Big Data Storage 3: New Storage Technologies

Shiow-yang Wu (吳秀陽)

CSIE, NDHU, Taiwan, ROC

Current Technologies Not Good Enough

- Data are growing fast. (175ZB/year by 2025)
- Traditional HDD technology is advancing slowly.
- Current storage technologies have limited capacities and lifespans

Storage Type	Capacity	Real Lifespan	Limitations
Magnetic Tape	GB ~ 15 TB	10 ~ 20 years	Sensible to humidity/temperature, susceptible to wear and tear
CD & DVD	GB	2 ~ 5 years	highly susceptible to damage (scratch, …)
Blue-Ray(BD)	25~100 GB	5 ~ 20 years	susceptible to scratch, temperature, sunlight
M-Disc	≈DVD/BD	1000 years	Must have M-Disk drive to write
HDD	~ 32/40 TB	3 ~ 5 years	Power, noise, heat, fail mechanical parts
Flash storage	GB ~ 100 TB	3 ~ 5 years	Expensive, limited write cycles

3-2-1-1 Backup Rule

- Data loss and ransomware demands the 3-2-1-1 backup straregy.
- **3**: Keep three copies of you data
- **2**: Store your backups on two different types of media
- **1**: Keep one copy offsite (in the cloud or secure storage)
- **1**: Ensure one copy of your data is immutable (in write-once-read-many format that can't be altered or deleted)

High Backup and Maintenance Cost

- Backup Or Die!
- All storage technologies are not REALLY durable and sustainable.
- Periodic maintenance is equally important (for backup, recovery, and storage lifespan)
- Both are even more important for big data applications.

New Technologies Needed

- Much bigger, safer and cheaper.
- Much more durable and sustainable.
- Much more flexible and scalable
- Much smaller form factor

DNA Storage

- Paper: Dave Landsman and Karin Strauss. The DNA Data Storage Model. *IEEE Computer*, vol. 56, pp. 78-85, July 2023.
- Using synthetic DNA as a data storage medium.
- High volumetric data density
- Highly durable and sustainable
- Dramatically lower total cost of ownership

Biotechnology Advancement

- Decades of investment in molecular-level technologies for medical and life sciences
- We can now construct and read synthetic DNA, base by base(鹼基).
- DNA is made of **4** types of bases (Adenine, Cytosine, Guanine, and Thymine, or ACGT) linked covalently into a chain (a DNA strand)
- Store genetic information

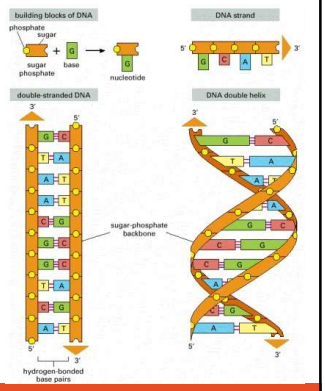

DNA Storage Idea

- Encode digital data into a sequence of bases (adenine, guanine, cytosine, and thymine, or AGCT).
- Write the sequence as a set of corresponding DNA molecules (synthesis).
- Store the molecules.
- Prepare them for reading (retrieval).
- Read them back as a sequence of bases (sequencing).
- Decode the original digital data.

DNA Storage System

- Encoding/decoding are performed in the electronic domain and translate bits to/from bases.
- Synthesis/sequencing are the interfaces to/from the molecular domain, creating and reading DNA sequences.

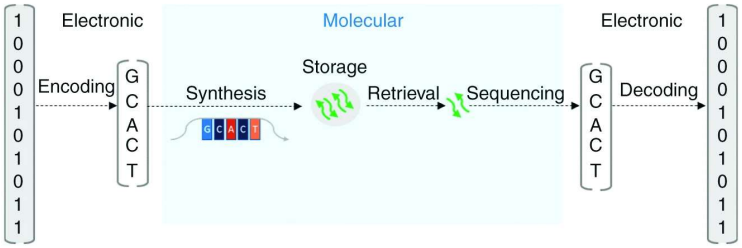

DAN Storage Advantages

- Can store 215 petabytes/gram
- Theoretical maximum density is 455 exabytes/gram
- **Advantages**:
  - Storage media (DNA molecules) are manufactured during write operations. Highly flexible and scalable.
  - Much higher volumetric density
  - DNA has unique error characteristics (eg. insertions/deletions may occur) which can be resolved in the encoding/decoding process
  - Last for thousands of years

Glass Storage

- Microsoft's Project Silica store data in quartz glass to hold 7TB in a DVD-size platter.
- Low cost, durable WORM (Write-Once-Read-Many) media
- EMF-proof (safe in electromagnetic field)
- Lifetimes of tens to hundreds of thousands of years
- Highly sustainable w/o costly cycle of periodically copying data to new media

Glass Storage Mechanism

- Ultrafast femtosecond lasers (飛秒雷射) to write
- Polarization-sensitive microscopy using regular light to read
- Impossible to accidentally overwrite data during read (not enough power to modify the glass)
- Much higher volumetric data densities than magnetic media.
- With a novel, low-power design for the media library through robotics and mechanics.

Silica Images

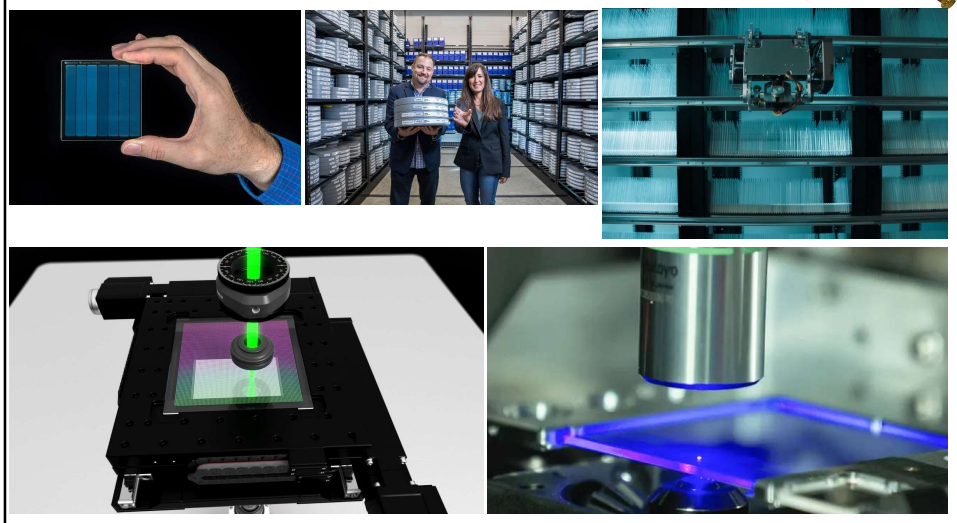

Silica References

- Patrick Anderson, et al. Project Silica: Towards Sustainable Cloud Archival Storage in Glass. *SOSP '23: Proceedings of the 29th Symposium on Operating Systems Principles*, October 2023, pp. 166–181.
- Project web site: https://www.microsoft.com/en-us/research/project/project-silica/

Next Generation of HDDs

- **OptiNAND** (WD): 50TB (2030)
- **UltraSMR** (WD): 28TB (2023)
- **Triple-Stage Actuator** (WD): 30TB+
- **HAMR** (Seagate): 32TB (2023), 40TB follows
- **MACH.2** (Seagate): 40TB+(2024), 50TB+(2025-6)
- **MAS-MAMR** (Toshiba): 30TB(2024)
- **NVMe** HDDs
- Heading toward 50TB HDD