Next-generation Magnetic Recording

Shingled magnetic recording is a low-cost way to increase the areal density of hard disk drives (it is low-cost in the *dollar* sense because it does not noticeably change a drive's manufacturing process).

The analogy that we see repeated in all SMR descriptions is that SMR tracks are overlapped "like shingles on a roof". Traditional HDD density gains have come by reducing both the reader and writer gap sizes. However, we have neared the physical limits of our ability to further reduce the *writer gap width* below its current size, while the *reader gap width* can be further reduced. SMR exports this asymmetry.

- Consecutive tracks are overlapped, but when the next downstream track is written, enough of the previous track is left un-obstructed so that its data can be read (the reader gap width is smaller than the writer gap width)

The result is that writes clobber downstream data. This adds limitations to the drive's support for

- random writes
- overwrites

Types of SMR

- **Host-managed SMR** drives expose important details about the drive's internal organization, giving the user control over data management decisions.
- **Drive-managed SMR** drives do not expose any details about their internal organization. Instead, a drive-managed SMR drive exposes the standard block interface, much like an SSD.

Both types of SMR require *someone* to implement a logic to manage data so that it is not lost. This is often done by maintaining a mapping from the logical blocks to physical blocks in a translation layer (like FTLs on SSDs). At a high level, there are two considerations for implementing an SMR translation solution:

- The drive layout
- The mapping scheme

Drive layout

- The disk is divided into contiguous regions of overlapping tracks called *bands*. The last track does not have any tracks that overlap it, and it is often called a *guard track*.
- The disk often has a region where random reads and writes are legal (no tracks are overlapped). This region is often used as a *persistent cache*.

Bands are typically *O(hundreds of MiB)*.

The persistent cache is typically *O(ones to tens of GiB)*

LBA -> PBA Mappings

Logical to physical block mappings can be *static* or *dynamic*.

- A static mapping requires that updates to a PBA are made by reading an entire band, modifying the band in memory, and writing the whole contents of the band as a unit.
- Dynamic mappings can take many forms. In my opinion, this is where the exciting research lies because there is no one "best" approach.

Drive-manged SMR drives typically employ a static mapping, and a persistent cache is used to buffer updates.

Questions

If you squint your eyes, the "SMR problem" is just like the "SSD problem". However, there are several differences that impact the designs that we can expect to employ.

1. What are the physical differences between SMR drives and SSDs?
2. What are the target use cases for SMR drives and SSDs?
3. What are the scales of SMR drives and SSDs (in terms of capacity)?
4. Are there concerns that an FTL must deal with that an SMR drive does not (or vice versa)?
5. What FTL approaches did we discuss that would work well on an SMR drive? Would not work well?
6. Where do you think HM-SMR drives will be used (if at all) moving forward? DM-SMR drives?