

1 Introduction

The buffer manager is responsible for managing pages in memory and receives page requests from the file and index manager. When pages are evicted from memory or new pages are read in to memory, the buffer manager communicates with the disk space manager to perform the required disk operations.

2 Buffer Pool

Memory is converted into a buffer pool by partitioning the space into frames that pages can be placed in. A buffer frame can hold the same amount of data as a page can (so a page fits perfectly into a frame). To efficiently track frames, the buffer manager allocates additional space in memory for a metadata table.

Frame ID	Page ID	Dirty Bit	Pin Count
0	5	1	3
1	3	0	1
2	10	1	0
3			

The table tracks 4 pieces of information:

1. **Frame ID** that is uniquely associated with a memory address
2. **Page ID** for determining which page a frame currently contains
3. **Dirty Bit** for verifying whether or not a page has been modified
4. **Pin Count** for tracking the number of requestors currently using a page

3 Handling Page Requests

When pages are requested from the buffer manager and the page already exists within memory, the page's pin count is incremented and the page's memory address is returned.

If the page does not exist in the buffer pool and there is still space, the next empty frame is found and the page is read into that frame. The page's pin count is set to 1 and the page's memory address is returned. In the case where the page does not exist and there are no empty frames left, a replacement policy must be used to determine which page to evict. The choice of replacement policy is heavily dependent on page access patterns and the optimal policy is chosen by counting page hits. A **page hit** is when a requested page can be found in memory without having to go to disk. Each **page miss** incurs an additional IO cost, so a good eviction policy is critical for performance. Additionally, if the evicted page has the dirty bit set, the page is written to disk to ensure that updates are persisted. The dirty bit is set to 1 if and when a page is written to with updates in memory. The dirty bit is set to 0 once the page is written back to disk.

Once the requestor completes its workload, it is responsible for telling the buffer manager to decrement the pin count associated with pages that it previously used.

4 LRU Replacement

A commonly used replacement policy is LRU (Least Recently Used). When new pages need to be read into a full buffer pool, the least recently used unpinned page (pin count = 0) is evicted. To track page usage, a last used column is added to the metadata table and measures the latest time at which a page's pin count is decremented.

Frame ID	Page ID	Dirty Bit	Pin Count	Last Used
0	5	1	3	20
1	3	0	1	32
2	10	1	0	40
3	6	0	0	25
4	1	0	1	15

Implementing LRU normally can be costly. The Clock policy provides an alternative implementation that efficiently approximates LRU using a ref bit (recently referenced) column in the metadata table and a clock hand variable to track the current frame in consideration.

Frame ID	Page ID	Dirty Bit	Pin Count	Ref Bit
0	5	1	3	1
1	3	0	1	1
2	10	1	0	1
3	6	0	0	1
4	1	0	1	0

Clock Hand
2

The Clock policy algorithm treats the metadata table as a circular list of frames. It sets the clock hand to the first unpinned frame upon start and sets the ref bit on each page's corresponding row to 1 when it is initially read into a frame. The policy works as follows when trying to evict:

- iterate through frames within the table, skipping pinned pages and wrapping around to frame 0 upon reaching the end, until the first unpinned frame with ref bit = 0 is found
- during each iteration, if the current frame's ref bit = 1, set the ref bit to 0 and move the clock hand to the next frame
- upon reaching a frame with ref bit = 0, evict the existing page (and write it to disk if the dirty bit is set; then set the dirty bit to 0), read in the new page, set the frame's ref bit to 1, and move the clock hand to the next frame


If accessing a page currently in the buffer pool, the clock policy sets the page's ref bit to 1 **without moving the clock hand**.

4.1 Sequential Scanning Performance - LRU

LRU performs well overall but performance suffers when a set of pages S , where $|S| > \text{buffer pool size}$, are accessed multiple times repeatedly.

To highlight this point, consider a 3 frame buffer pool using LRU and having the access pattern:

ABCDABCDABCDABCD



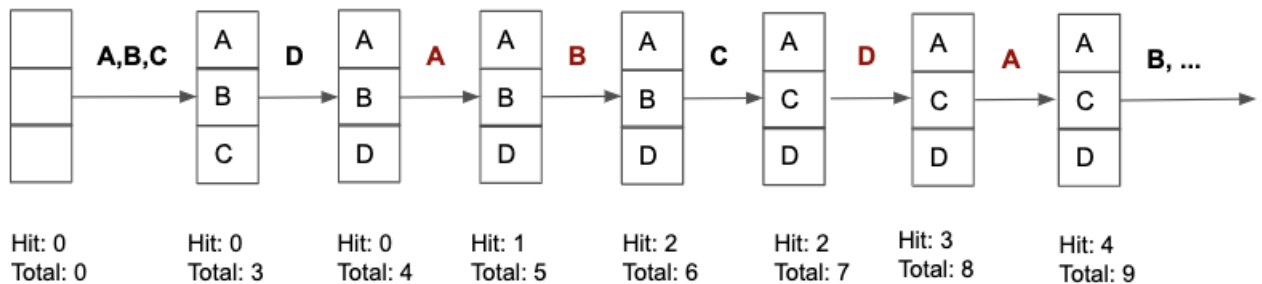
5 MRU Replacement

Another commonly used replacement policy is MRU (Most Recently Used). Instead of evicting the least recently used unpinned page, evict the most recently used unpinned page measured by when the page's pin count was last decremented.

5.1 Sequential Scanning Performance - MRU

At first it might seem counter-intuitive to use this policy but consider the scenario where a 3 frame buffer pool using MRU has the access pattern:

ABCDABCDABCDABCD



Clearly, MRU far outperforms LRU in terms of page hit rate whenever a sequential flooding access pattern occurs.