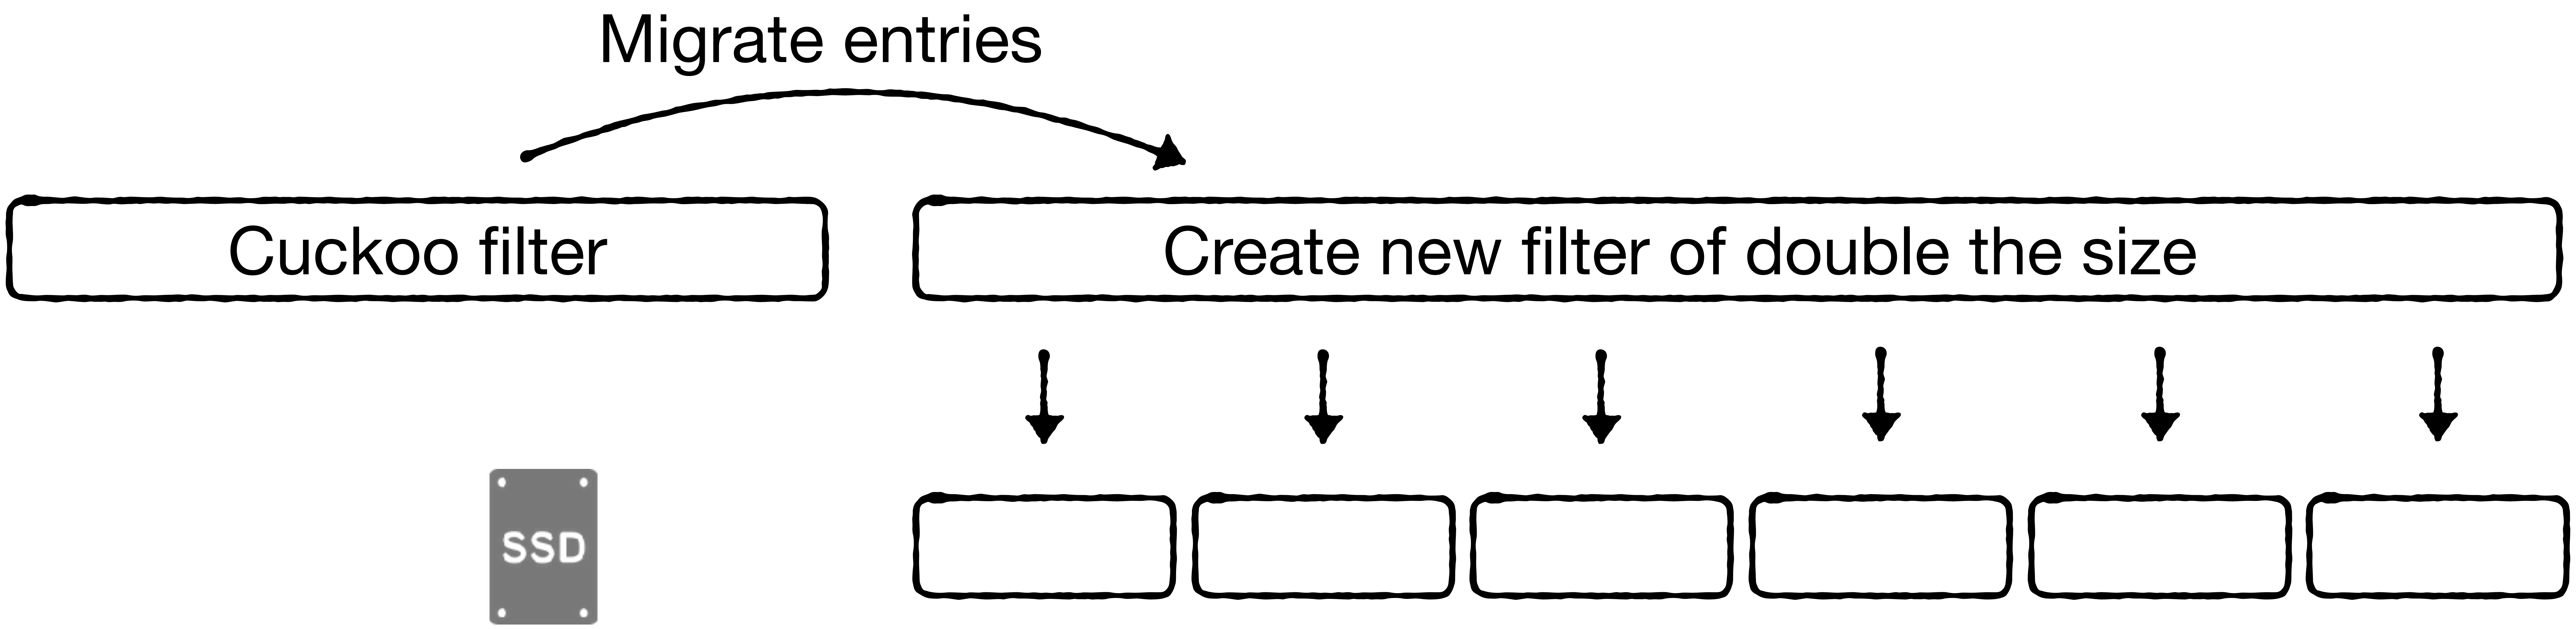

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. Comment on any trade-offs or downsides.

Challenge: we do not have the full keys to rehash. We only have fingerprints.

Approach: view fingerprint and first bucket address as components of the same hash

$$\textbf{hash(X)} = \overbrace{\textbf{01001}}^{\text{address}}\ \overbrace{\textbf{001110}}^{\text{fingerprint}}$$

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. Comment on any trade-offs or downsides.

Challenge: we do not have the full keys to rehash. We only have fingerprints.

Approach: view fingerprint and first bucket address as components of the same hash

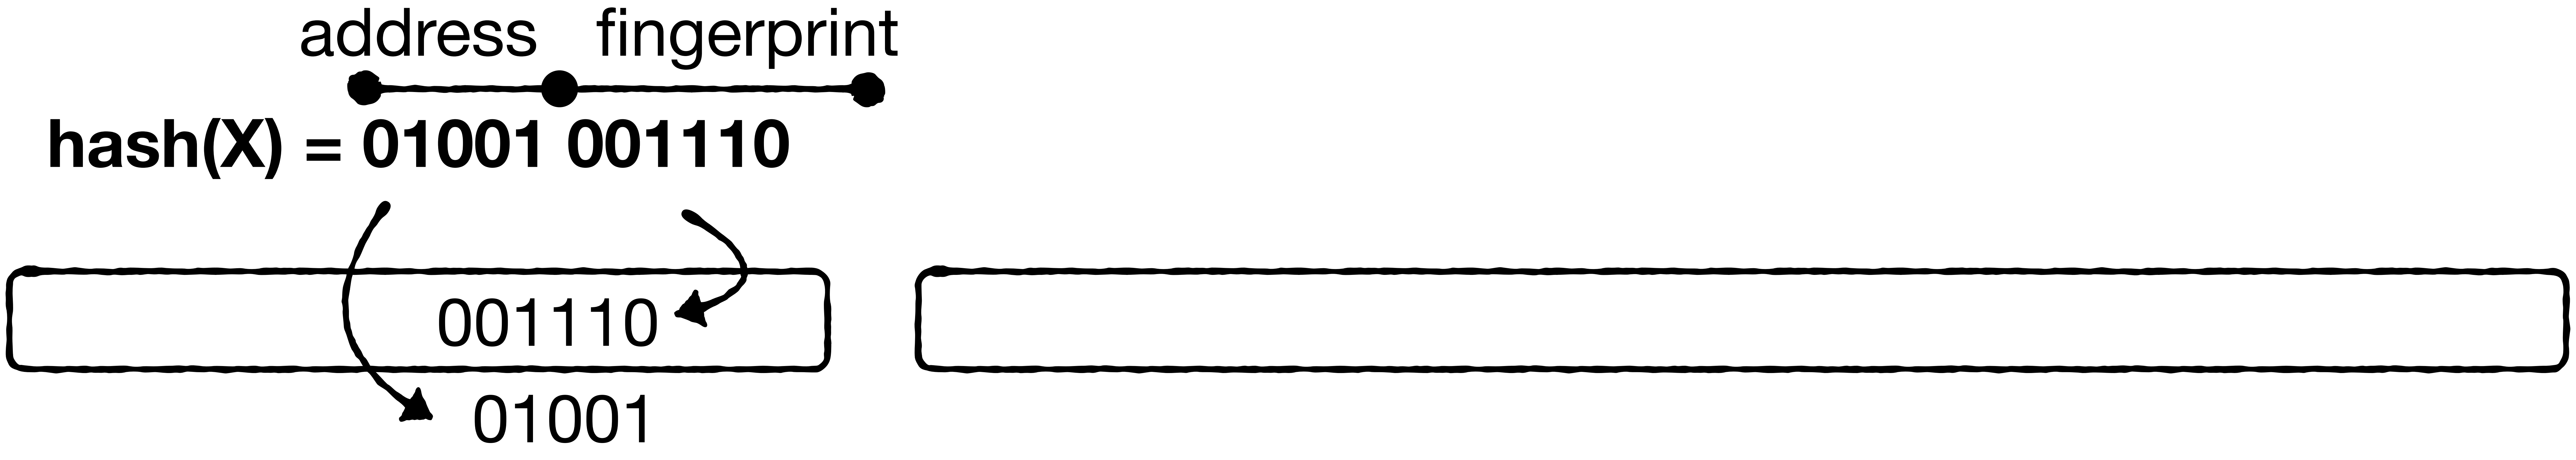

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. Comment on any trade-offs or downsides.

To migrate, transfer one bit from fingerprint to address

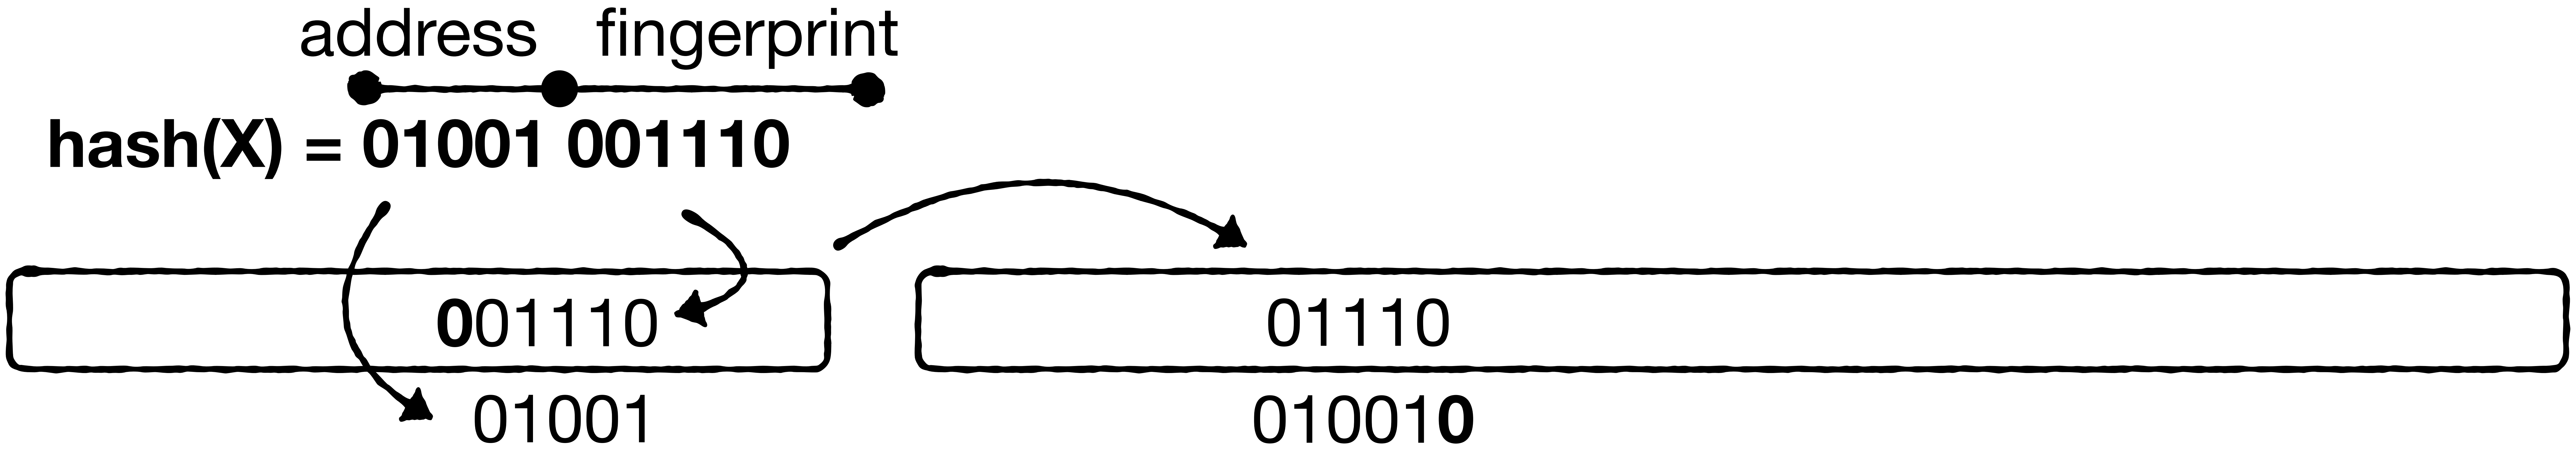

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. Comment on any trade-offs or downsides.

To migrate, transfer one bit from fingerprint to address

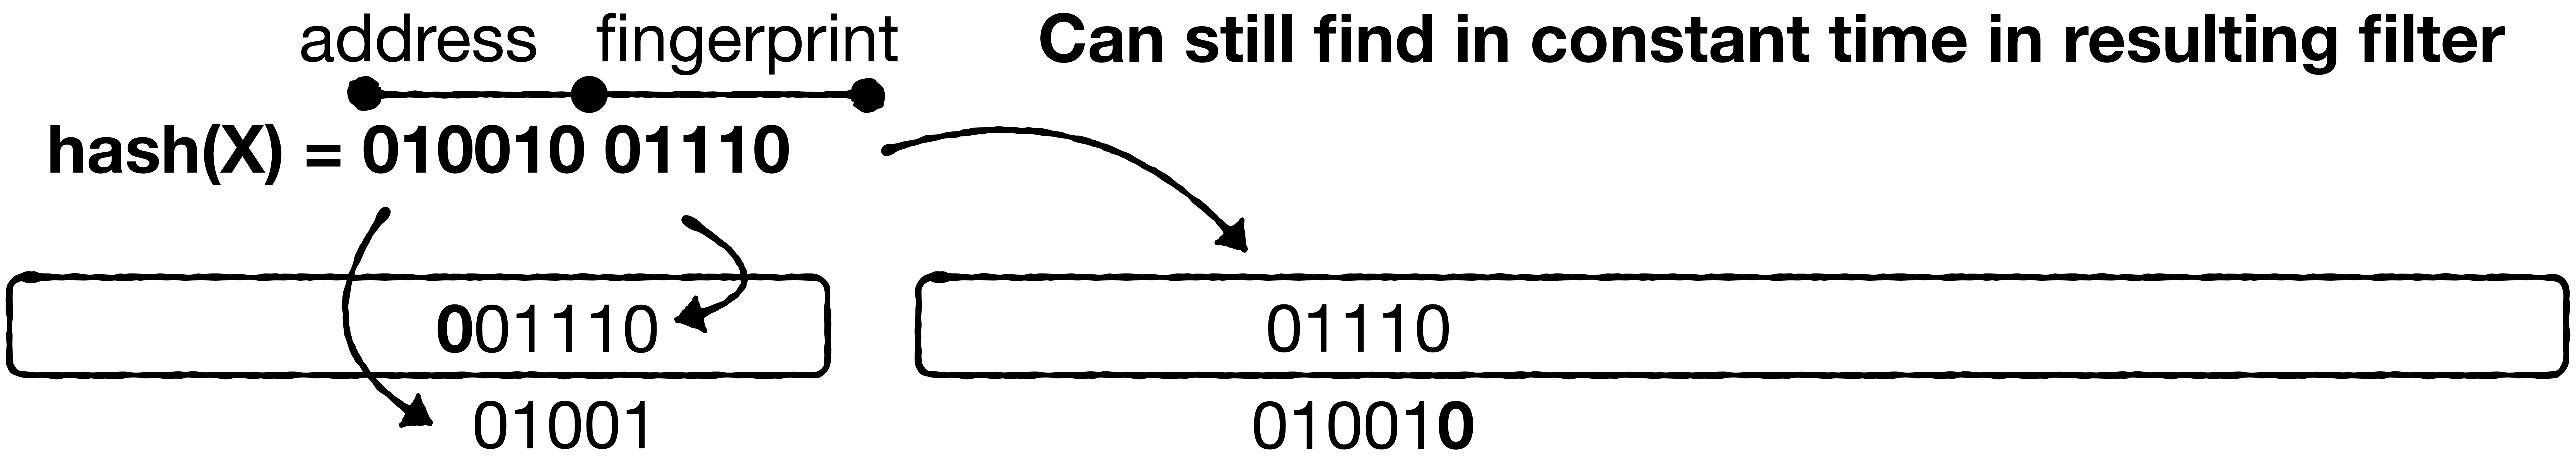

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. Comment on any trade-offs or downsides.

Complication: in a cuckoo filter an entry can be in one of two buckets, the canonical address and the alternative address. Only the canonical address should be viewed as a part of the original hash.

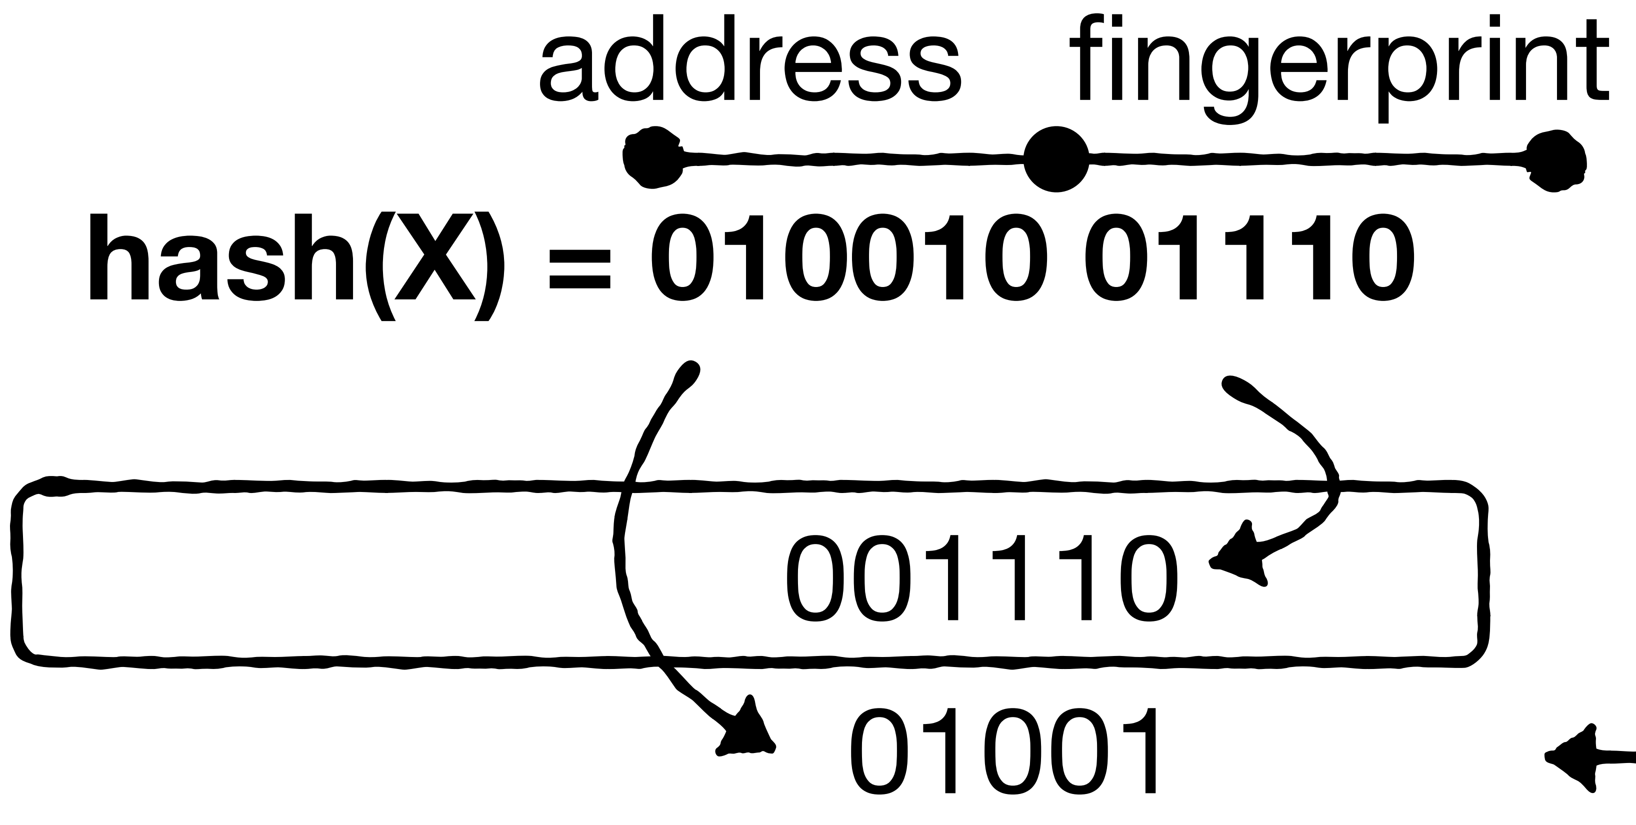

Alternative address = Canonical address XOR h(fingerprint)

Canonical or alternative address?

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. Comment on any trade-offs or downsides.

Complication: in a cuckoo filter an entry can be in one of two buckets, the canonical address and the alternative address. Only the canonical address should be viewed as a part of the original hash.

Solution: add a bit to indicate whether the current address is canonical or alternative. If alternative, switch to canonical via XOR.

0 001110

01001

Question 1 - Expanding Cuckoo Filters

A Cuckoo filter is allocated with a fixed capacity. As a circular log grows, we must expand its cuckoo filter to map more data. Devise a Cuckoo filter expansion algorithm. Ideally, this algorithm should maintain constant time performance and not have to read any data from storage. **Comment on any trade-offs or downsides.**

Every time we double the data size, we lose one bit from all fingerprints, meaning the false positive rate doubles. Hence, the false positive rate as we expand is: $O(N \cdot 2^{-M+3})$

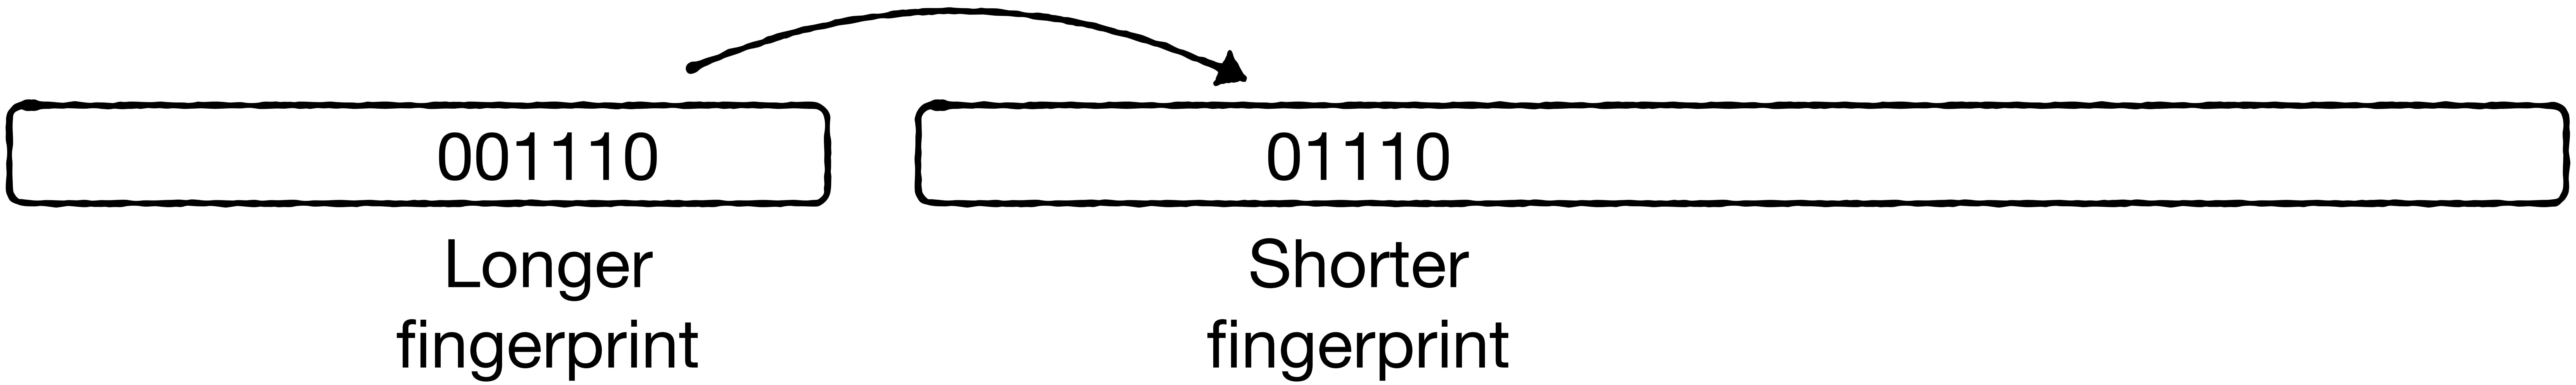

Question 2 - Cuckoo Filter for LSM-tree

As we have seen, the expected worst case query cost over Bloom filters for a basic LSM-tree is O(L + M), where L is the number of levels and M is the number of bits per entry. (Assume only unique entries in the tree).

(A) How can we employ a cuckoo filter to achieve constant time?

(B) What are the implications on the false positive rate and memory footprint? Any downsides compared to plain Bloom filters?

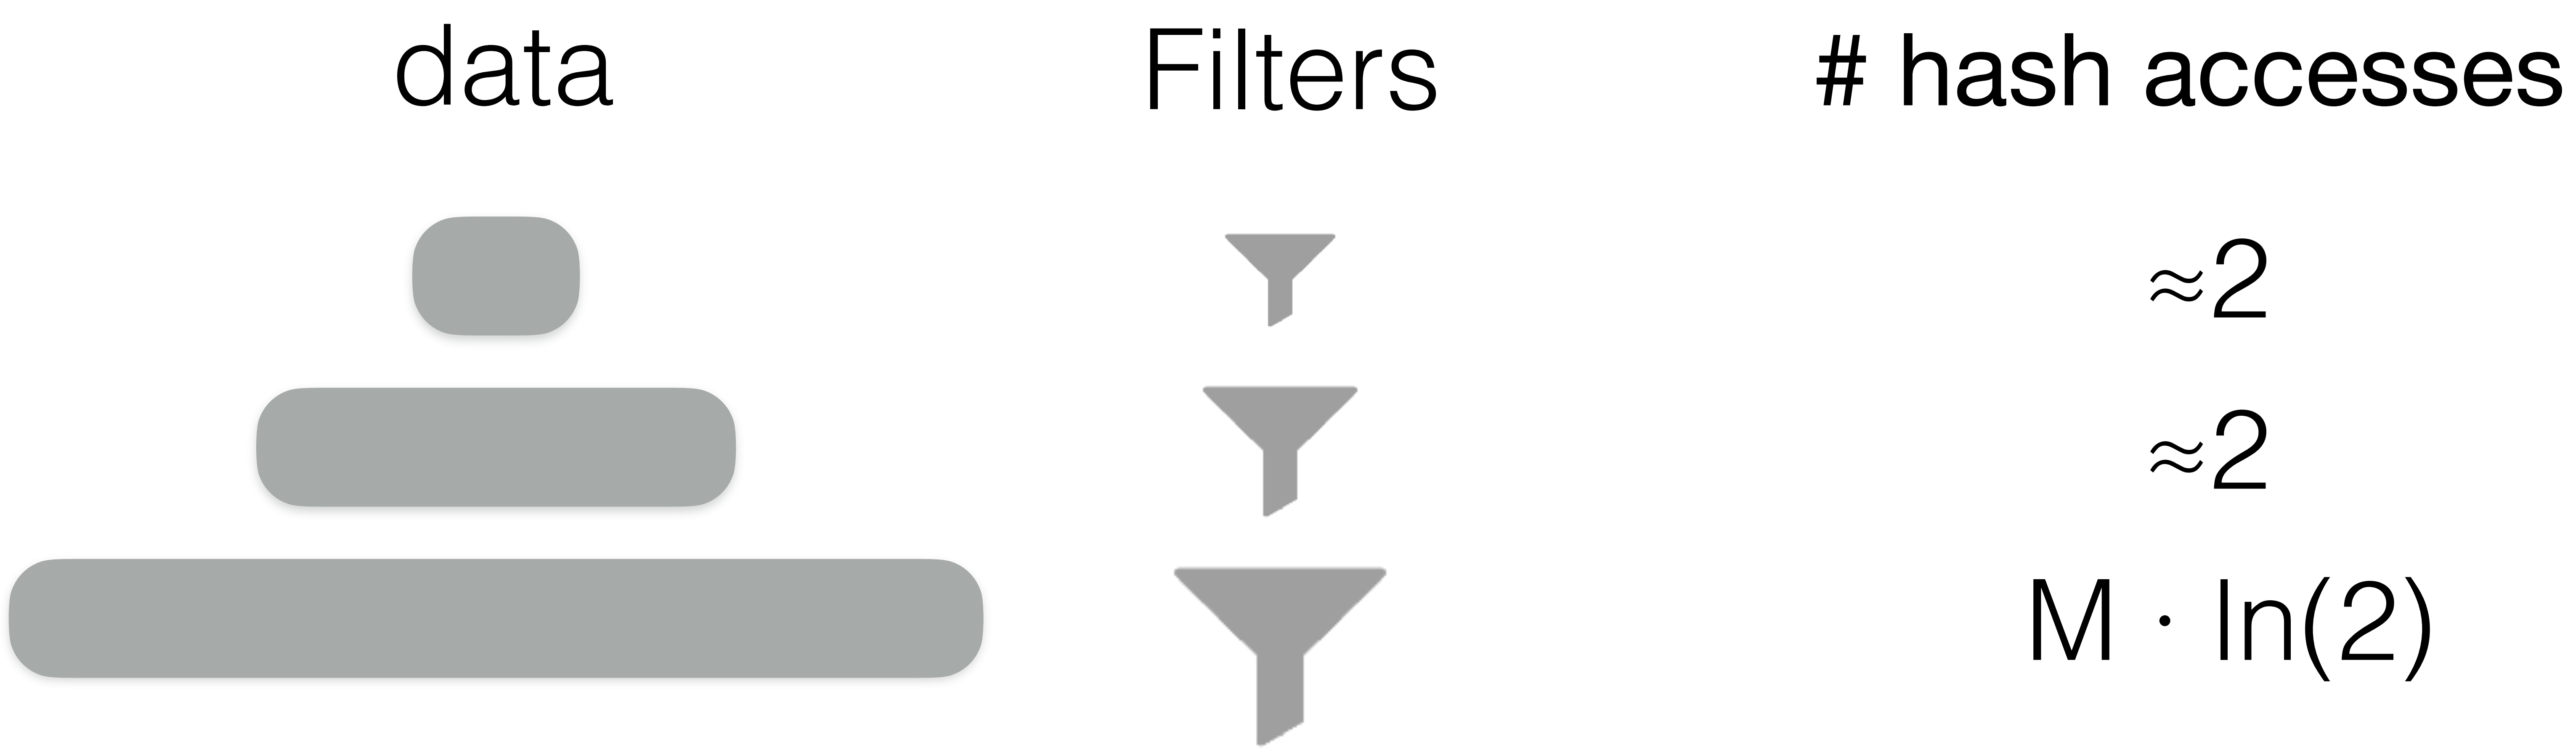

Question 2 - Cuckoo Filter for LSM-tree

As we have seen, the expected worst case query cost over Bloom filters for a basic LSM-tree is O(L + M), where L is the number of levels and M is the number of bits per entry. (Assume only unique entries in the tree).

(A) How can we employ a cuckoo filter to achieve constant time?

(B) What are the implications on the false positive rate and memory footprint? Any downsides compared to plain Bloom filters?

Question 2 - Cuckoo Filter for LSM-tree

As we have seen, the expected worst case query cost over Bloom filters for a basic LSM-tree is O(L + M), where L is the number of levels and M is the number of bits per entry. (Assume only unique entries in the tree).

(A) How can we employ a cuckoo filter to achieve constant time?

(B) What are the implications on the false positive rate and memory footprint? Any downsides compared to plain Bloom filters?

Unified Cuckoo filter

	Unified Cuckoo filter
Filter accesses:	$O(1)$
False positive rate:	$O(2^{-M+3})$
Memory (bits/entry)	$O(M + \log_2(L))$
Construction:	$O(L)$

FP(X), 2
FP(Y), 3
…

<M bit Fingerprint, level>

Question 2 - Cuckoo Filter for LSM-tree

As we have seen, the expected worst case query cost over Bloom filters for a basic LSM-tree is O(L + M), where L is the number of levels and M is the number of bits per entry. (Assume only unique entries in the tree).

(A) How can we employ a cuckoo filter to achieve constant time?
(B) What are the implications on the false positive rate and memory footprint? Any downsides compared to plain Bloom filters?

	Unified Cuckoo filter	**With Bloom filters**
Filter accesses:	$O(1)$	$O(M+L)$
False positive rate:	$O(2^{-M+3})$	$O(L \cdot 2^{-M \cdot \ln(2)})$
Memory (bits/entry)	$O(M + \log_2(L))$	$O(M)$
Construction:	$O(L)$	$O(L \cdot M)$

Question 2 - Cuckoo Filter for LSM-tree

As we have seen, the expected worst case query cost over Bloom filters for a basic LSM-tree is O(L + M), where L is the number of levels and M is the number of bits per entry. (Assume only unique entries in the tree).

(A) How can we employ a cuckoo filter to achieve constant time?

(B) What are the implications on the false positive rate and memory footprint? Any downsides compared to plain Bloom filters?

	Unified Cuckoo filter	**With Bloom filters**	**Monkey**
Filter accesses:	$O(1)$	$O(M+L)$	$O(M+L)$
False positive rate:	$O(2^{-M+3})$	$O(L \cdot 2^{-M \cdot \ln(2)})$	$O(2^{-M \cdot \ln(2)})$
Memory (bits/entry)	$O(M + \log_2(L))$	$O(M)$	$O(M)$
Construction:	$O(L)$	$O(L \cdot M)$	$O(L \cdot (L + M))$

(This memory analysis here only account for the filters and not the fence pointers (internal nodes) being stored in memory)

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

(B) Estimate write-amplification assuming perfect hot/cold data separation.

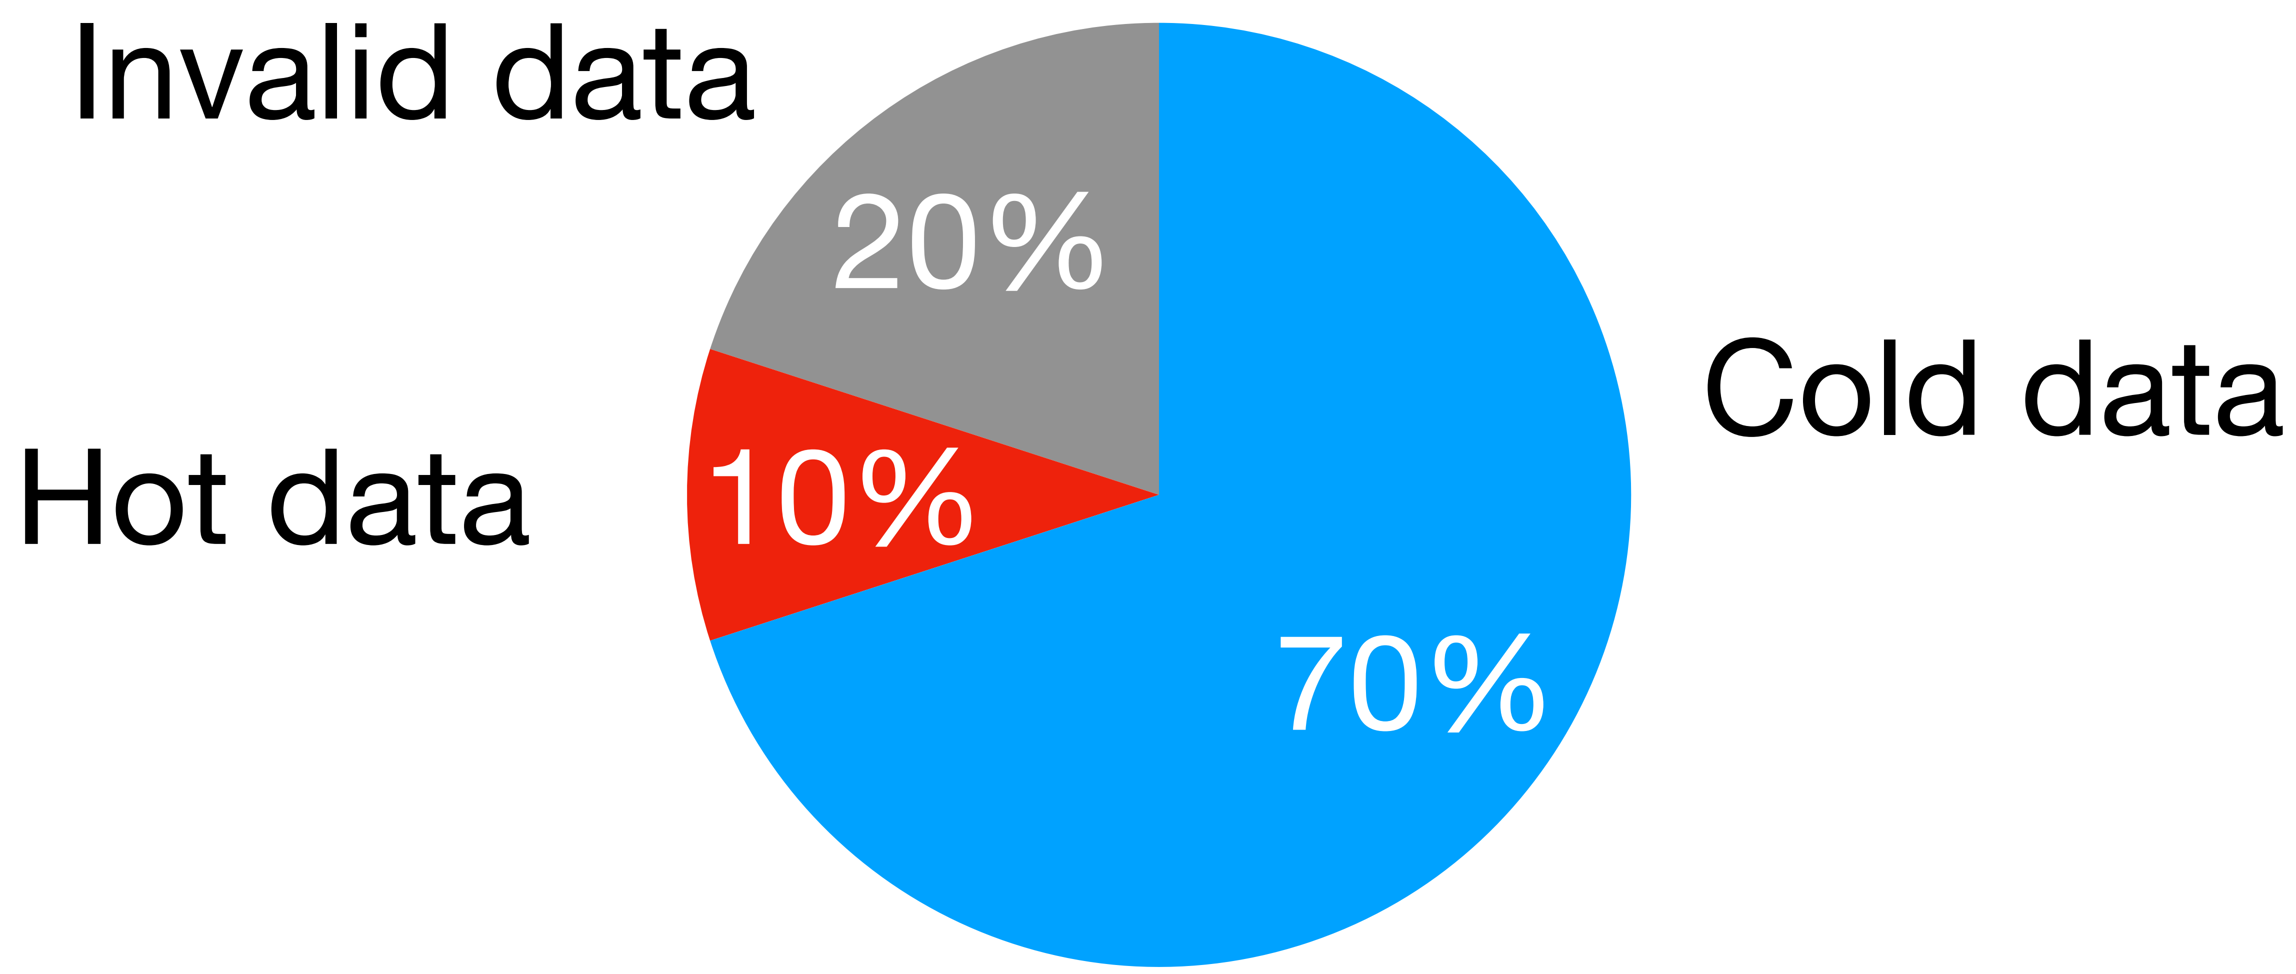

Garbage-Collection Write-Amplification

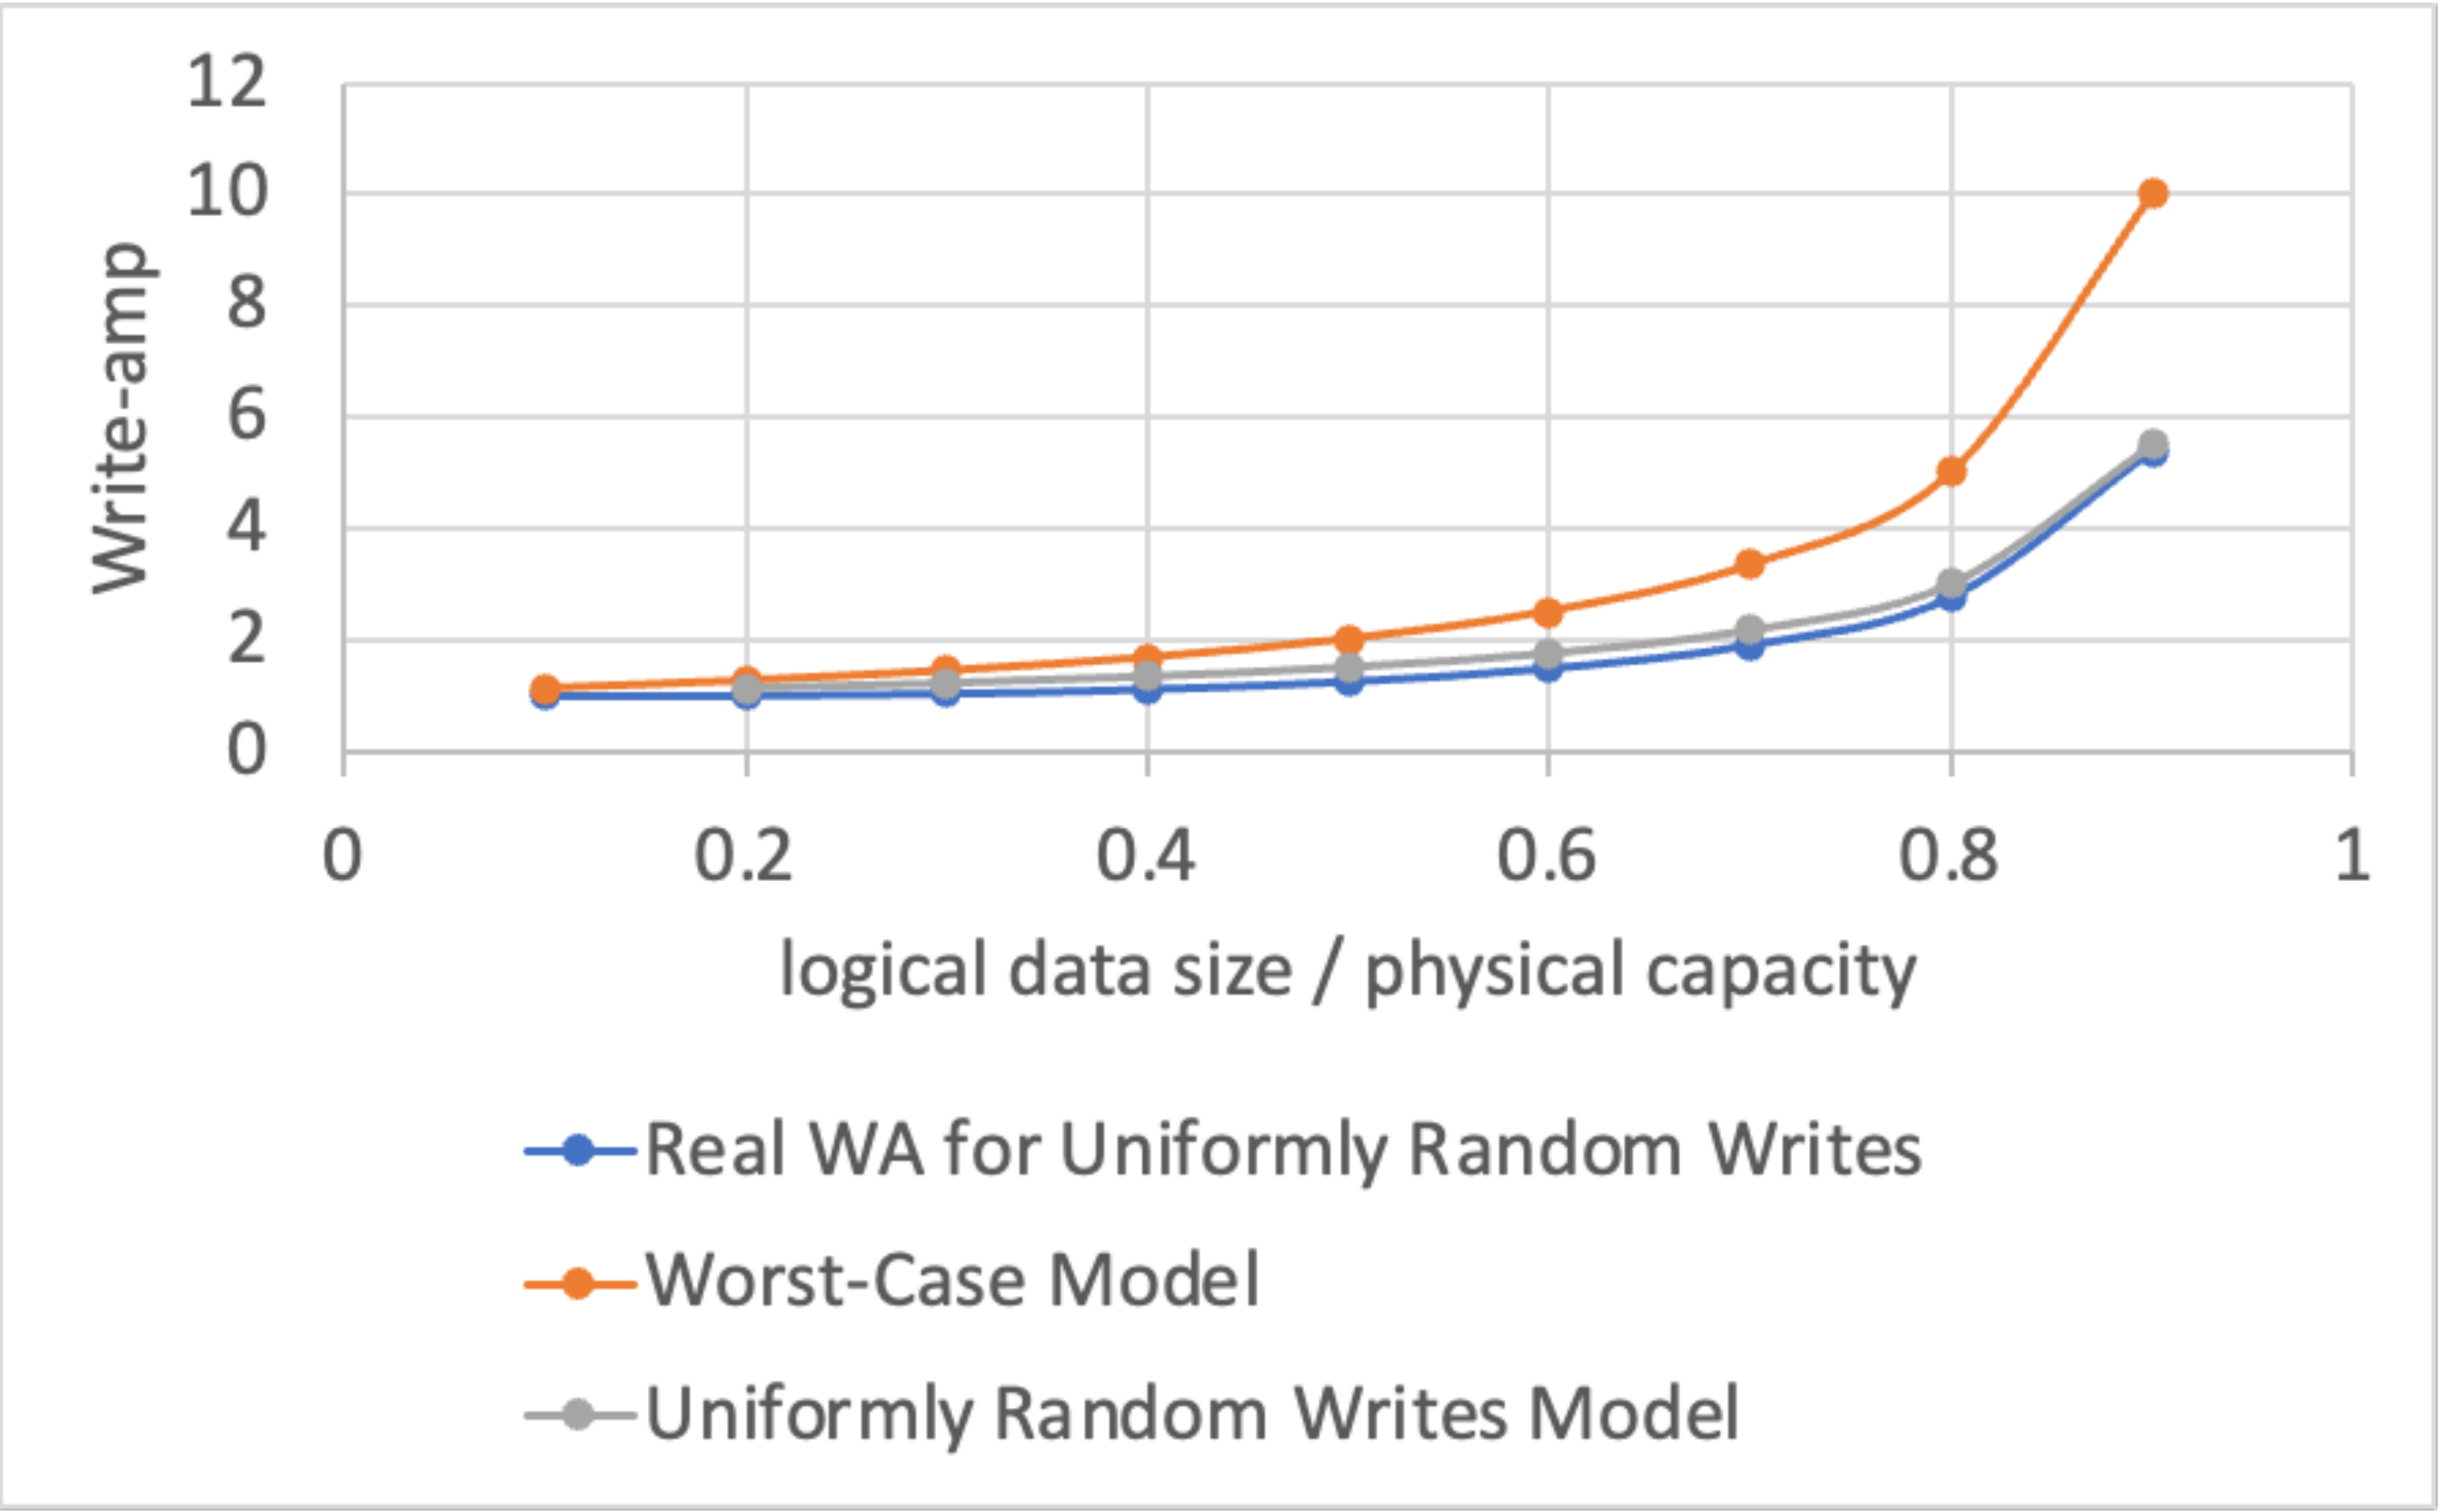

$$1 + \frac{L/P}{1 - L/P}$$ **Worst case**

$$1 + \frac{1}{2} \cdot \frac{L/P}{1 - L/P}$$ **Uniformly random**

L = logical data size

P = physical data size

Garbage-Collection Write-Amplification

$$1 + \frac{L/P}{1 - L/P} \quad = 1 + \frac{L}{P - L}$$

Worst case

$$1 + \frac{1}{2} \cdot \frac{L/P}{1 - L/P} \quad = 1 + \frac{1}{2} \cdot \frac{L}{P - L}$$

Uniformly random

L = logical data size

P = physical data size

Garbage-Collection Write-Amplification

$$1 + \frac{L/P}{1 - L/P} \quad = 1 + \frac{L}{P - L} \quad = 1 + \frac{L}{O}$$

Worst case

$$1 + \frac{1}{2} \cdot \frac{L/P}{1 - L/P} \quad = 1 + \frac{1}{2} \cdot \frac{L}{P - L} \quad = 1 + \frac{1}{2} \cdot \frac{L}{O}$$

Uniformly random

L = logical data size

P = physical data size

O = Overprovisioned space (P-L)

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

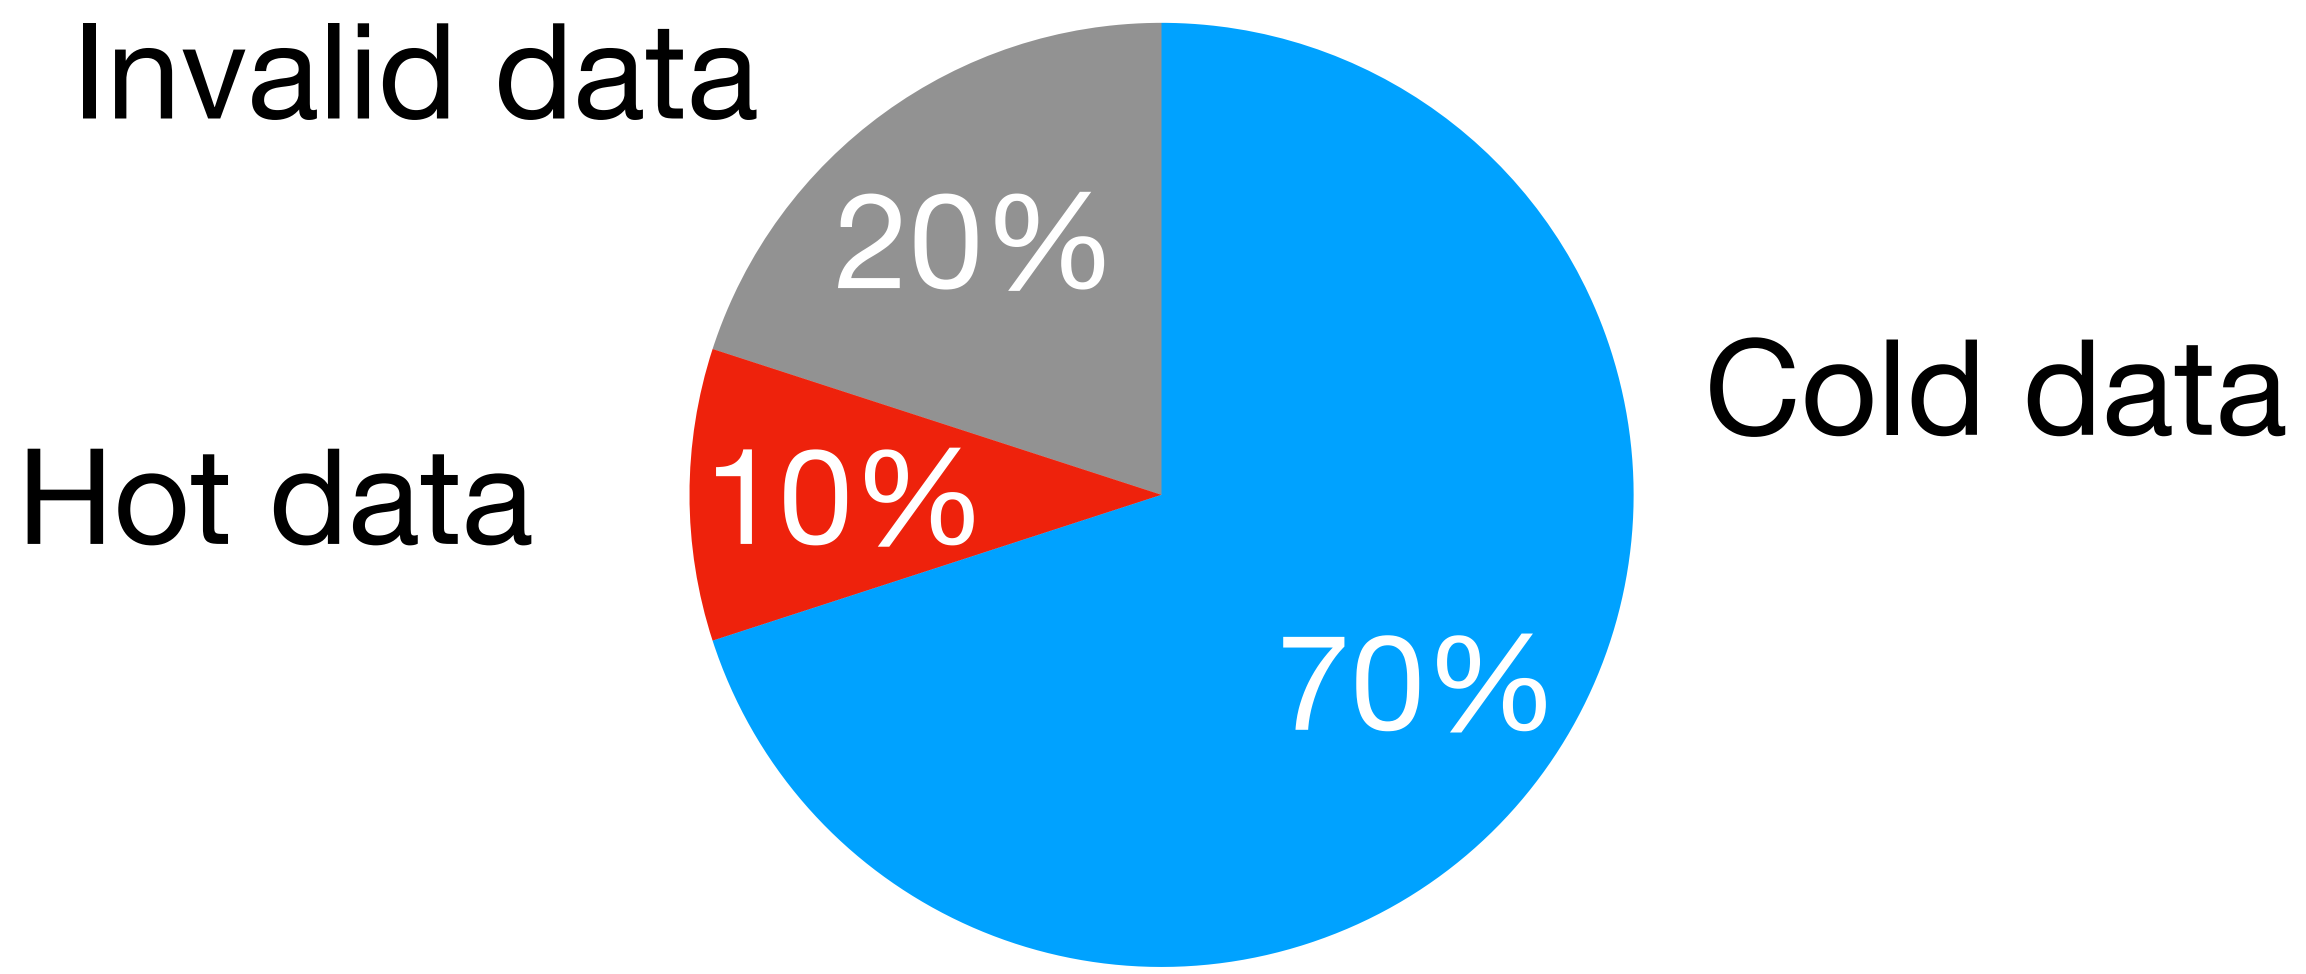

Worst case WA: $= 1 + \frac{L}{O}$

Uniformly random: $= 1 + \frac{1}{2} \cdot \frac{L}{O}$

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

Let C=0.7, H=0.1 and O=0.2

Worst case WA: $= 1 + \frac{L}{O}$

Uniformly random: $= 1 + \frac{1}{2} \cdot \frac{L}{O}$

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

Let C=0.7, H=0.1 and O=0.2

In worst-case, same amount of live data in each area

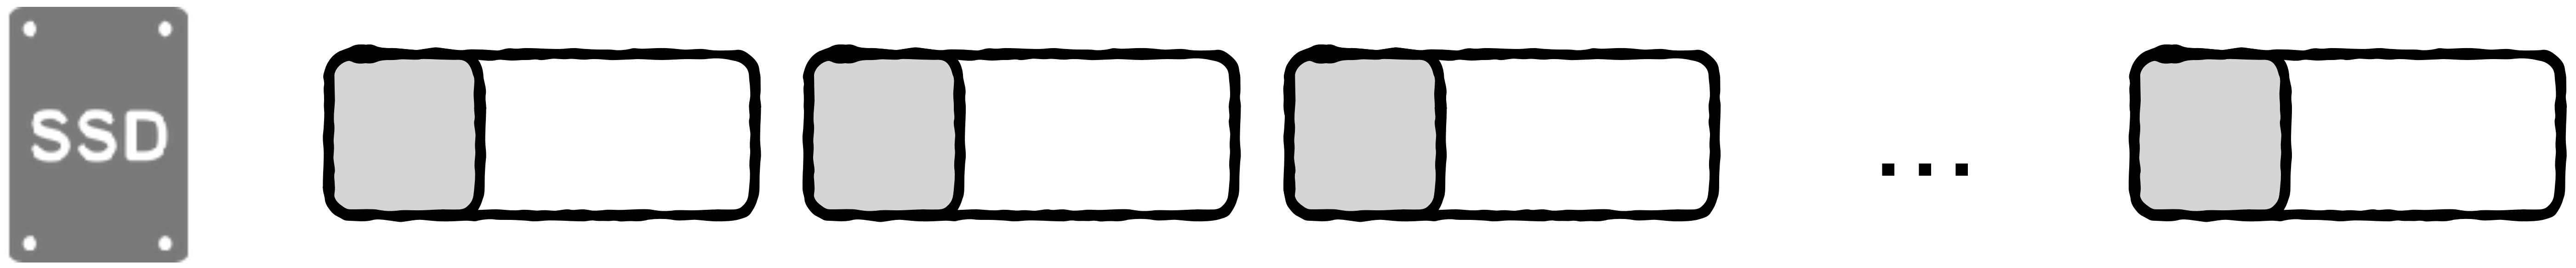

Worst case WA: $= 1 + \frac{L}{O}$

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

Let C=0.7, H=0.1 and O=0.2

In worst-case, same amount of live data in each area

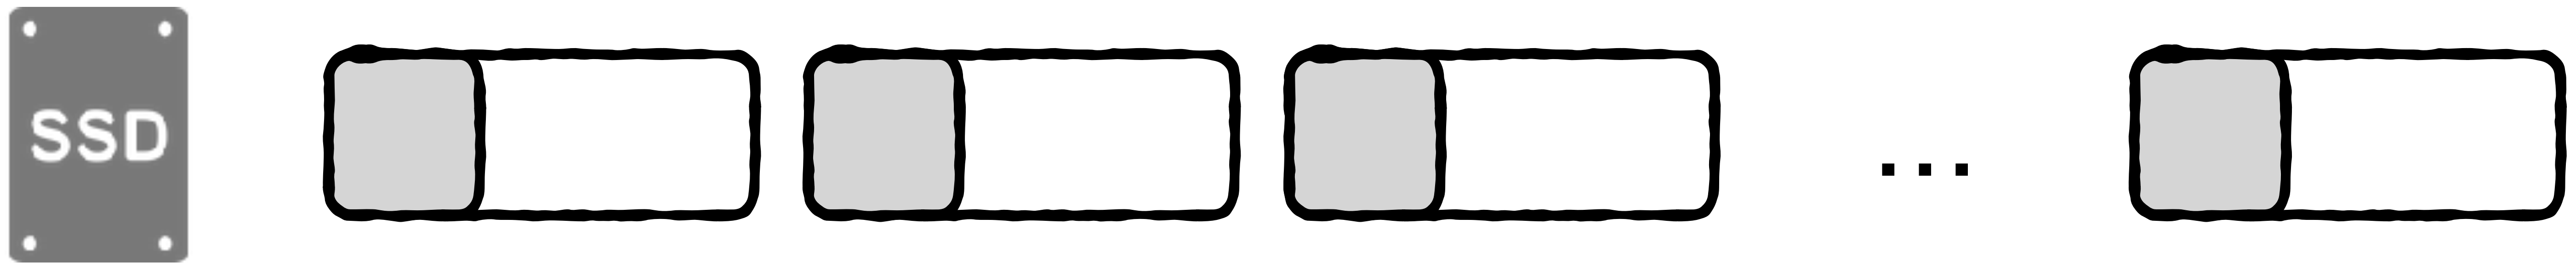

Upper bound: $$= 1 + \frac{L}{O} = 1 + \frac{H + C}{O} = 1 + \frac{0.8}{0.2} = 5$$

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

Let C=0.7, H=0.1 and O=0.2

For lower bound, let's use our uniform workload distribution estimation.

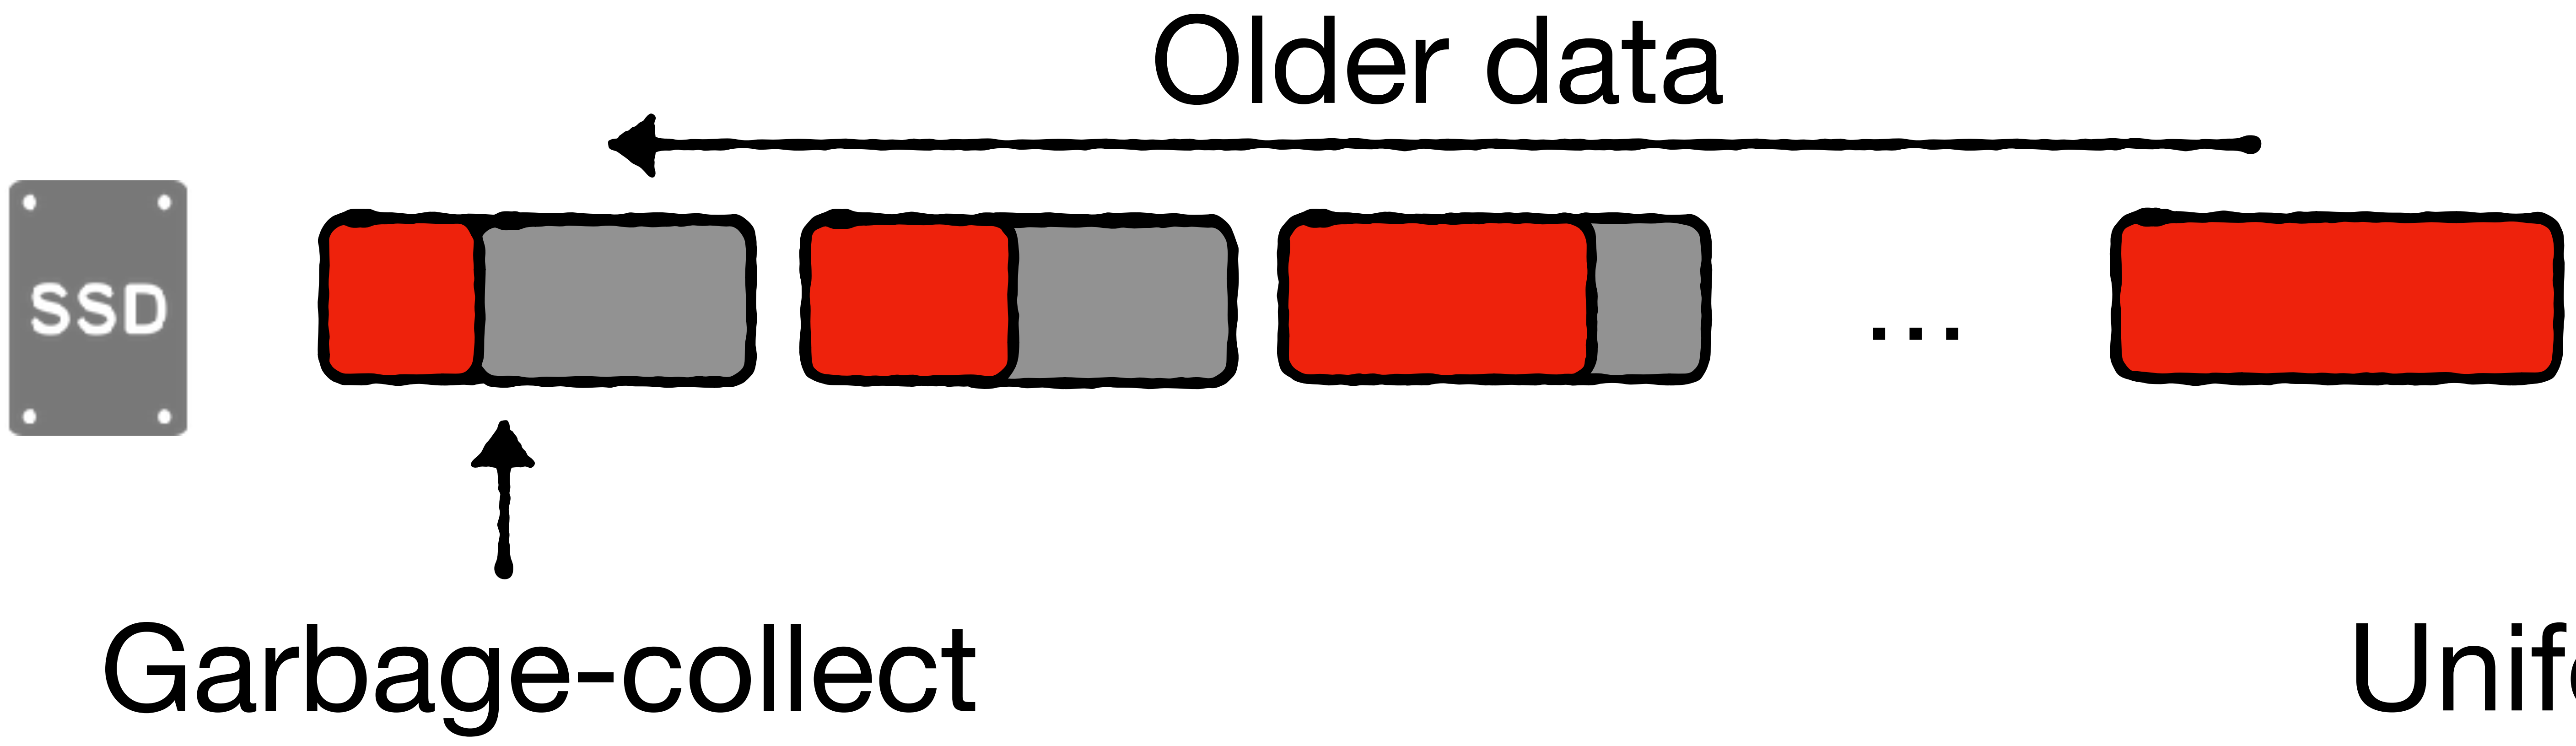

Uniformly random: $= 1 + \frac{1}{2} \cdot \frac{L}{O}$

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

Let C=0.7, H=0.1 and O=0.2

For lower bound, let's use our uniform workload distribution estimation.

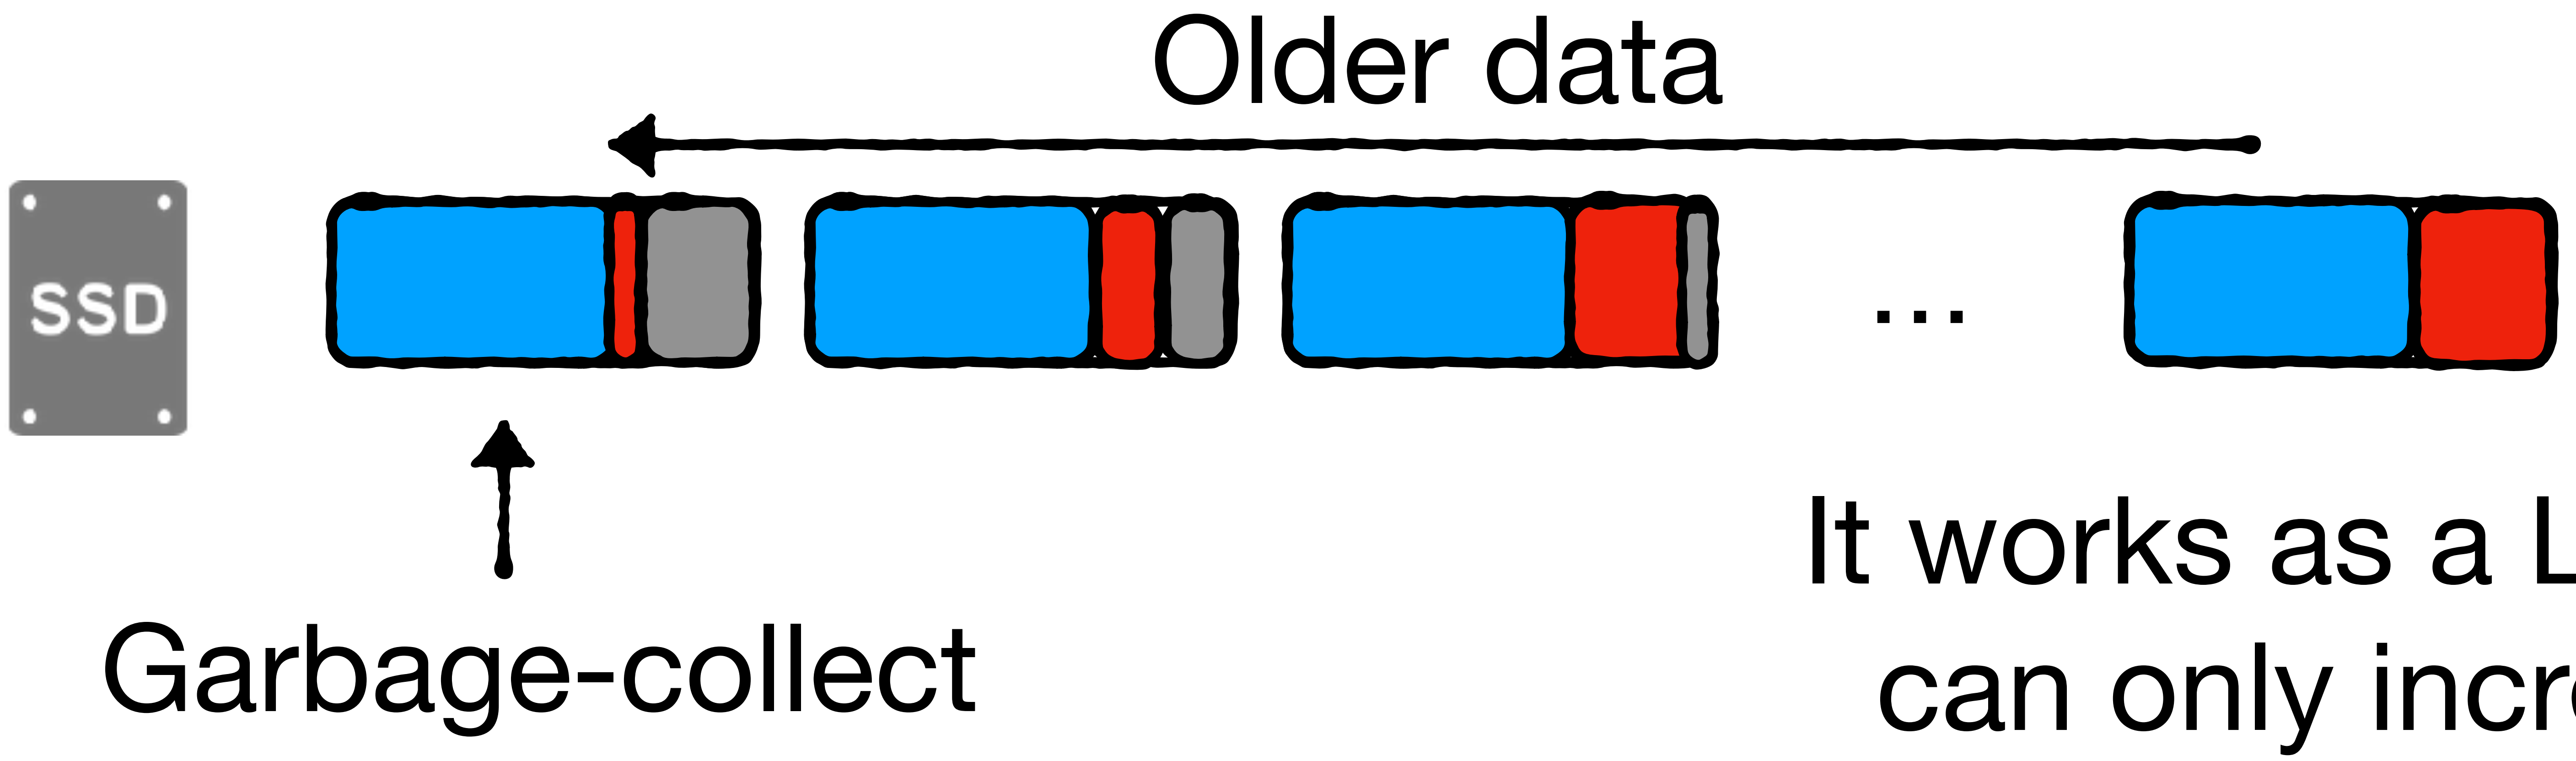

It works as a Lower bound since static data can only increase fraction of valid data in areas we garbage-collect

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

Let C=0.7, H=0.1 and O=0.2

For lower bound, let's use our uniform workload distribution estimation.

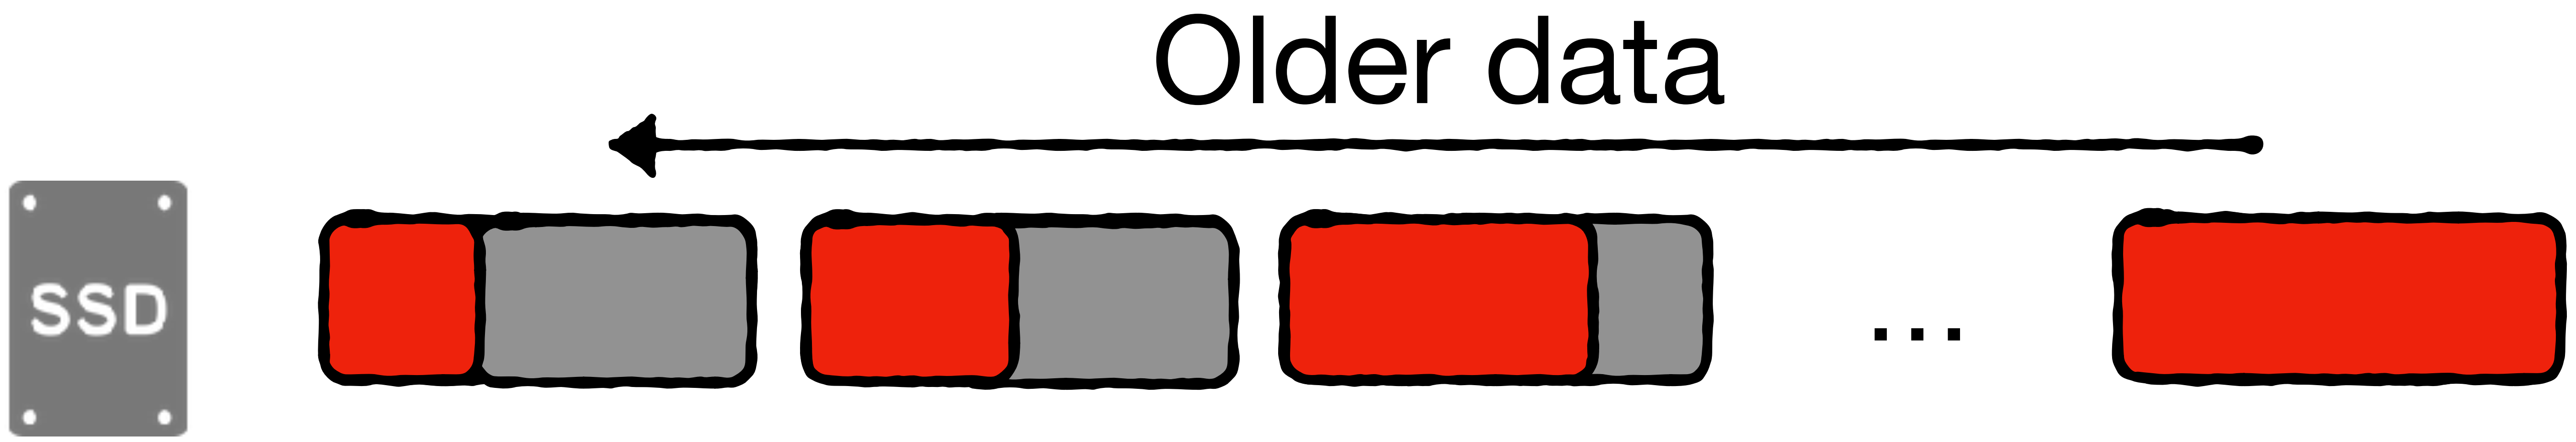

Lower bound: $$= 1 + \frac{1}{2} \cdot \frac{L}{O} = 1 + \frac{1}{2} \cdot \frac{H + C}{O} = 1 + \frac{1}{2} \cdot \frac{0.8}{0.2} = 3$$

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

(B) Estimate write-amplification assuming perfect hot/cold data separation.

Let C=0.7, H=0.1 and O=0.2

For hot/cold separation estimation, assume all over-provisioned space is applied on hot areas.

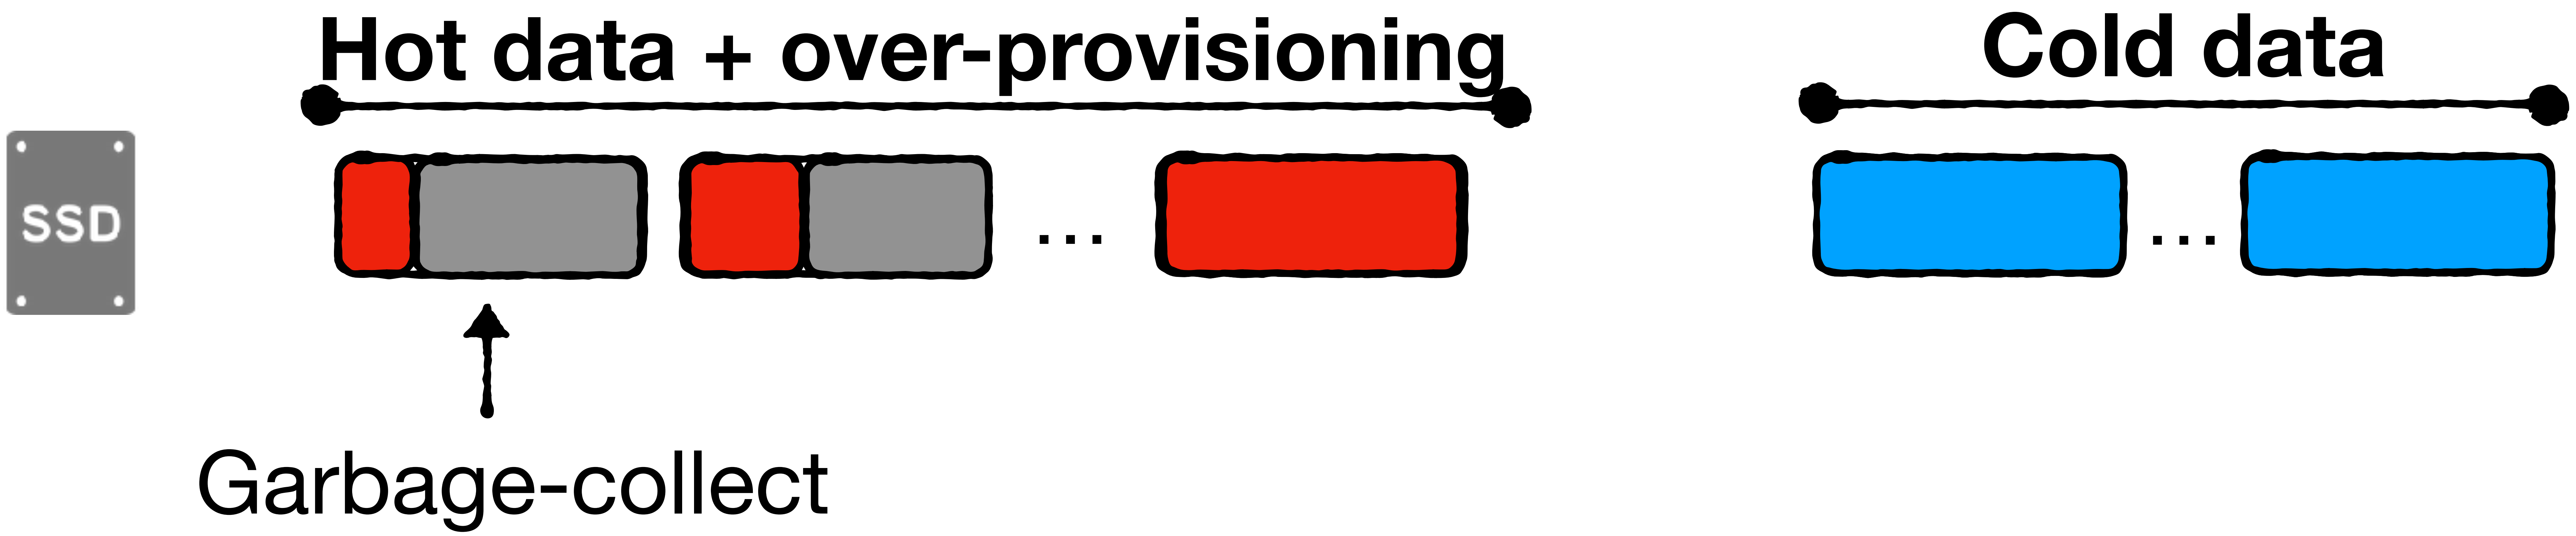

Question 3 - Hot/Cold Data Separation

Consider a circular log where the physical capacity consists of 70% static data (never updated), 10% hot data, and 20% over-provisioning.

(A) Estimate a lower bound and upper bound for write-amplification assuming no hot/cold data separation.

(B) Estimate write-amplification assuming perfect hot/cold data separation.

Let C=0.7, H=0.1 and O=0.2

For hot/cold separation estimation, assume all over-provisioned space is applied on hot areas.

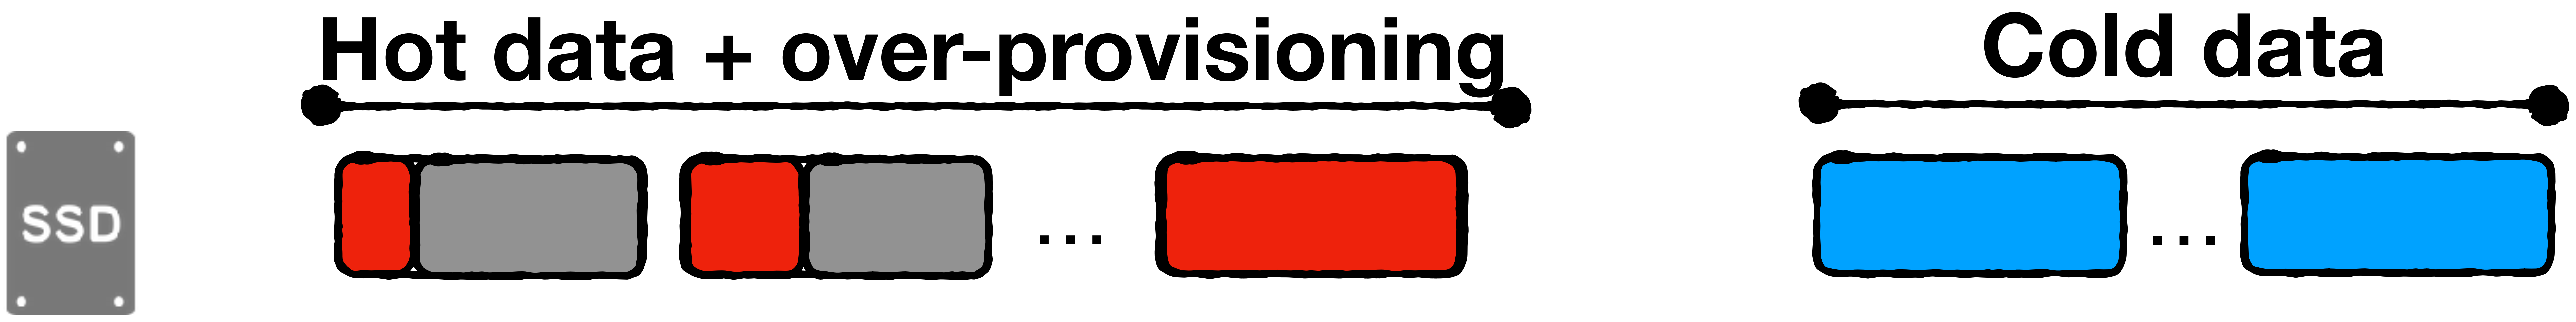

Estimation: $= 1 + \frac{1}{2} \cdot \frac{L}{O} = 1 + \frac{1}{2} \cdot \frac{H}{O} = 1 + \frac{1}{2} \cdot \frac{0.1}{0.2} = 1.25$