

# GARBAGE COLLECTION

---

SAURABH KADEKODI  
RAJAT KATEJA  
TIANYUAN DING

### PROBLEM STATEMENT

- ▶ To implement efficient garbage collection in Peloton

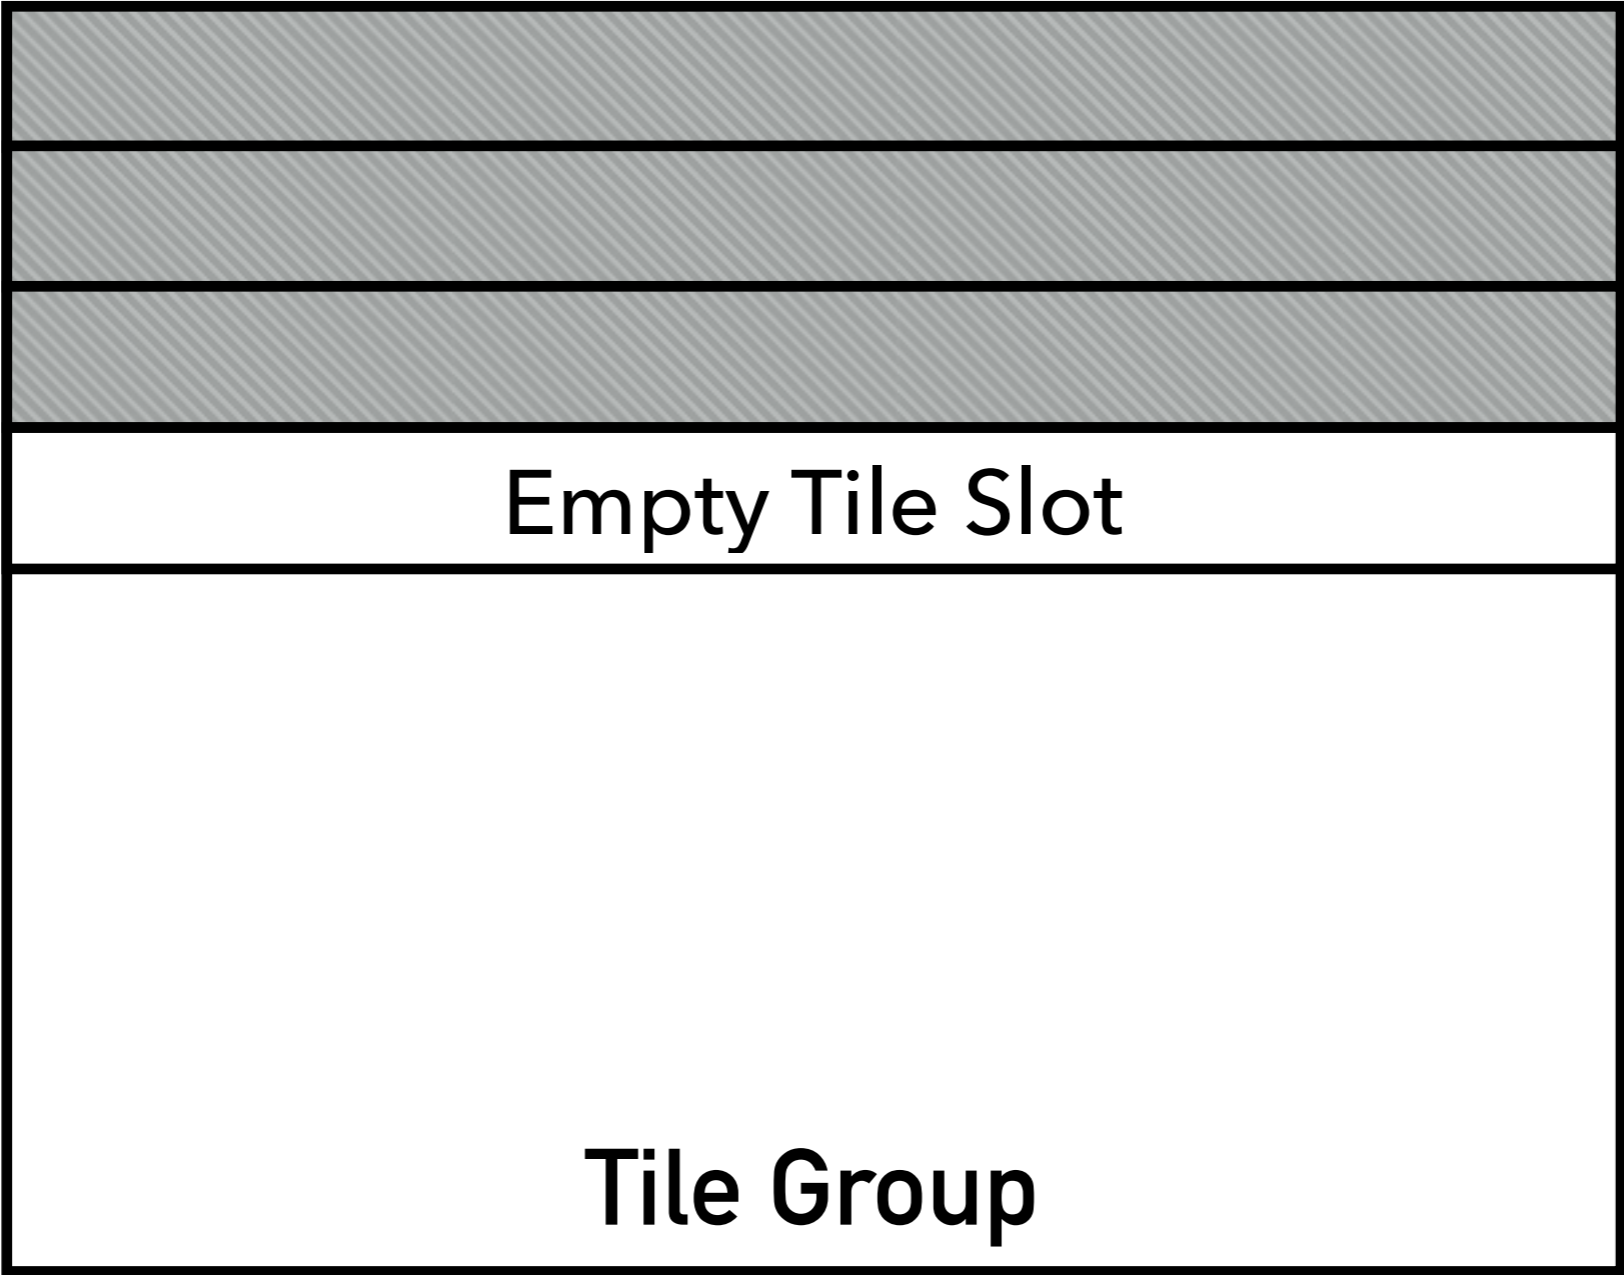
### GOALS (REVISED)

- ▶ 75% - implement basic tuple recycling using vacuum
- ▶ 100% - implement epoch based co-operative gc
- ▶ 112.5% - lock-free implementations, GetMemoryFootprint()
- ▶ 125% - DDL garbage collection (deferred after discussion)

# WITHOUT GC

Add new tuple



Txn



# WITHOUT GC

Add new tuple

Txn






Tile Group



# WITHOUT GC

Update tuple 3

Txn

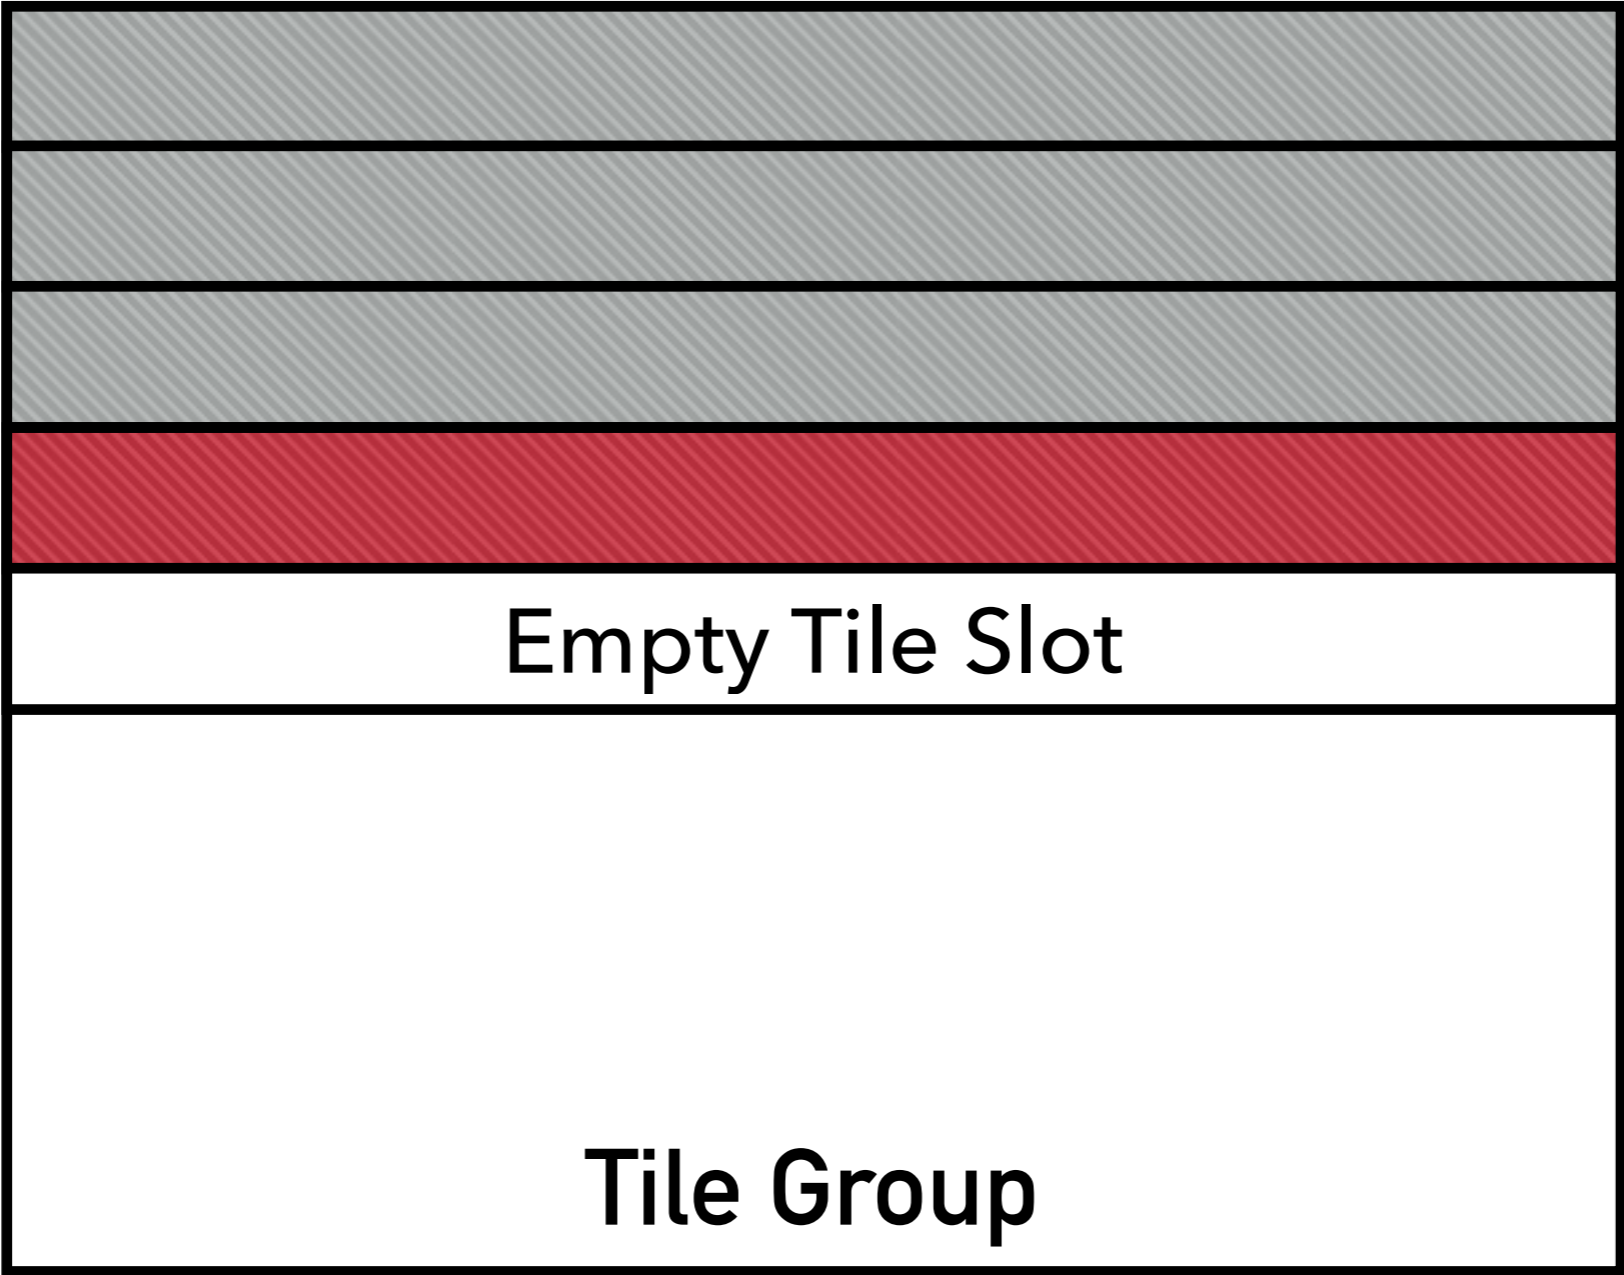


Tile Group

# WITHOUT GC

Update tuple 3


Txn



# WITHOUT GC

## Update tuple 3



Txn



# WITHOUT GC

Update tuple 3

Txn

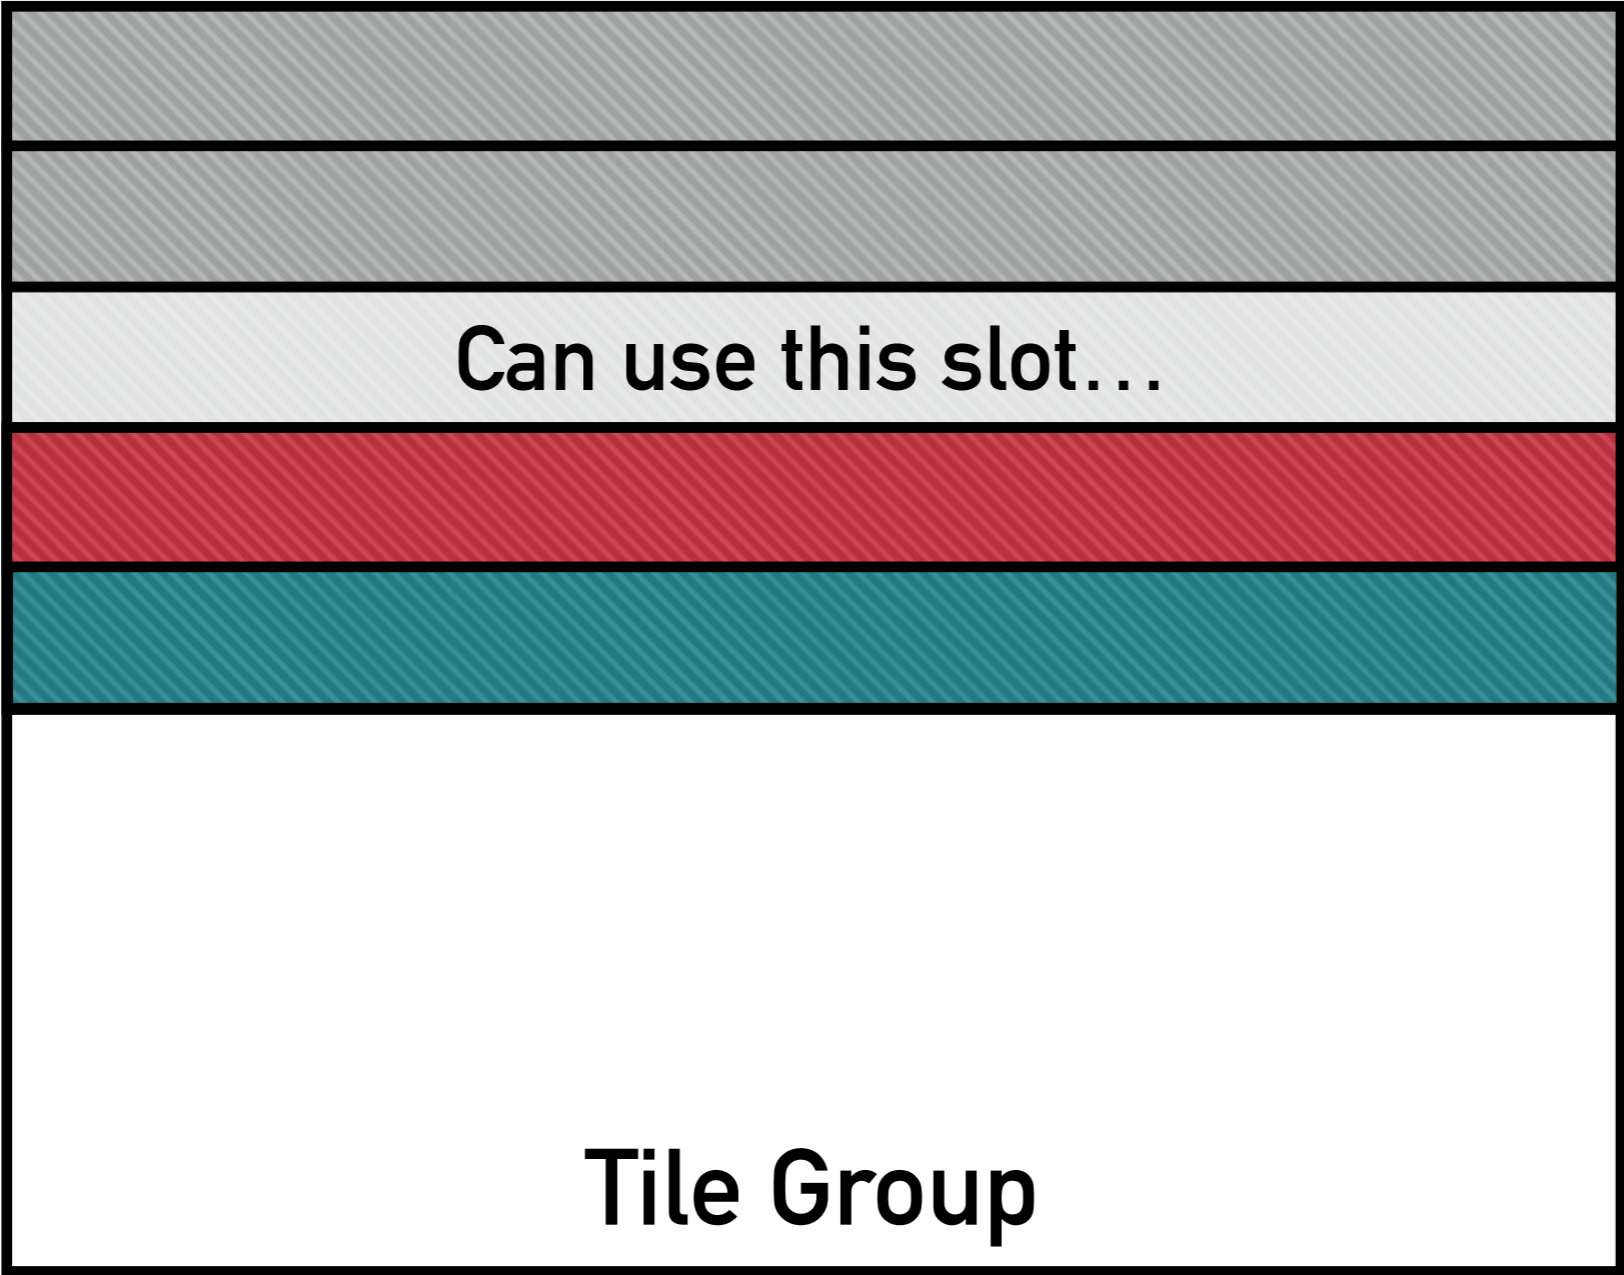





# WITHOUT GC

Update tuple 1

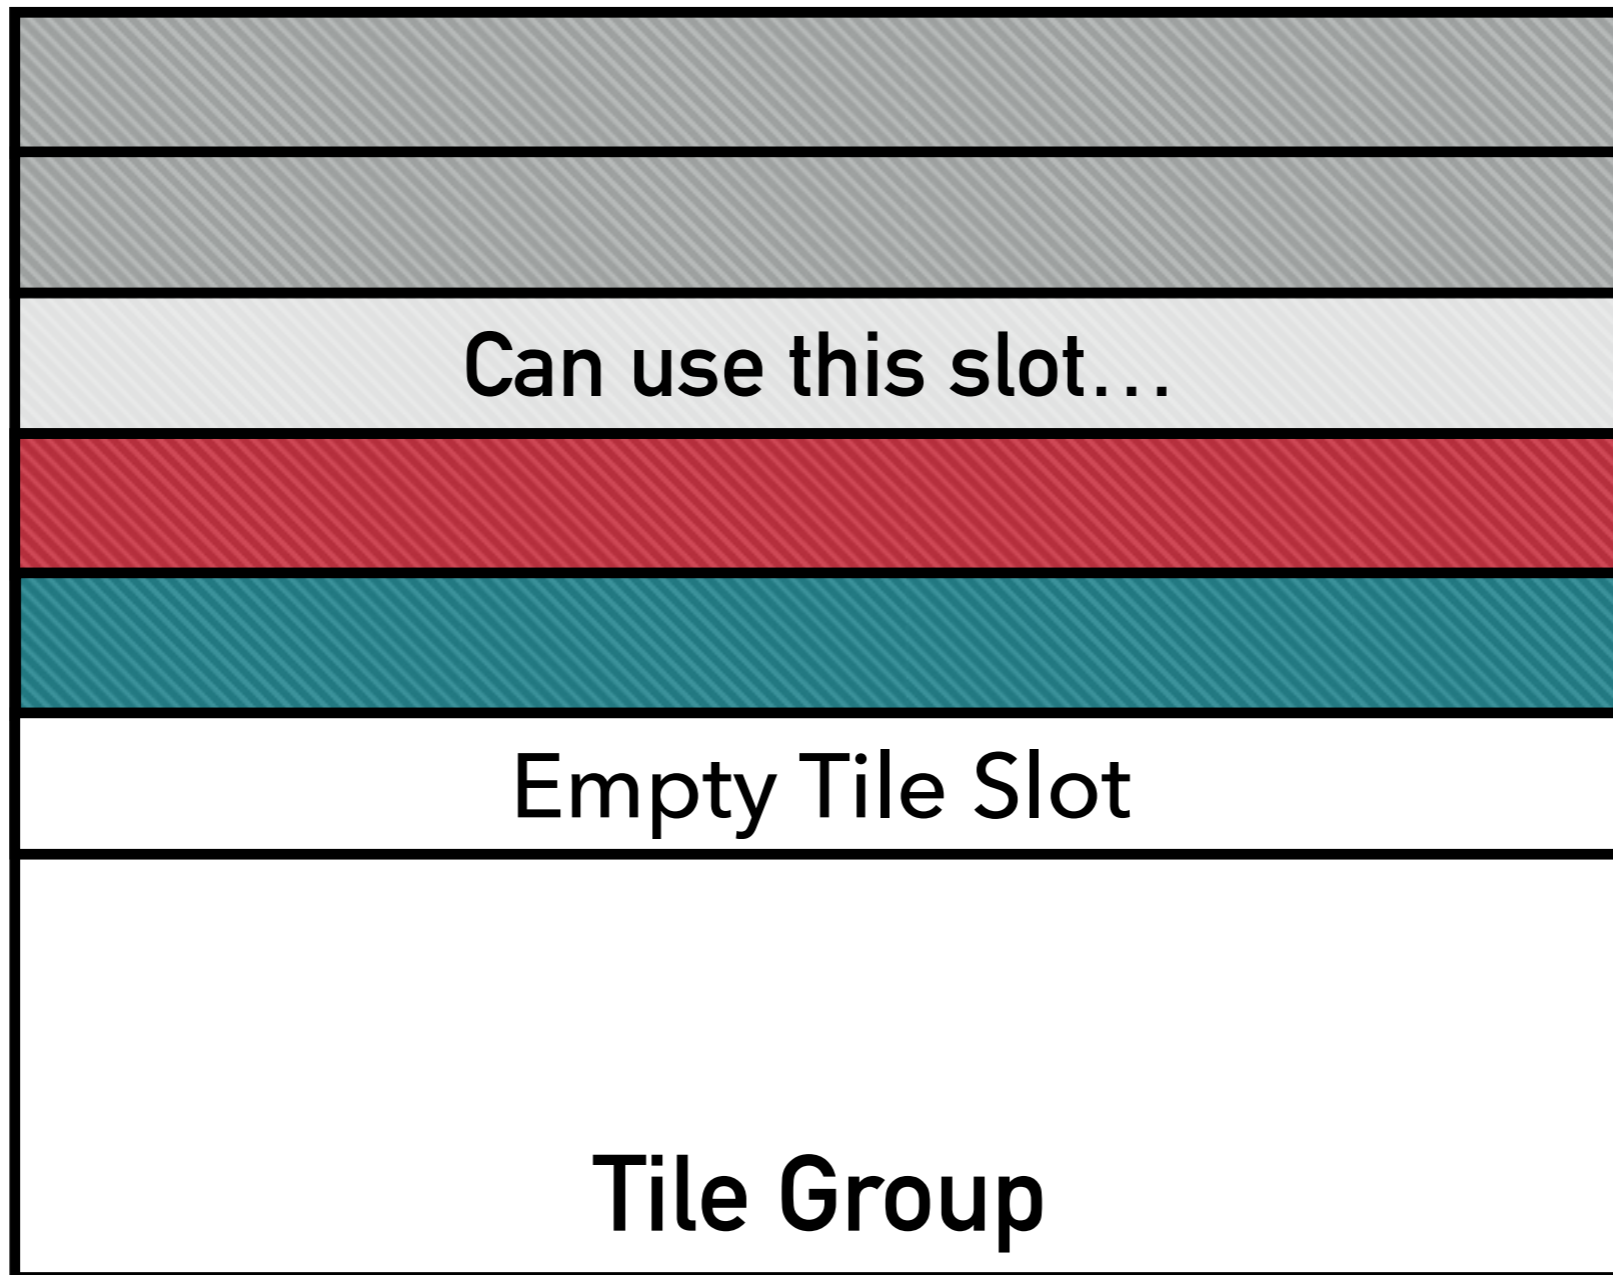
Txn



# WITHOUT GC

Update tuple 1

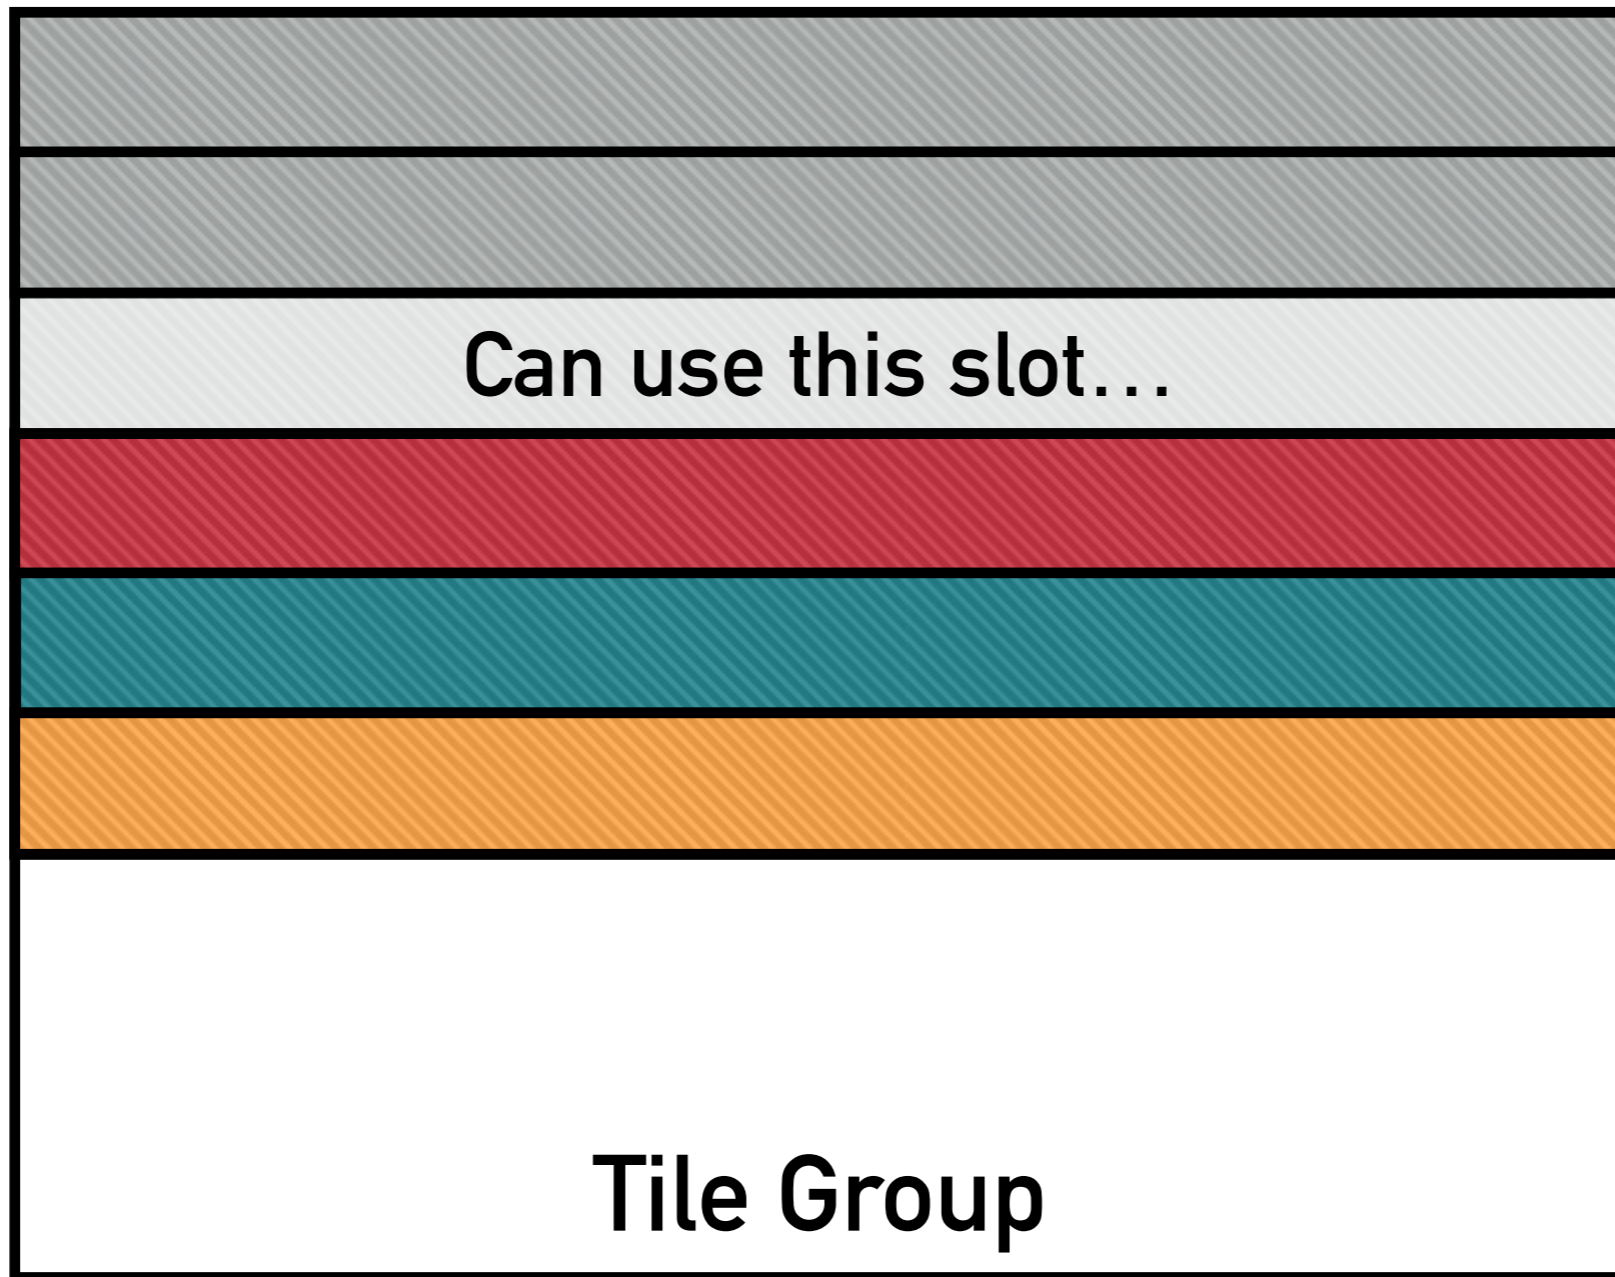
Txn



# WITHOUT GC

Update tuple 1

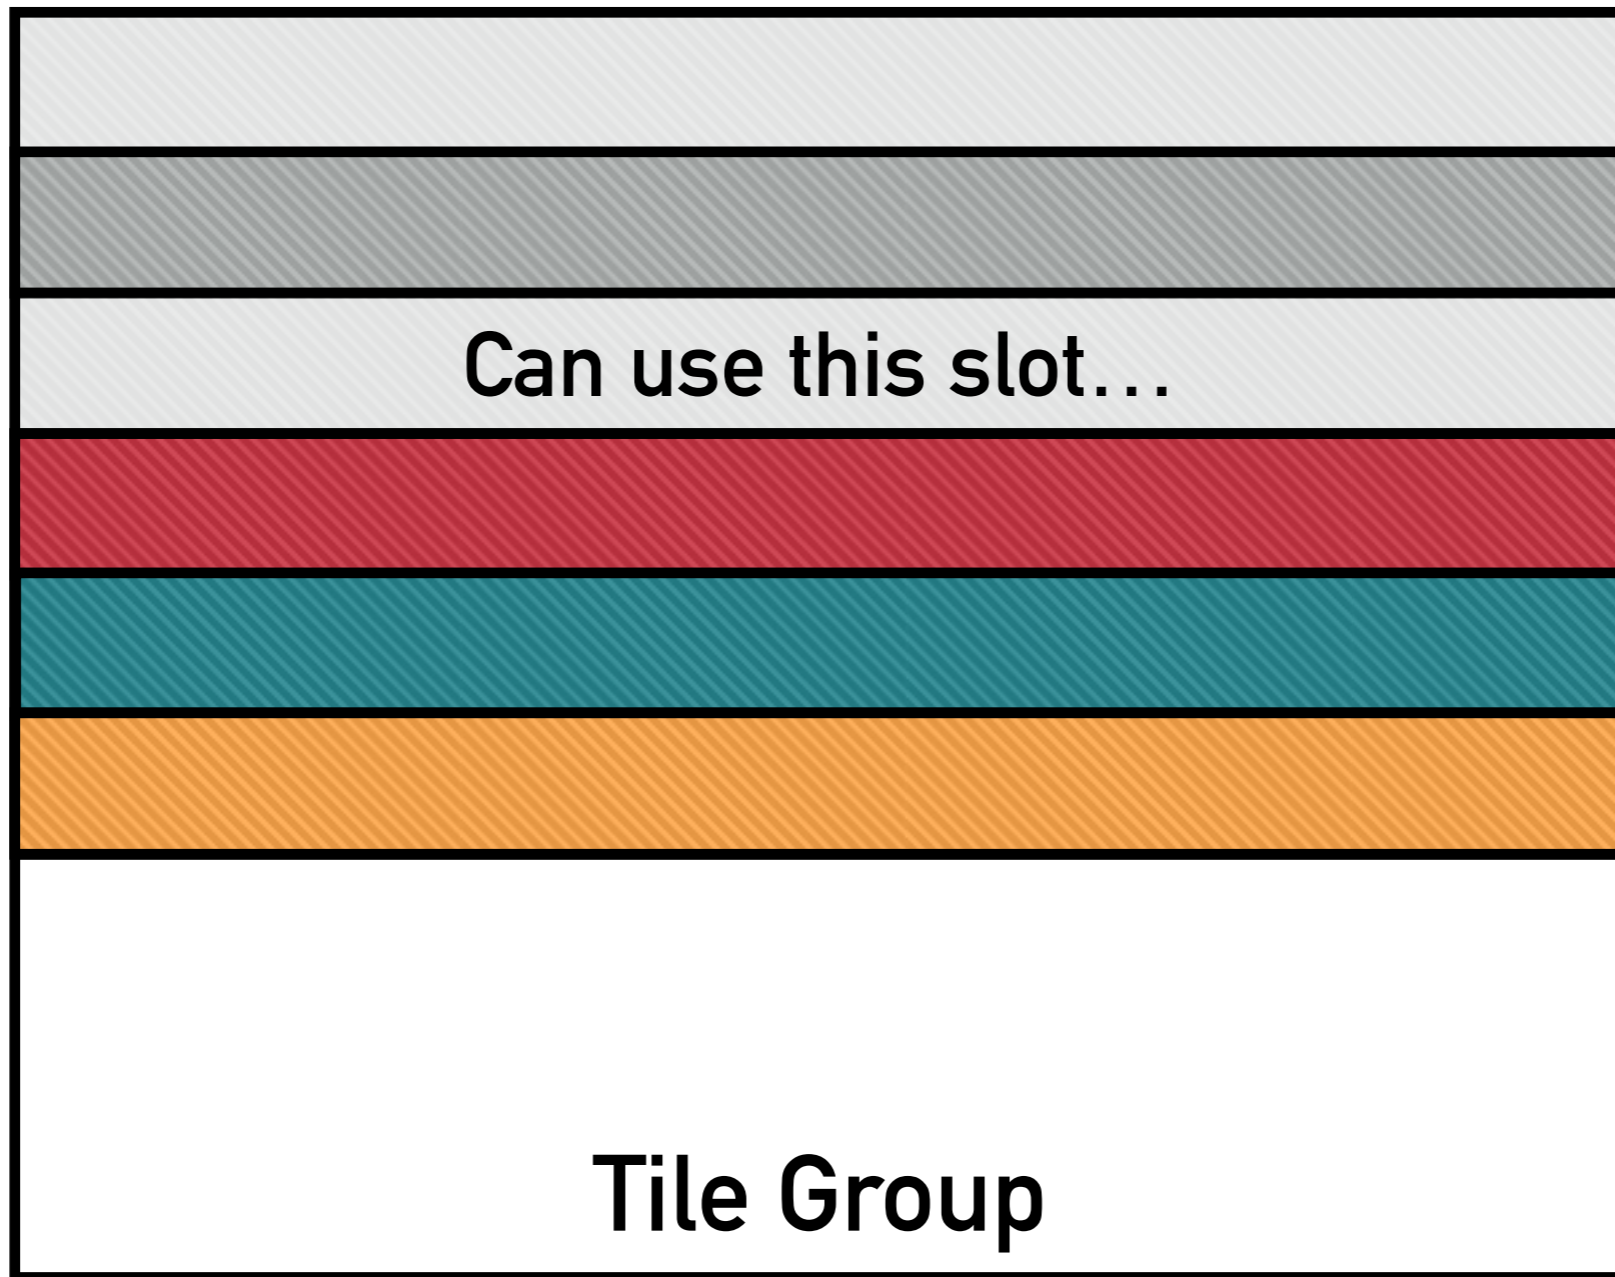
Txn



# WITHOUT GC

Update tuple 1

Txn

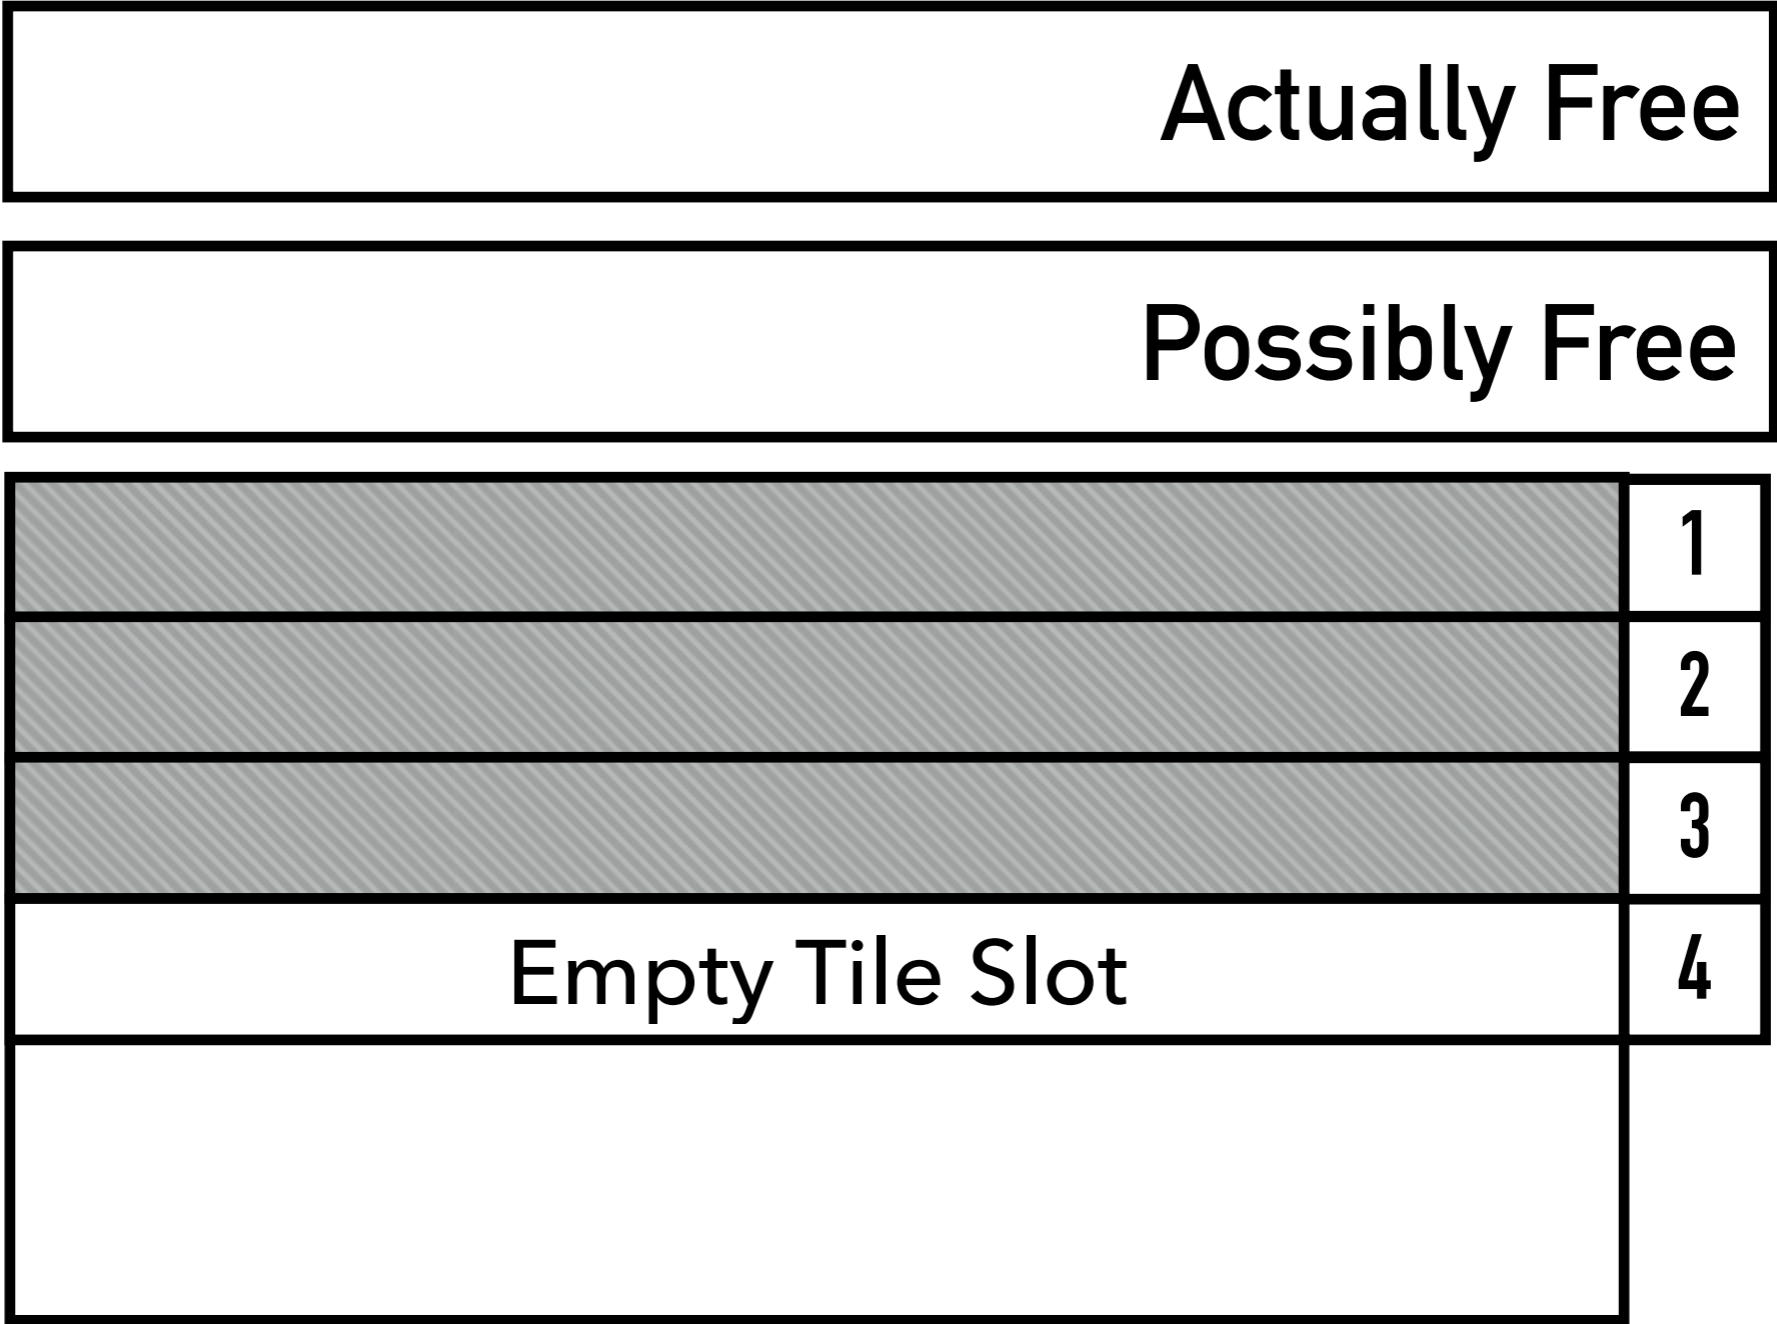


Unused Slots

Tile Group

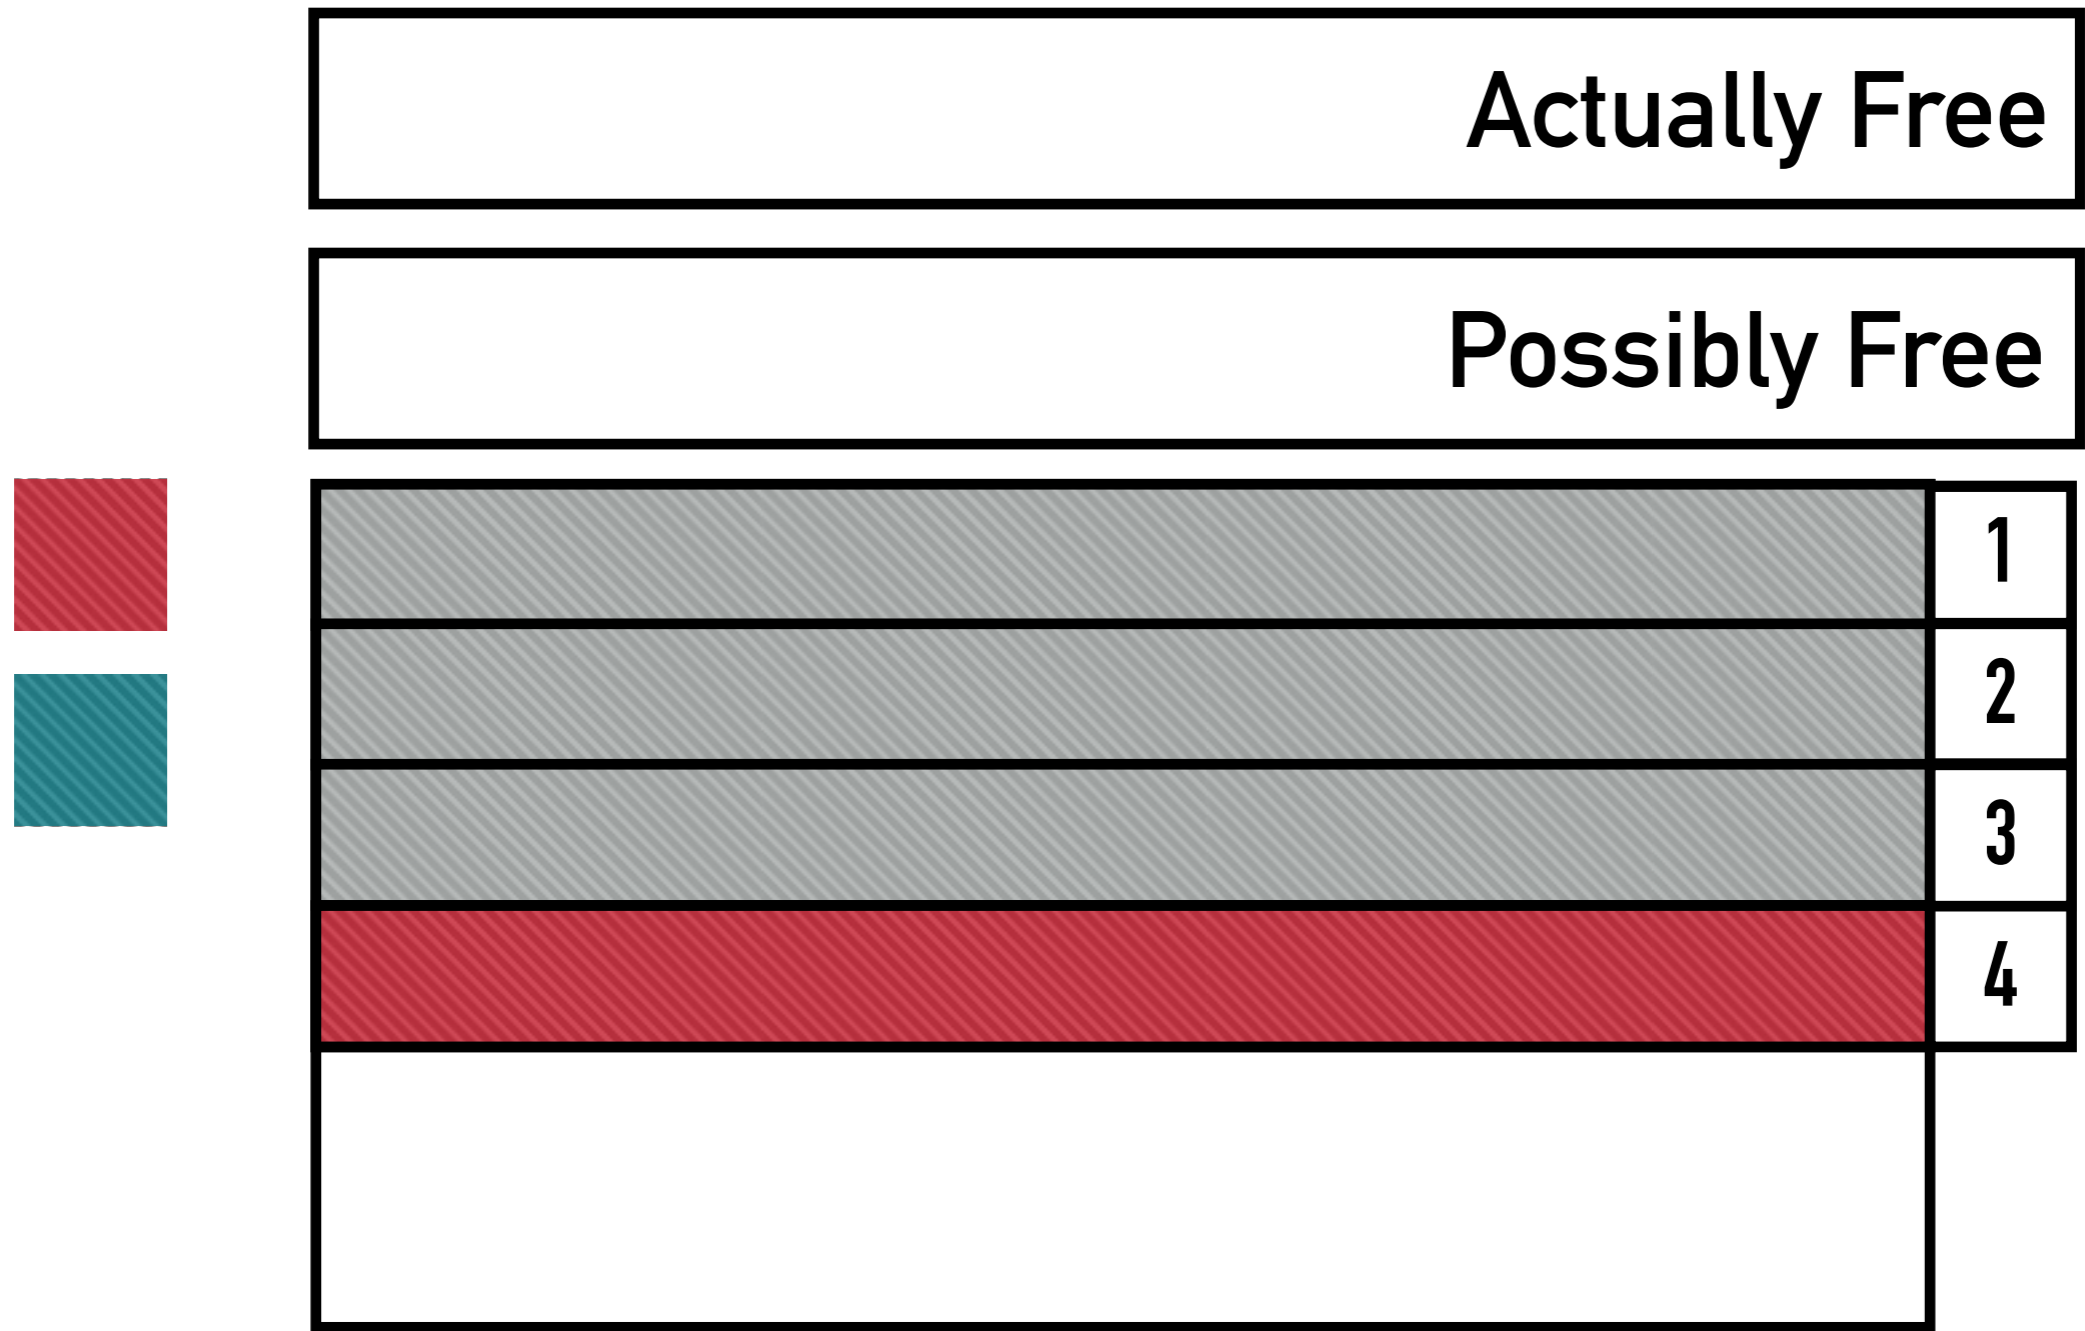


WITH GC



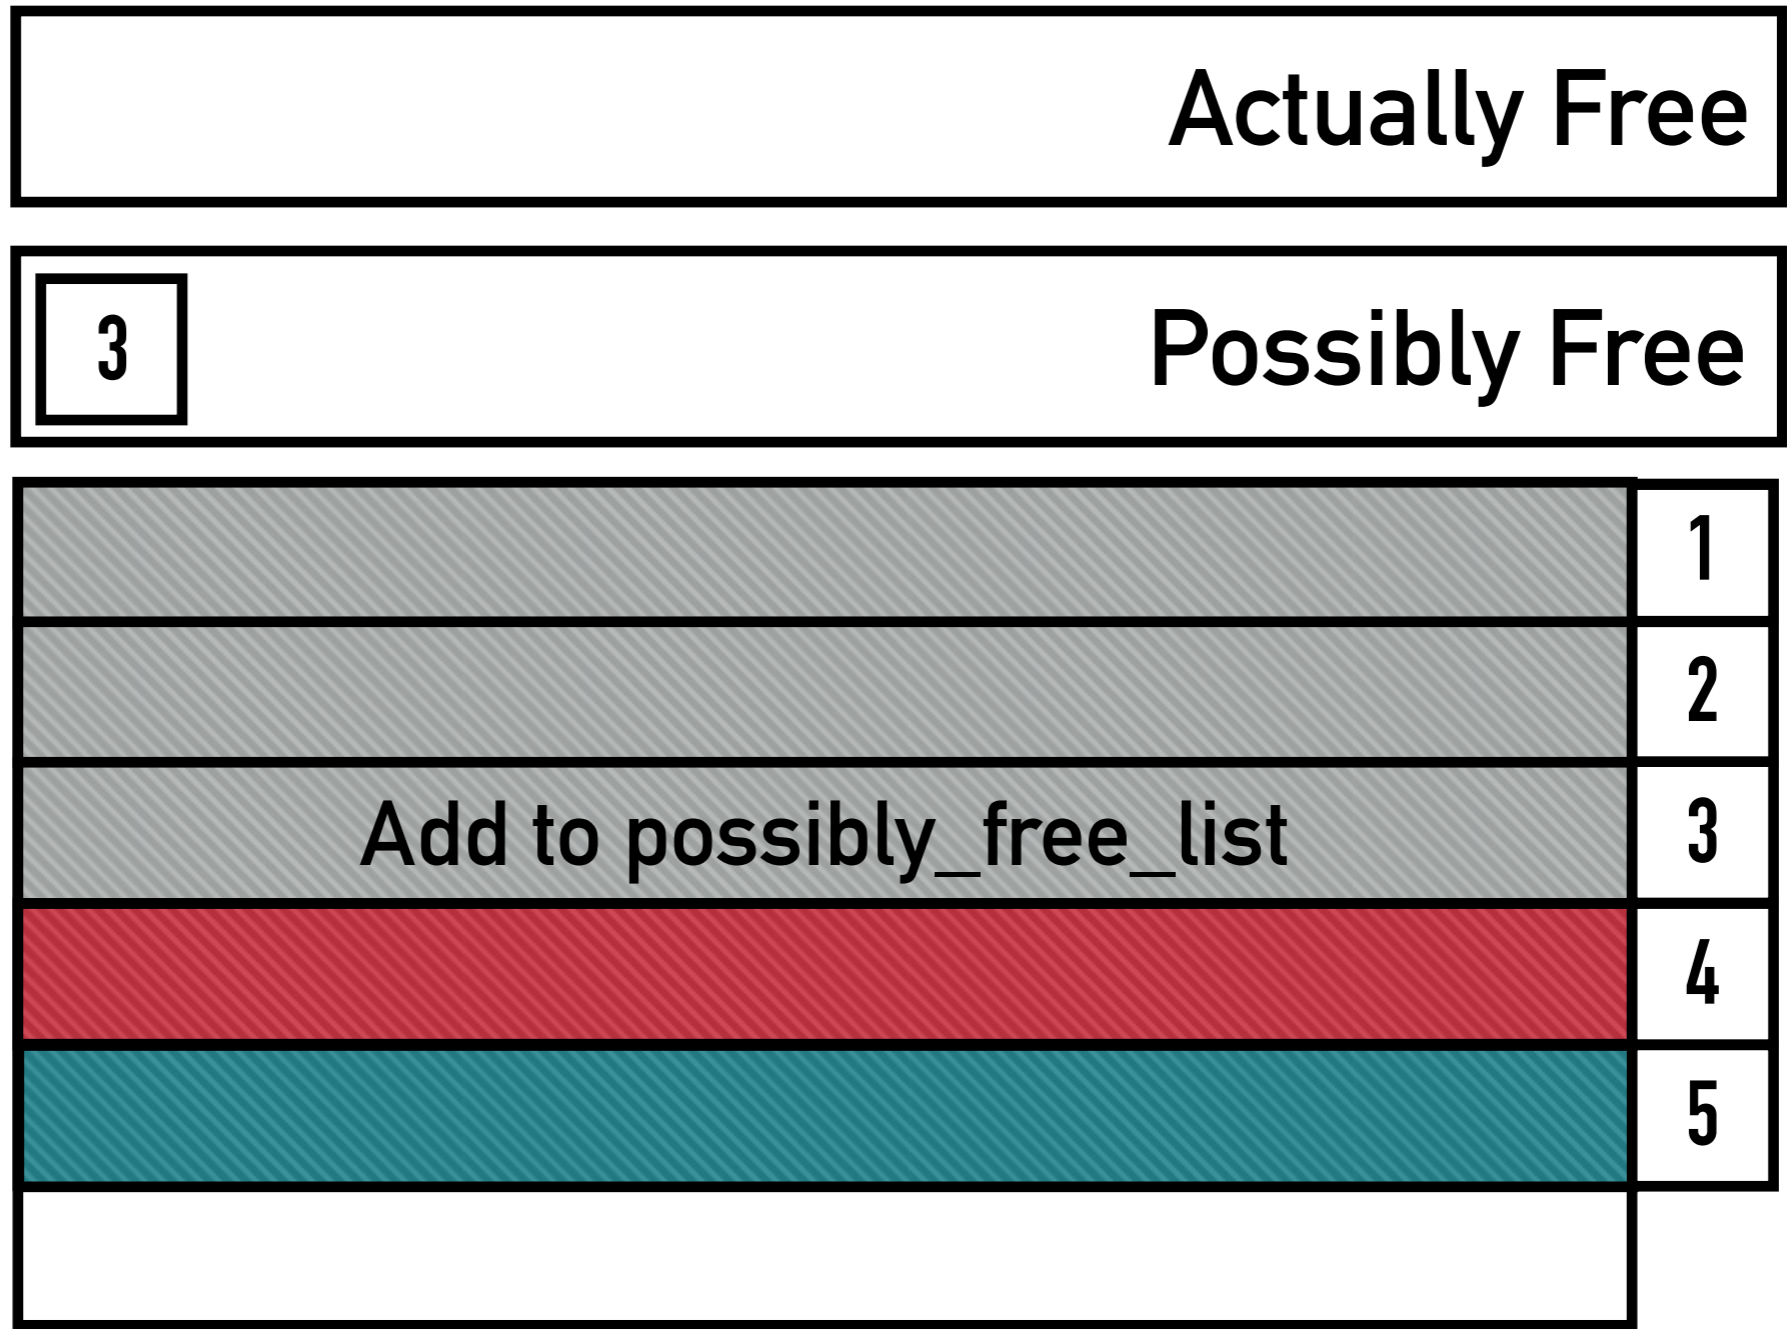
Tile Group

# WITH GC



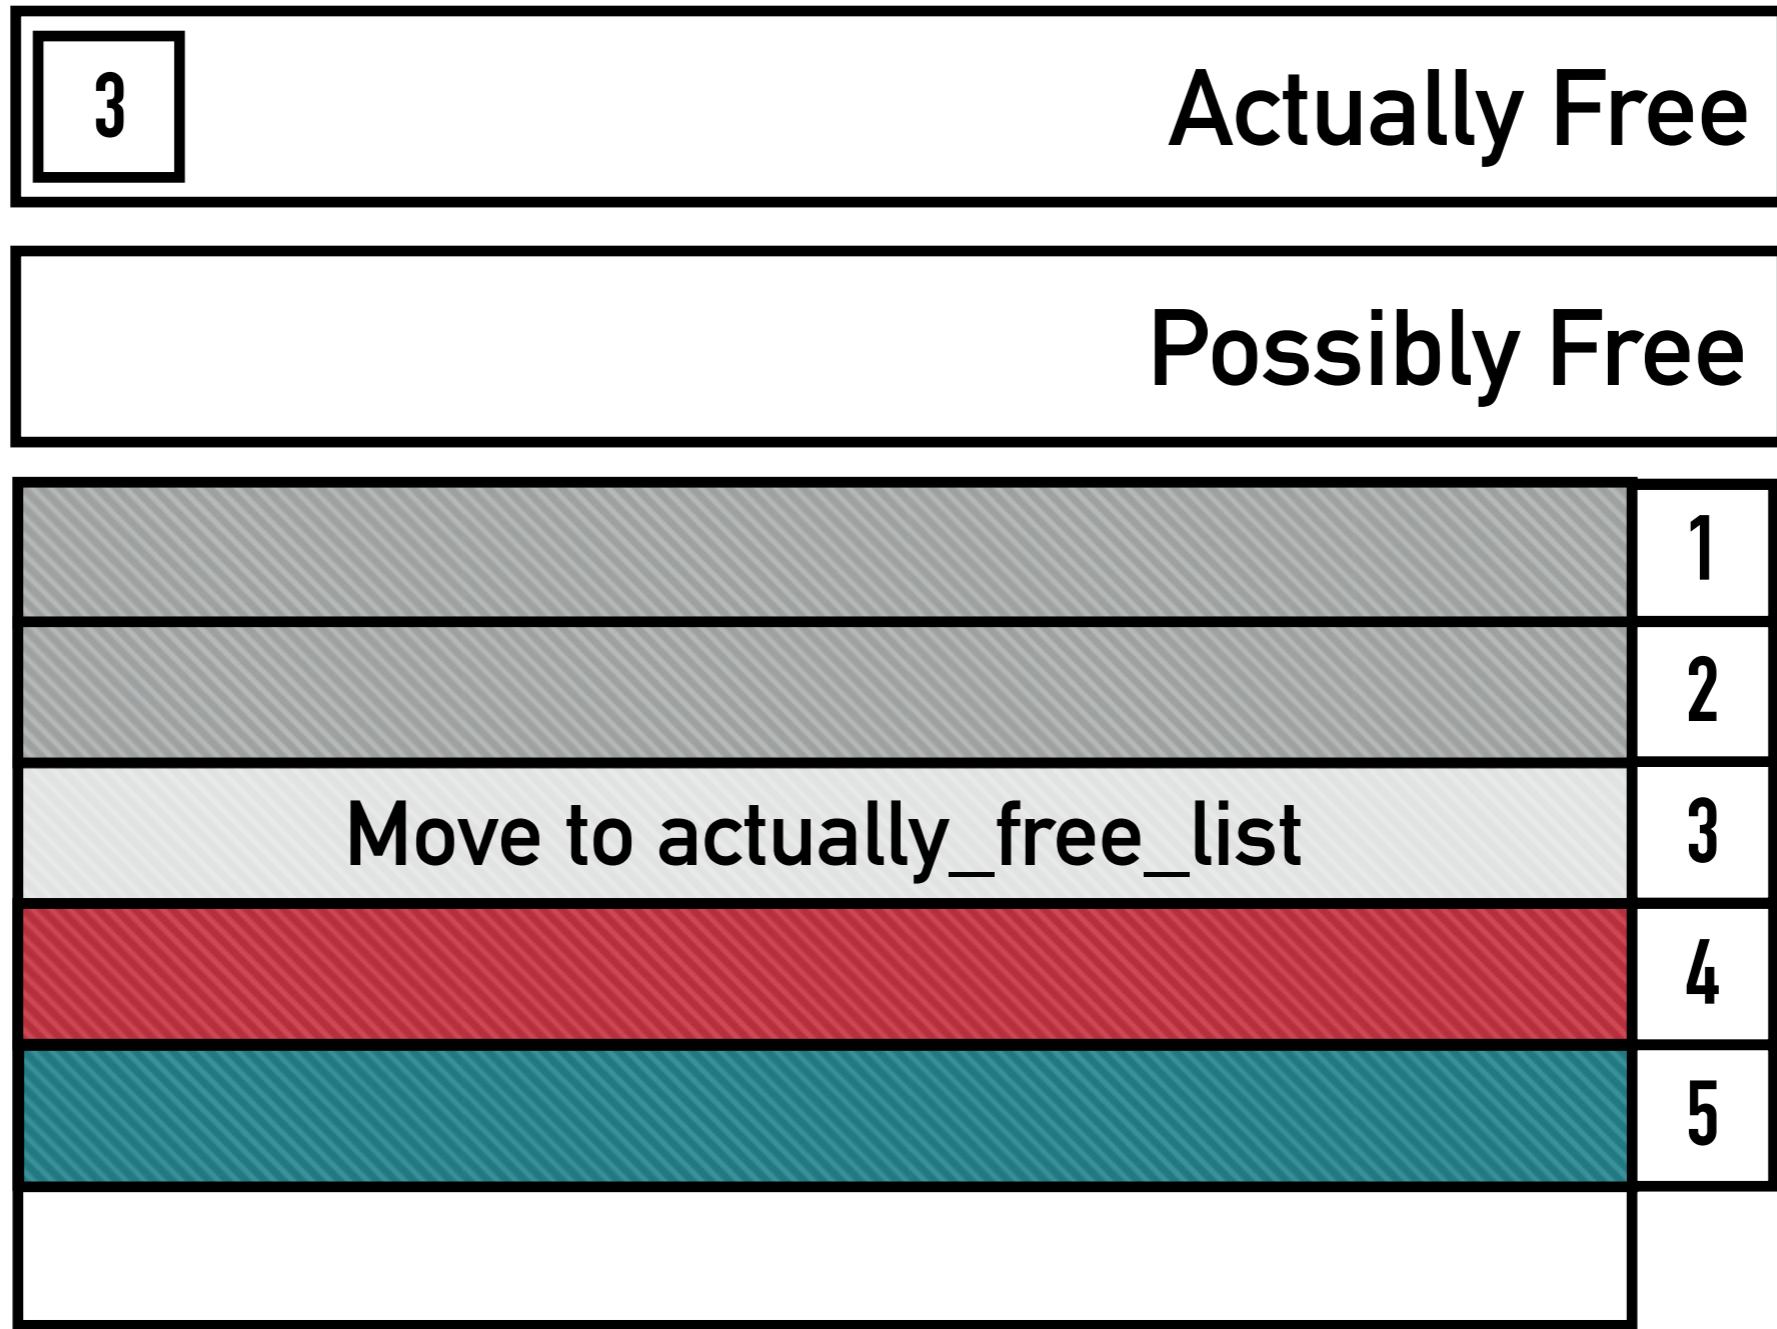
Tile Group

# WITH GC



Tile Group


WITH GC



Tile Group

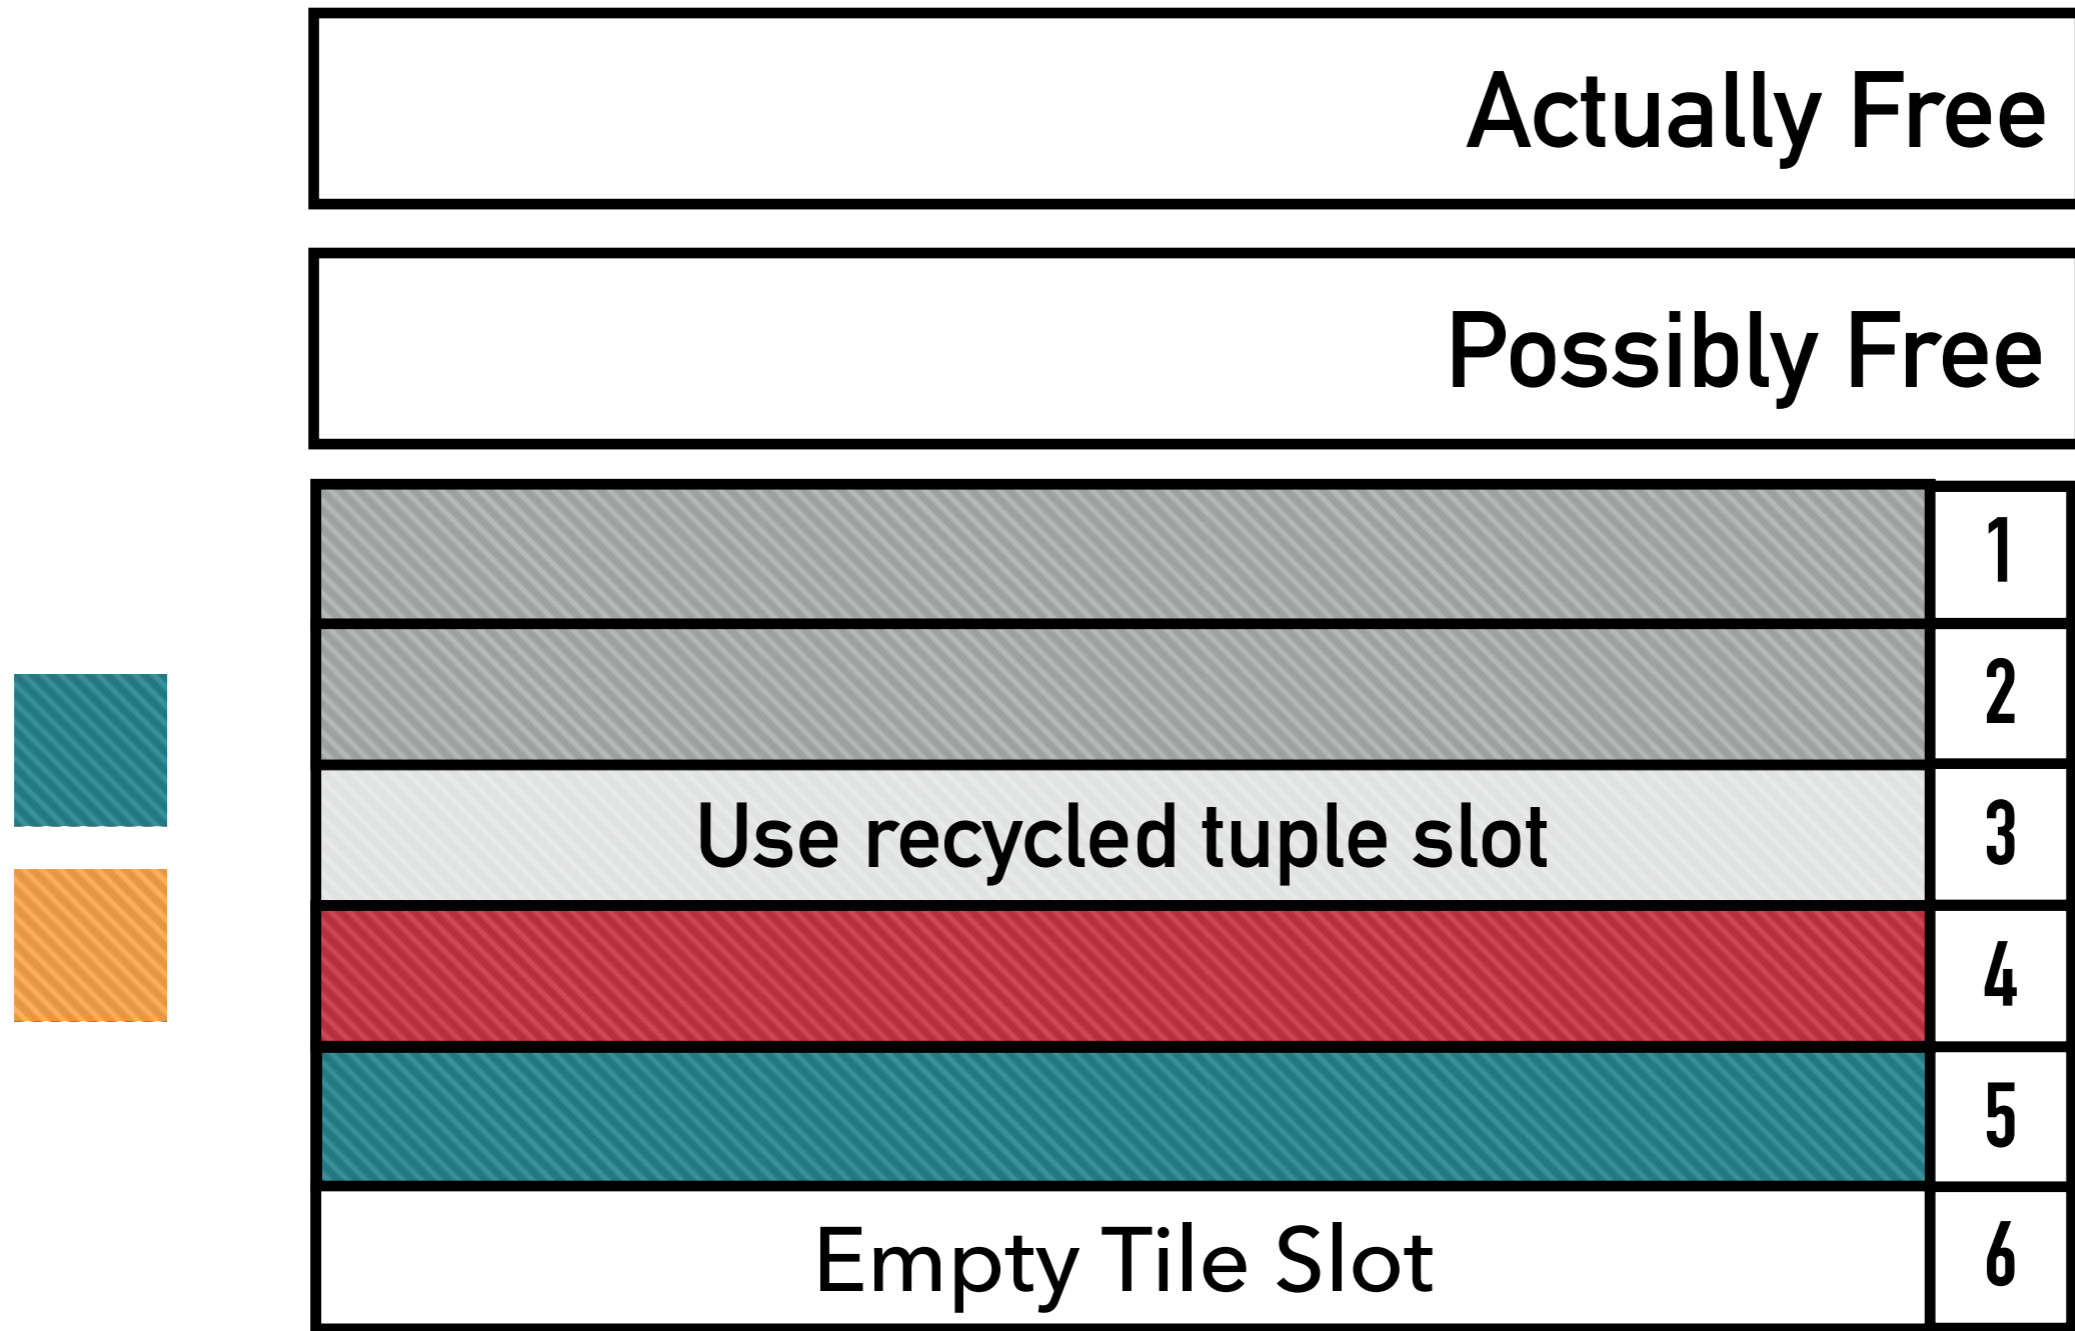


WITH GC



Tile Group

WITH GC



Tile Group

WITH GC

Actually Free


1 Possibly Free



	1
	2
	3
	4
	5
Empty Tile Slot	6

Tile Group

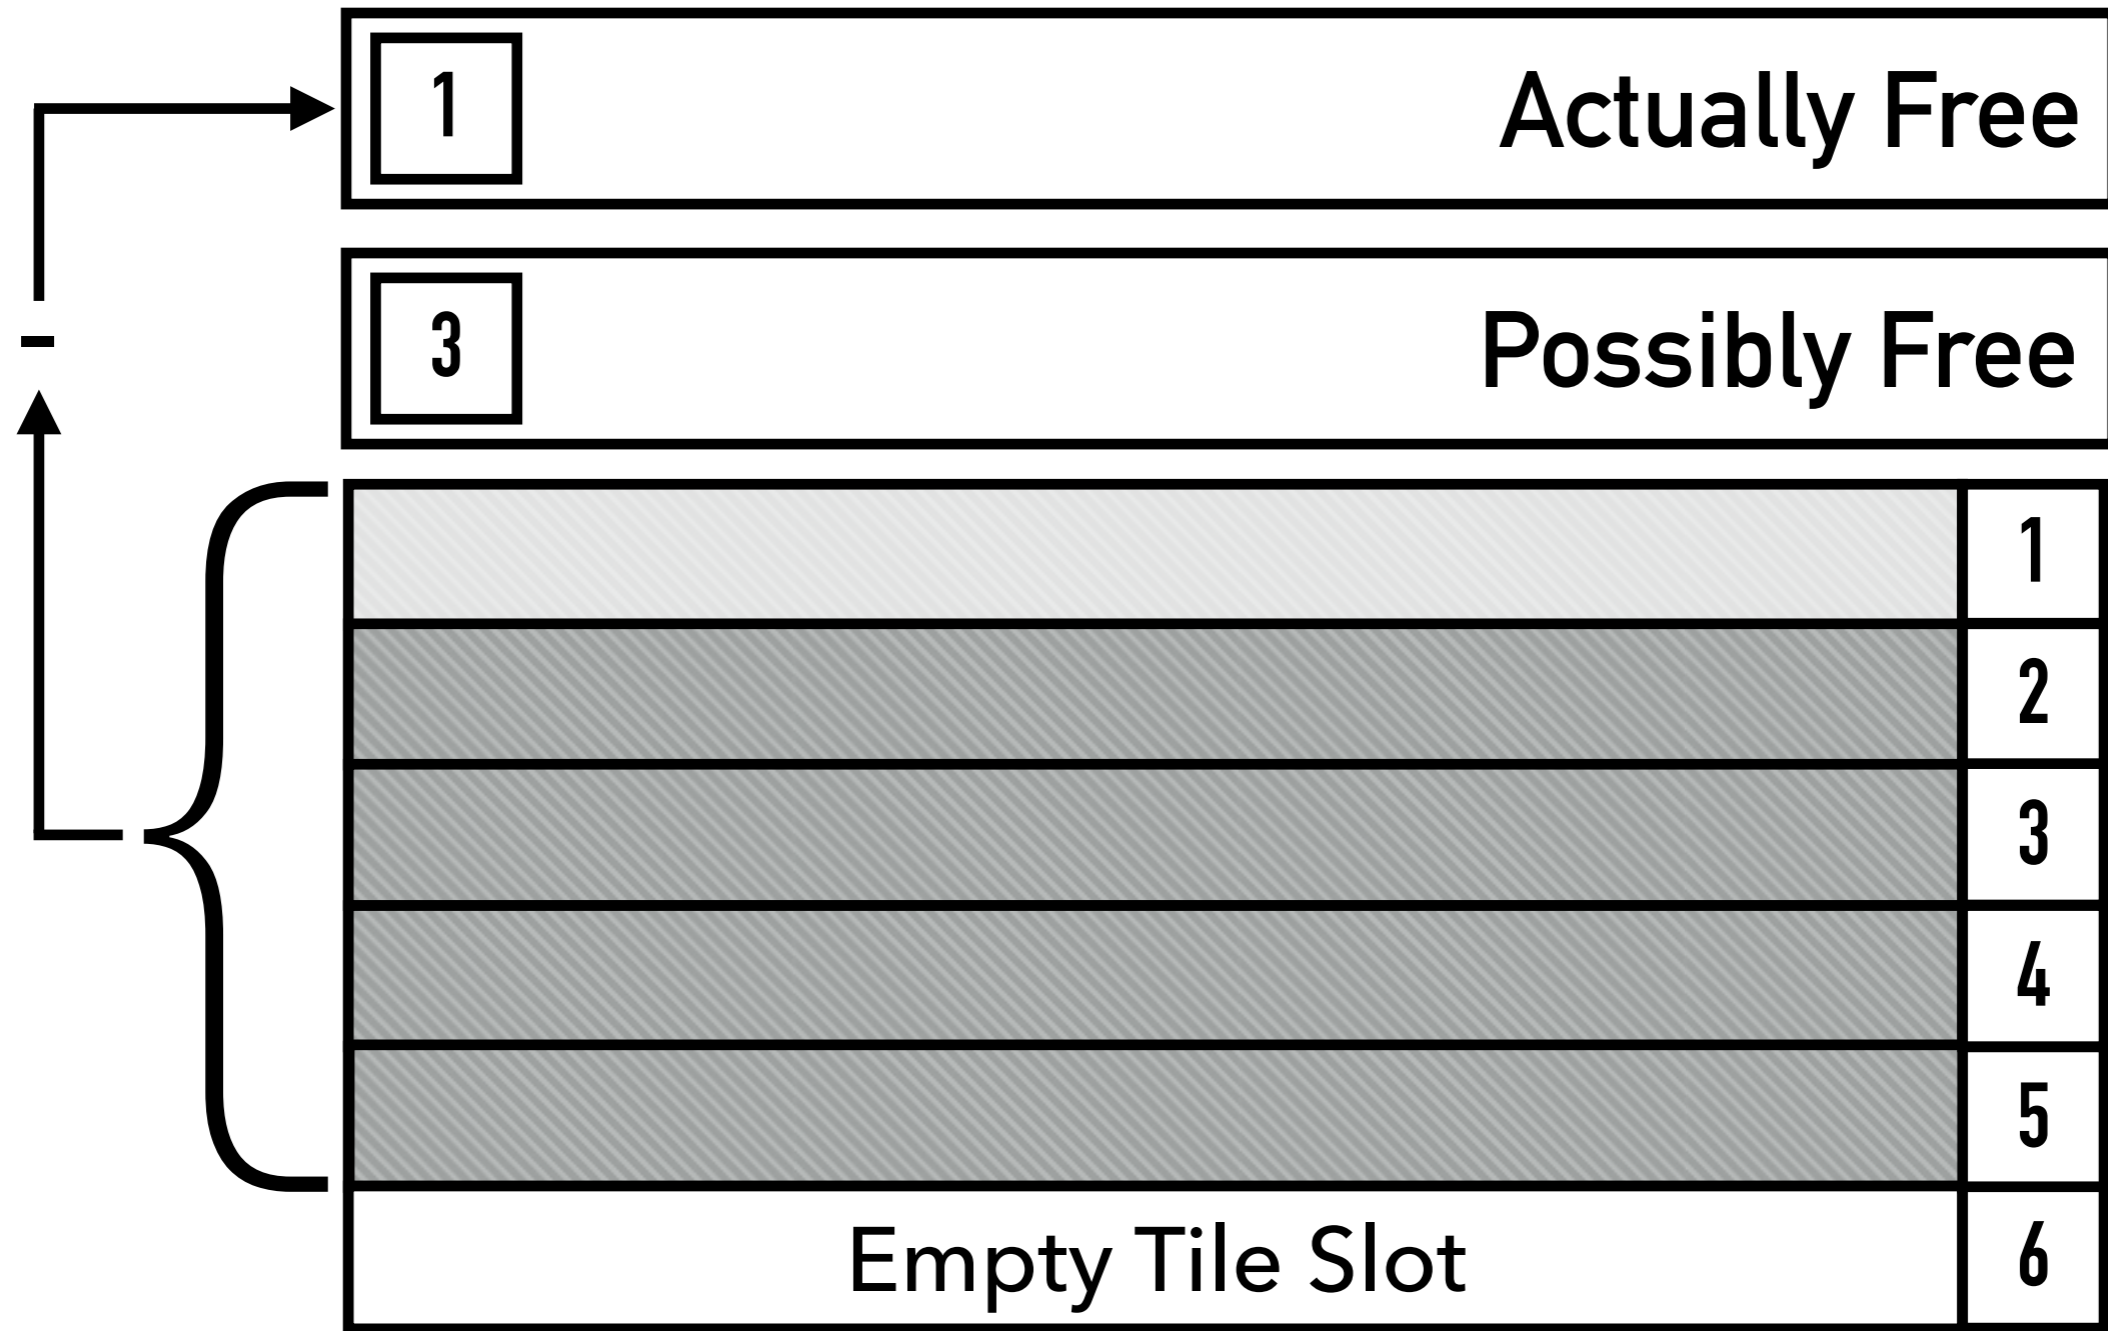
# WITH GC



**Tile Group**



# MEMORY FOOTPRINT



Tile Group

### GC MODES

- ▶ Off (default)
- ▶ Vacuum
- ▶ Naïve co-operative
- ▶ Epoch based co-operative

### WORKLOAD

- ▶ YCSB - 80% updates, 20% reads, 10 terminals, 100000 tuples and 100 sec runtime
- ▶ YCSB - 50% updates, 50% inserts, 10 terminals, 100000 tuples and 100 sec runtime

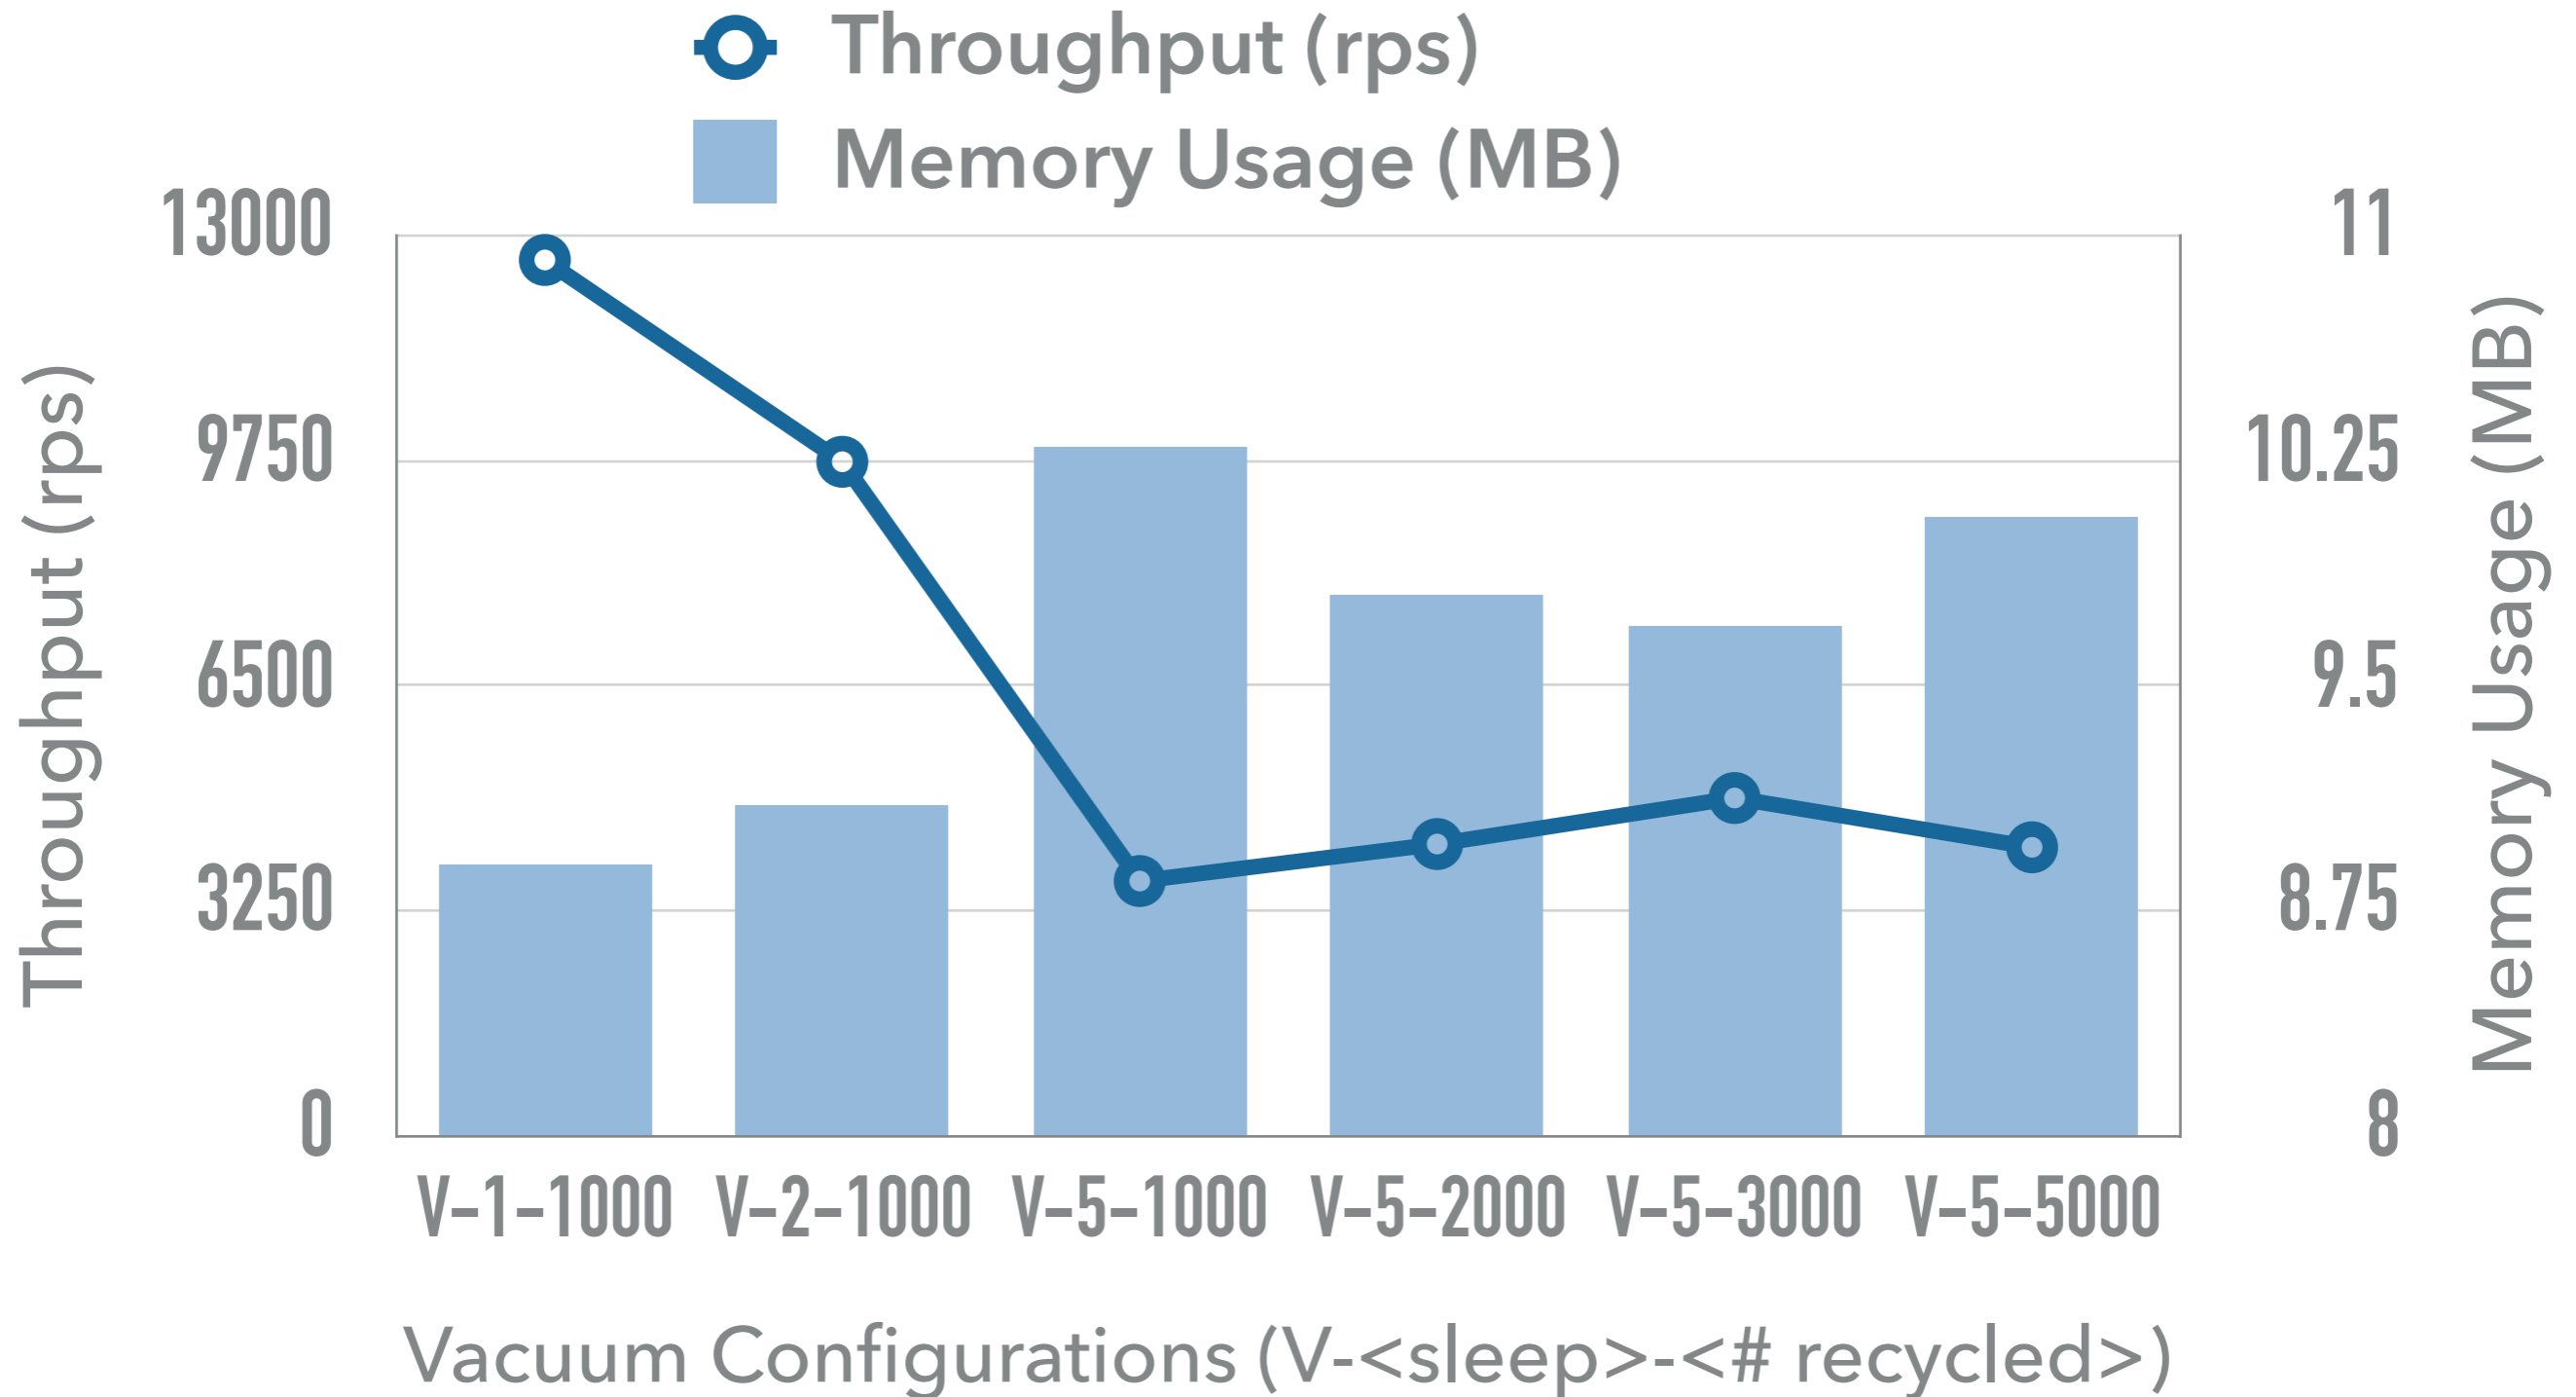
### VACUUM

- ▶ Separate thread started as Peloton bootstraps
- ▶ Periodically moves elements from `possibly_free_list` to `actually_free_list`



### VACUUM TRADEOFFS

- ▶ Vacuum thread period vs GC efficiency
- ▶ # Recycled per vacuum invocation vs GC efficiency

# VACUUM TRADEOFFS MEASURED





## NAÏVE CO-OPERATIVE (COOPERATIVE)







# NAÏVE CO-OPERATIVE (COOPERATIVE)



**PERFORM GC**



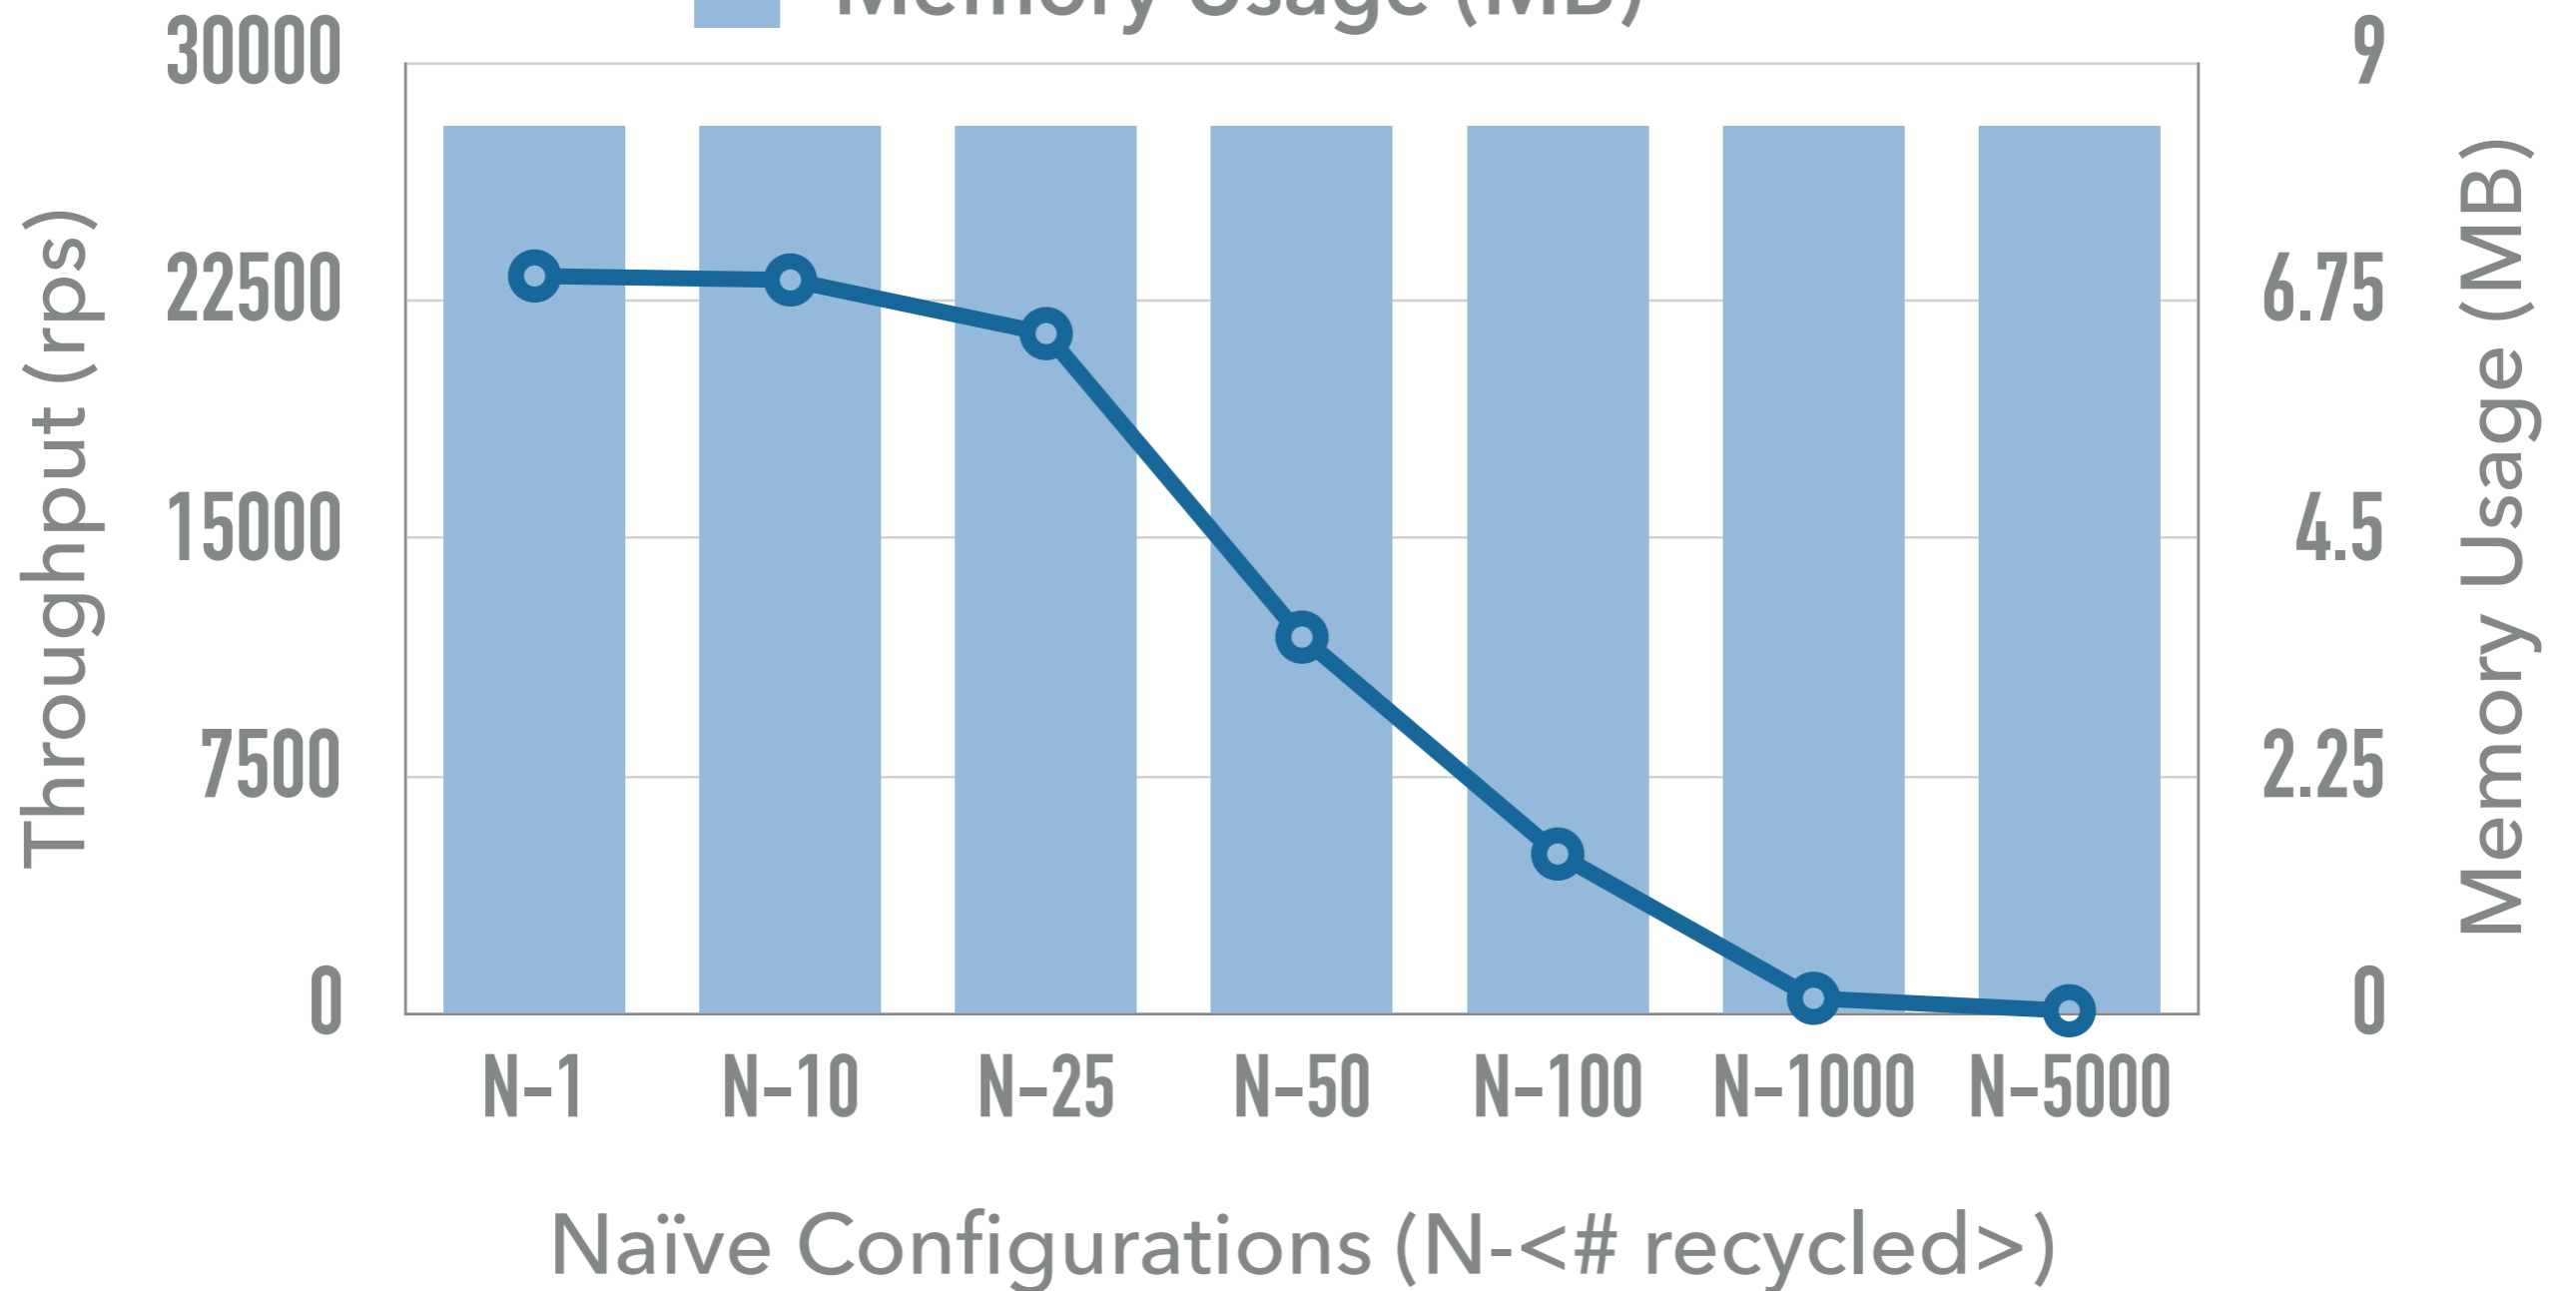
# NAÏVE CO-OPERATIVE (COOPERATIVE)





**PERFORM GC**

# COOPERATIVE TRADEOFFS MEASURED



○ Throughput (rps)  
■ Memory Usage (MB)



## EPOCH CO-OPERATIVE (EPOCH)




# EPOCH CO-OPERATIVE (EPOCH)







# EPOCH CO-OPERATIVE (EPOCH)




# EPOCH CO-OPERATIVE (EPOCH)



# EPOCH CO-OPERATIVE (EPOCH)



# EPOCH CO-OPERATIVE (EPOCH)




# EPOCH CO-OPERATIVE (EPOCH)

join




leave  
e  
e  
e






# EPOCH CO-OPERATIVE (EPOCH)

join



leave

e




# EPOCH CO-OPERATIVE (EPOCH)



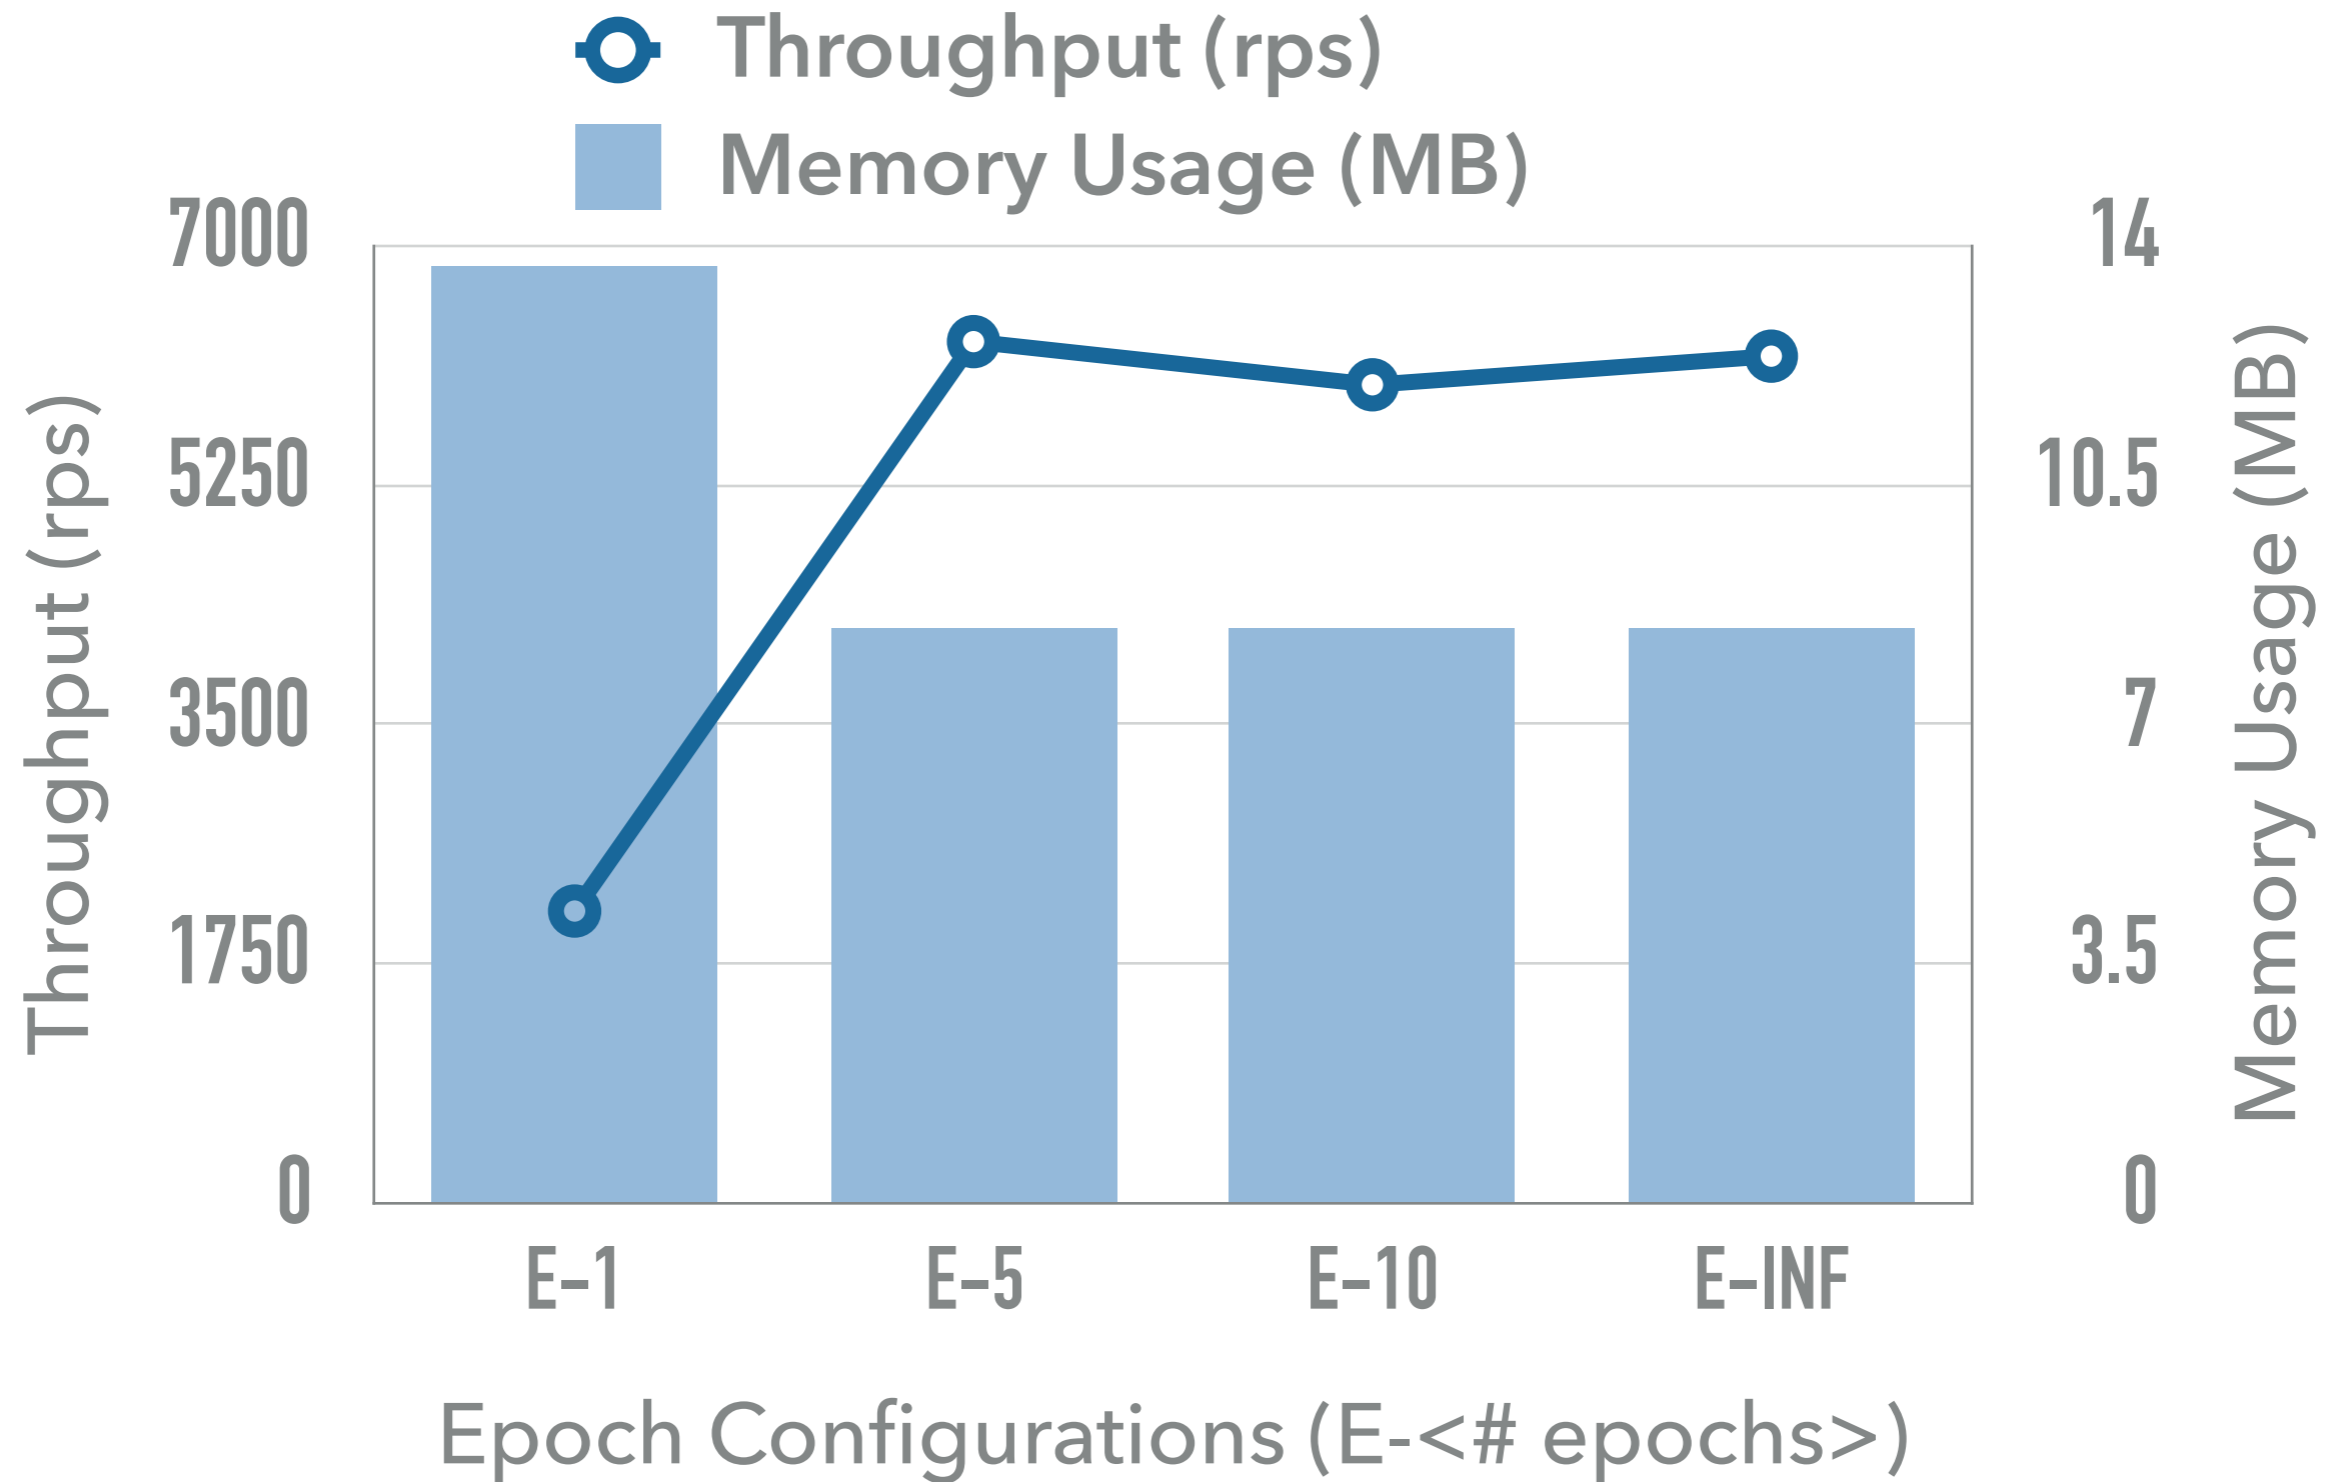
join

leave

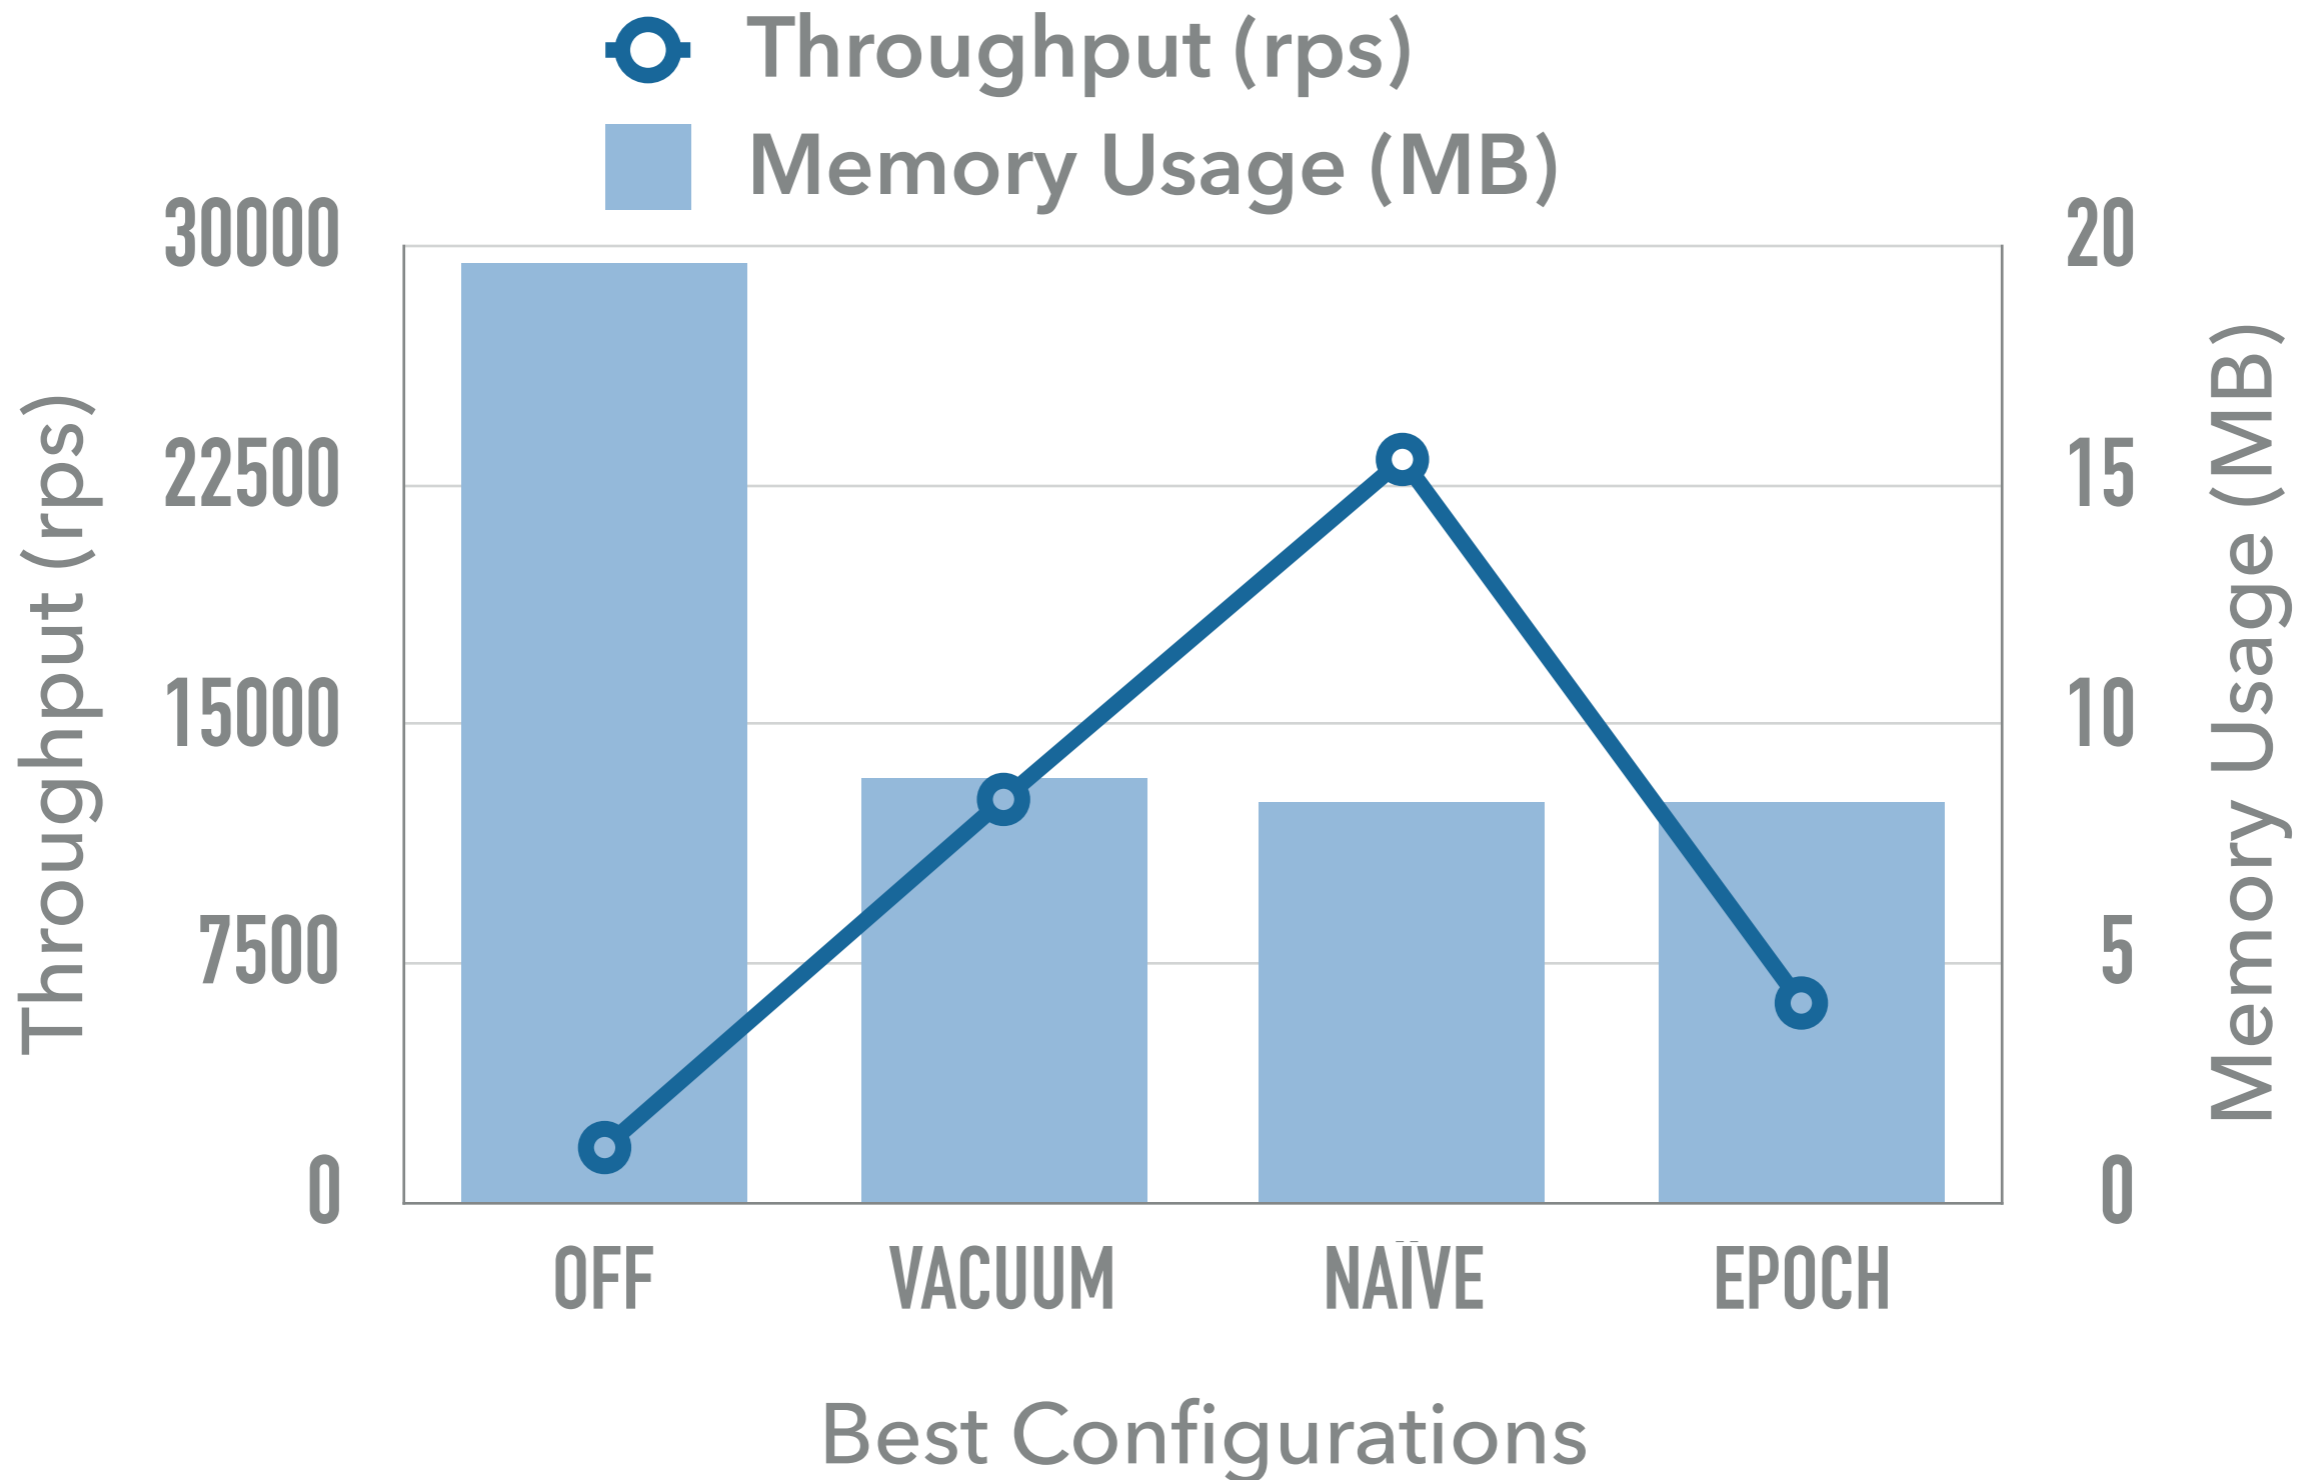
e



# EPOCH TRADEOFFS MEASURED

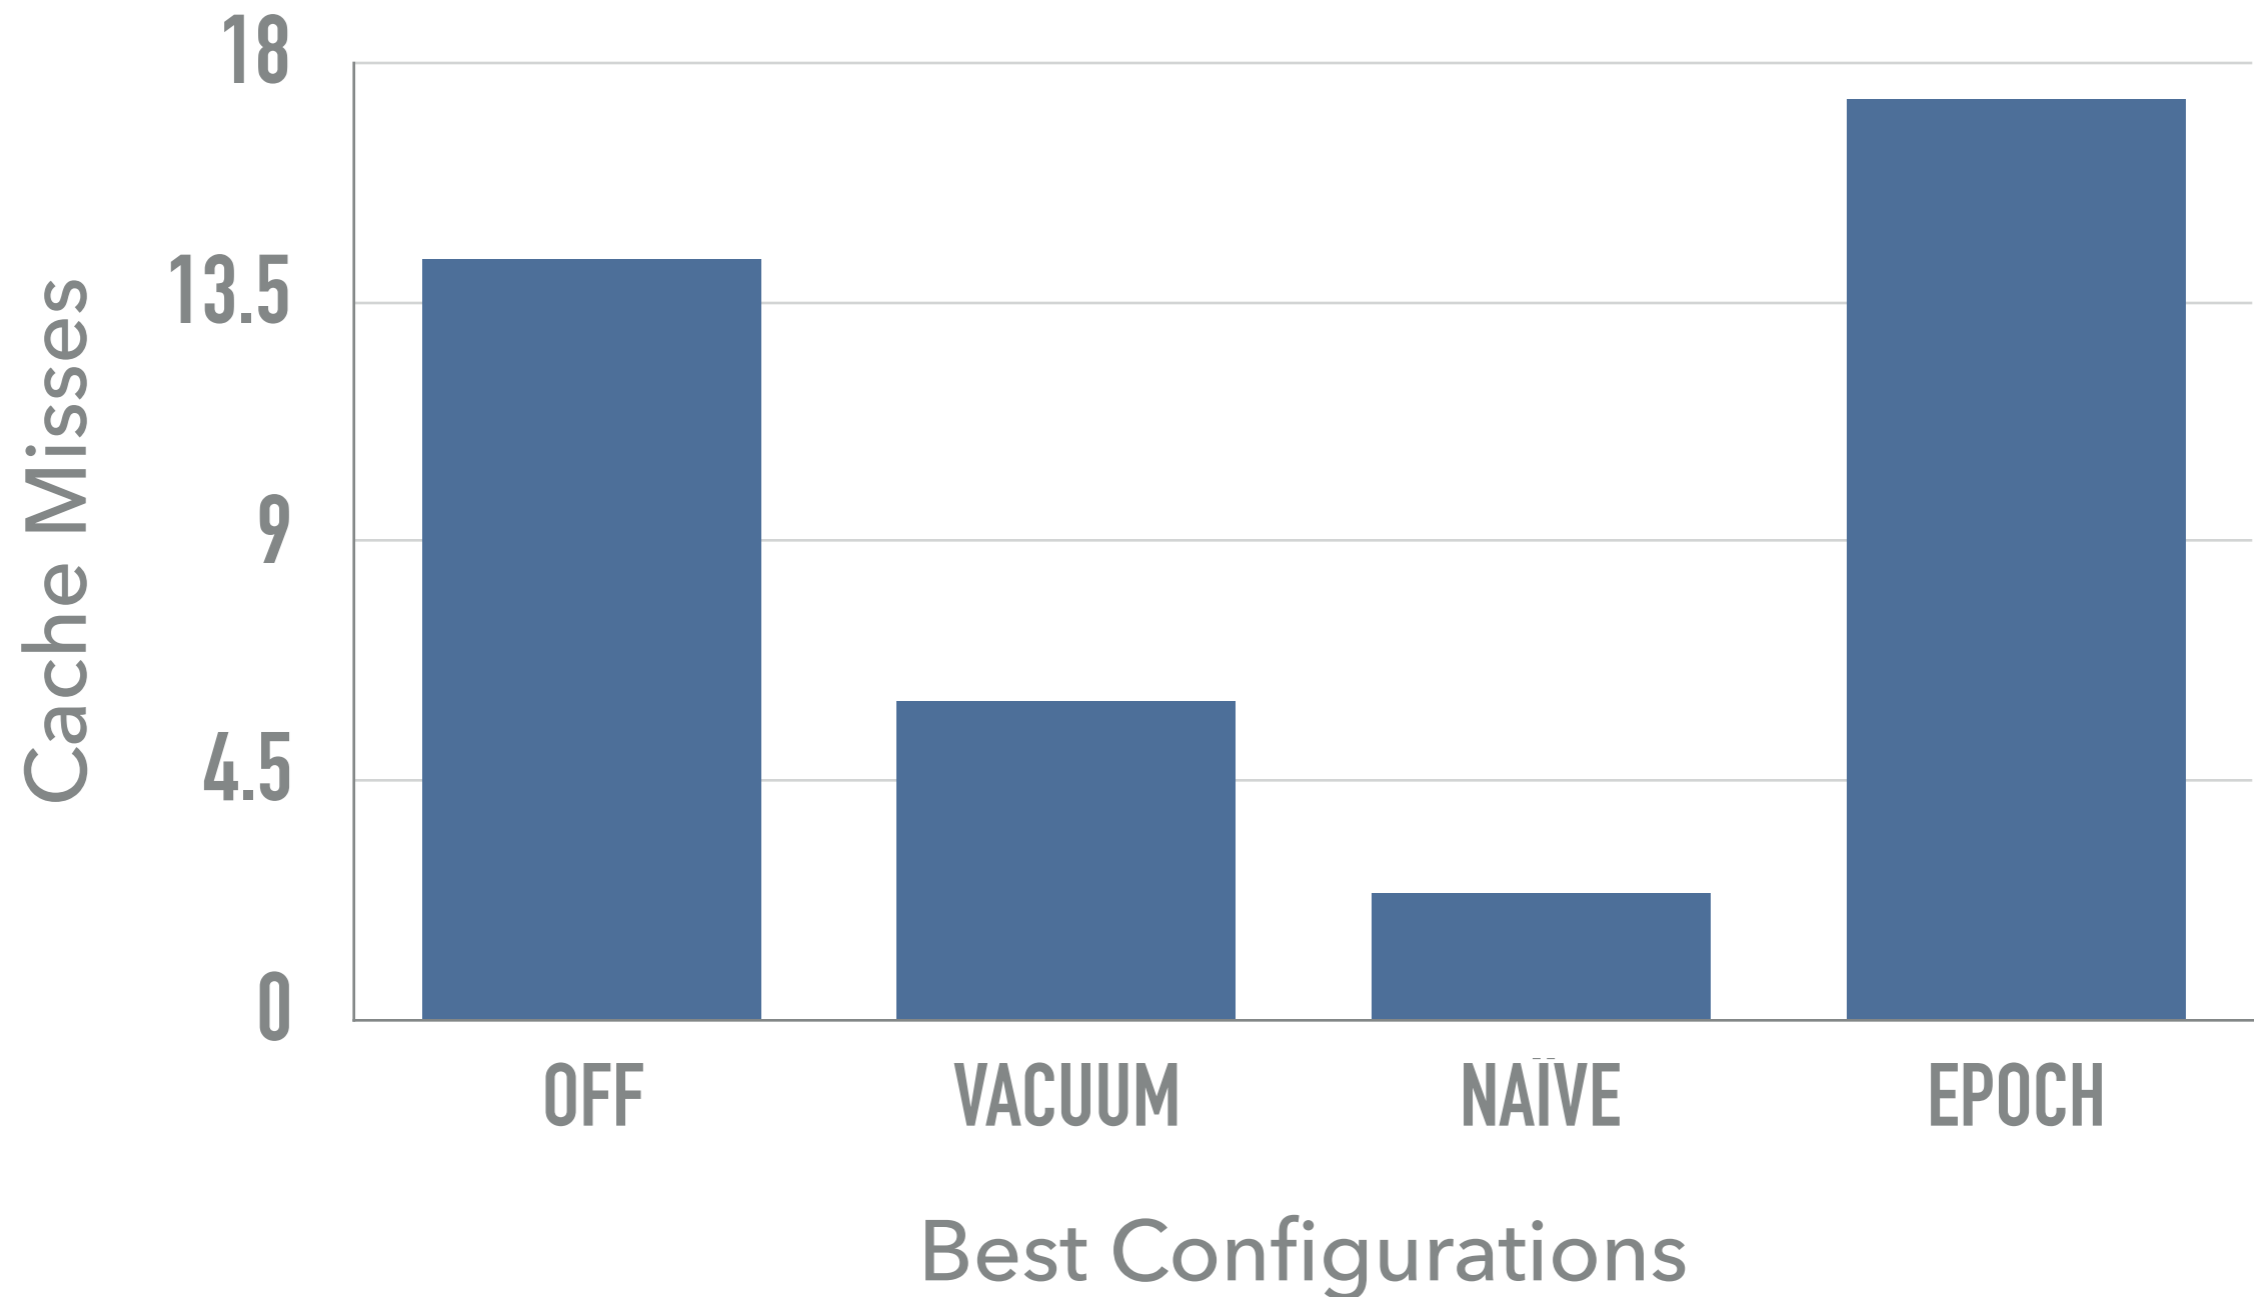


# OVERALL GC COMPARISON



# CACHE MISSES MATTER AND THEY DON'T!

■ Cache Misses (%)





### ENTIRE TABLE TRUNCATIONS

- ▶ Possibly\_free\_list may potentially end up with an absurdly large number of free slots
- ▶ Has to be handled as a special case

### RECYCLING ACROSS TILE GROUPS

- ▶ Current recycling is at table granularity - i.e. across tile groups
- ▶ Different tile group schemas in the same table may become problematic
- ▶ Tradeoff: recycling granularity vs # tuples recycled



**QUESTIONS?**

**Thank You . . .**