

Hashing

We will continue from where we left off in the last class before the university closure. In particular, we will look at one more static hashing technique - Cuckoo Hashing and three dynamic hashing techniques of different complexity: chained, extendible and linear hashing (not to be confused with linear probe hashing).

Reading Material: Chapter 11 of Ramakrishnan and Gerhke

As you go through the material, answer all the relevant assessment questions. The deadline for submission of this assessment is Tuesday, March 24 @ 1:00 pm. You can update your answers as often as you want. It is not a timed assessment.

Assessment: There are **3** assessments in GradeScope linked to this lesson: Cuckoo Hashing, Extendible Hashing and Linear Hashing.

Cuckoo Hashing

Watch the video lecture from 43:02 to 47:19

<https://www.youtube.com/embed/j5X1bBPuIAk?start=2582&end=2839>

In case you were wondering, the unique name comes from how the hashing scheme mimics the cuckoo chick, which pushes out unhatched eggs or even other chicks from the nest to make room for itself. Here is how this looks:



Now answer the Cuckoo Hashing Assessment on GradeScope.

Dynamic Hashing Schemes

Why we need dynamic hashing?

Watch the video from 47:22 to 48:48

(<https://www.youtube.com/embed/j5X1bBPuIAk?start=2842&end=2928>)

Chained Hashing

Watch the video lecture from 48:50 to 50:18

(<https://www.youtube.com/embed/j5X1bBPuIAk?start=2930&end=3018>)

Extendible Hashing

Watch the video lecture from 50:38 to 1:00:08

(<https://www.youtube.com/embed/j5X1bBPuIAk?start=3038&end=3608>)

Now answer the Extendible Hashing Assessment on Grade Scope.

Linear Hashing

Reading Material: Read the following 4-pager to understand how linear hashing works.

Zhang D., Manolopoulos Y., Theodoridis Y., Tsotras V.J. (2009) Linear Hashing. In: LIU L., ÖZSU M.T. (eds) Encyclopedia of Database Systems. Springer, Boston, MA
(https://doi.org/10.1007/978-0-387-39940-9_742)

Congrats on completing this unit! If you have questions/comments regarding this unit please write a Piazza post.

Buffer Pool Manager

The buffer pool is an in-memory cache of pages read from disk. The DBMS often just mallocs a large memory cache for the buffer pool on startup that is of fixed size. You can often configure the size of this space. This space is usually organized as an array of frames. On a page request, an exact copy is placed into one of these frames from disk. The Buffer Pool Manager manages the placement and replacement of pages in the buffer.

Reading Material: Chapter 9.4 of Ramakrishnan and Gerhke

As you go through the material, answer all the relevant assessment questions. The deadline for submission of this assessment is Tuesday, March 24 @ 1:00 pm. You can update your answers as often as you want. It is not a timed assessment.

Assessment: Complete the buffer pool manager assessment in GradeScope.

[Buffer Pool Manager Assessment](#)

The Job of a Buffer Manager

(<https://www.youtube.com/embed/qOO0HGVTtoYA?feature=oembed>)

Moving pages in and out of the Buffer Pool

(<https://www.youtube.com/embed/6gHKFYgsWLQ?feature=oembed>)

Dealing with Page Updates: Dirty Pages

(<https://www.youtube.com/embed/Ddltpilw47E?feature=oembed>)

Buffer Page Meta Data

(<https://www.youtube.com/embed/e - SCTcl70?feature=oembed>)

To summarize the meta-data maintained by the buffer pool:

- **Page Table:** In-memory hash table that keeps track of pages that are currently in memory. It maps page ids to frame locations in the buffer pool.
- **Dirty-flag:** When a page is modified by the DBMS threads, this flag is set. A dirty page must be written back to disk.
- **Pin Counter:** This tracks the number of threads that are currently accessing that page (either reading or modifying it). A thread has to increment the counter before they access the page. If a page's count is greater than zero, then we cannot evict that page from memory.

Now you can complete questions 1 and 2 of the Buffer Pool Management Assessment on Gradescope.

Making Room: Page Replacement Overview

A replacement policy is an algorithm that the DBMS implements that makes a decision on which pages to evict from buffer pool when it needs space.

There are several criteria that guide the choice of the policy to use including:

- Correctness - You don't want to evict pinned pages!
- Accuracy - You shouldn't evict pages that you will likely read again soon.
- Speed - Deciding which page to evict should be quick for interactive query performance: you don't want to use an inefficient strategy to decide which pages to replace.
- Meta-data overhead - Some policies like LRU keep a timestamp that tracks when every page was accessed. This overhead is one of the reasons behind its approximation: the clock. You want to strike a balance between how much information you keep and its ability to improve the replacement policy. You are also constrained by how much space you have to store this metadata.

<https://www.youtube.com/embed/NHp0nGW3aAg?feature=oembed>

We will now examine two main policies.

LRU

https://www.youtube.com/embed/Uj_UOxZwBNQ?feature=oembed

LRU Approximation: The Clock

<https://www.youtube.com/embed/gjLywr7GbLU?feature=oembed>

Problematic Access Patterns: LRU & Sequential Flooding

<https://www.youtube.com/embed/cESKTI12Ulg?feature=oembed>

MRU

https://www.youtube.com/embed/bwTrNL_xKdw?feature=oembed

Now you can complete questions 3, 4, 5, and 6 of the Buffer Pool Management Assessment on Gradescope.

Prefetching & Why not just rely on the OS cache?

<https://www.youtube.com/embed/QNUKYIZybUI?feature=oembed>

Now you can complete question 7 of the Buffer Pool Management Assessment on Gradescope.

For the fun of it

If you would like to understand how caching (buffer pools) work in Postgres and how they interact with OS file caches, read the following blog post: <https://madosudanan.com/blog/understanding-postgres-caching-in-depth/>

It also describes a number of Postgres commands that can help you debug/improve your query performance if you suspect poor buffering to be the cause of the problem!

Congrats on completing this unit! If you have questions/comments regarding this unit please write a Piazza post.