

1 Buffer Management

(a) Fill in the following tables for the given buffer replacement policies. You have 4 buffer pages, with the access pattern **A B C D A F A D G D G E D F**

Least Recently Used (LRU)

A			*		*							F	Hit Rate $\frac{6}{14}$
	B			F							E		
		C					G		*				
			D			*		*			*		

Most Recently Used (MRU)

A			*	F	A								Hit Rate $\frac{2}{14}$
	B												
		C											
			D			*	G	D	G	E	D	F	

Clock ("second chance LRU")

A			*	F								D	Hit Rate $\frac{4}{14}$
	B				A							F	
		C					G		*				
			D			*		*		E			

Red frame outline is where the clock hand points at the end of the access
 Yellow frame outline is where the clock hand considered ejecting a page, but the bit was one
 Red box means the second chance bit is zero
 Green box means the second chance bit is one
 * means a hit

(b) Fill in the following tables for the given buffer replacement policies. You have 4 buffer pages, with the access pattern **A, B, C, D, A, F (remains pinned), D, G, D, unpin F, G, E, D, F**. Remember that unpinning does not contribute to the hit count!

Least Recently Used w/ Pinning

A				*								E			
	B				F	x	x	x	x	x				*	
		C						G			*				
			D				*		*				*		
$\frac{6}{13}$															

Most Recently Used w/ Pinning

A				*	F	x	x	x	x	x	G	E			
	B														
		C													
			D				*	G	D				*	F	
$\frac{3}{13}$															

* = hit
X = pinned

Remember that unpinning doesn't contribute to the hit count

(c) Is MRU ever better than LRU?

Yes; MRU prevents sequential flooding during sequential scans

(d) Why would we use a clock replacement policy over LRU?

Efficiency (approximation of LRU; don't need to maintain entire ordering)

(e) Why would it be useful for a database management system to implement its own buffer replacement policy? Why shouldn't we just rely on the operating system?

The database management system knows its data access patterns, which allows it to optimize its buffer replacement policy for each case

2 Relational Algebra

Consider the schema:

- Songs(SONG_ID, song_name, album_id, weeks_in_top_40)
- Artists(ARTIST_ID, artist_name, first_yr_active)
- Albums(ALBUM_ID, album_name, artist_id, yr_released, genre)

Write relational algebra expressions for the following queries:

(a) Find the names of the artists who have albums with a genre of either 'pop' or 'rock'.

π artists.artist_name (σ albums.genre = 'pop' \vee albums.genre = 'rock' (Artists \bowtie Albums))

(b) Find the names of the artists who have albums of genre 'pop' and 'rock'.

Solution 1:

π artists.artist_name ((σ albums.genre = 'pop' Albums) \bowtie Artists) \cap
 π artists.artist_name ((σ albums.genre = 'rock' Albums) \bowtie Artists)

Solution 2:

$\pi_{\text{artist_name}}$ (
Artists \bowtie (
 $\pi_{\text{artist_id}}$ ($\sigma_{\text{genre} = \text{'pop'}}$ (Albums)) \cap
 $\pi_{\text{artist_id}}$ ($\sigma_{\text{genre} = \text{'rock'}}$ (Albums))
)
)

(c) Find the id of the artists who have albums of genre 'pop' or have spent over 10 weeks in the top 40.

$\pi_{\text{artist_id}}(\sigma_{\text{genre} = \text{'pop'}}(\text{Albums})) \cup$

$\pi_{\text{artist_id}}(\sigma_{\text{weeks_in_top_40} > 10}(\text{Albums} \bowtie \text{Songs}))$

(d) Find the names of the artists who do not have any albums.

$\pi \text{ artists.artist_name } (\text{Artists} \bowtie ((\pi \text{ artists.artist_id } \text{Artists}) - (\pi \text{ albums.artist_id } \text{Albums})))$