Design of HDD and SSD hybrid architecture.

Overview

- Advantages and disadvantages of both HDD and SSD.
- Motivation behind proposed architecture.
- Detailed design of proposed hybrid architecture.
- Conclusion.
- Future work.

Overview

- Advantages and disadvantages of both HDD and SSD.
- Motivation behind proposed architecture.
- Detailed design of proposed hybrid architecture.
- Conclusion.
- Future work.

Advantages of SSD over HDD

SSD	HDD
It has an array of semiconductors organized using ICs.	It has various mechanical moving parts like spindle, stepper motors, rotating disks and read-write heads.
Because of absence of mechanical parts, data transfer to and from becomes very fast and easy.	Data transfer rates are slower as compared to SSD.
Lighter and portable.	Heavier and very large storage capacity HDDs are not very portable.

Presence of only semiconductors make SSDs durable and quieter.	Presence of mechanical moving parts makes the HDD unreliable and noisy.
Random access time is very much lower than the HDDs.	Random access time is quite greater than SSD. Because alignment of mechanical moving parts take substantial amount of time and lowers the access speed.

Advantages of HDD over SSD

SSD	HDD
For any overwrite, the data must be first erased and then the writing can be done on the pages of the block.	In-place of data concept works with HDD and hence no need of erasure operation before write operation.
Cost/bit is very heavy.	Cost/bit is very less as compared to HDD.

Overview

- Advantages and disadvantages of both HDD and SSD.
- Motivation behind proposed architecture.
- Detailed design of proposed hybrid architecture.
- Conclusion.
- Future work.

Motivation

- By considering the advantages and disadvantages of the storage devices, designing hybrid architecture can be an excellent way to gain maximum benefits.
- Rather than completely moving HDD and replacing with SSD, SSD is introduced as one of the layers in memory hierarchy.
- One of a motivation behind the proposed architecture is some basics of SSD and HDD.

Some basics of SSD and HDD

- HDDs give excellent performance when used in sequential writing mode.
- Many of the writes seen by block devices are overwrites of a small popular blocks. E.g. desktop workloads.
- Performance of SSDs degrade when too many number of writes performed on SSD. Hence, to increase the lifetime of an SSD, less number of writes need to be performed.

(contin)

- SSDs work excellent for sequential as well as random readings.
- Cost of HDDs are very much lesser as compared to SSDs.

All above reasons form a motivation behind the proposed hybrid architecture.

Overview

- Advantages and disadvantages of both HDD and SSD.
- Motivation behind proposed architecture.
- Detailed design of proposed hybrid architecture.
- Conclusion.
- Future work.

Overview of design

- For write operation:

  1. Disk cache is the embedded memory in a HDD which acts as buffer.

  2. All writes are appended to this disk cache in a log structured manner.

  3. This data is eventually flushed to SSD periodically.

(contin)

- For read operation:

  1. If data is not found from a storage device of one memory hierarchy level, then it is called as a MISS and search is continued in another storage device present in next level. This is called as look up penalty.

  2. Proposed architecture performs a simultaneous search in both of the storage devices i.e. HDD as well as SSD and tries to completely avoid any look up penalty.

  3. Such a technique is called as “look aside” technique.

Detail design description

- For write operation:

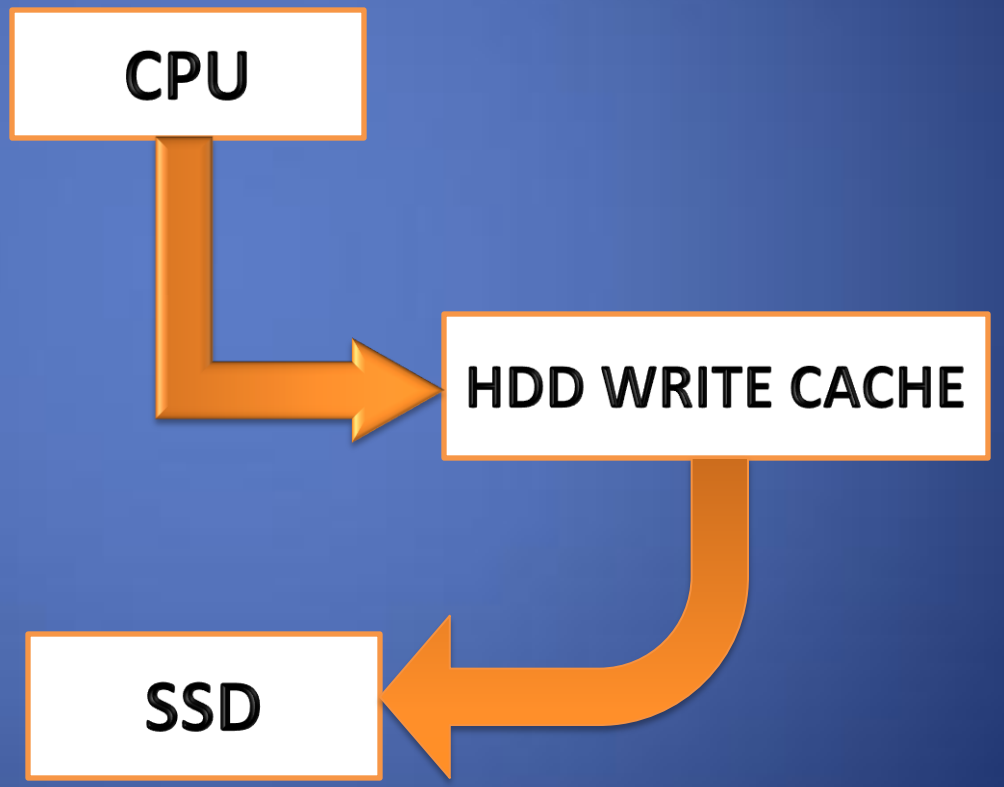

- The proposed architecture uses a hard disk as a persistent write cache.
- The writes are never performed directly on the SSD. All writes are logged sequentially on the HDD write cache and eventually migrated to the SSD, preferably before reading the written block.
- Using HDD as a write cache to coalesce overwrites can reduce the number of writes that are seen by SSD.
- Writes on write cache are in a sequential manner and hence HDD works faster.

- For read operation:

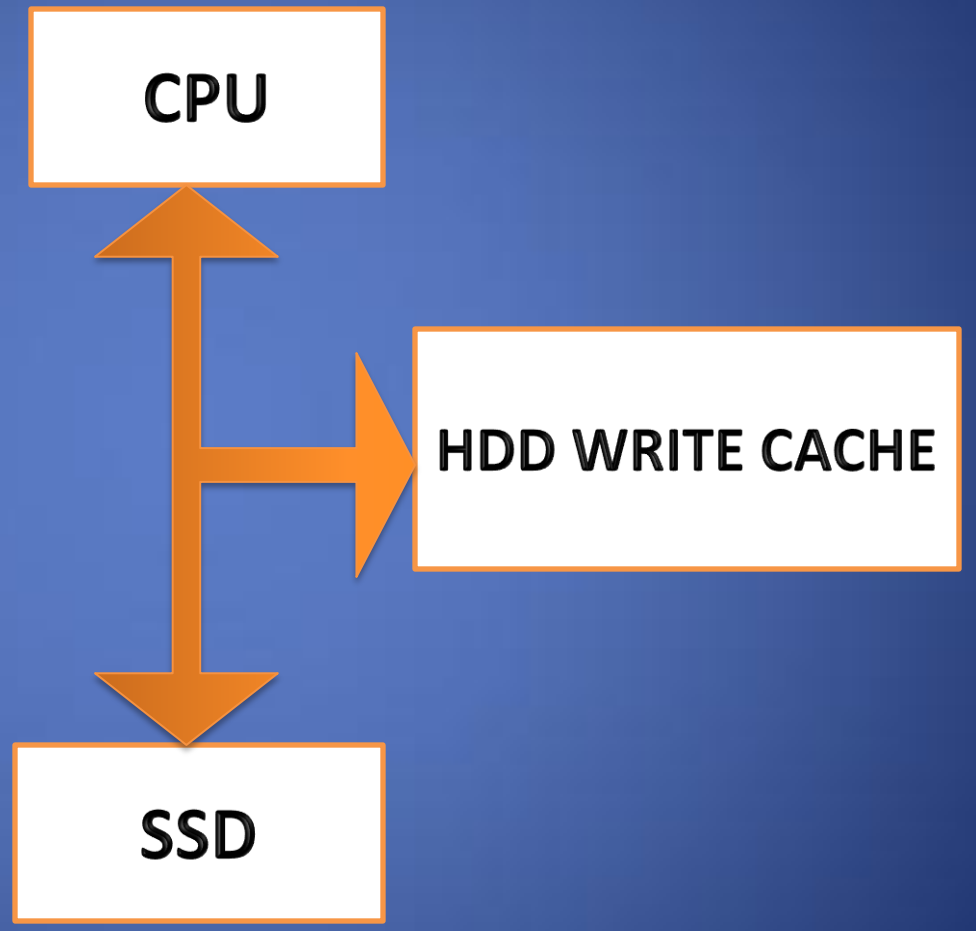

- To avoid look up penalties, CPU performs a search on both SSD as well as HDD.
- If the data is found in SSD, then excellent performance can be achieved.
- If data is found in HDD, then it will suffer from slower and costlier reads called as read penalties.
- To avoid read penalties, the data must be flushed to SSD periodically preferably it is read from the HDD. (contin)

- Look aside technique works best for single processor system.
- But as CPU will be always using a bus, the other bus masters have to wait for it to get free. Hence this technique gives low performance when many other processors are there in the system and want to use bus.

Triggers

For the system to give an excellent performance, read penalties should be least. To have least read penalties, migration policies need to be implemented efficiently.

- Timeout triggers: This trigger fires if certain time elapses without migration. Example, the timeout value of 10 minutes will ensure that no write older than 10 minutes will be performed.
- Read-threshold trigger: This trigger fires when

(contin)

the measured read penalty since the last migration goes beyond a threshold.

- Migration-size trigger: This trigger fires when the total size of migrating data exceeds a certain size.

The act of migration is very quick and simple; data is simply read sequentially from the HDD log and written to the SSD.

Overview

- Advantages and disadvantages of both HDD and SSD.
- Motivation behind proposed architecture.
- Detailed design of proposed hybrid architecture.
- Conclusion.
- Future work.

Conclusion

The design of the proposed hybrid architecture is motivated by hardware characteristics. In the proposed architecture, HDD serves as the cache for SSD. Reading is done simultaneously to avoid the look up penalties. The careful evaluation of the design proposal shows that it has the potential to exploit SSD performance and also improve its lifetime significantly.

Overview

- Advantages and disadvantages of both HDD and SSD.
- Motivation behind proposed architecture.
- Detailed design of proposed hybrid architecture.
- Conclusion.
- Future work.

Future work

- Different types of SSDs can be used instead of using only SLC based SSDs.
- “Look through” technique can be implemented in multiprocessor environment.

Thank you.