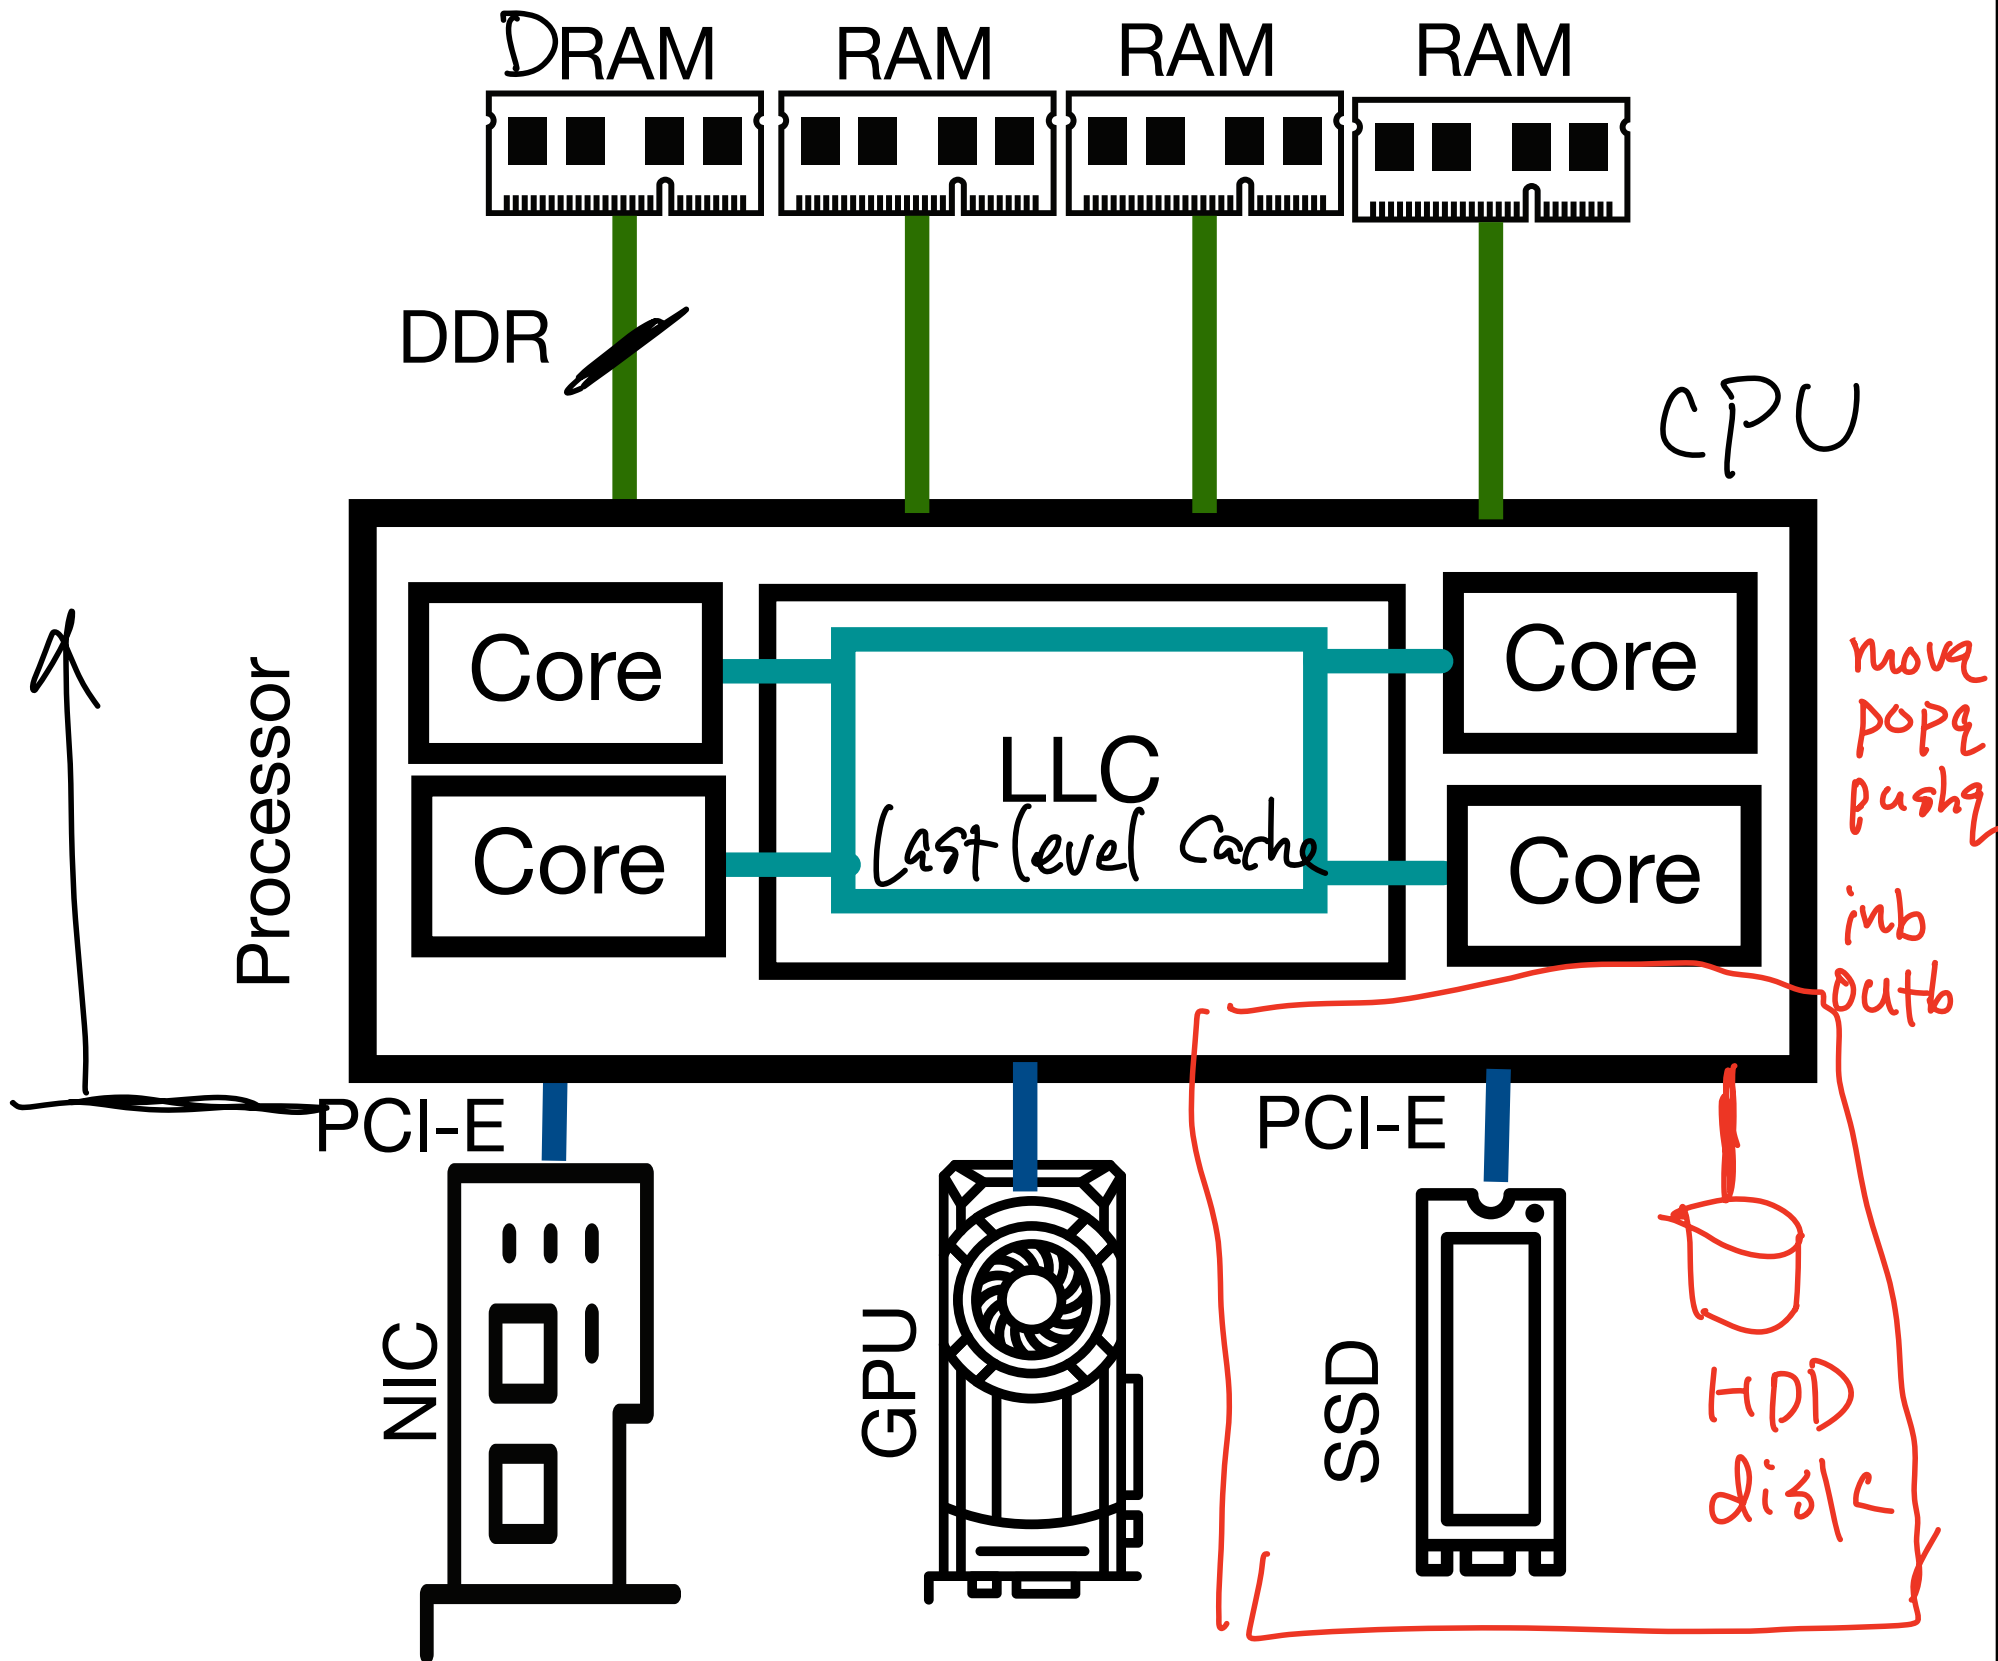


Machine



Hard disk drive organization






Fig 5.11, 5.12 from <http://www.ccs.neu.edu/~pjd/x600-book-v0903.pdf>



A Flash-based SSD (logical diagram)




Fig 44.3 from <https://pages.cs.wisc.edu/~remzi/OSTEP/file-ssd.pdf>

wear-out

Q: What to do to prevent/mitigate the wear-out?

1. I/O architecture and storage devices
2. Intro to file systems
3. Files

- Fancy storage devices
 - Low latency? persistent memory
 - long-term storage? project silica
- Intro to FS

Q: What does FS do?


- read pdf from fs
- deleting files
- Creating files

① provide persistency

② name a seq of bytes (files)

③ give a way mapping from user-friendly names to named bytes (dirs, namespace)

• Q: Where are FSes implemented?




- Files Q: what is a file?
 - user: a seq of bytes.
 - FS: a collection of blocks.

file:

$\langle \text{file}, \text{offset} \rangle \rightarrow \text{block address}$

per-file: [SOMETHING] (inode) \rightarrow file mapping




offset: $4096 + 1 \Rightarrow$ block # 100

- FS parameters
 - small files vs. large files
 - disk utilization (metadata, fragmentation)
 - access patterns
 - sequential access vs. random access

Q: How will you design this SOMETHING?


Candidates:

- contiguous allocation (extent-based)



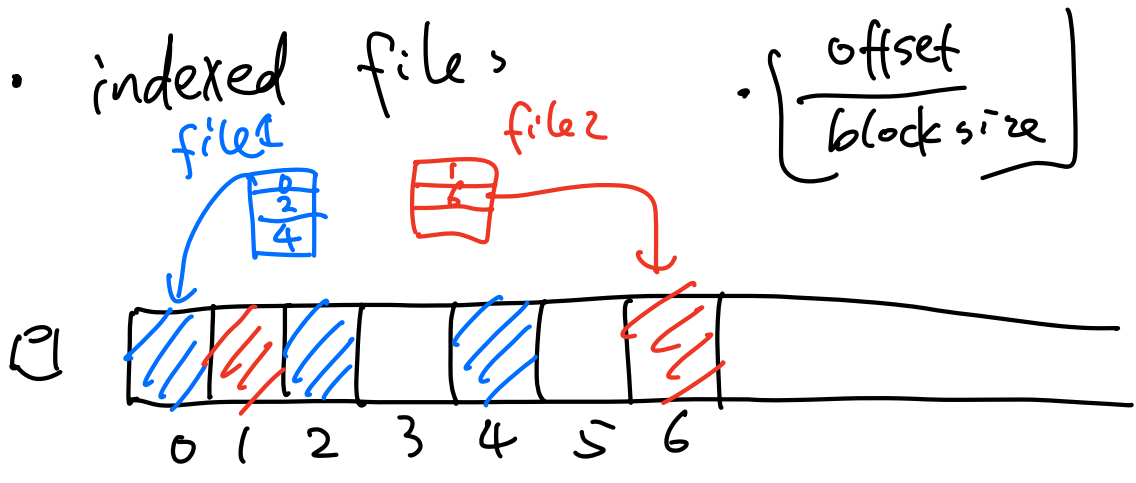
$(base, size)$ \uparrow file1 file2 file3
 $\left[\frac{offset}{block\ size} \right] + base$

• linked-base files.



• indexed files

$$\frac{4GB}{4KB} = 2^{20}$$



• proposal

