Holographic Optical Storage for the Cloud?

Grace Brennan, Nathanael Cheriere, Jiaqi Chu, Jannes Gladrow, Douglas Kelly, Giorgio Maltese, Dushyanth Narayanan, Greg O'Shea, Alan Sanders, Xingbo Wu, Mengyang Yang, Benn Thomsen Microsoft Research

benn.thomsen@microsoft.com

Abstract: We assess the performance and energy efficiency of an end-to-end Holographic optical storage system to determine if this technology can cost effectively solve the access rate challenges in online cloud storage applications.

1. Introduction

Today the bulk of online data in the cloud is stored on Hard Disk Drives (HDD) as these provide the lowest cost storage with acceptable access rates. However, as hard disk drive capacity continues to scale, albeit more slowly than in the past, the data access rates per unit of stored data, measured in IOPS/TB, are declining. The access rates provided by the latest generation of HDDs are already too low to support all online workloads in the cloud. This is necessitating the deployment of more expensive hybrid storage solutions that use a combination of HDD and Solid State Drives (SSD) to achieve the required access rate performance [1].

Holographic optical storage has long been touted as a technology that could provide cost effective high-capacity storage with fast access rates enabled by the inherent parallelism that optical systems offer. Whilst there have been many impressive demonstrations of the storage density that rewritable holographic optical storage can achieve [2] very little has been done to assess whether the combination of storage density, access rate and energy efficiency performance is sufficient to build a cost-effective cloud storage system. In rewritable holographic storage media this assessment is made even more challenging by the fact that the stored data within the media is degraded by each write or read operation. For a holographic storage system to be competitive with the incumbent storage technologies it needs to have: sufficient storage density for the media to be cost effective, a spatial multiplexing scheme that does not rely on mechanical movement of the media to give fast access over a large capacity device and an energy efficient scheme for managing the data degradation. To evaluate the performance for online cloud storage we built and evaluated an end-to-end holographic storage system with these features.

We show that by jointly optimizing the media properties and the write and read optical energy profiles used to access the data in combination with a machine learning based decoder to enable 2 bit/symbol we were able to obtain a net storage density of 9.6 GB/cm3 an increase of 1.8× over previous work [2]. Furthermore, we developed a zone-based garbage collection and data refresh scheme that handles the data degradation of rewritable holographic media in an energy efficient manner achieving a net optical energy of 16nJ/bit at the media. We also propose a spatial multiplexing scheme that provides fast access without mechanical movement. However, we find that the overall system suffers from low energy efficiency and low overall device storage density due to the trade-offs between the media efficiency, spatial multiplexing and the refresh scheme. Therefore, we argue that Holographic optical storage is still not a cost-effective online storage technology for today's cloud storage applications.

2. Media performance, storage density and energy efficiency

The end-to-end holographic storage system testbed, shown in fig. 1, maps the storage workload requests onto a zone based holographic storage system and manages the data lifetime in the media. We use rewritable Iron doped Lithium Niobate (Fe:LiNb0₃) media to store data as optical holograms. The data is written into the media as pages within a zone using angle multiplexing in the 90° configuration [3] at a wavelength of 532nm. To scale the capacity across multiple zones we developed a spatial multiplexing concept, described in section 3, to provide fast access across zones. The individual zones are erased for rewriting with exposure to UV light. In a holographic optical storage system with rewritable media write and read operations to the same zone result in the degradation of all data within the zone. Whilst for a single operation this degradation is small and can be ignored for a viable storage system this needs to be managed. We manage the zone life cycle by optimizing the optical energy profile that is used to write each page into the zone and the energy for each page readout operation in a zone to jointly minimize the required energy and data degradation. Before the data degradation within a zone reaches the point at which the pages are unreadable the valid pages are read out and written into a fresh zone and then the zone is erased and reclaimed as shown in fig. 1(b). This additional refresh overhead is also included in our energy and performance assessment.

To optimize the write and read energy profiles it is necessary to characterize the media write sensitivity and erasure properties over the entire "zone life cycle" – all the writes and reads between one UV erase of a zone and the next. To do this we use the testbed, shown in fig.1(a), to write 500-1000 pages into a single zone and then read these back 10,000s of times. We used these media characterization measurements to evaluate different media compositions (Fe doping and annealing, to adjust Fe2+ concentration) and extended the typically used exponential media erasure model to a stretched exponential model [4] which gave a better fit to the experimental data and resulted in energy performance gains. The best media composition has high sensitivity and low erasure, but typically there is a trade-off between the two. The best media overall was chosen by combining the media model with an endto-end optimizer to find the best combination of storage density and energy efficiency. We also used this framework to determine the optimum Fe doping level and annealing to improve the media performance. The optimizer was then used to generate write and read energy profiles to maximize the storage density. Using these energy profiles with the angle multiplexing scheme we were able to store 705 pages at net page size of 73kB in a storage volume of 5.4mm3 $(1.6 \times 1.6 \times 2.1 \text{ mm})$ to give a net density of 9.6GB/cm3 (note: the net density includes all the overheads e.g. error correction, the raw density here was 16.8 GB/cm3). The data pages were encoded using 4PAM and decoded using machine learning based UNET type decoder [6]. We achieved an optical energy efficiency of 16nJ/bit using a 0.03% Fe doped crystal with non-optimal annealing with an Fe2+ concentration of 4.11×1017 cm-3. We then used the experimental results from this and characterization of other media compositions with the media model to estimate that the upper bound with an optimally doped and annealed media is around 40 GB/cm3 equating to 5% of the so called $1/\lambda 3$ limit. However, even with the optimizations that we have presented the required 16 nJ/bit optical energy at the media is still significantly higher than the wall-plug energy of the competing technologies: HDD 3-5nJ/bit and SSD 1nJ/bit. Wall-plug energy for the holographic system will be even larger depending on: the additional optical losses e.g. in the spatial MUX, the laser electrical to optical efficiency and power consumption of SLM and camera.



Fig. 1. (a) End-to-end holographic optical storage system, consisting of: media management, angle multiplexing within a single zone and spatial multiplexing across zones, with machine learning based decode. (b) Managed zone life cycle with optimized write and read energy profiles and refresh.

3. Spatial Multiplexing

To amortize the cost of the optical components (laser, spatial light modulator and camera) we need to multiplex them over many zones. The key challenge is to do this without physically moving the media: the mass of the media makes physical movement incompatible with high access rates. We developed new active optical routing networks, shown in fig. 2(e, f), that use relay imaging, waveguides and solid state switching to transport 2D images to different zones in the media with fast (~1ms) switching times. The network concepts consist of four key components: (b) a 4-lens image relay system for lossless propagation of images through 90° bends and through the media, square waveguides for compact transport over longer distances, (c) liquid crystal polarization switches for fast switching to actively route the images through the network and (f) a spatial coherent receiver to recover the transmitted data and digitally compensate for optical distortions that occur in the network.

In contrast to 4f optical systems (fig. 2(a)) typically used for holographic recording, where the physical extent of the image plane is much larger than that of the Fourier plane in the storage media, in this system we use a 4-lens system (fig.2(b)) that balances the physical extent of the image and Fourier planes to maximize the fill factor in both the network and the media when multiplexing across adjacent zones. The 4-lens scheme is also used to route the light around bends without loss of information. With a fill factor optimized 4-lens design and a raw data page size of 65kB, (N2= 512×512 , p=6µm symbol pitch, 2bit/symbol) the optimal cross-sectional width (w=Np) of the routing



Fig. 2. Spatial multiplexing. (a) Typical 4f image relay. (b) 4-lens balanced image relay. (c) Fast polarization switch. (d) Coherent receiver (e) Low loss binary tree network. (f) Compact linear network.

To scale to this number of zones whilst minimizing the overall optical insertion loss we developed a binary tree routing architecture shown in fig. 2(e) where the loss scales logarithmically with the number of zones (Z). The loss in this network is dominated by the 95% transmission of the polarization switch. For the paths in this network that are longer than the maximum 4-lens relay length (L=2Np2/ λ) 70mm we used 3mm square cross-section waveguides to route the images. Whilst the waveguides provide a compact means of transport they distort the transmitted images, to undo the distortions at the receiver we used a homodyne coherent receiver [5] (fig. 2(d)) to measure the received wavefronts. The data was then digitally recovered using a differentiable mode-based waveguide model to invert the channel, combined with a UNET that was jointly trained using machine learning to decode the data images [6]. We demonstrated the transport of data pages over 75mm waveguides with a net page size of 46kB [6].

However, we find that there is a fundamental trade-off between path loss and physical volume of the binary network. The three networks required for the data write, the camera readout, and the reference beam require a volume of approximately $6\log_2(Z)$ the media volume. This makes the design uncompetitive on rack density, e.g. at 1024 zones the optics of the network is 60x the media volume. To address this issue, we developed a simpler and volume-optimized linear network (fig. 2(f)) where the volume overhead is 3x regardless of zone count, allowing for competitive rack densities. However, optical path loss is linear in zone count, and beyond 150 zones the higher laser cost due to the increased power requirements outweighs the gains from density. For our achievable page and zone sizes this gives 150GB of raw storage per set of optics, insufficient to amortize the optics cost.

- [1] A. Chatzieleftheriou, I. Stefanovici, D. Narayanan, et al., "Could cloud storage be disrupted in the next decade?" HotStorage 20 (2020).
- [2] G. W. Burr, et al., "Volume holographic data storage at an areal density of 250 gigapixels/in²," Optics Letters 26, 444-446 (2002).
- [3] Hans J. Coufal, Demetri Psaltis, Glenn T. Sincerbox, Holographic Data Storage, (Springer, 2000)
- [4] Nee, I., et al., "Role of iron in lithium-niobate crystals for the dark-storage time of holograms," J. Appl. Phys. 88, 4283-4286 (2000).
- [5] A. C. Urness, W. L. Wilson and M.R. Ayres, "Homodyne detection of holographic memory systems," Optical Data Storage 9201 (2014).
- [6] J. Lim, J. Gladrow, D. Kelly, et al. "High-bandwidth Close-Range Information Transport through Light Pipes," arXiv:2301.06496 (2023).