



THE UNIVERSITY  
*of* EDINBURGH

# Advanced Database Systems

Spring 2026

---

## Q&A Session 1

# ADMINISTRIVIA

Coursework was released last week

Start early – several things you can already begin working on

Ask questions on Piazza

Q&A sessions

Like office hours

You can ask questions about material & provide feedback

Each Q&A session includes a practice worksheet available on Learn

# FILES, PAGES, RECORDS

Tables stored as **logical files** consisting of **pages**, each containing a collection of **records**

**File** (corresponds to a table)

**Page** (many per file)

**Record** (many per page)

The unit of access to physical disk is the page

**1 I/O** = read or write 1 page

# PAGE BASICS

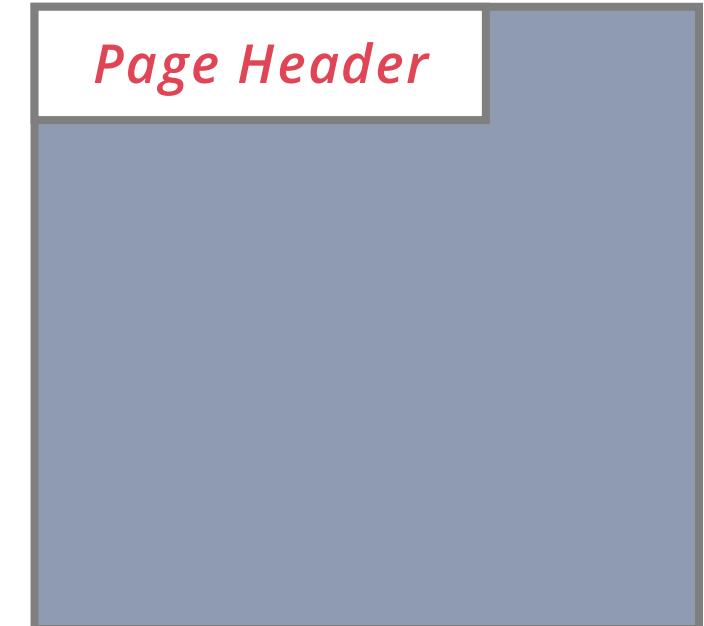
The **page header** keeps track of the records in the page

The page header may contain fields such as:

- Number of records in the page

- Pointer to segment of free space in the page

- Bitmap indicating which parts of the page are in use

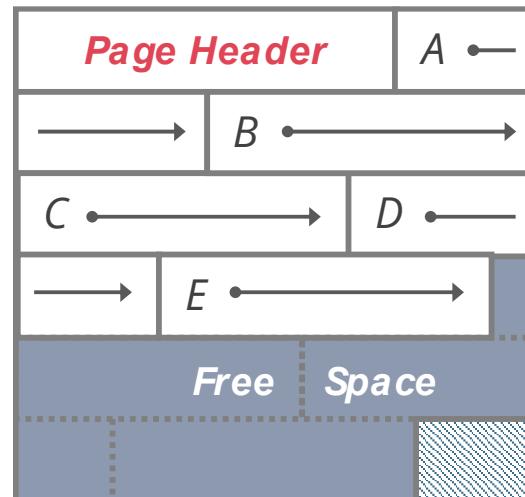


Page

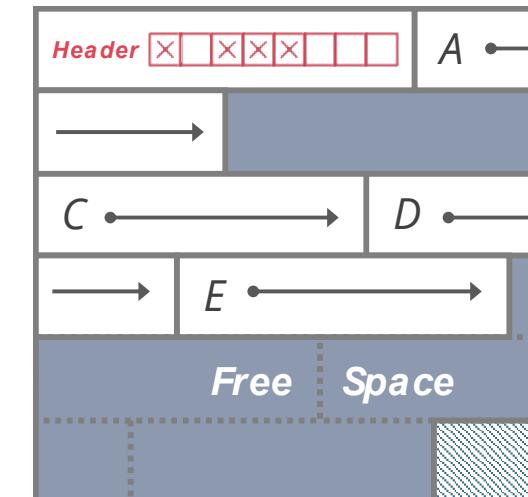
# FIXED-LENGTH RECORDS

**Fixed-length records** = record lengths are fixed and field lengths are consistent

Packed Records: no gaps between records, record ID is location in page



Unpacked Records: allow gaps between records, use a bitmap to keep track of where the gaps are



# VARIABLE-LENGTH RECORDS

**Variable-length records** may not have fixed & consistent field lengths

We can store variable length length records with an array of field offsets:



Each record contains a **record header**

Variable length fields are placed *after* fixed length fields

Record header stores **field offset** (where variable length field ends)

# QUESTION 1

Consider the following relation:

Assume record header stores only pointers (4B) to variable-length fields

```
CREATE TABLE Customer (
    customer_id INTEGER PRIMARY KEY,
    age INTEGER NOT NULL,
    name VARCHAR(10) NOT NULL,
    address VARCHAR(20) NOT NULL
)
```

Record header size = ???

Min record size = ???

Max record size = ???

# QUESTION 1, PART 2

Consider the following relation:

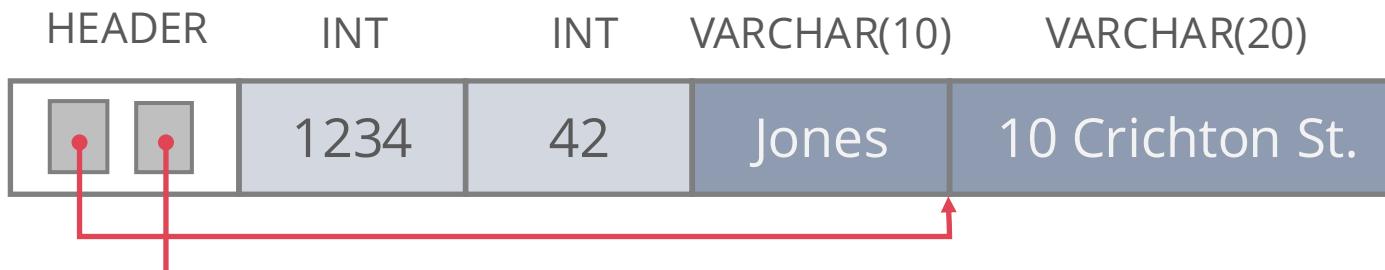
Assume record header stores only pointers (4B) to variable-length fields

```
CREATE TABLE Customer (
    customer_id INTEGER PRIMARY KEY,
    age INTEGER NOT NULL,
    name VARCHAR(10) NOT NULL,
    address VARCHAR(20) NOT NULL
)
```

Record header size = **8**

Min record size = **16**

Max record size = **46**



# SLOTTED PAGES

Most common layout scheme is called **slotted pages**

Slot directory maps “slots” to the records’ starting position offsets

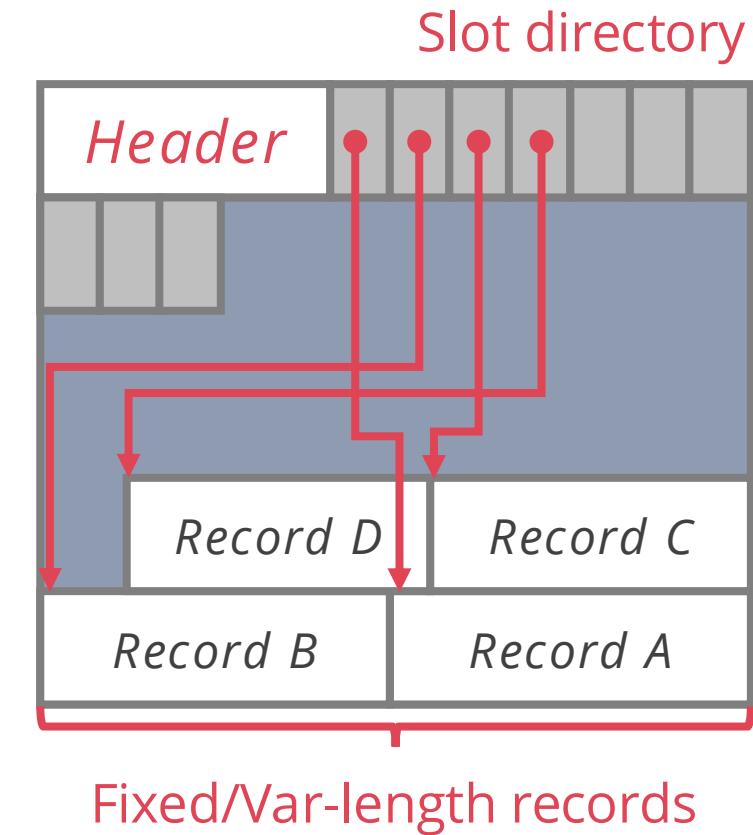
Record ID = (page ID, slot ID)

Header keeps track of:

The number of used slots

The offset of the last slot used

Records stored at the end of page



# QUESTION 2

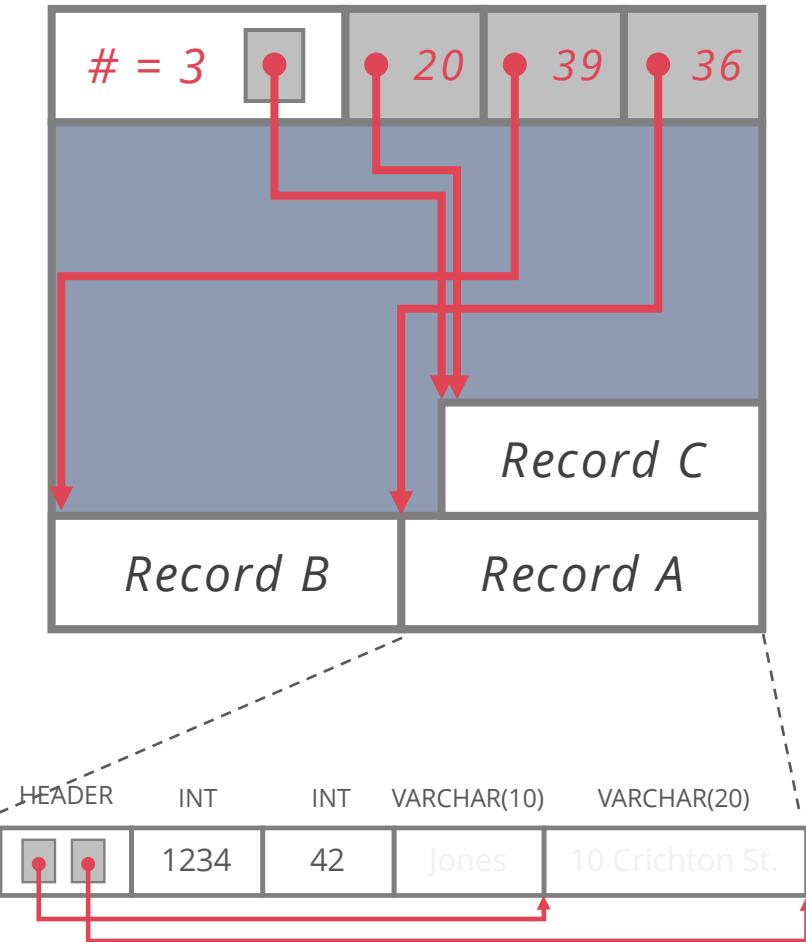
Suppose the Customer relation is stored using a slotted page layout

**Page header** stores the number of records and a pointer to free space

**Directory slot** stores a pointer and length

**Page size** is 8KB

Max number of records = ???



# QUESTION 2, PART 2

Suppose the Customer relation is stored using a slotted page layout

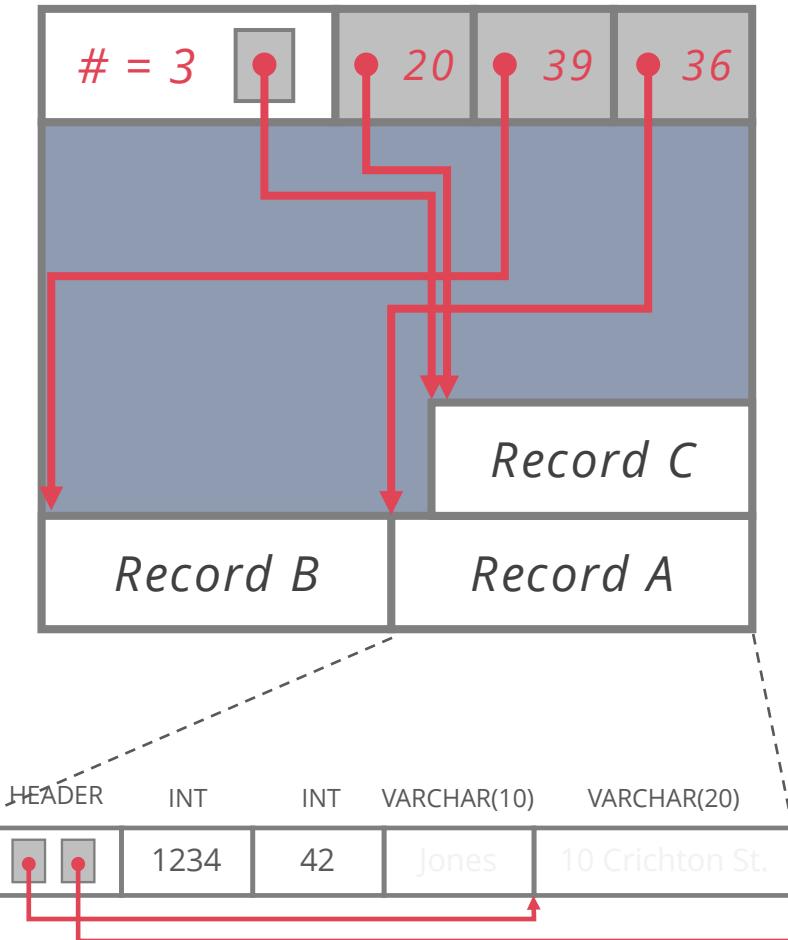
Page header stores the number of records and a pointer to free space (**4B + 4B**)

Directory slot stores a pointer and length (**4B + 4B**)

Page size is 8KB

Max number of records

$$\begin{aligned}
 &= (\text{page size} - \text{header size}) / (\text{min record size} + \text{slot size}) \\
 &= (8192 - 8) / (16 + 8) = \text{341 records}
 \end{aligned}$$



# BUFFER MANAGEMENT

# BUFFER MANAGER

Layer that manages which pages are loaded in memory

Controls when pages are read from & written to disk

When no space in memory, decides what page to **evict**

Decision process is the **page replacement policy**

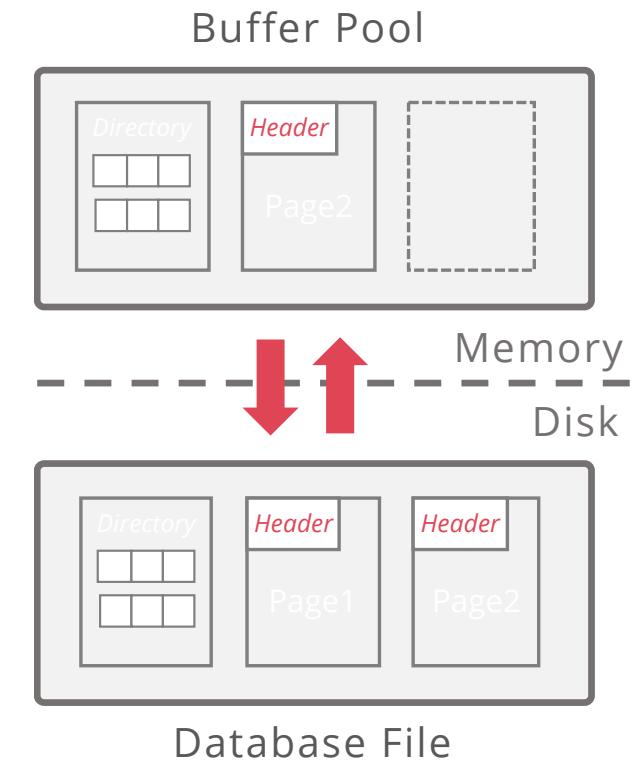
Big impact on I/Os depending on access pattern

Common policies:

LRU (Least Recently Used)

MRU (Most Recently Used)

Clock

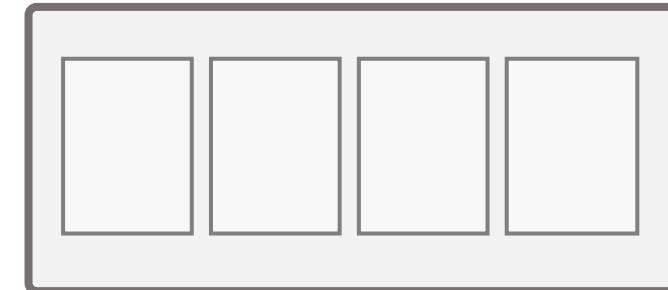


# QUESTION 3

Page access sequence:

A B C D E B A D C A E C

Buffer Pool



Most Recently Used (MRU)

Buffer hits = ???

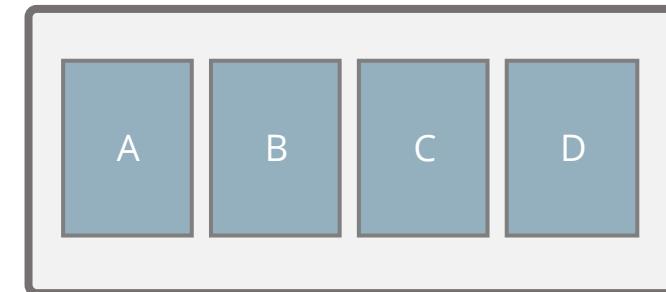
# QUESTION 3 – AFTER 4 REQUESTS

Page access sequence:

A B C D E B A D C A E C



Buffer Pool



Most Recently Used (MRU)

Buffer hits = **0**

A (miss), B (miss), C (miss), D (miss)

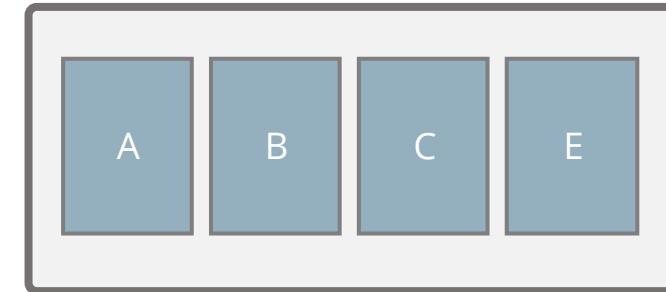
# QUESTION 3 – AFTER 7 REQUESTS

Page access sequence:

A B C D E B A D C A E C



Buffer Pool



Most Recently Used (MRU)

Buffer hits = **2**

A (miss), B (miss), C (miss), D (miss), E (miss, D out), **B (hit)**, **A (hit)**

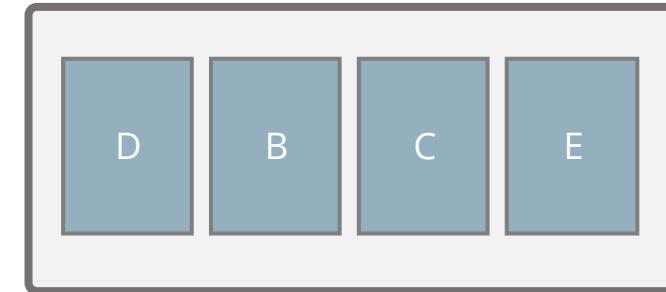
# QUESTION 3 – AFTER 10 REQUESTS

Page access sequence:

A B C D E B A D C A E C



Buffer Pool



Most Recently Used (MRU)

Buffer hits = **3**

A (miss), B (miss), C (miss), D (miss), E (miss, D out), **B (hit)**, **A (hit)**,

D (miss, A out), **C (hit)**

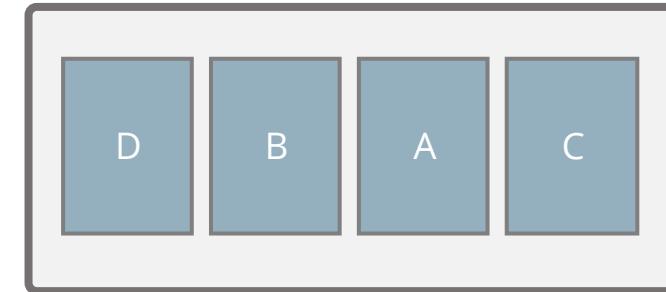
# QUESTION 3 – AFTER 12 REQUESTS

Page access sequence:

A B C D E B A D C A E C



Buffer Pool



Most Recently Used (MRU)

Buffer hits = **4**

A (miss), B (miss), C (miss), D (miss), E (miss, D out), **B (hit)**, **A (hit)**,

D (miss, A out), **C (hit)**, A (miss, C out), **E (hit)**, C (miss, E out)

# CLOCK

Efficient approximation of LRU

Arrange frames in a circle (like numbers on a clock)

Advance **clock hand** around the clock to find pages to evict

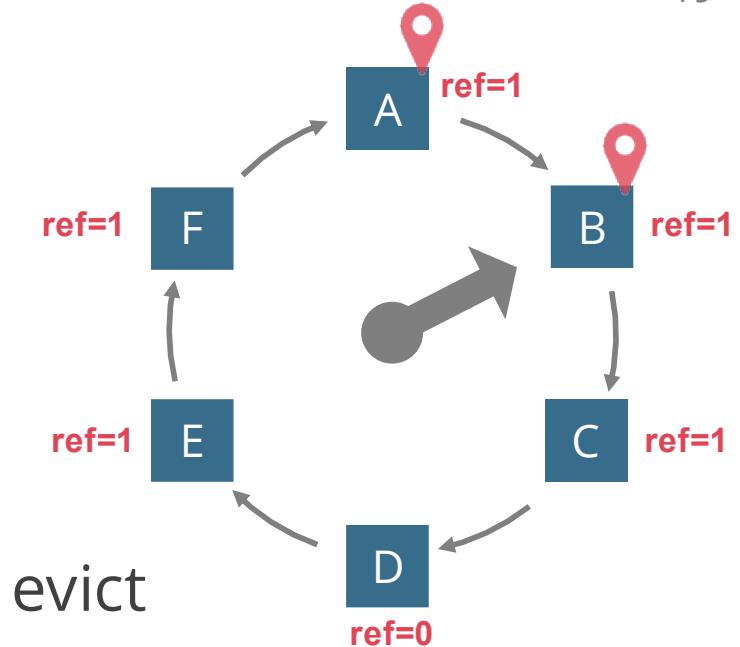
Only do this if you need to evict a page

To make this approximate least recently *used* (rather than least recently *loaded*): add a **reference bit** to each frame

Set to 1 on load/hit, 0 if clock hand passes the frame and the frame is unpinned

Evict unpinned frame if clock hand reaches it and bit = 0

(bit = 0 means less recently used than those with bit = 1)

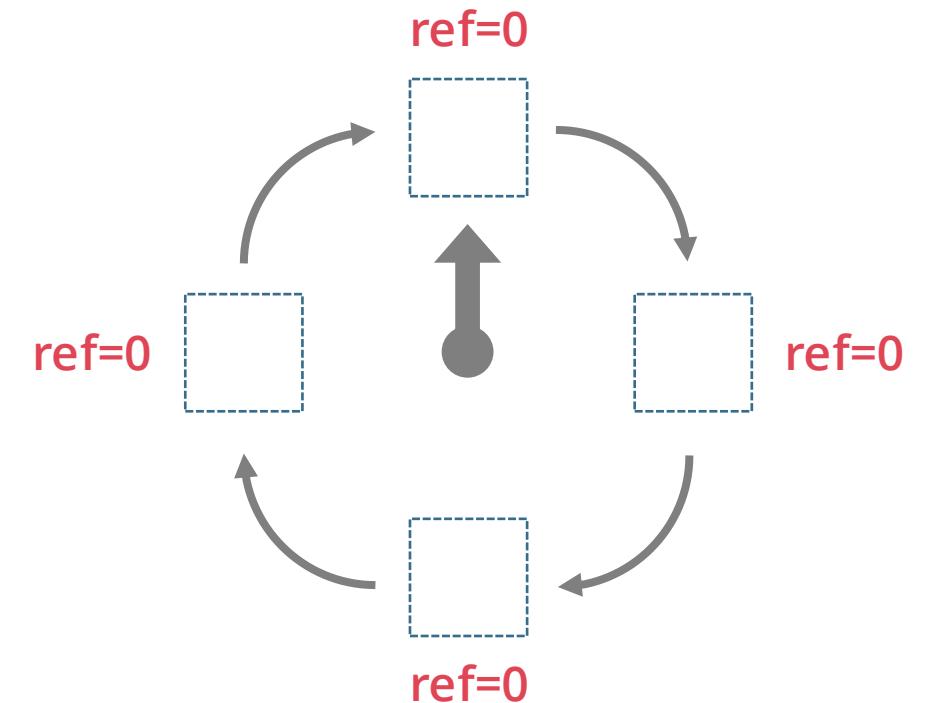


# QUESTION 4

Page access sequence:

A B C D E B A D C A E C

Assume pages are immediately unpinned  
after being pinned



Buffer hits = ???

# QUESTION 4, PART 2

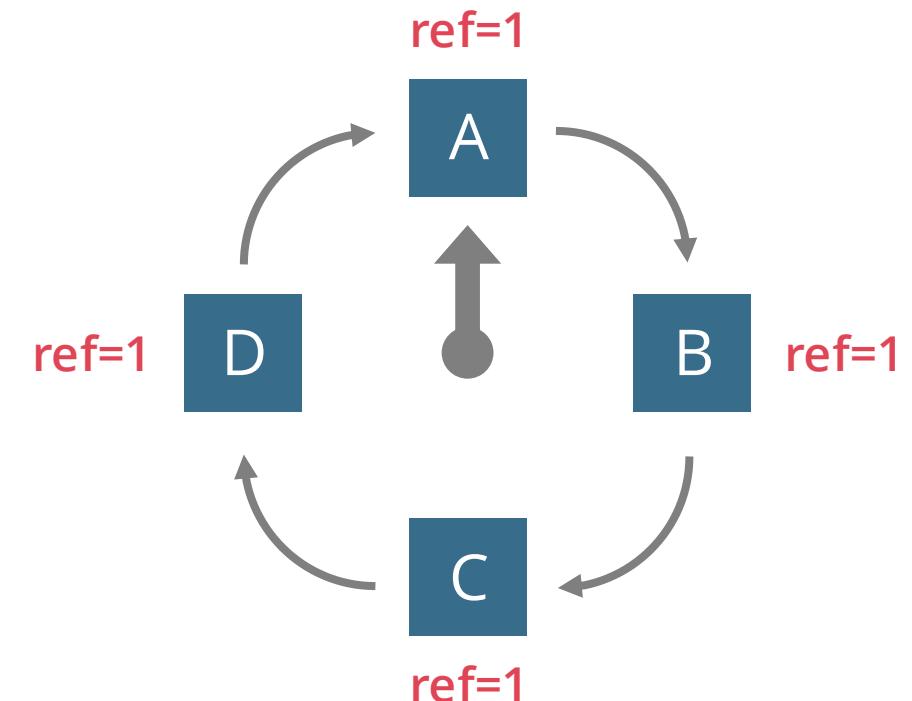
Page access sequence:

A B C D E B A D C A E C



Pages A, B, C, D populate the buffer pool

The clock hand stays still



Buffer hits (so far) = 0

# QUESTION 4, PART 3

Page access sequence:

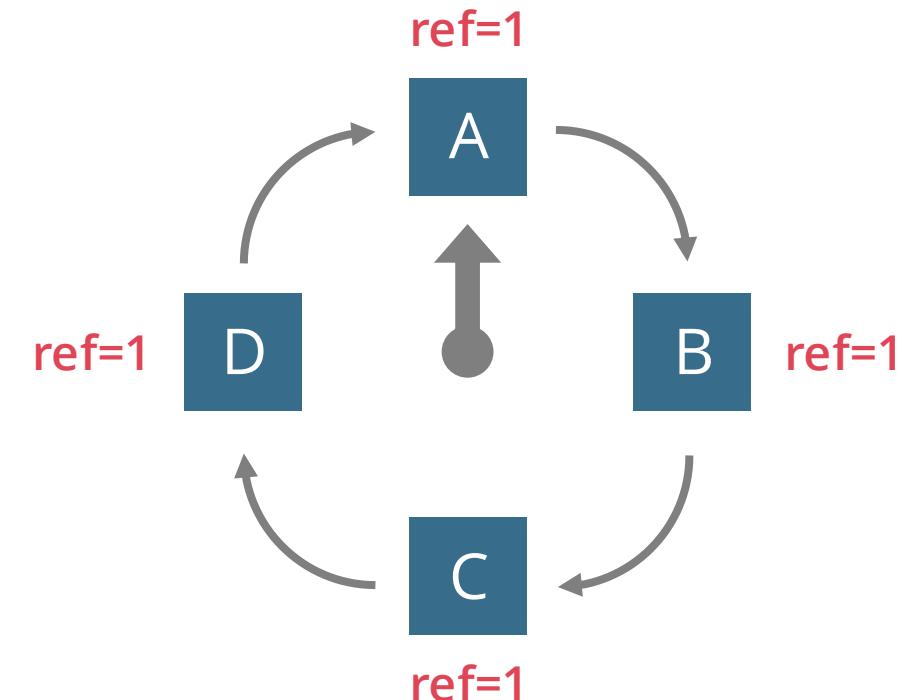
A B C D E B A D C A E C



Page E not present  $\Rightarrow$  buffer miss!

Find first frame with ref = 0

If ref = 1, unset it and move the hand



Buffer hits (so far) = **0**

# QUESTION 4, PART 4

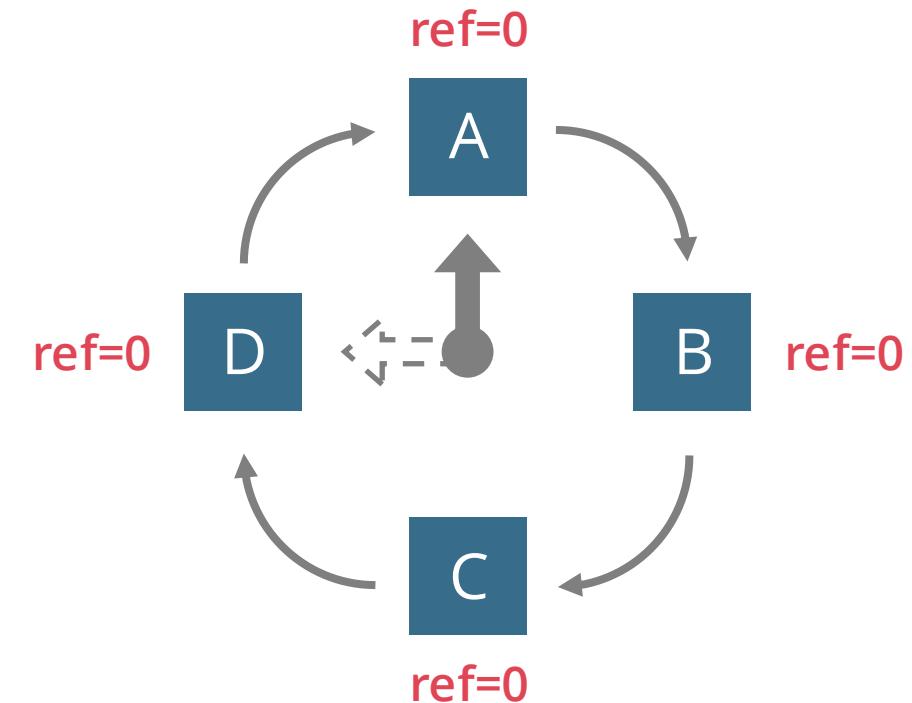
Page access sequence:

A B C D E B A D C A E C



Resets bits of A, B, C, D while moving the hand

First frame with ref = 0 holds A



Buffer hits (so far) = **0**

# QUESTION 4, PART 5

Page access sequence:

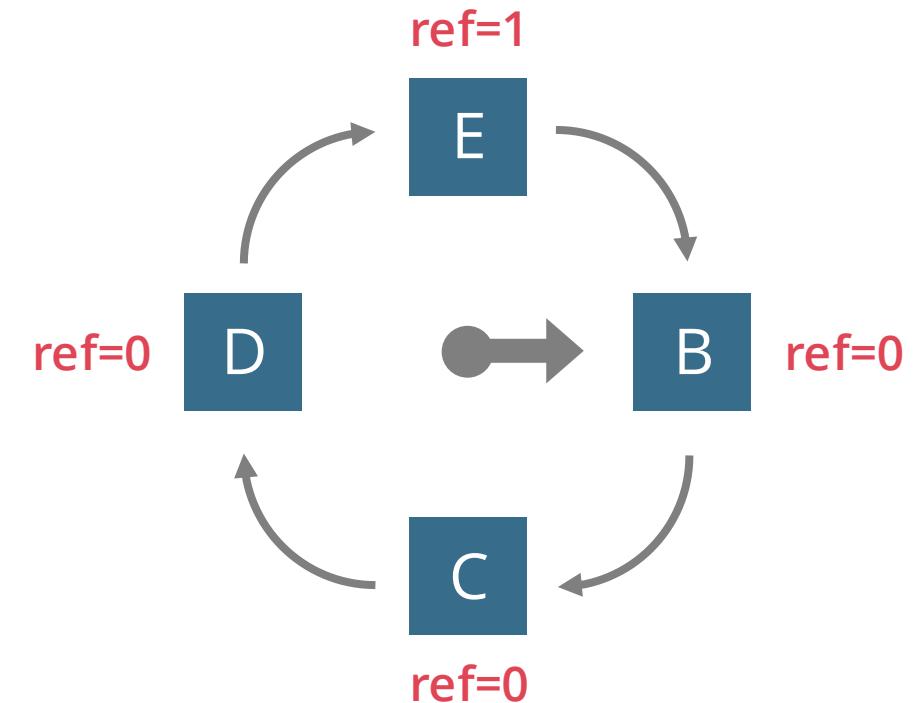
A B C D E B A D C A E C



Resets bits of A, B, C, D while moving the hand

First frame with  $\text{ref} = 0$  holds A

Replace A with E, set reference bit, move the hand



Buffer hits (so far) = **0**

# QUESTION 4, PART 6

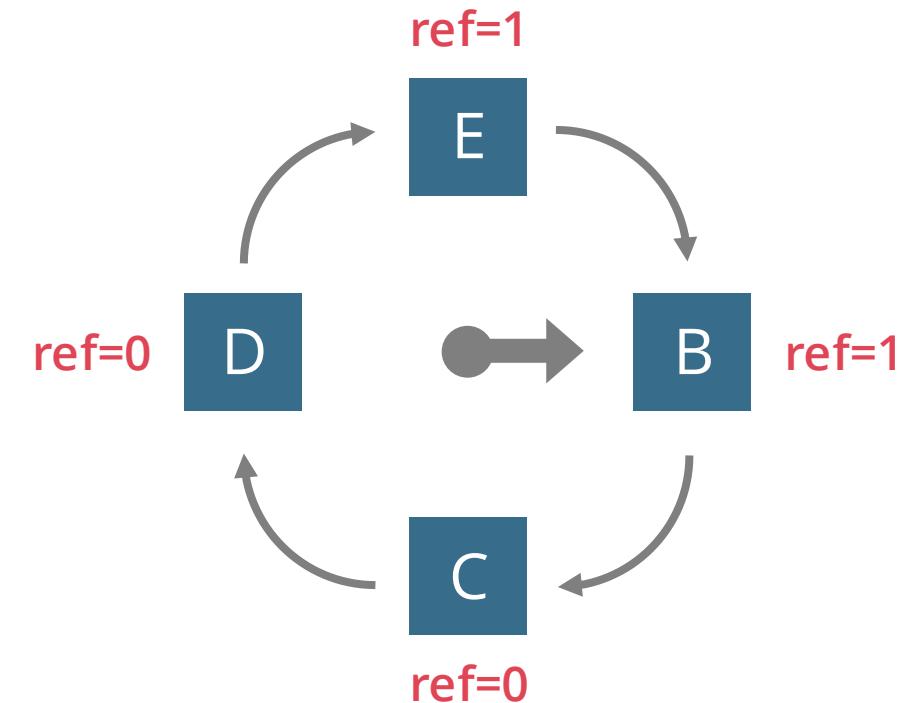
Page access sequence:

A B C D E B A D C A E C



Page B is present  $\Rightarrow$  buffer hit!

Set reference bit



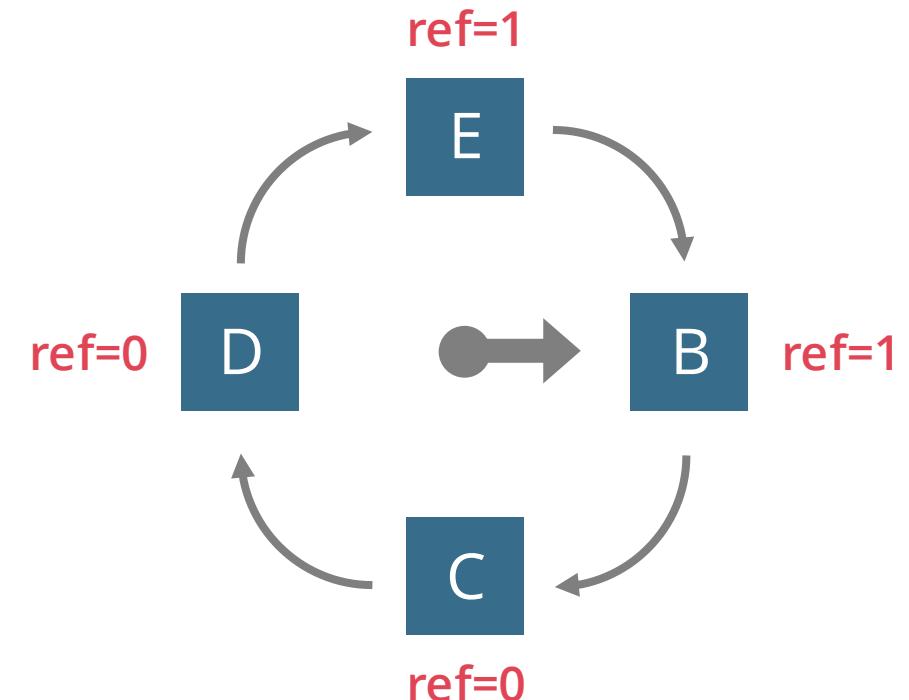
Buffer hits (so far) = 1

# QUESTION 4, PART 7

## Page access sequence:

A B C D E B A D C A E C

Page A not present  $\Rightarrow$  buffer miss!



Buffer hits (so far) = 1

# QUESTION 4, PART 8

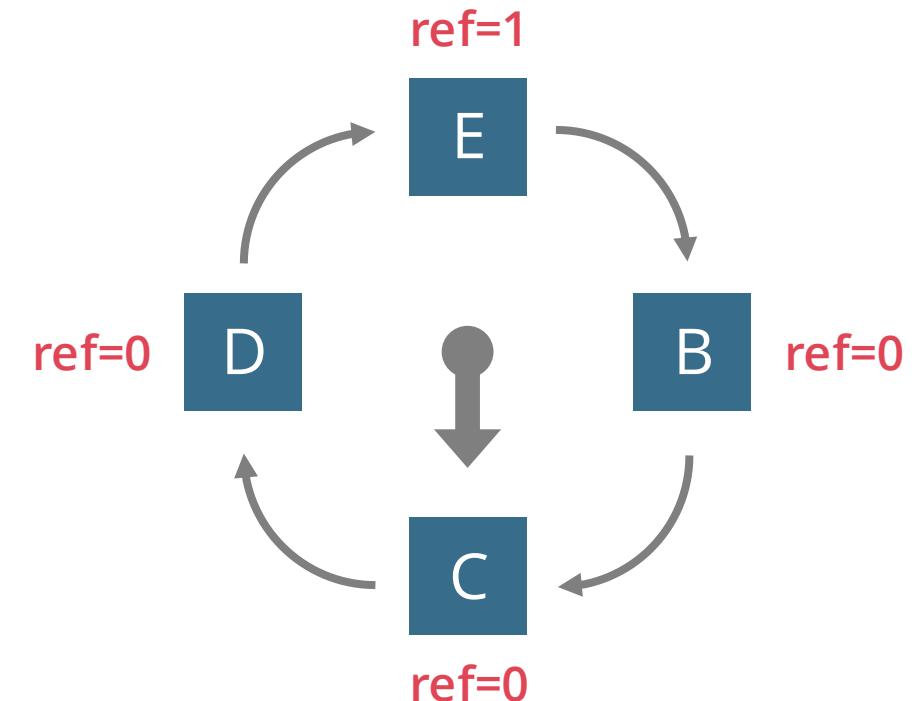
Page access sequence:

A B C D E B A D C A E C



Page A not present  $\Rightarrow$  buffer miss!

Unset refence bit for B, move the hand



Buffer hits (so far) = 1

# QUESTION 4, PART 9

Page access sequence:

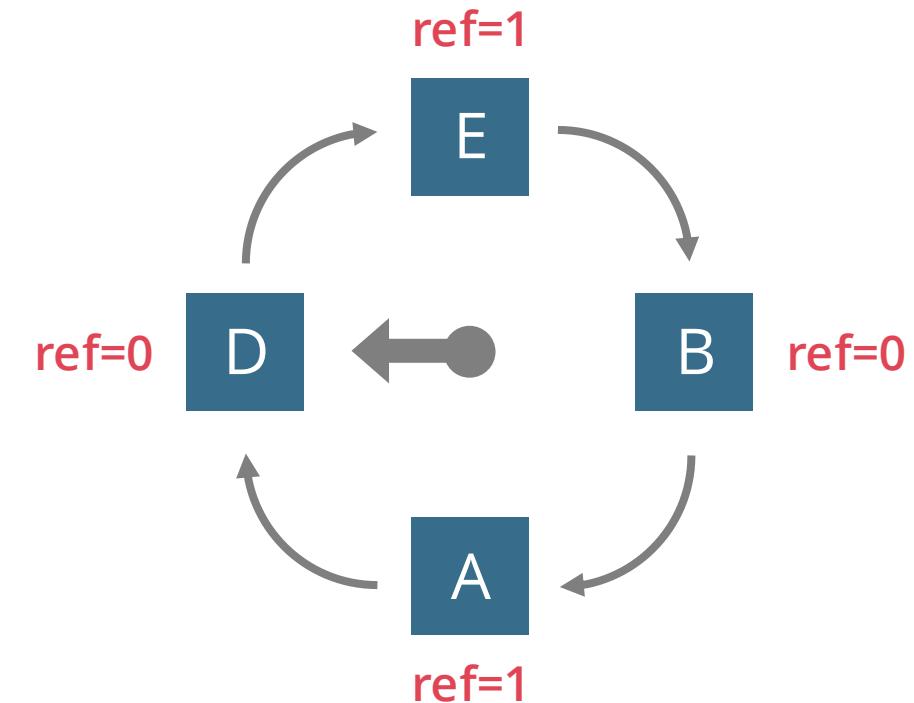
A B C D E B A D C A E C



Page A not present  $\Rightarrow$  buffer miss!

Unset refence bit for B, move the hand

Replace C with A, set refence bit, move the hand



Buffer hits (so far) = 1

# QUESTION 4, PART 10

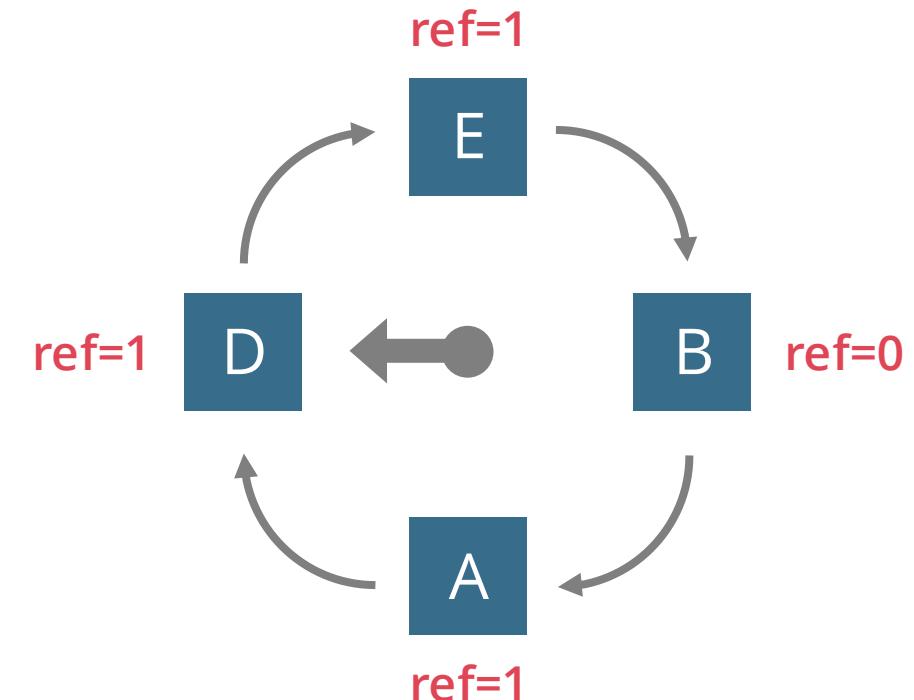
Page access sequence:

A B C D E B A D C A E C

↑

Page D is present  $\Rightarrow$  buffer hit!

Set reference bit



Buffer hits (so far) = 2

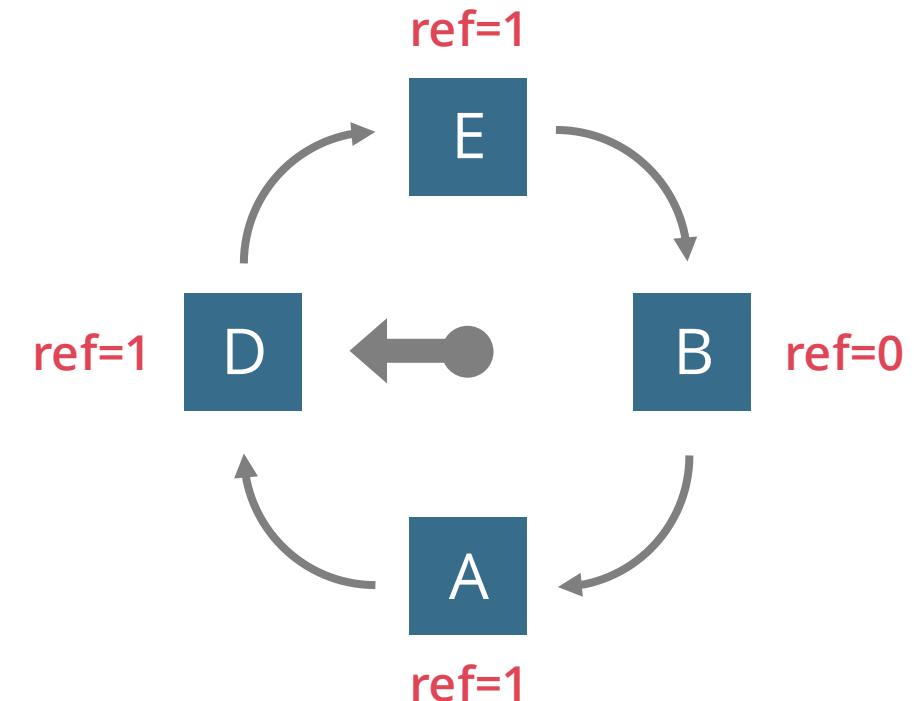
# QUESTION 4, PART 11

Page access sequence:

A B C D E B A D C A E C



Page C is not present  $\Rightarrow$  buffer miss!



Buffer hits (so far) = 2

# QUESTION 4, PART 12

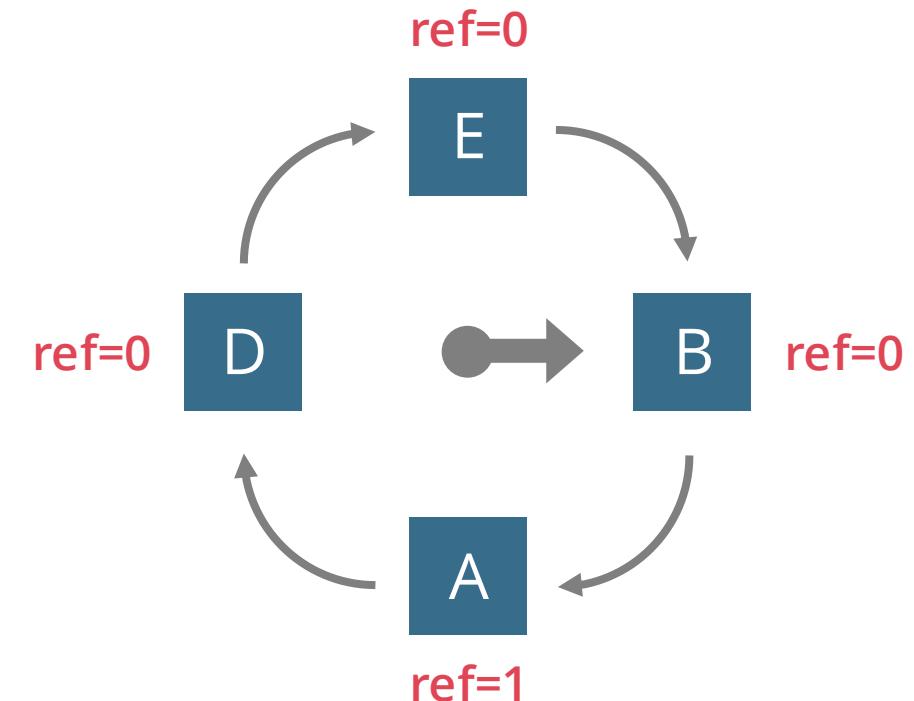
Page access sequence:

A B C D E B A D C A E C



Page C is not present  $\Rightarrow$  buffer miss!

Unset ref bits for D & E, move the hand to B



Buffer hits (so far) = 2

# QUESTION 4, PART 13

Page access sequence:

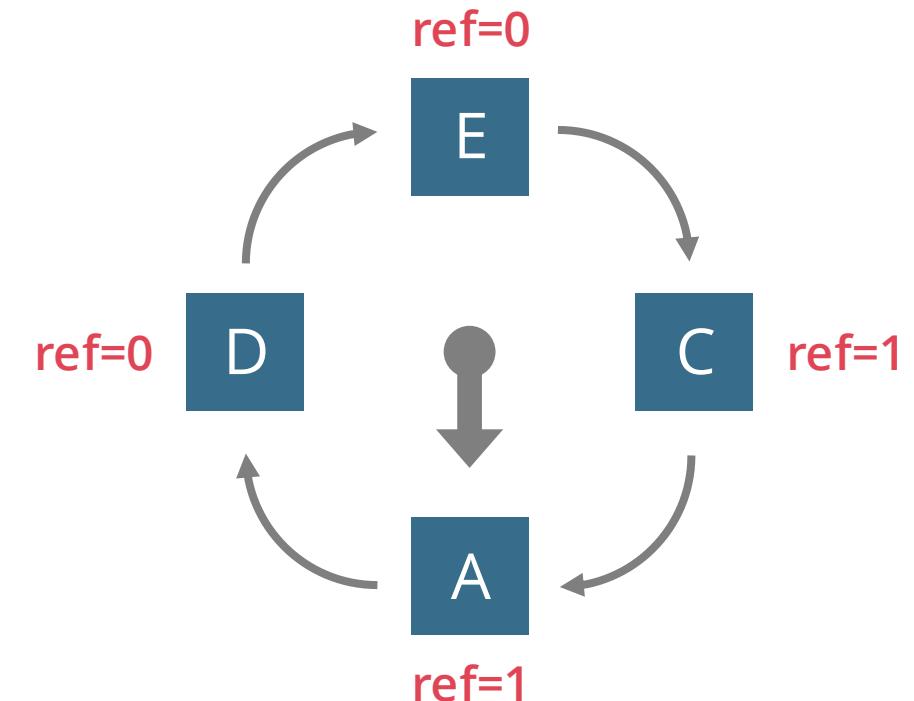
A B C D E B A D C A E C



Page C is not present  $\Rightarrow$  buffer miss!

Unset ref bits for D & E, move the hand to B

Replace B with C, set reference bit, move the hand



Buffer hits (so far) = 2

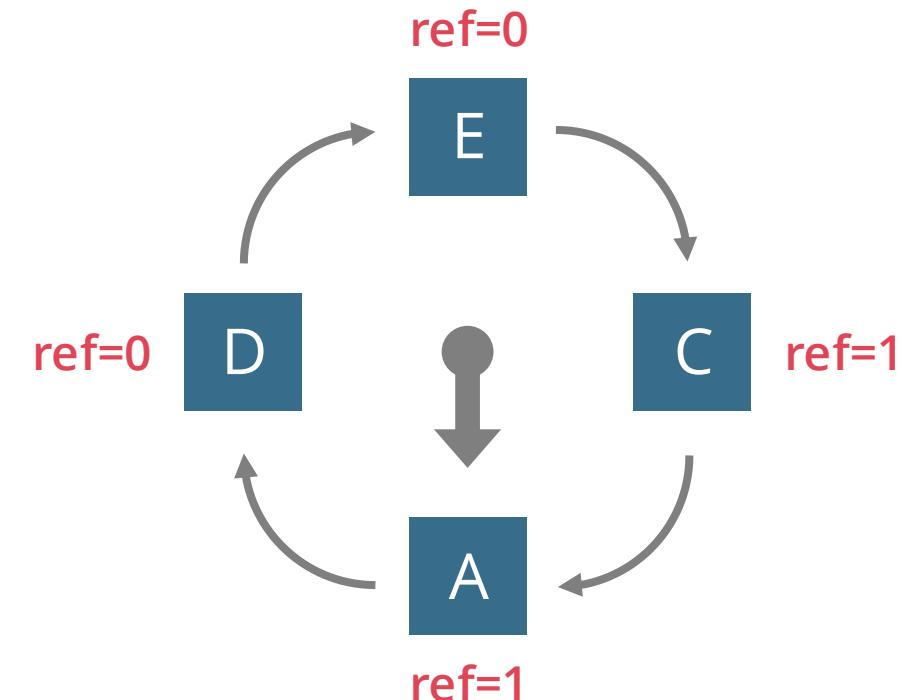
# QUESTION 4, PART 14

Page access sequence:

A B C D E B A D C A E C



Pages A, E, C are present  $\Rightarrow$  buffer hits!



Buffer hits (so far) = 2

# QUESTION 4, PART 15

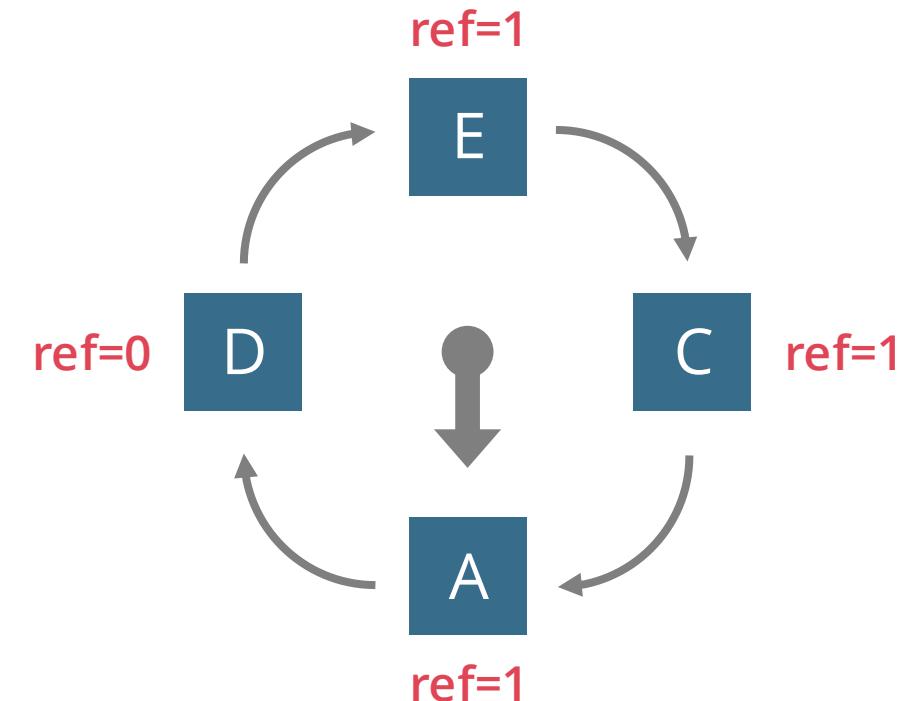
Page access sequence:

A B C D E B A D C A E C



Pages A, E, C are present  $\Rightarrow$  buffer hits!

Set their reference bits



Buffer hits = 5

# POSTGRESQL – BUFFER POOL DEMO

## Purpose

What the PostgreSQL buffer pool (*shared\_buffers*) is

How sequential scans use a ring buffer to avoid cache pollution

How index-based access uses the normal buffer pool

## Key ideas

PostgreSQL stores data in 8 KB pages

Pages are cached in *shared\_buffers*

Different access patterns use the cache differently

## Try it out

`buffer_demo.sql` is available on Learn → Practice Worksheets

Requires PostgreSQL installed locally

Run the script step by step and observe buffer behaviour

# POSTGRESQL – SLOTTED PAGES DEMO

## Purpose

- How tables are stored as slotted pages
- How inserts/deletes create dead space
- How **VACUUM** and **VACUUM FULL** reclaim space

## Key ideas

- Pages contain headers, pointers, tuples, and free space
- Dead tuples persist until cleaned

## Try it out

- `page_demo.sql` is available on Learn → Practice Worksheets
- Run the script step by step and observe space usage before and after VACUUM