Locating records by the DBMS

- Fact: pointers (addresses) are part (stored) of records.
- (This is not typical in tuples of a relation)
- But: this is common for tuples of an Object Oriented DBMS. (Object Relational DBMS).
- Pointers are also used in index files.
- The DBMS has a management system for pointers: needs

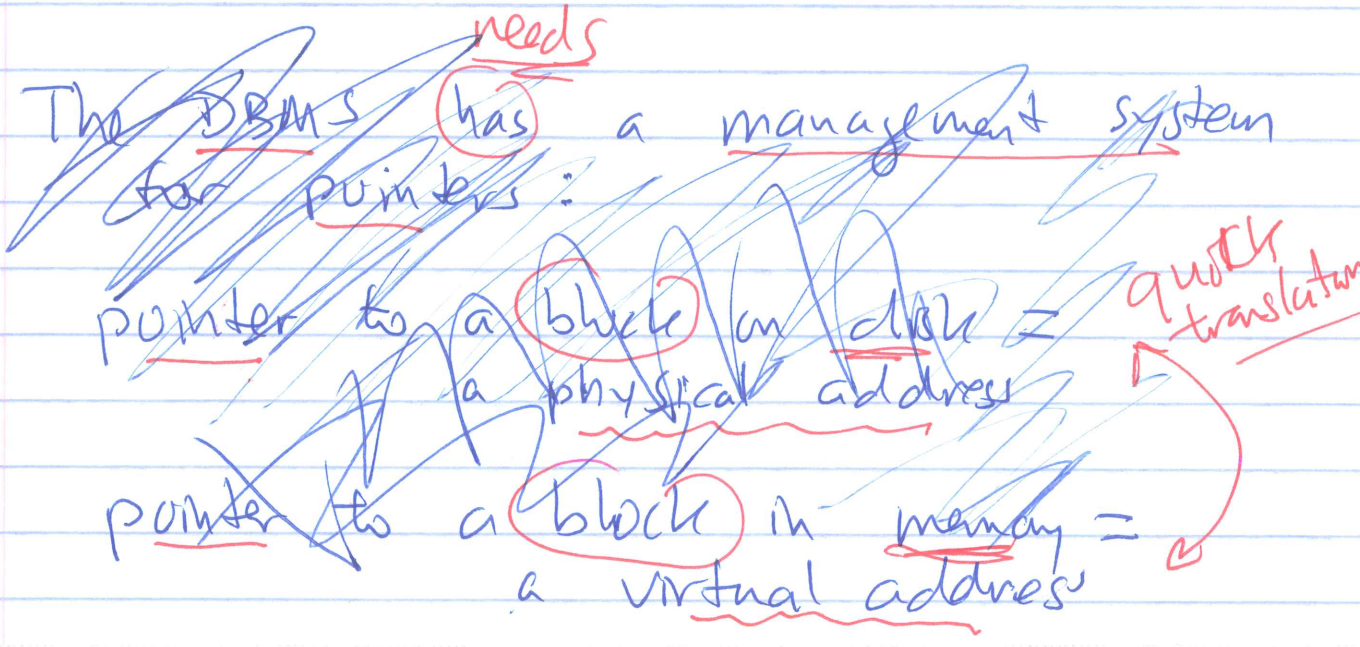

- Recall that:

A database object (block/record) is identified by:

(1) Database address of object (= logical/physical address) when object is on Disk.

(2) A virtual addr. in memory when object is read/stored in memory.

(The object is still on disk, but the one on disk is NOT used until the in-memory copy is written back to disk).

- The DBMS has/needs a management system to convert:

Database address ⟷ virtual address

Observation:

(1) When a database object (block/record) is on disk, we can ONLY refer to the object using a database address.

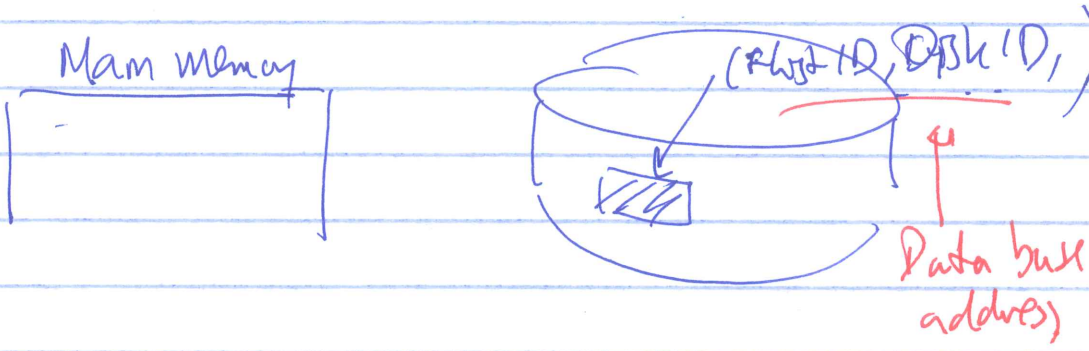

(2) When a database object (block/record) is in memory:

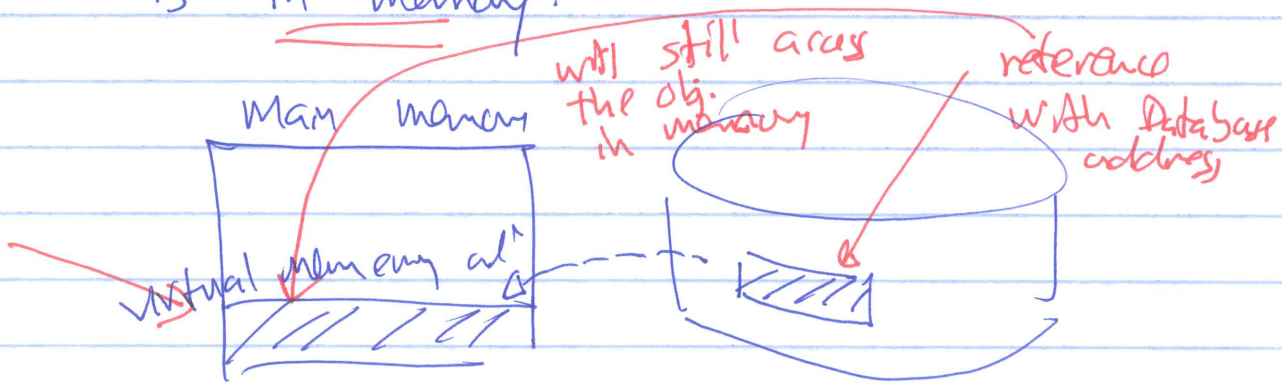

(A) we must use the object in memory.

(B) but: we can refer (identify) this object by:
- virtual memory addr
- Database address.

- So: when the DBMS uses a Database Address to reference an object that is in memory, we must:

  translate the Database Address → to a (virtual) memory addr !!!

- Translation table:

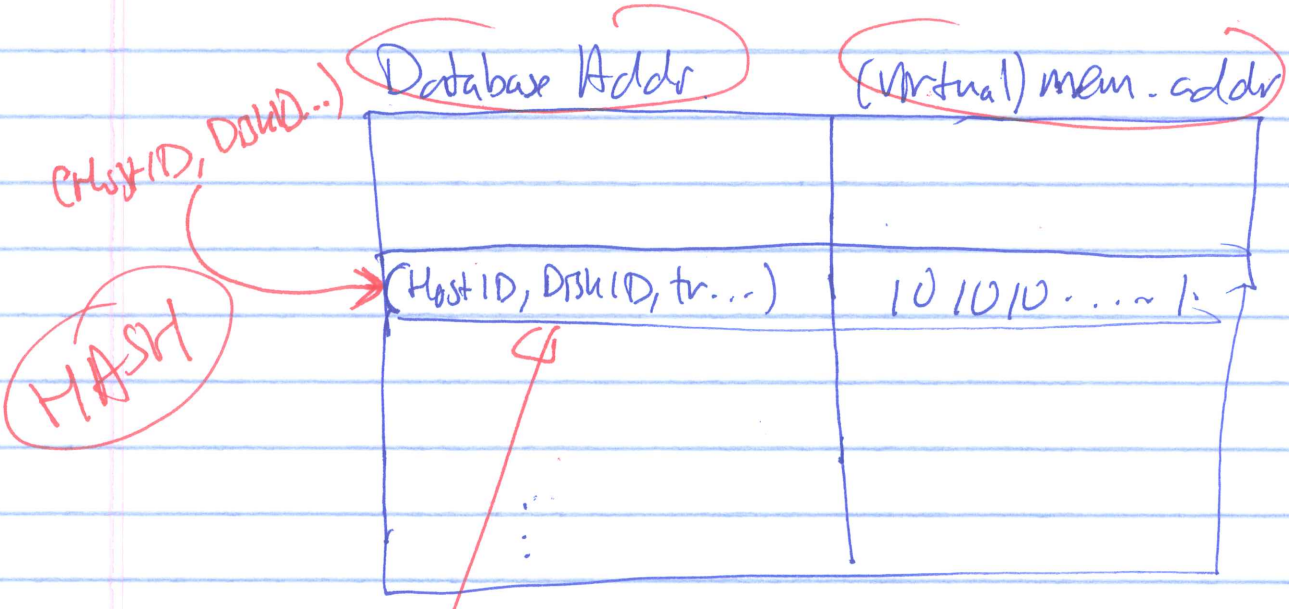

Naive Access to Database Objects.

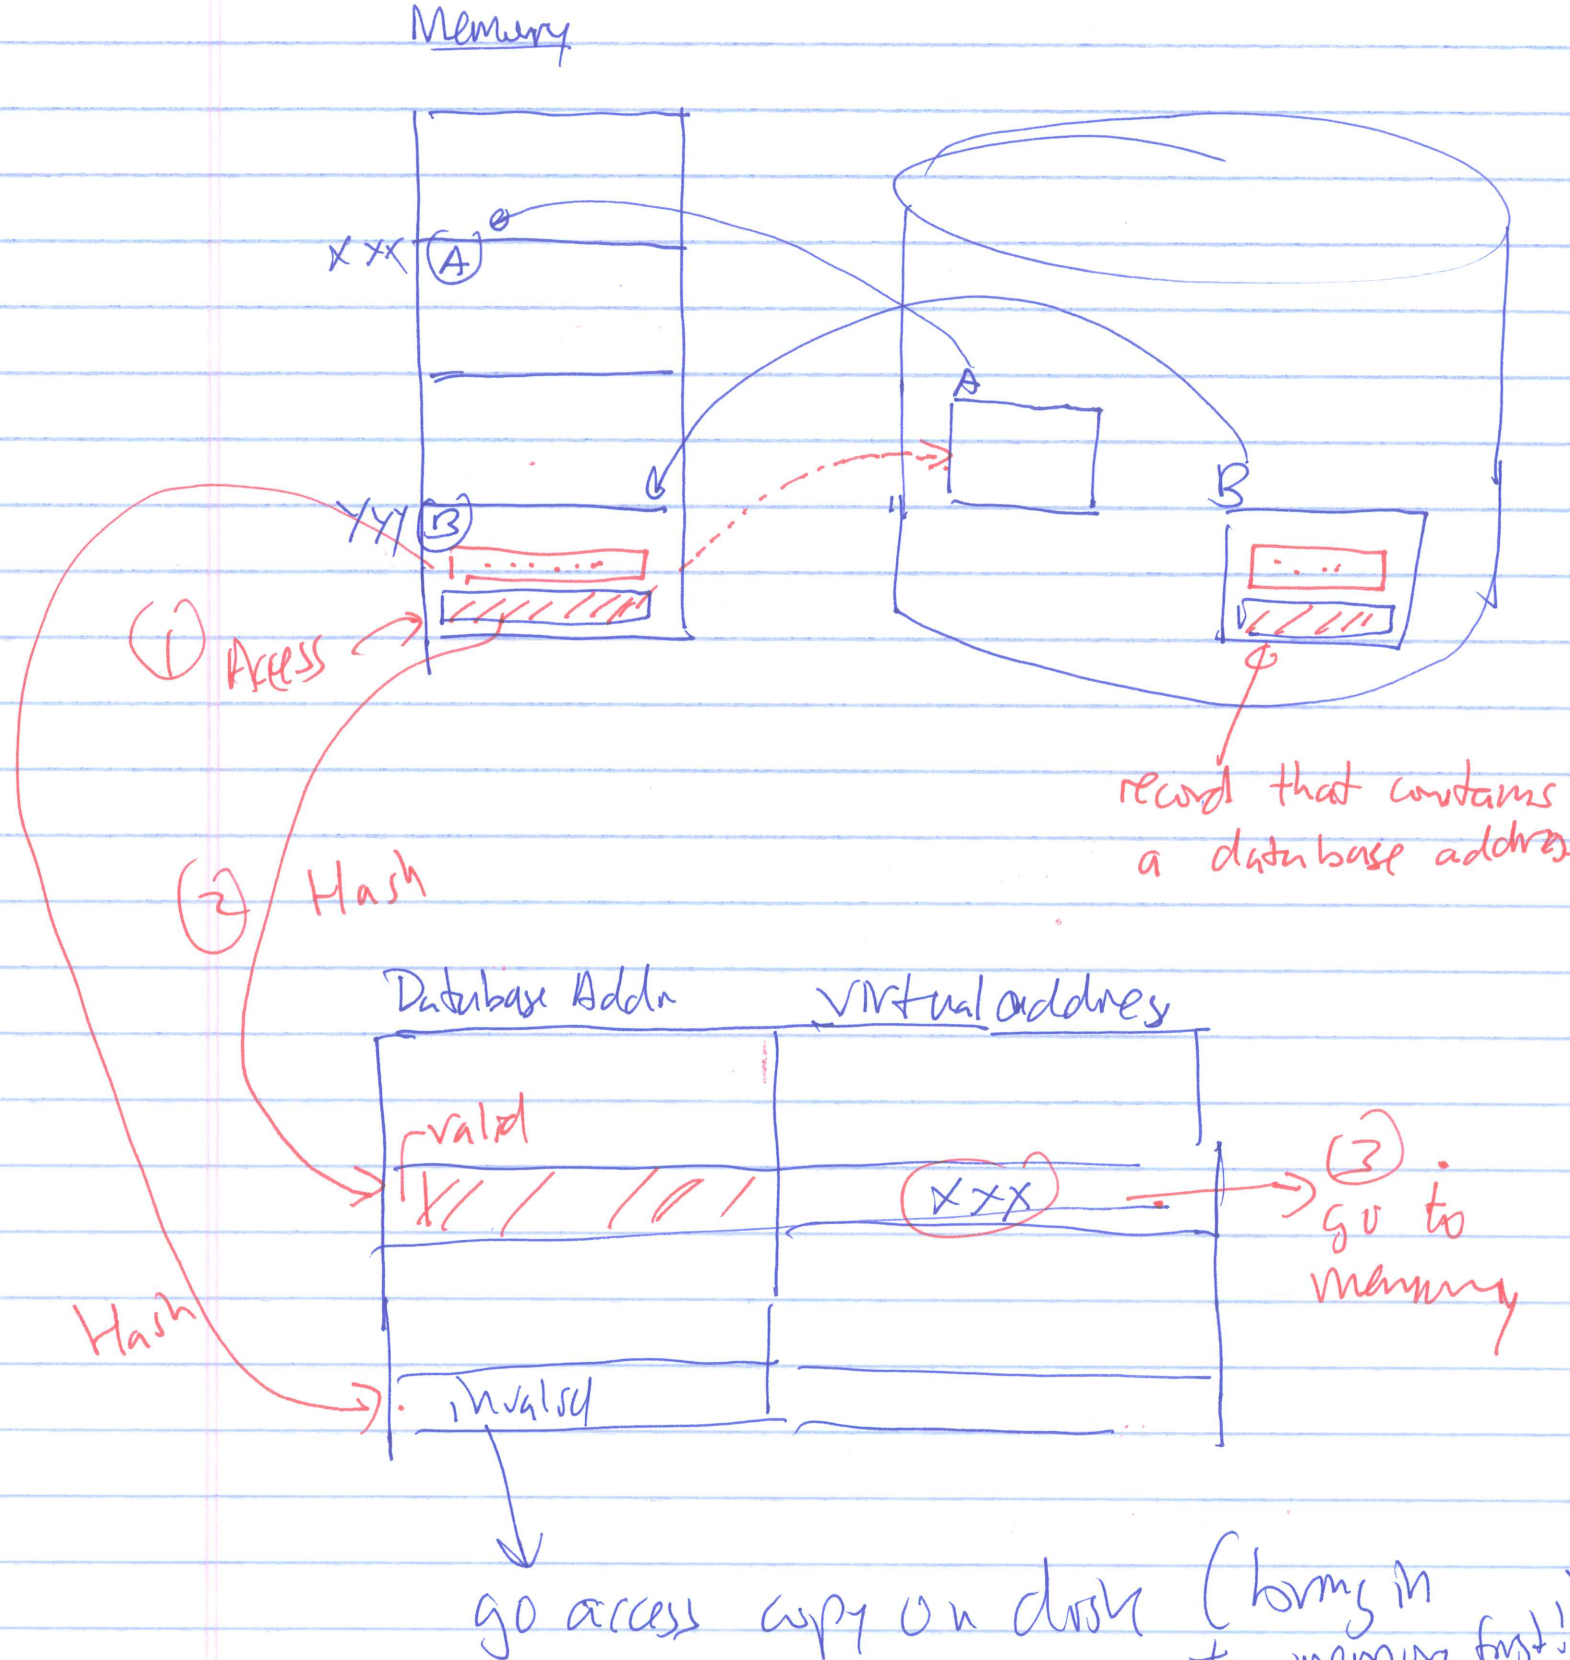

- Advantage: (1) easy to implement
  (only need a hashing table)
  (2) records/blocks are not "pinned" in memory (later)

Disadvantage: slow...

Faster Access: Trick (pointer swizzling)

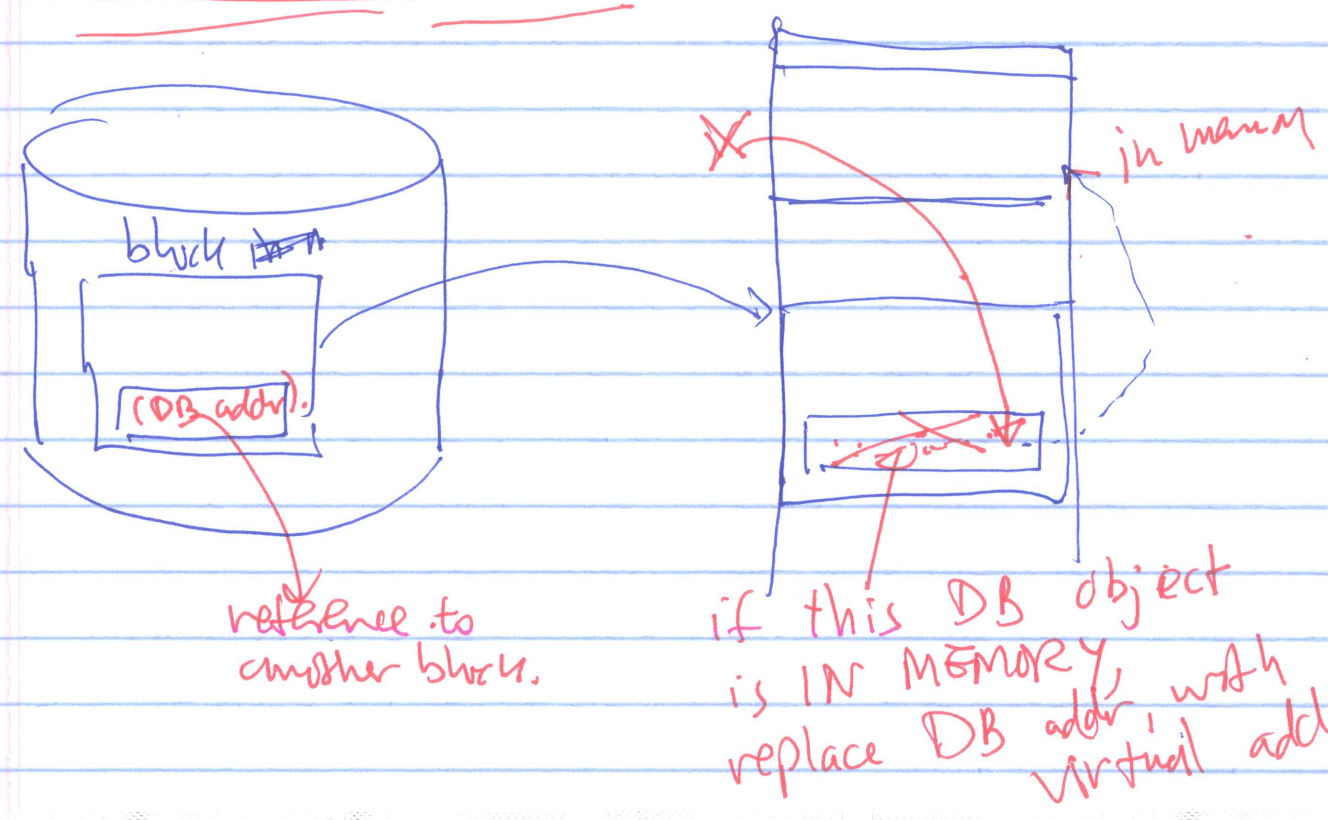

Pointer swizzling:

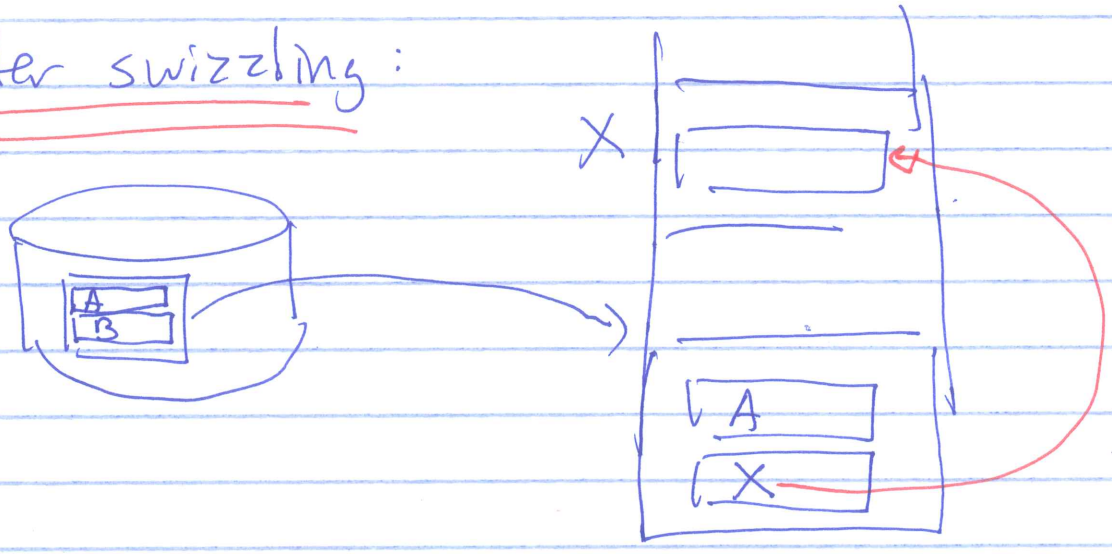

When we move a block from disk to main memory, the pointers (= Database addresses) in the block may be "swizzled"

//

translated (using the Hash table/Map from Database Address ↓ virtual address.

Implementation:

The ~~data~~ record field for a DB addr. is expanded with 1 bit:

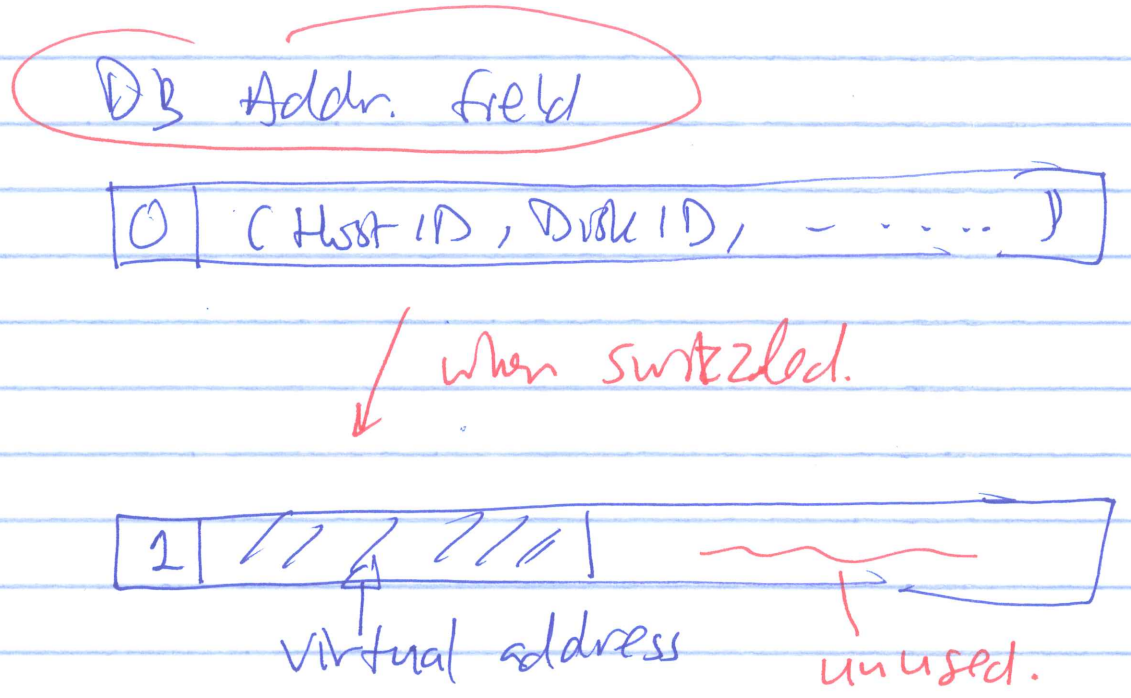

Example:

Disk

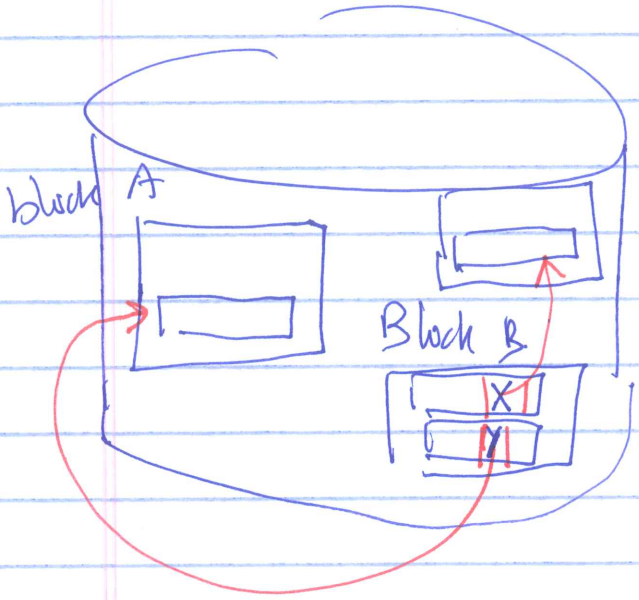

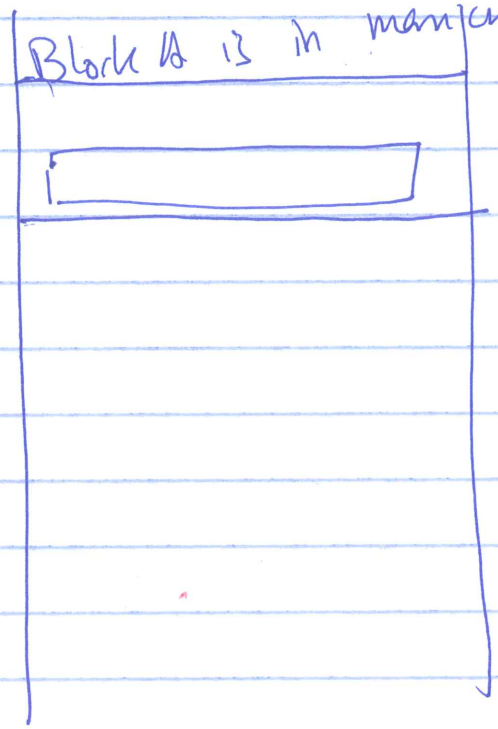

When we read in Block B:

translation table

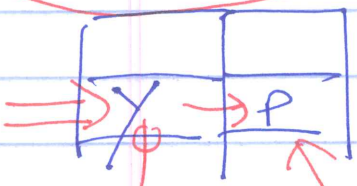

DB addr

virt. mem. addr.

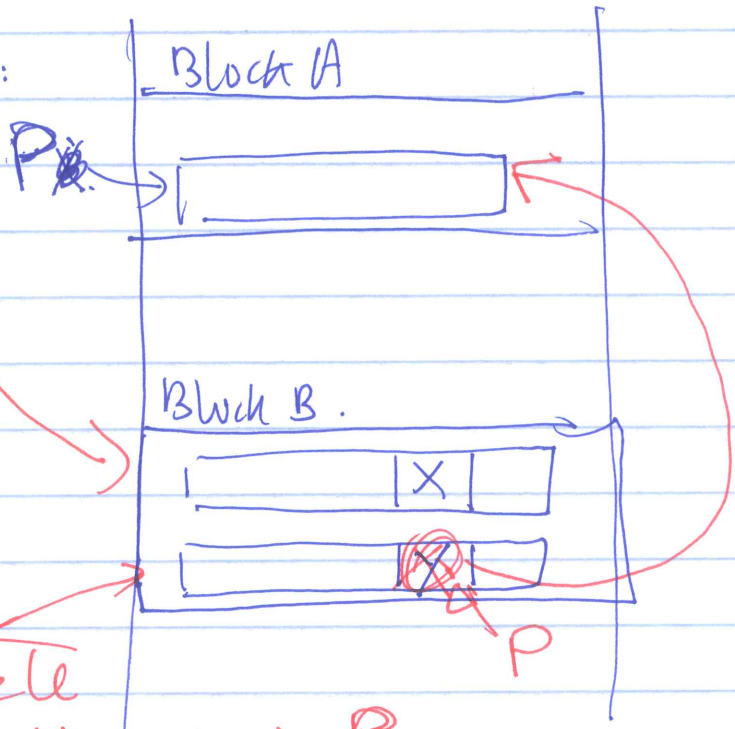

We swizzle the DB addr Y → P

Swizzling Techniques:

(1) Automatic Swizzling

- Swizzle all DB addr → mem. addr inside a block when the block is FIRST read in to memory.

(2) On Demand swizzling

- Leave all DB addr. unswizzled when block read in.

- When DB addr. is used. (first time)

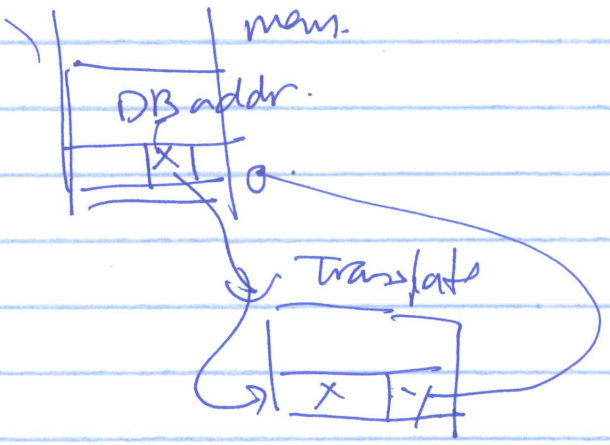

Swizzle (so next accesses will be fast).

Implementing auto matic swizzling:

- Information required:

  Knowledge on the location of every DB address type field in a block.

  (1) Block holds records of one schema

  ⇒ Schema will tell us the location of the DB addr. fields

  (2) Block used for indexes (later)

  - Structure is known
  - DB addr's can be located

  (3) Structure unknown ...

  → put a list of offsets in Block header that contain locations of DB addresses.

Implementing on-demand swizzling:

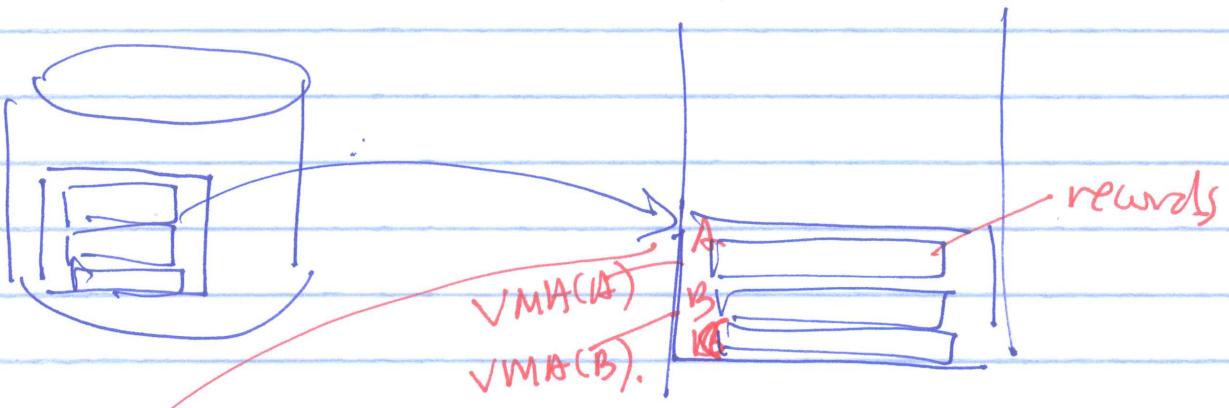

When we read in a block:

Enter translation entries for every record in the block:

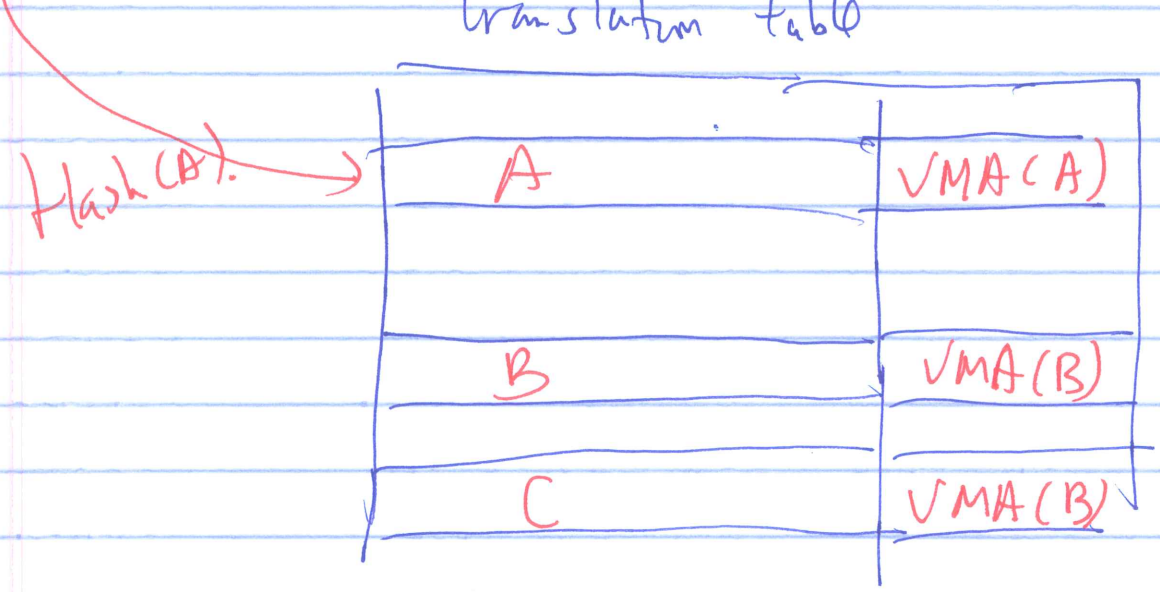

Problem created by Pointer Swizzling

(1) Problem: the block in memory is changed:

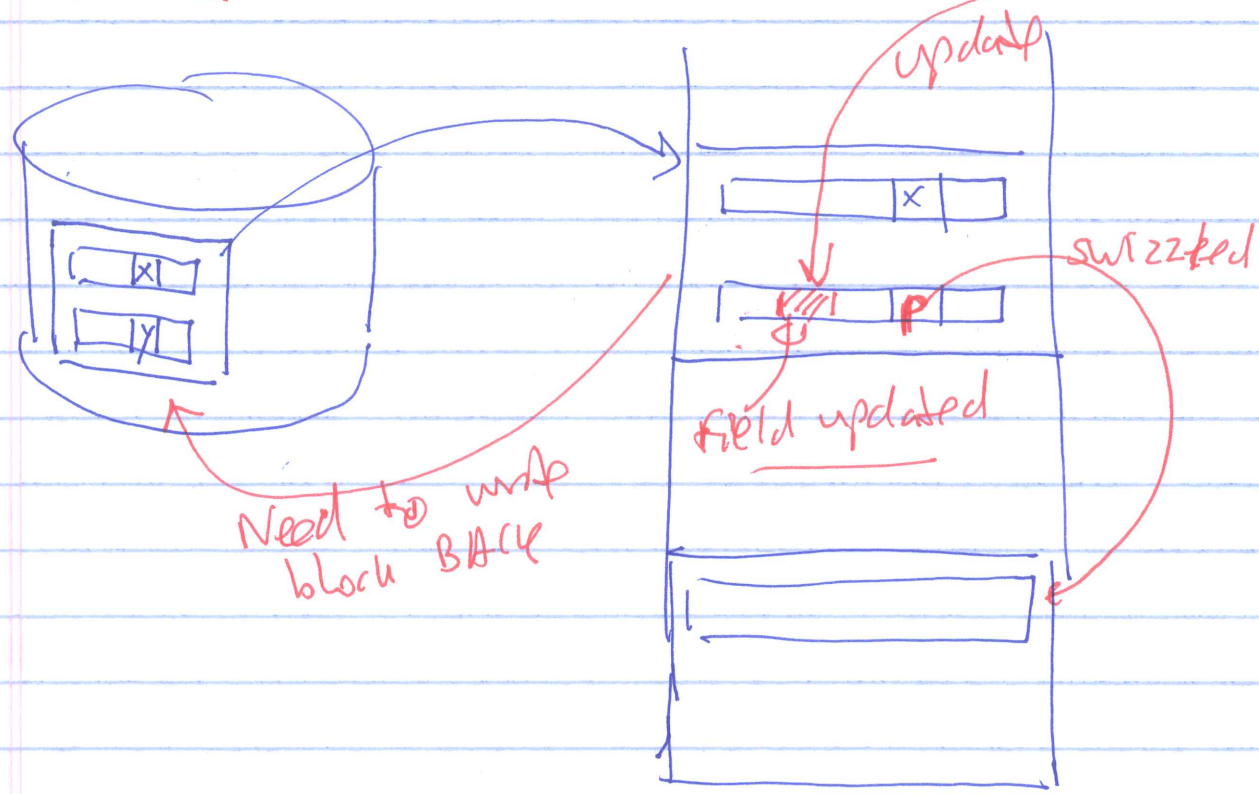

⇒ We must UNSWIZZLE the Virtual addr. p ⇒ DB addr. y BEFORE writing the block Back to the DISK !!!

Naive Solution:

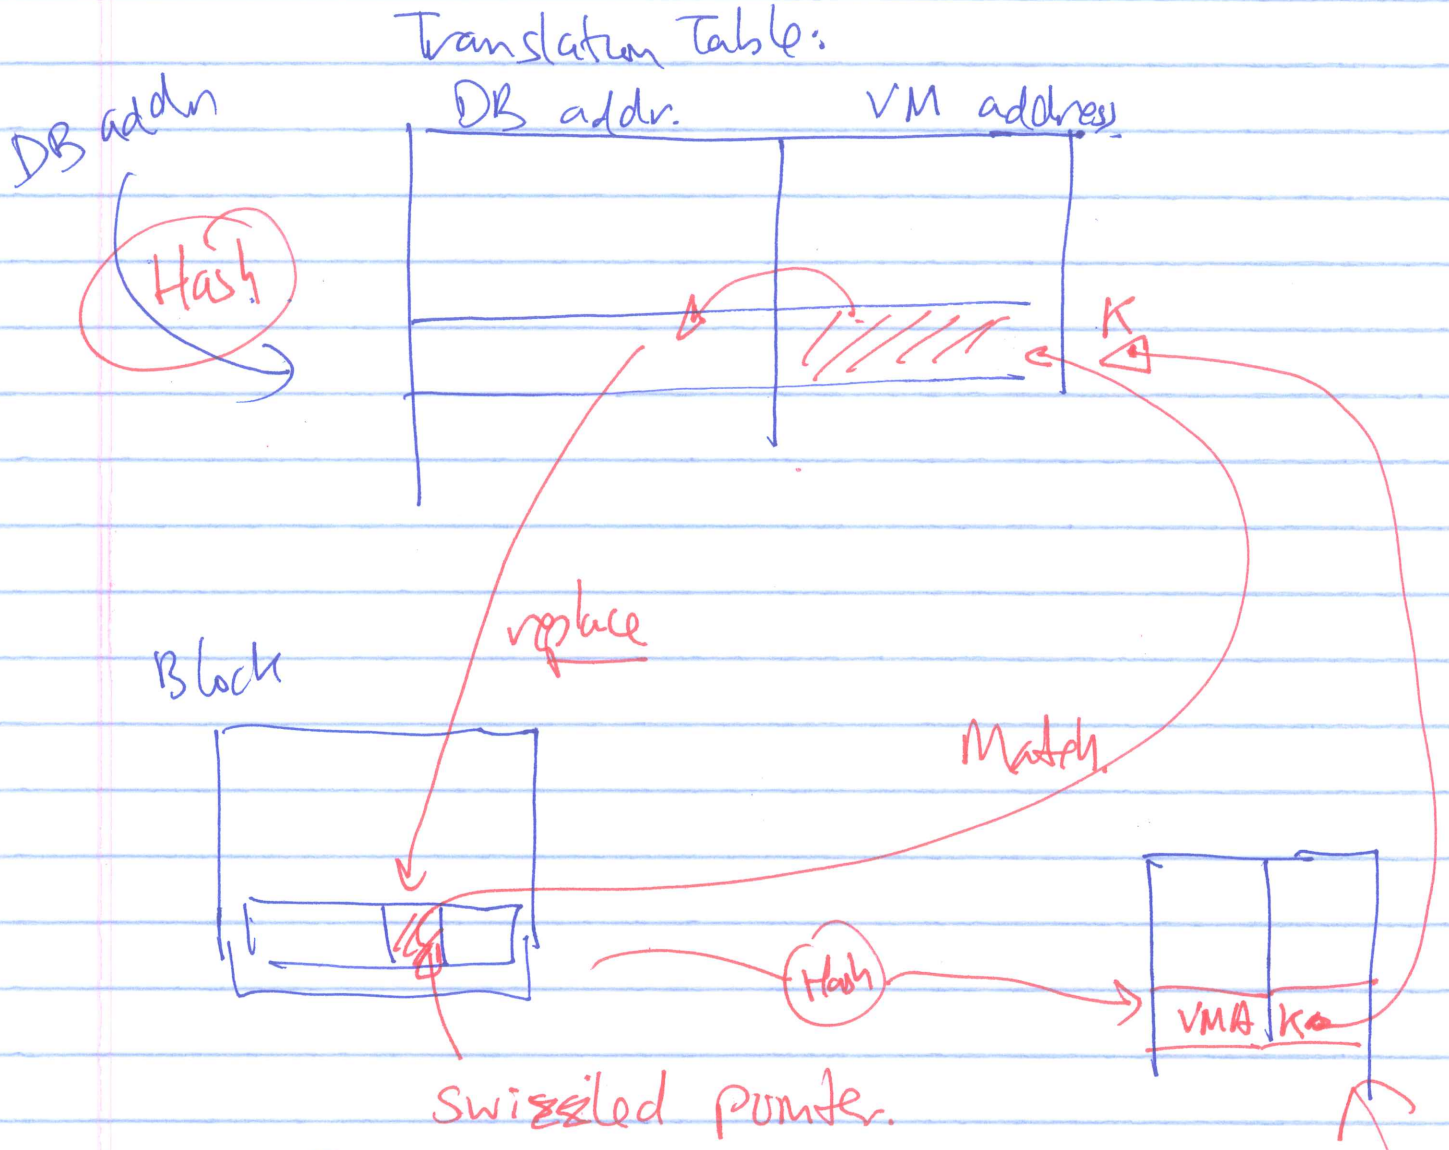

Problem: requires a search to find entry!

Faster ~~Possible~~ solution: Add an index (Hashing).

based on virtual memory addresses.

Summary: faster technique to unswizzle a VM address => DB address

(1) Use a hash map:

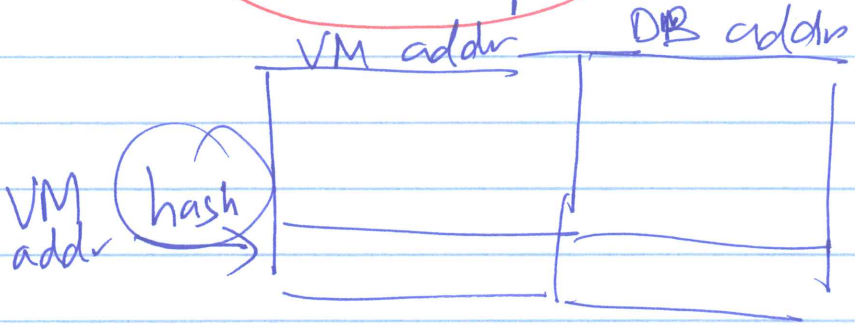

In Chapter 14, the book will discuss an indexing data structure to ~~how to~~ implement this search faster.

pp: ???

Problem created by pointer swizzling

(2) Problem: a block (that the DBMS wants to remove from memory) cannot be removed from memory if:

"pinned"

some records in block if referenced by swizzled pointers in other (in memory) blocks.

Example:

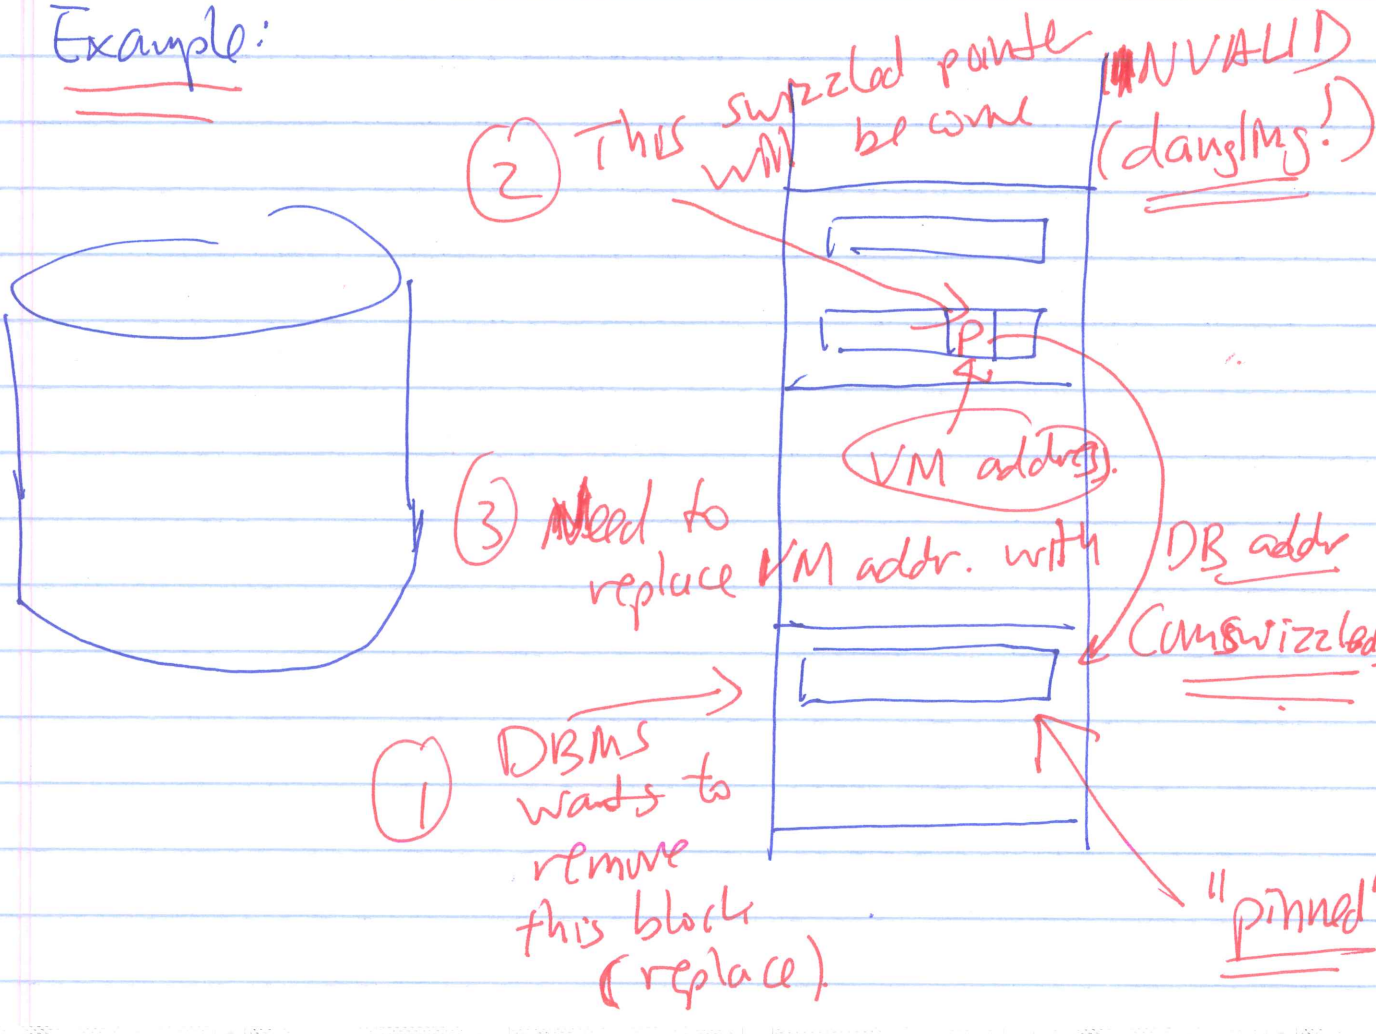

How to unpin a block:

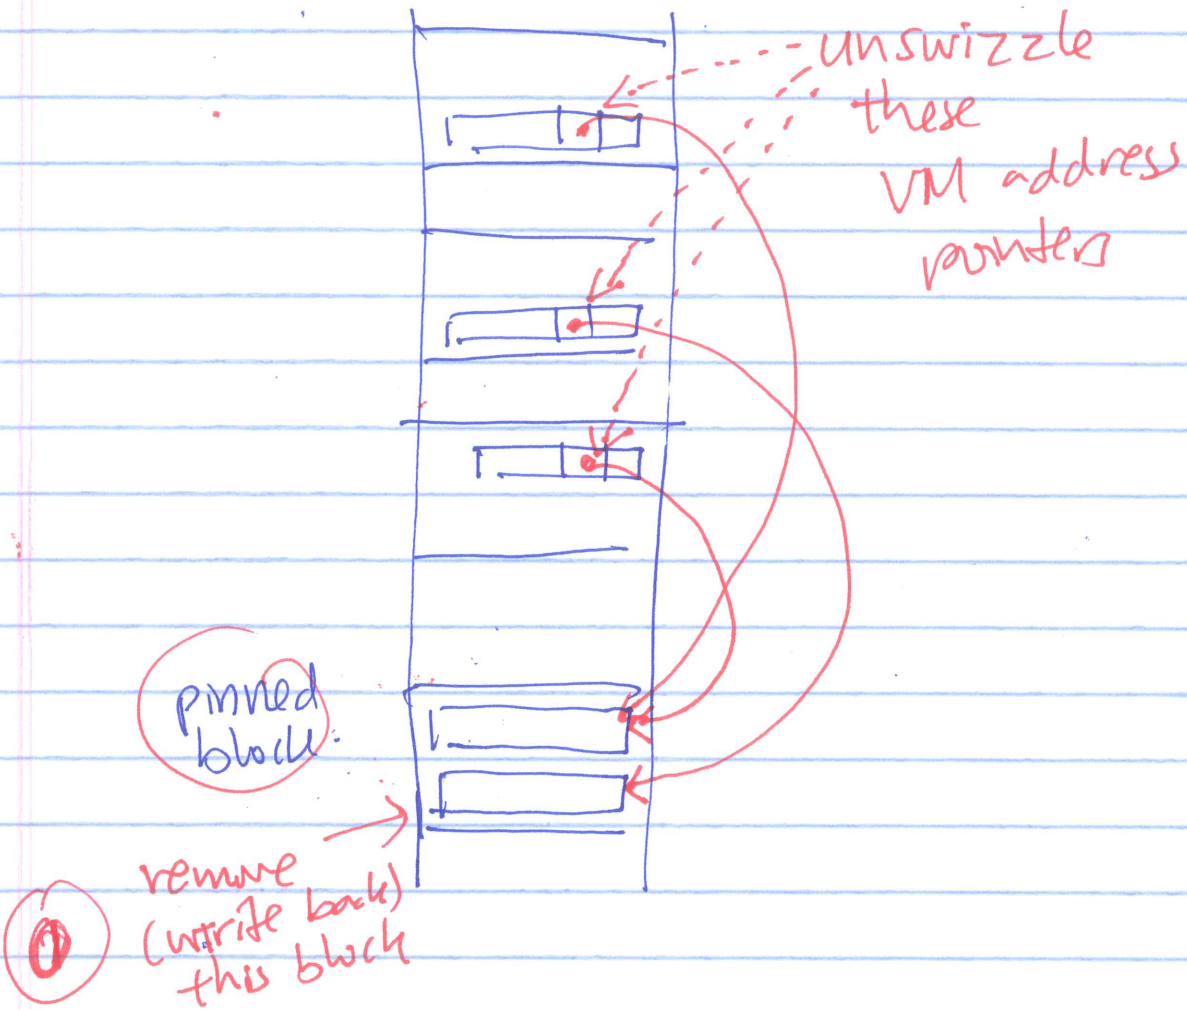

Conclussion:

write back.

When the DBMS wants to release/remove a block from memory:

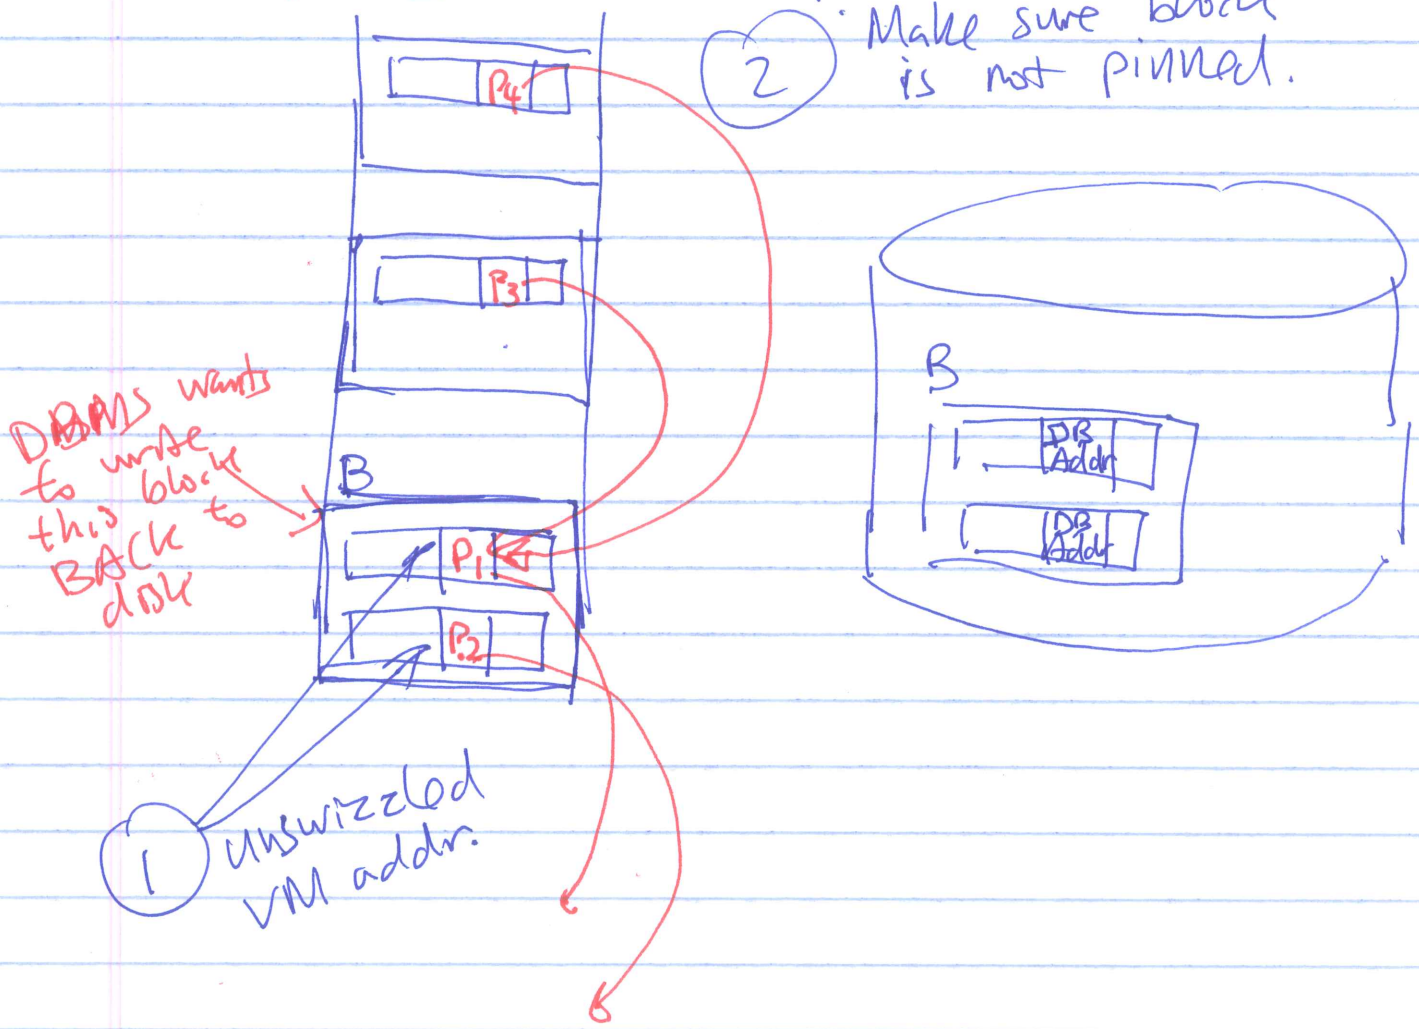

Implementing unpin efficiently

(1) Keep a (linked) list of

memory addresses

of swizzled addresses.

(2) How to use list:

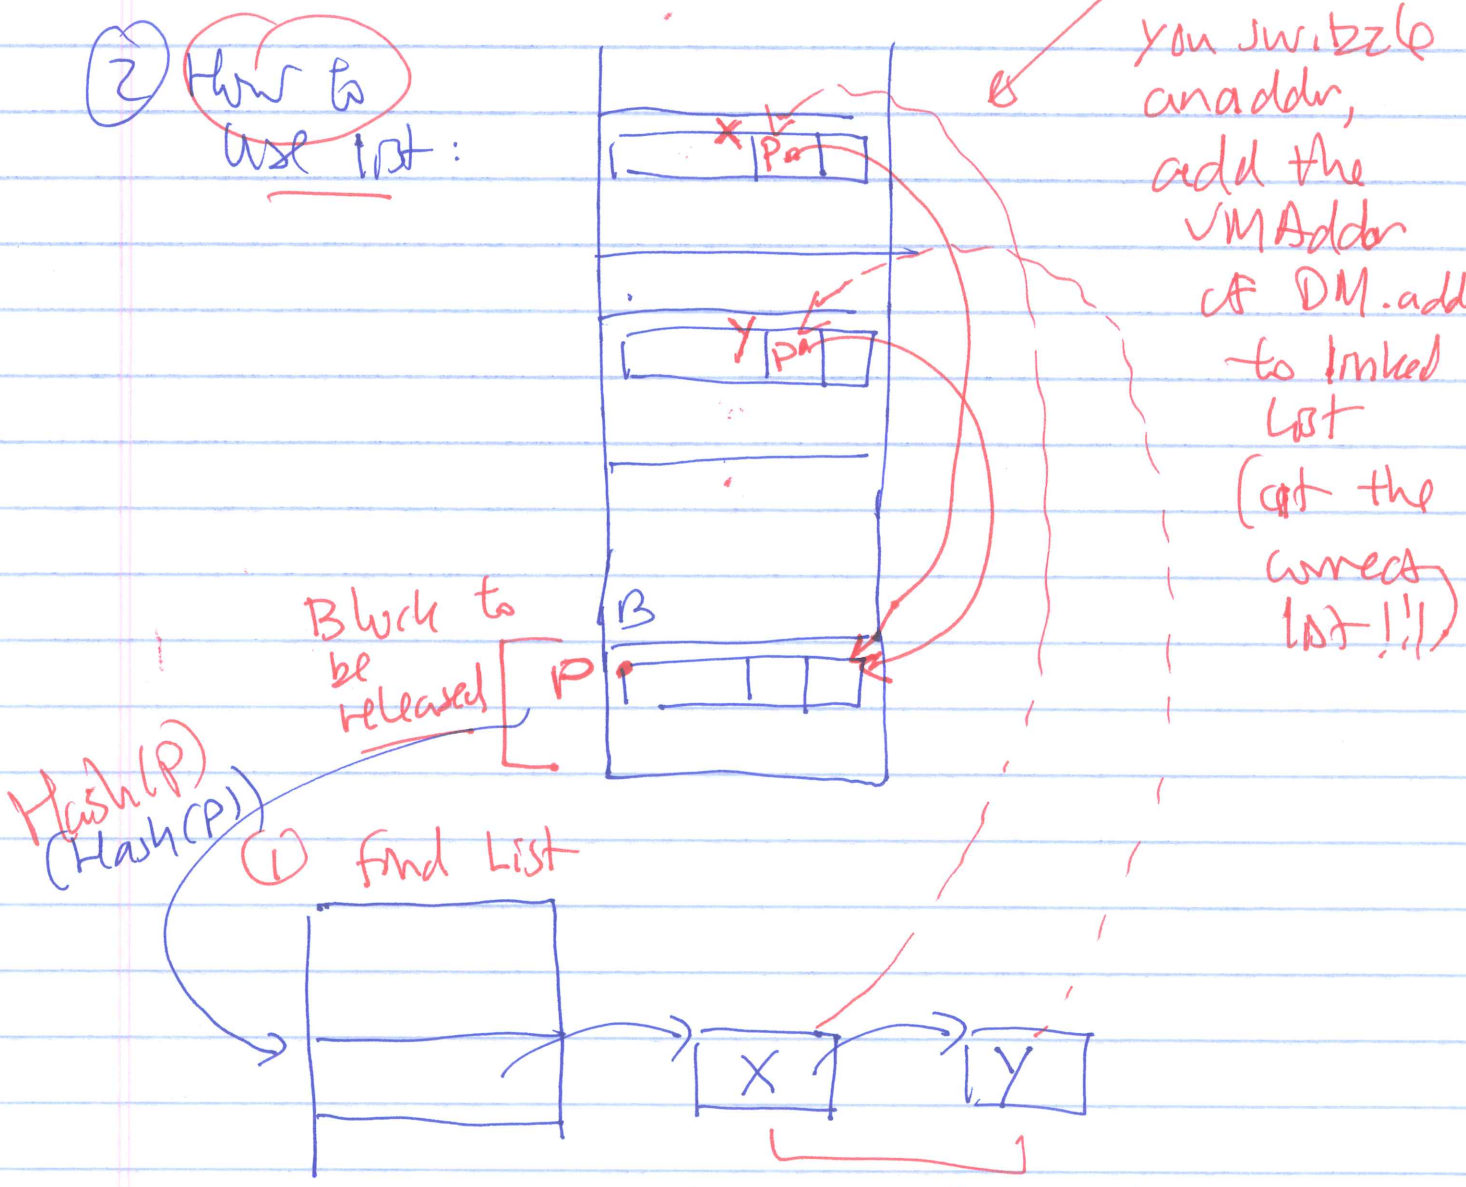

More efficient implementation of the Linked List:

- often: the DB address is much larger than a VM address:

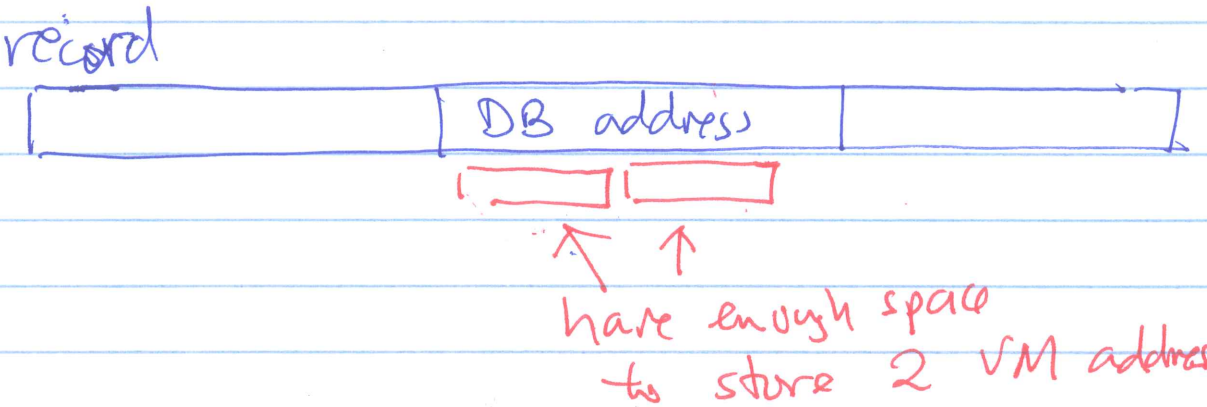

- We can construct the Linked list of swizzled address using the space that store the DB address.

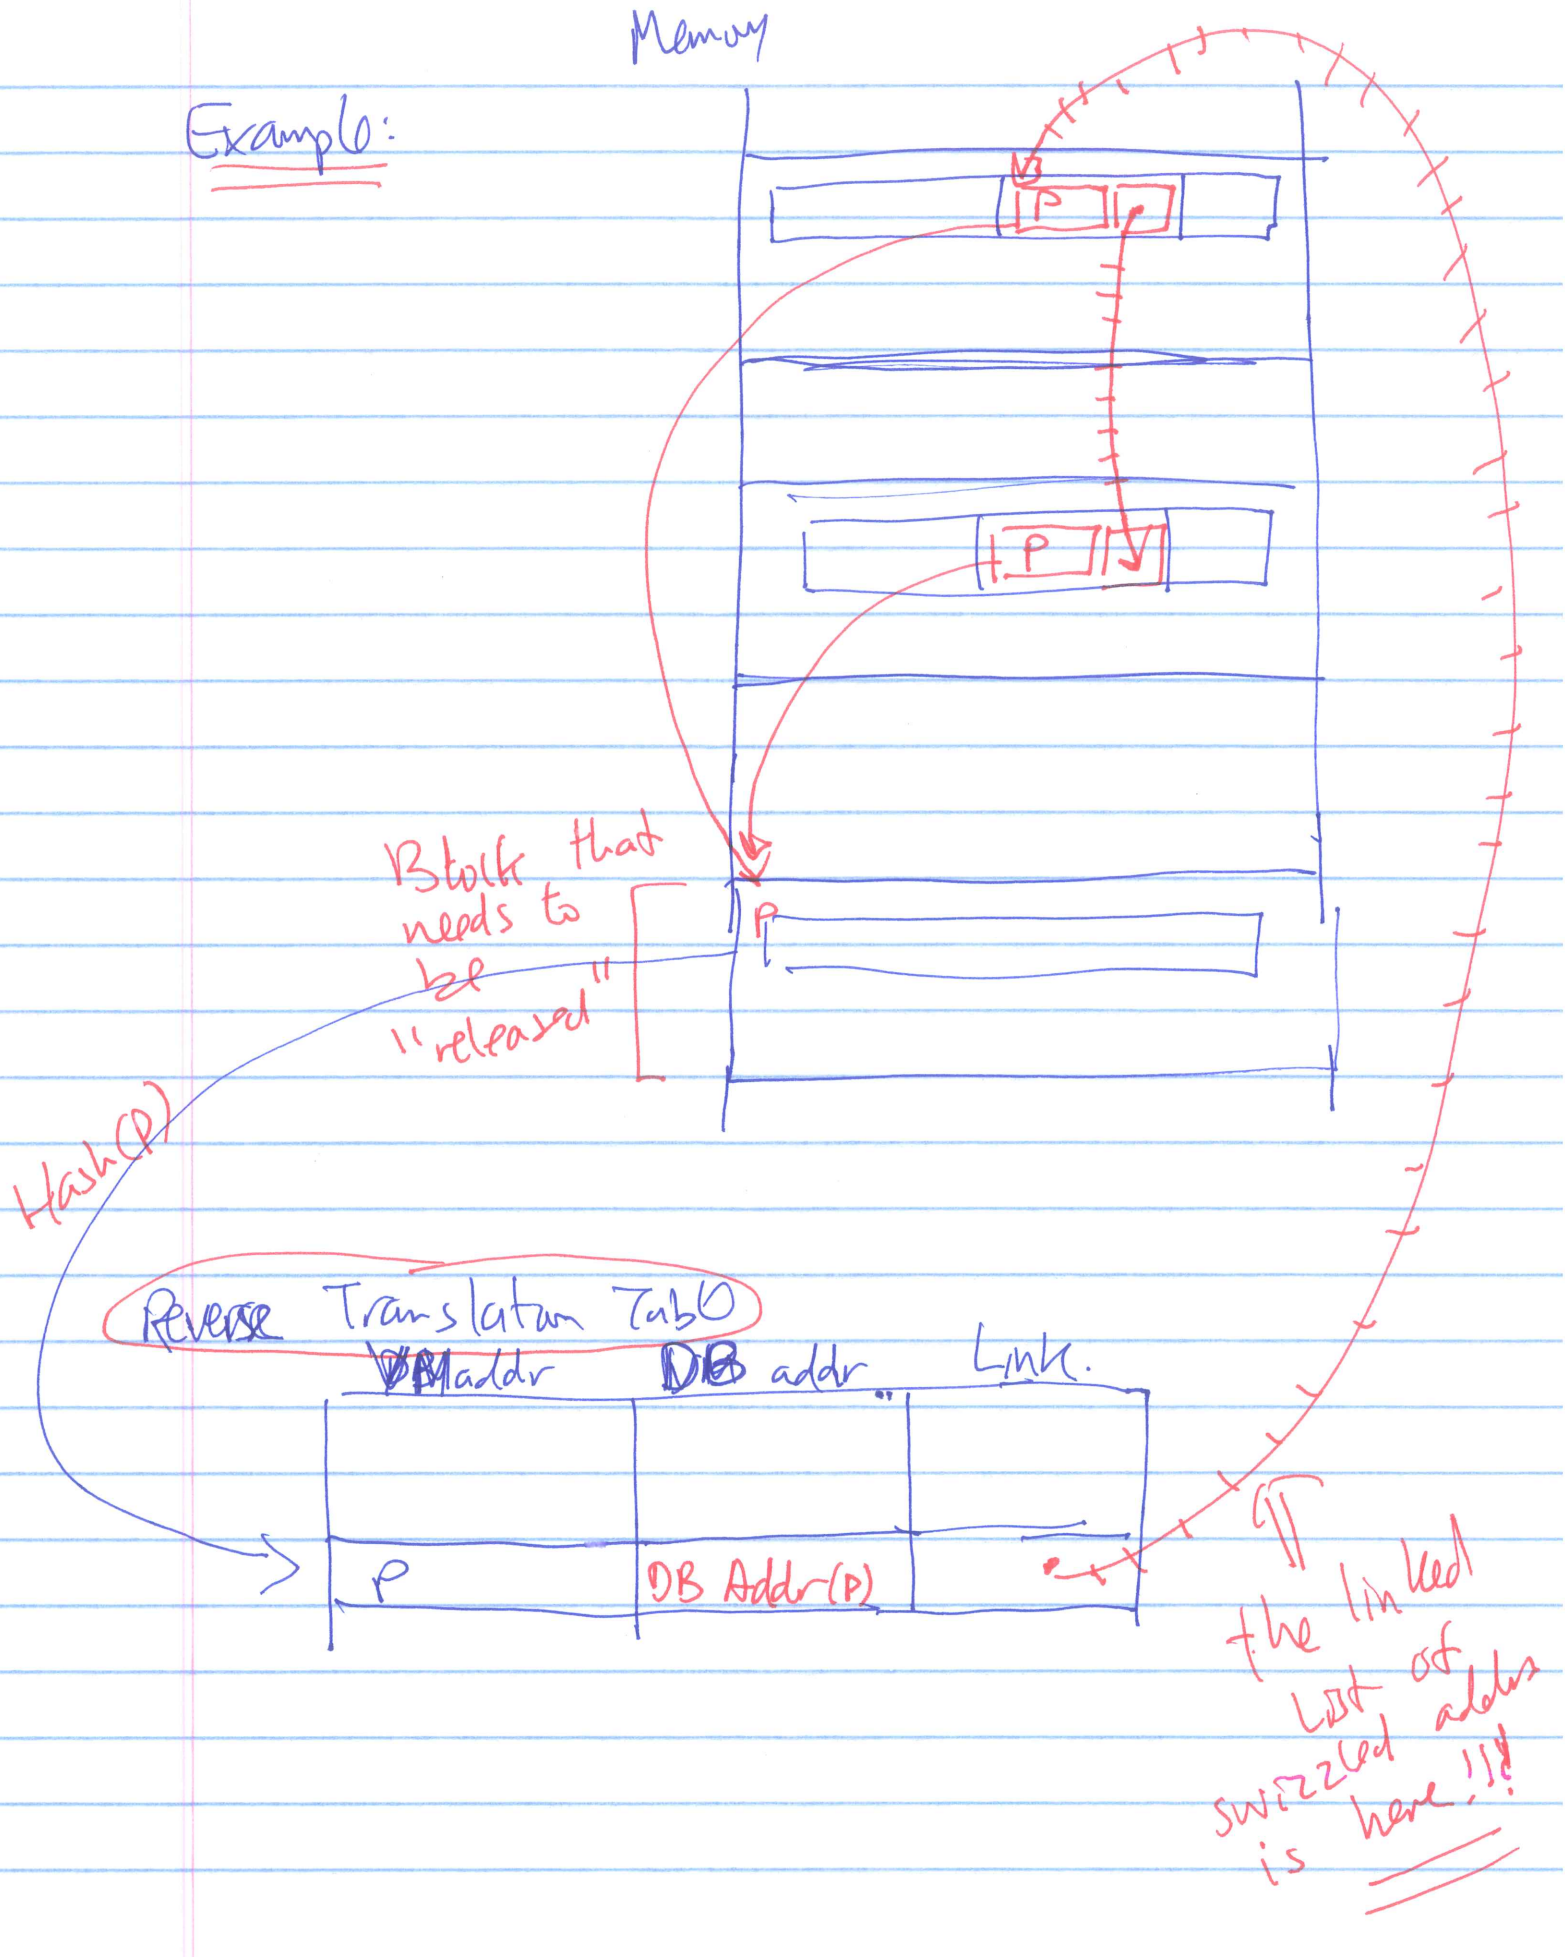
Memory
Example:
P
P
P
Block that needs to be "released"
Hash(P)
Reverse Translation Table
DB addr
DB addr
Link.
P
DB Addr(P)
the linked list of swizzled addrs is here!!!