Future Generation Computer Systems ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Bridging data-capacity gap in big data storage

Wasim Ahmad Bhat

Department of Computer Sciences, University of Kashmir, India

HIGHLIGHTS

- Introduces the working principles of three emerging storage technologies, i.e., Optical storage, DNA storage and Holographic storage.
- Evaluates the advances received in storage density, throughput and lifetime of these three emerging storage technologies.
- Quantitatively compares these advances with the trends and advances in current storage technologies like HDD, SSD and Tape.
- Discusses the implications of adopting these emerging storage technologies, evaluates their prospects, and highlights the challenges.

ARTICLE INFO

Article history:
Received 30 June 2017
Received in revised form 15 November 2017
Accepted 31 December 2017
Available online xxxx

Keywords:
Big data
Data-capacity gap
Optical storage
DNA storage
Holographic storage
Magnetic storage

ABSTRACT

Big data is aggressive in its production, and with the merger of Cloud computing and IoT, the huge volumes of data generated are increasingly challenging the storage capacity of data centres. This has led to a growing data-capacity gap in big data storage. Unfortunately, the limitations faced by current storage technologies have severely handicapped their potential to meet the storage demand of big data. Consequently, storage technologies with higher storage density, throughput and lifetime have been researched to overcome this gap. In this paper, we first introduce the working principles of three such emerging storage technologies, and justify their inclusion in the study based on the tremendous advances received by them in the recent past. These storage technologies include *Optical data storage*, *DNA data storage* & *Holographic data storage*. We then evaluate the recent advances received in storage density, throughput and lifetime of these emerging storage technologies, and compare them with the trends and advances in prevailing storage technologies. We finally discuss the implications of their adoption, evaluate their prospects, and highlight the challenges faced by them to bridge the data-capacity gap in big data storage.

1. Introduction

Big Data is large in volume, complex in structure, and aggressive in its production. The value promised by big data has been envisioned across all the fields. However, the challenges faced by the technology to deliver that promise are still being discussed, and the solutions are yet to be finalized [1]. At the same time, integration of Internet of Things (IoT) with Cloud computing is eminent [2]. IoT is expected to grow to 30 billion units by 2020 making them one of the main sources of big data [3]. Consequently, it will have a massive impact on volume, velocity, and variety (and other axes) of big data. Data centres that are mainly based on magnetic storage technology, with petabyte (PB) and even exabyte (EB) capacities have proved to be the core platforms for cloud computing and big data storage [4]. However, there is a growing gap between the volume of digital data being created and the extent of available storage capacities. As an example, in a report prepared by International Data Corporation (IDC) [5], it is forecast that the amount of data generated globally will reach 44 zettabytes (ZBs) in year 2020. This forecast is based on the estimation that the information generated worldwide doubles every two years. IDC's report also noted that the rate of production of data continues to outpace the growth of storage capacity. In 2013, the available storage capacity could hold just 33% of the digital universe, and by 2020, it will be able to store less than 15% [6]. Recent estimate of IDC suggests that 13 ZBs of 44 ZBs generated in 2020 will be critical, and should be stored. However, since the storage capacity available at that time will be able to hold only 15% of 44 ZBs, a minimum *data-capacity gap* of over 6 ZBs (which is nearly double all the data produced in 2013) is expected in year 2020.

Prevailing storage technologies are increasingly challenged by their limited storage density and throughput as well as the shortcomings associated with energy consumption, capacity footprint, lifetime, and other like features. Accordingly, storage technologies with greater storage densities, higher throughput, lower energy consumption and longer lifetimes are in high demand to support

E-mail address: wab.cs@uok.edu.in.

https://doi.org/10.1016/j.future.2017.12.066
0167-739X/

big data centres. Even though many next-generation mass storage technologies have been actively researched [7], three emerging storage technologies have received tremendous advances in the recent past, and are emerging as the next-generation storage technologies for big data storage.

The purpose of this paper is to evaluate the recent advances received by these three emerging storage technologies, i.e. *Optical data storage*, *DNA data storage* & *Holographic data storage*, for their potential to bridge the data-capacity gap in big data storage. The evaluation criteria adopted for the study include the parameters that are critical to any emerging storage technology for its ability to overcome the data-capacity gap. These include storage density, throughput and lifetime. The paper also discusses current trends and advances in storage density, throughput and lifetime of current storage technologies, and quantitatively compares them with those of the emerging storage technologies. The paper further investigates the implications of the adoption of these emerging storage technologies, evaluates their prospects, and highlights the challenges faced by them.

The rest of this paper is organized as follows. Section 2 introduces the working principle of each of these emerging storage technologies and justifies their inclusion in the study. Section 3 evaluates and compares the advances in storage density of the three emerging storage technologies and current storage technologies, while as Section 4 does the same for throughput, and Section 5 for lifetime. Section 6 highlights the implications of the adoption of these emerging storage technologies, and Section 7 discusses their prospects and challenges. Finally, Section 8 discusses basic features of other emerging storage technologies, and Section 9 presents the conclusion.

2. Preliminaries

This section introduces the working principle of Optical, DNA and Holographic storage technologies, and based on the technological-breakthroughs that are relevant to big data centres justifies their inclusion in this study.

2.1. Optical data storage

Optical data storage (ODS) first emerged as compact discs (CDs) in 1980s, and soon its green features, high storage capacity and high energy efficiency became apparent. The technology continued to progress thereafter to deliver higher capacity discs in the form of digital video discs (DVDs) and Blu-ray discs. Unfortunately, these optical discs can only record information in a diffraction-limited region within a layer beneath the surface of the disc. Therefore, their maximum capacity is limited to a few tens of gigabytes (GBs). The classical diffraction limit imposes a limit on the size of the optical spot that records digital bits. The optical spot size is directly proportional to wavelength (λ) of the Laser beam, and inversely proportional to the effective numerical aperture (NA) of the optical head. This implies that the storage density of an optical disc is proportional to $(NA/\lambda)^2$ [8]. Fig. 1 shows how these parameters affected the storage capacity of three generations of optical discs. Over a period of last three decades, a lot of research and development efforts have been invested in exploiting the volume of recording media for volumetric or multilayer memories [9–11]. Unfortunately, the diffraction limit here also limits the theoretical storage capacity to a few terabytes (TBs) per DVD-sized disk [12–14]. This storage density is not enough to overcome the data-capacity gap in big data storage.

Nevertheless, studies based on the new photonic principles have led to the development of artificial materials with negative refractive indices [15–17], nano-optical circuits [18,19], nanoscale light-emitting sources [20,21], imaging beyond the diffraction limit [22–24] and super-resolution optical lithography [25,26]. These studies have been disruptive and innovative in their approach, and have demonstrated confinement of light–matter interactions to the nanometre scale. This has paved the way towards breaking or circumventing the diffraction barrier, and increasing storage capacity tremendously by using entirely new nanophotonic approaches.

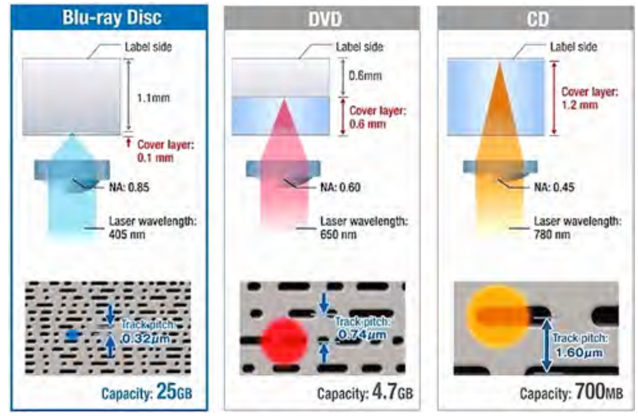

Fig. 1. Basic parameters that affect the storage capacity in ODS.
Source: http://www.cd-info.com/blu-ray/.

2.2. DNA data storage

DNA was identified and advocated as the promising ultra-dense storage technology early in 1960s, and since then, DNA data storage (DDS) has been researched. Studies have demonstrated that DNA has the potential to act as a huge-capacity and longterm digital storage medium for three main reasons; it is incredibly dense (it can store one bit per base, and a base is only a few atoms large), it is volumetric rather than planar, and it is incredibly stable (DNA can survive for hundreds of thousands of years). It is so dense that a human body that typically contains 100 trillion cells is able to store approximately 150 ZBs of data in its DNA.

DNA consists of four types of nucleotides: adenine (A), cytosine (C), guanine (G), and thymine (T). A DNA strand, or *oligonucleotide*, is a linear sequence of these nucleotides. Two single strands can bind to each other and form a double helix if they are complementary: A in one strand aligns with T in the other, and likewise for C and G. Arbitrary single-strand DNA sequences can be synthesized chemically, nucleotide by nucleotide [27,28]. The probability that a nucleotide binds to an existing partial strand at each step of the process is as high as 99%, and is commonly referred to as *coupling efficiency* [29]. On the other hand, DNA polymerase enzymes are used for sequencing purposes, and involves creating a complement strand, using fluorescent nucleotides, against the strand of interest, and the process is commonly referred to as *sequencing by synthesis* [29].

For the purpose of storing binary data in DNA, each base pair can be used to represent 2 bits. Consequently, the method encodes binary data in base 4, which converts a string of n binary bits into a string of $n/2$ quaternary digits. These digits can then be mapped to DNA nucleotide pairs. Therefore, 4 different base pairs (i.e. AT, GC, TA, CG) can encode 4 quaternary digits (i.e. 00, 01, 10, 11). Consequently, following encoding is possible: $00 \rightarrow AT$, $01 \rightarrow GC$, $10 \rightarrow TA$ and $11 \rightarrow CG$. As an example, to store a binary string 00011011, the DNA sequence generated is ATGCTACG. Conversely, while reading the data stored in DNA, nucleotide pairs are decoded into quaternary digits and then translated into binary data. The

Please cite this article in press as: W.A. Bhat, Bridging data-capacity gap in big data storage, Future Generation Computer Systems (2018), https://doi.org/10.1016/j.future.2017.12.066.

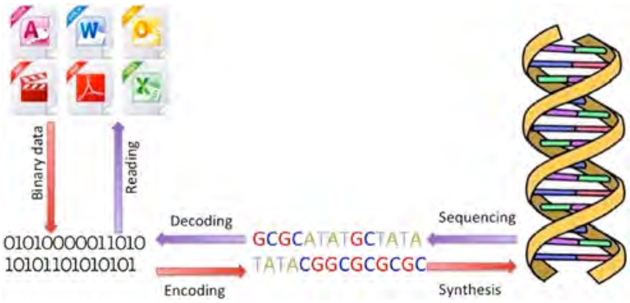

Fig. 2. Basic process of storing and retrieving data in DDS.

basic process of storing and retrieving data in DNA is illustrated in Fig. 2.

Unfortunately, DNA synthesis and sequencing is highly vulnerable to a wide variety of errors (substitutions, insertions, and deletions of nucleotides), with error rates of the order of 1% per nucleotide. Also, nucleotide sequences tend to degrade while stored, compromising data integrity. Nevertheless, encoding schemes that can tolerate errors have been studied. On the other hand, DNA synthesis and sequencing is a time consuming task that hinders throughput. Fortunately, tremendous research efforts have been invested to speed up the process of synthesis and sequencing, and the results are promising.

2.3. Holographic data storage

Holographic data storage (HDS) is based on the idea of volumetric storage in which information is stored throughout the volume of the media and not just on its surface. Magnetic storage technology and conventional-optical storage technology record information bits on the surface of the recording media as distinct magnetic and optical changes respectively. Unfortunately, physical phenomena limit the size of these *bits* on the surface. Too small in size can compromise the integrity of recorded information as storing and retrieving information becomes too difficult. Hence, this restricts the storage density of these technologies to a certain limit. By being volumetric, HDS overcomes this conventional storage capacity limit.

HDS employs multiplexing for recording multiple holograms in the same volume of media. Light from a Laser source is split into two beams — Signal beam that is passed through a *Spatial Light Modulator* for encoding binary data, and a Reference beam that recombines with the modulated Signal beam before both enter the recording media. An optical interference pattern is produced by the overlapped beams that gets imprinted as a hologram in the recording media. Each hologram has a unique volume address that is provided by the Reference beam. Therefore, to record more holograms in the same media, Reference beam properties like angle of incidence, wavelength, phase, and so on, are changed. While retrieving the recorded data same Reference beam that was used for recording a particular hologram is used. Therefore, each Signal beam can be retrieved from its corresponding hologram without any interference from the rest of the holograms. This ensures integrity of the recorded data. The process of recording to and reading from HDS is illustrated in Fig. 3.

The idea of HDS is not new and was conceived decades ago. The research efforts in HDS started in early 1960s [31], and continued through 1970s [32]. Unfortunately, due to unavailability of suitable recording media, and lack of precise and tunable input–output devices, these efforts faced many technical challenges. Nevertheless, with the emergence of low-cost enabling technologies, significant results from long-standing research efforts, and progress in holographic recording materials and multiplexing techniques [33], holographic storage has made significant progress.

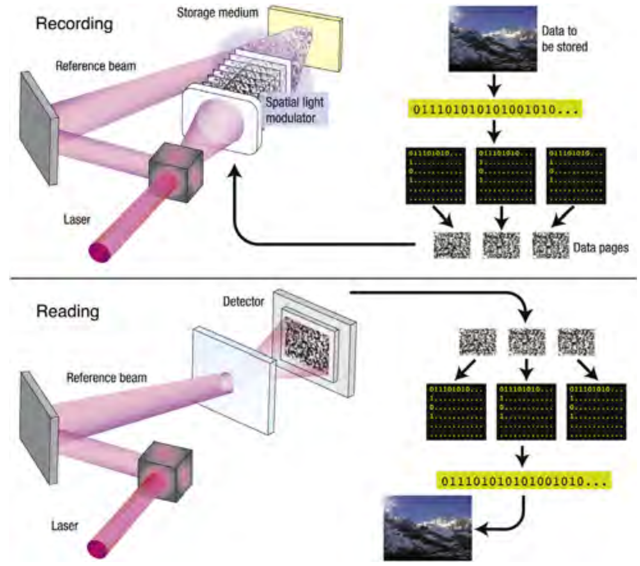

Fig. 3. Recording and reading process of HDS [30].

3. Advances in storage density

Storage density is the main attribute of any storage technology that primarily decides the prospects of its adoption in big data centres. Considering the huge and growing data-capacity gap in big data storage, sustained and significant progress in storage density of a storage technology is necessary to overcome the gap. This section qualitatively evaluates and quantitatively compares the recent advances in storage density of the emerging and current storage technologies.

3.1. Optical data storage

ODS exploits multiple-dimensions of light–matter interaction at nanometre scale to offer huge storage density. Advances in nanophotonics remain the cornerstone of this dramatic increase in storage density of ODS. Nanophotonics can facilitate either encoding of information in physical dimensions, such as those defined by the frequency and polarization parameters of the writing beam, or by achieving 3D super-resolution recording [14]. These approaches hold significant potential to break the conventional storage capacity limit. Nanophotonics allows for sharp colour and polarization selectivity [34]. These characteristics make nanoparticles suitable for the implementation of spectrally encoded memory. As a result, three-colour spectral encoding by using gold nanorods of various sizes has been demonstrated [35]. Similarly, by means of selective excitation by different polarization states of a writing beam, polarization anisotropy can be created in nanocomposite materials [36,37]. When such polarization selectivity is combined with the sharp spectral selectivity of nanophotonics, gold nanorods can enable information recording in five dimensions, encoded across three wavelengths and two polarization states resulting into a huge capacity of 1.6 TB per disc (which is roughly an areal density of 1 Tbpsi) [35,14]. Other possible dimensions that can be exploited include the intensity of the optical beam [38] and the angular momentum [39].

Another approach to enhance storage density is inspired by a diffraction-unlimited far-field imaging to develop Super-resolution Photoinduction Inhibition Nanolithography (SPIN) [40]. SPIN can overcome the diffraction barrier and achieve 3D super-resolved writing. The development of 3D super-resolved writing

Please cite this article in press as: W.A. Bhat, Bridging data-capacity gap in big data storage, Future Generation Computer Systems (2018), https://doi.org/10.1016/j.future.2017.12.066.

methods promise to further reduce the effective size of optical spots to below 50 nm. This enhances the storage capacity of an optical disc significantly to approach or even surpass the PB scale areal density in the near future [14]. Combining super-resolution techniques and multiplexing in the physical dimensions can result into optical discs with hundreds of PBs capacity. This results into a huge areal density of the order of tens of Pbpsi that far exceeds the areal density of magnetic and flash storage technology. These discs when aggregated as an array to the EB (or even ZB) scale, can dramatically alleviate the infrastructure and operational costs of big data centres besides bridging the data-capacity gap in big data storage.

3.2. DNA data storage

DDS exploits the peculiar characteristics of a DNA molecule to encode information bits into *nucleotides* to yield an extremely huge storage density. The encoding schemes employed for the process determine the storage density of DDS. Theoretically, DNA can encode two bits per nucleotide yielding a storage density of 455 EB/gram of single-stranded DNA. The theoretical limit of volumetric density of DDS is above 10 EB/mm^3. Researchers have successfully stored and retrieved data in DNA using different DNA encoding schemes [41–44]. These efforts have encoded and decoded English text, mathematical equations, Latin text and simple musical notations to DNA [45]. However, DNA encoding faces several practical limitations. Using a novel encoding scheme that encodes one bit per base (A or C for zero, G or T for one) and employs next-generation DNA synthesis and sequencing technologies, storage density of 700 TB/gram (which is approximately a volumetric density of 5.5 Petabits/mm^3) has been achieved [46]. Unfortunately, the process of DNA synthesis and sequencing is vulnerable to a wide variety of errors. Encoding schemes with error-detection (and error-correction) techniques that focus on physical errors like substitution, insertion and deletion of *oligonucleotides* [47,48] have been proposed. Huffman code with error correction techniques have been used to achieve a storage density of 2.2 PB/gram of DNA that equals an areal density of around 100 Tbpsi [49]. However, all these achievements have failed to deliver full recovery of information, and have realized less than 50% of the Shannon information capacity of DNA molecules. Nevertheless, an encoding scheme equipped with Reed Solomon error-correction technique was able to achieve a huge storage density of 25 PB/gram of DNA (which is around an areal density of 1 Pbpsi) with 100% recovery of information, and exploited 62% of the Shannon information capacity of DNA molecules [50]. In a more recent work, using Fountain codes, researchers were successfully able to achieve 86% of the Shannon information capacity of DNA molecules, 100% information recovery, and an extremely huge storage density of 214 PB/gram of DNA that translates into an extremely huge areal density that approaches 100 Pbpsi [51]. Apart from this, techniques that allow rewritable data storage in DNA using strand exchange between site specific DNA sequences and other possible approaches have been demonstrated [52,53].

The technology offers tremendous storage density that besides bridging the data-capacity gap, suffices all the capacity requirements of big data storage for all the times to come and significantly reduces infrastructure and operational costs of big data centres. Furthermore, the technology is destined to advance further; the cost of DNA synthesis and sequencing are declining exponentially (as shown in Fig. 4), and it is believed that DNA storage will become cost-effective in coming years for the data to be archived for at least 50 years [54].

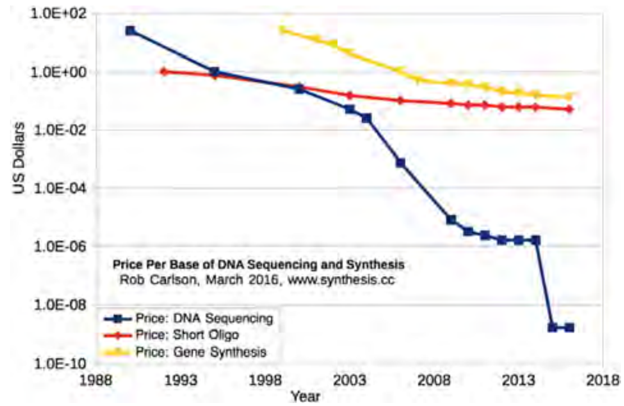

Fig. 4. Growth in cost-drop of DNA Sequencing and Synthesis.

3.3. Holographic data storage

HDS exploits multiplexing techniques to record hundreds of data pages at the same location in the recording media. These multiplexing techniques are at the core of HDS to offer high storage density. The theoretical limits for volumetric density of holographic storage is around tens of Terabits/cm^3 [55]. Many techniques of holographic multiplexing have been proposed [56–58]. Angle multiplexing roughly delivers a user capacity of 700 GB on a 120 mm disk with some other assumptions in place [59]. Another technique, namely Dynamic Aperture Holographic multiplexing [60], improves storage density offered by HDS by greatly increasing the scan range available for Angle multiplexing [61]. Dynamic Aperture multiplexing dynamically alters the portions of the available angular aperture used for the Signal and Reference beams. Holograms are recorded in order of ascending Reference beam angle, shrinking the Signal aperture as the Reference beam scans from left to right. With this geometry, the number of holograms multiplexed increase from 192 to 1178. Because the size of the data page shrinks as the Reference angle increases, the amount of data per hologram diminishes as well. However, the total amount of data recorded increases substantially from a maximum user capacity of 700 GB to TB, which an improvement of over 240% [62].

Another parameter that can be exploited for storage density is the phase of the Reference beam to record another hologram for the same Reference beam as long as a phase difference of 90° is maintained. Accordingly, Phase Quadrature Holographic Multiplexing (PQHM) promises to double the storage density [63]. A medium 1.5 mm thick and a bulk index of refraction of 1.5, and a recording wavelength of 0.4 μm, the areal density of the HDS is approximately 12 Tbit/in^2. By increasing the material thickness, and index of refraction, it is possible to build devices and media with areal densities of 40 Tbit/in^2. PQHM doubles (or triples) the achievable addressable space limit [64].

Yet another determinant of storage density is the refractive index contrast and thickness of recording media as they determine the number of holograms multiplexed. Research efforts so far have concluded that only photopolymers are commercially viable materials for HDS. A Two-Chemistry approach for designing such polymers was introduced by Bell Laboratories, Lucent Technologies, and InPhase Technologies [65]. A major advance over the Two-Chemistry approach was introduced by Akonia Holographics that serves the requirement of Dynamic Aperture Multiplexing [66].

HDS has the potential to bridge the data-capacity gap in big data storage as further advances are made in the technology. Since, the technology is highly competitive with current and emerging storage technologies on cost/TB basis, it is expected to significantly cut-down the infrastructure and operational costs of big data centres.

Please cite this article in press as: W.A. Bhat, Bridging data-capacity gap in big data storage, Future Generation Computer Systems (2018), https://doi.org/10.1016/j.future.2017.12.066.

3.4. Current storage technologies

Magnetic disk drives, commonly known as hard disk drives (HDD), rely on progress in recording technology to deliver advances in storage density. Perpendicular magnetic recording (PMR) technology currently used by magnetic disks is more than a decade old technology and should have been replaced by a better technology, like Heat assisted magnetic recording (HAMR) a long time ago [67]. Unfortunately, deploying HAMR was more difficult and expensive than expected. This delay has significantly affected growth in storage density of magnetic drives by bringing down areal density growth rate from 100%/yr to 10–15%/yr. Current industry roadmap predicts that HAMR will be deployed in late 2018. However, even with HAMR in place, the density growth is not expected to be more than 20–30%/yr. Industry believes that another technology namely Bit-patterned media (BPM), which is expected to be deployed after 2020, is the answer to the declining storage density growth of magnetic disks. Unfortunately, adoption of BPM is expected to be more difficult and expensive than that of HAMR, and thus is expected to face same delays. The technical difficulties and economic factors have affected the growth of areal density of magnetic disks to such an extent that the technology can neither fulfil the capacity demand nor be economical for big data storage. HDDs have recently touched an areal density of 1 Tbpsi, and with HAMR and BPM, they are expected to reach a maximum areal density of 10 Tbpsi in year 2025. In addition to this, *superparamagnetism* sets an upper limit on areal density of magnetic disks, limiting it to well under 100 Tbpsi. Unfortunately, even reaching halfway to this areal density is difficult. Besides, studies have shown that this limit could be encountered sometime between 2020 to 2040 depending upon the growth in areal density.

In contrast, progress in areal density of solid-state drive (SSD) technology has been impressive during the last decade. The areal density of SSDs has recently superseded that of magnetic disks with an areal density of 1.69 Tbpsi. Shrinking the cell size to increase storage density has been a traditional approach. Unfortunately, smaller cells are less reliable and support fewer write cycles only. Recent approach includes stacking layers of Flash one on top of the other while having larger cells. This approach significantly increases storage density and write endurance, and is estimated to achieve a maximum areal density of 10 Tbpsi in 2020. However, stacking more and more layers is becoming difficult, and it is currently expected that 64-layers will be the limit. Also, adding layers adds processing steps, and hence cost. Each byte of Flash contains about 50 times as much capital as a byte of magnetic disk. Hence, competing with magnetic disks on cost/TB will be challenging for SSDs. Therefore, it is expected that SSDs will be restricted to work in tandem with HDD for performance gains in applications [68,69]. Furthermore, studies have shown that all SSD technologies would be approaching their technological limits before magnetic disk technology does. At the other end of the spectrum, historically magnetic tape, open standard of which is known as Linear Tape Open (LTO), has been highly reliable and cost effective technology. However, storage density of LTO has not progressed to an extent so as to compete with other current storage technologies, let alone the emerging storage technologies as shown in Fig. 5.

It is evident from Fig. 5 that the areal densities of current storage technologies are far smaller than those of the emerging storage technologies. Also, DNA and Optical storage technology has crossed the Petabit/in^2 scale, and thus can undoubtedly overcome the data-capacity gap in big data centres. Unfortunately, HDS has recently touched Terabit/in^2 scale, like HDD and SSD. However, the potential for further improvements in areal density of HDS and its cost-effectiveness supersedes that of HDD and SSD. Also, further progress in areal density of HDD and SSD is becoming more difficult and they are not expected to reach even near the Petabit/in^2 scale in the foreseeable future.

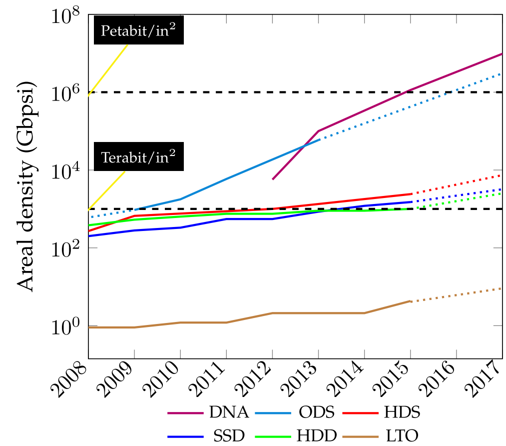

Fig. 5. Progress in areal density of current and emerging storage technologies.

4. Advances in throughput

Big data centres run wide-variety of applications which are typically characterized by I/O intensive operations. These I/O operations access large datasets and are mostly random and read-intensive in nature. To deliver real-time responses, storage subsystem of a big data centre must offer high transfer rates and low latencies. This section qualitatively evaluates and quantitatively compares the recent advances in throughput of the emerging and current storage technologies.

4.1. Optical data storage

Conventional ODS device such as a typical Blu-ray disc has an average transfer rate of 20 MBps, which is roughly 5 times lesser than the average transfer rate of a typical magnetic disk. Both transfer rates are not acceptable for a big data centre. As an example, accessing 1 PB of information residing on such magnetic disks will take 100 days while as the same process will take $1\frac{1}{2}$ year for Blu-ray discs. To overcome this problem, progress in ODS has resulted into various optical parallelism methods for generating multifocal arrays in single Laser shots [70–72]. Among these methods, Debye-based Fourier transformation method [73] enables the formation of diffraction-limited multifocal spots using a high-NA objective wherein each focal spot can be dynamically programmed [14]. If few hundred such focal spots are used, the method promises to deliver throughput as high as 10 GBps, and access time of the order of tens of milliseconds. This implies that the time to access 1 PB of information will dramatically reduce to only a few days, rendering the technology relevant in the era of big data. Furthermore, the method can generate not only in-plane multifocal arrays but also high-quality 3D multifocal arrays for volumetric parallel writing [74], and it can even be integrated with other physical parameters to increase the throughput beyond tens of GBps [14].

Furthermore, in an array, each disc unit has the ability to read and write in optical parallelism to maximize the throughput. This implies that such an arrangement will allow an overall throughput of a big data centre to far exceed TBps scale.

4.2. DNA data storage

DDS reading and writing transfer rate is around 400 Bps [75] and is not currently competitive with any existing or emerging storage technology; so is the access time which is of the order of tens of hours. This pitfall of the technology is partly due to the limitations of the semiconductor technology that renders synthesis

Please cite this article in press as: W.A. Bhat, Bridging data-capacity gap in big data storage, Future Generation Computer Systems (2018), https://doi.org/10.1016/j.future.2017.12.066.

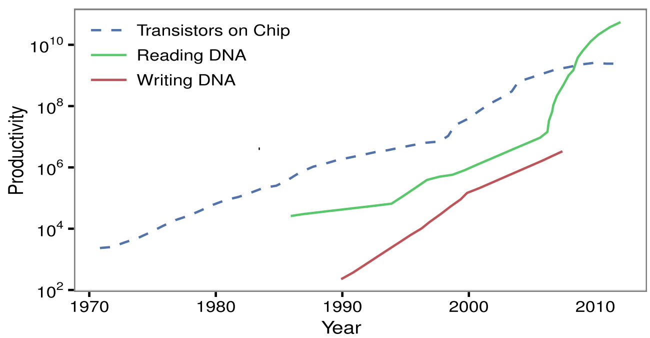

Fig. 6. Carlson's curves: Moore's law vs DNA synthesis and sequencing.

and sequencing a highly time consuming task, and partly due to the encoding/decoding schemes employed for synthesis/sequencing. However, synthesis and sequencing technology has witnessed exponential growth over time. Comparing Moore's law with the advances in throughput of synthesis and sequencing technology reveals that sequencing technology has been growing faster than Moore's law as shown in Fig. 6 [29]. In fact, recent advances in microfluidics and labs-on-a-chip have made synthesizing and sequencing DNA an everyday task [76].

Other improvement in this direction has been the demonstration of Random-Access DDS [53]. Earlier, DNA storage technology provided access to large blocks only, and thus it required sequencing of entire DNA pool to read even a single byte from the storage system. This sequential access results in overall read latency that is much longer than write latency. Using a polymerase chain reaction (PCR) to amplify only the desired data, read latency can be significantly reduced and sequencing entire DNA pool can be avoided. Apparently, it seems that the technology is not feasible for big data centres due to its low throughput. However, the advocates of the technology argue that it is even currently suitable for read-only, high-latency applications of big data centres, like longterm archival of data, which interestingly is the major contributor of data deluge in this era of big data. And, further progress in the technology will extend its feasibility for other applications of big data centres.

4.3. Holographic data storage

HDS recording and reading processes incur different set of steps each, and hence have different transfer rates, with recording process being latent compared to reading process. Advances by Akonia Holographics over the past few years have demonstrated transfer rates between 200 MBps and 300 MBps, and access times of <50 ms [60]. Improvements in recording process is the result of advances in media exposure time and mechanical positioning time. Significant improvement in media exposure time has been achieved with better media sensitivity and higher Laser power. In contrast, mechanical positioning that involves galvo step-and-settle and media step-and-settle. This process is mechanical in nature, and thus much improvements cannot be expected. Nevertheless, the galvo step-and-settle times have improved from 700 to 100 μs by simply decreasing the mirror size.

On the other hand, reading transfer rate is dependent upon media diffraction efficiency, movement times, Laser power, and camera sensitivity. Progressive improvements in movement times, Laser power and diffraction efficiency are expected. Among these factors, camera sensitivity significantly improves both recording and reading transfer rate. Since, all the pixels are read or recorded simultaneously during a single exposure, increasing the number of pixels per page dramatically increases the transfer rate for both reading and recording. With such advances, recording and reading transfer rates are expected to improve further from 300 MBps to 400 MBps. This allows HDS to closely compete with the transfer rates of SSDs and win the competition for prospective storage technology of big data centres together with cost/TB parameter.

Table 1
Transfer rates and access times of the emerging and current storage technologies.

Technology	Transfer rate	Access time
ODS	10 GBps	10s ms
SSD	250–500 MBps	∼0.1 ms
HDS	300–400 MBps	<50 ms
HDD	100–200 MBps	10s ms
LTO	Under 240 MBps	minutes
DDS	400 Bps	10s hrs

4.4. Current storage technologies

Magnetic disk drives are mechanical at heart. Primarily, seek-time and rotational-latency of a magnetic disk determine its latency. Improvements in access-time have been the result of progressive improvement in rotational-speed that has hit the performance wall at 22,000 rpm. Also, physical limits imposed on rotational speed has limited the access-time of magnetic disk to the order of tens of milliseconds, and further improvements are proving to be more challenging. Besides access-time, transfer rate of a magnetic disk has ceased to improve further. Transfer rate of magnetic disk is a function of rotational speed and areal density. Unfortunately, limitations in both parameters have restricted magnetic disk transfer rate to around 200 MBps. Not surprisingly, SSDs offer better access time of around 0.1 ms, and transfer rate of 250–500 MBps, as they do not involve any mechanical motions. Surprisingly, LTO has a better transfer rate of 240 MBps than magnetic disk, though being mechanical in nature. However, LTO suffers from huge latency, which is of the order of minutes.

Table 1 compares transfer rates and access times of the emerging and current storage technologies. The table clearly indicates that a lot of improvement in transfer rate and access time of DDS is needed so as to be better than, and comparable with, current and emerging storage technologies, respectively. In contrast, the transfer rate of ODS is far ahead of the rest of the technologies and fulfils the throughput requirements of big data centres. Also, it is worth noting that the transfer rate of HDS is comparable with those of SSDs, and further improvements in throughput promised by HDS will not only allow them to overshadow SSDs but will also allow them to compete with ODS.

5. Comments on lifetime

Lifetime of a storage technology points towards the period of the time for which the technology can reliably store and retrieve data, and thus decides the frequency of recurrent investment into procuring new storage units. In addition, operational and maintenance costs of storage technologies vary greatly. Big data centres prefer storage technologies that have longer lifetimes and incur lesser operational and maintenance costs. This section compares the lifetime of (and operational and maintenance costs incurred by) the emerging and current storage technologies.

5.1. Optical data storage

ODS devices use Laser beams to engrave information bits into optical media. The imprinted data can maintain stability and readability for 100–1000 years. ODS devices are also power proportional, and thus consume energy only during operation. Energy consumed by a single ODS unit for a period of 50 years is around

35 kWh, which is significantly small as compared to 4480 kWh energy consumed by a single HDD unit for the same period of time [14]. Also, the inscribed writing can sustain temperatures of up to 1000 °C. ODS devices have touched huge areal density of Pbpsi scale, and when aggregated to EB or ZB scale, the footprint of a big data centre is significantly reduced. It is estimated that enhancing the lifetime of ODS devices can result into about 96% savings in overall expenditures incurred by big data centres on account of energy consumption, maintenance, operational and infrastructure costs.

5.2. DNA data storage

DDS employs DNA to encode information. DNA is a naturally occurring information molecule that has an observed half-life of over 500 years in harsh environments. A study has shown that information encoded in DNA can be recovered without errors after treating DNA in silica at 70 °C for one week [77]. The study claims that it is thermally equivalent to storing information in DNA in central Europe for 2000 years. The energy consumed by DDS largely depends on the energy consumed by DNA sequencers and synthesizers which is progressively improving. Since, DDS offers extremely huge areal density of around hundred Pbpsi, the footprint of a big data centre is extremely small, and the infrastructure, operational and maintenance costs are significantly less.

5.3. Holographic data storage

HDS devices encode data pages in holographic media, which is an important factor in determining many features of HDS including lifetime. New breakthroughs in HDS by Akonia Holographics have resulted into polymer based recording media with more than 50 years of lifetime [60]. The durability of the media to sustain atmospheric conditions of 80 °C and 95% relative humidity has been also demonstrated. Even though areal density of HDS is within Tbpsi scale, the polymer media that can be packaged between two plastic substrates (which are 700 μm thick) significantly reduces the footprint of a big data centre besides reducing the infrastructure, operational and maintenance costs.

5.4. Current storage technologies

Magnetic disk drives have a shorter lifetime of about 5–7 years before they become unreliable. This implies that the investment in procuring new storage units is significantly frequent. In addition, magnetic disks are not power proportional, i.e., they consume energy even in idle state. The energy consumed by a magnetic disk unit is huge and is about 4480 kwhr for a period of 50 years. Since they require replacement after every 5–7 years, frequent data migration also adds huge energy consumption costs. They also require an operating temperature of 5–50 °C to maintain data stability and reliability, and thus, they add significantly to the cooling costs of a big data centre. To increase the capacity of a data centre to PB scale, thousands of magnetic disks are assembled. This results into huge footprint of a big data centre besides contributing significantly to its infrastructure costs.

In contrast, SSDs have a better lifetime of 10 years, each unit consumes about $\frac{1}{3}$rd of the energy consumed by a magnetic disk unit, are power proportional, result into smaller footprint of a data centre, and have a wider operating temperature range of 0–70 °C. However, they are 50 times as expensive as magnetic disks. For the same criteria, magnetic tapes are better than the two. LTO has a lifetime of 30 years, each unit consumes $\frac{1}{10}$th of the energy consumed by a magnetic disk unit, are power proportional, result into small data centre footprint, and are much cheaper than magnetic disks.

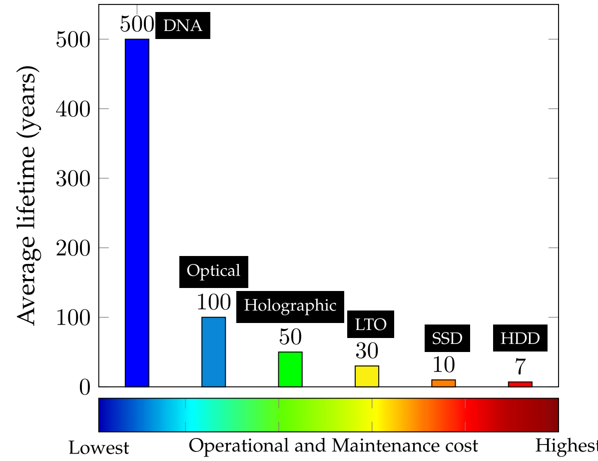

Fig. 7. Average lifetime of and overhead incurred by the emerging and current storage technologies.

Unfortunately, no current storage technology matches the lifetime of, and operational, maintenance and infrastructure costs incurred by, the emerging storage technologies as shown in Fig. 7. There is a strong correlation between the lifetime of a storage technology and the operational, maintenance and infrastructure costs incurred by it. DDS is at one end of the spectrum that offers highest lifetime and lowest overhead. At the other end, HDD suffers from the lowest lifetime and highest overhead.

6. Implications of adoption

Even though the emerging storage technologies promise to fulfil all the storage requirements of a big data centre to overcome the data-capacity gap, adopting them do have certain implications. This section discusses these implications.

6.1. System level disruptive changes are expected

Emerging storage technologies can confront the limitations of existing IT infrastructure, and may necessitate many disruptive system level changes. As an example, holographic recording is very data sensitive, i.e., it requires data streaming as it is not yet appropriate for partial recordings. If the data buffer is emptied during a recording process, the recording fails and the media is left useless like earlier CDs and DVDs. This puts the interfacing system under pressure to deliver an uninterrupted and high-speed input–output channel. Coping up with such a demand can push other components of the system beyond their current performance capabilities. This can have a cascading effect and can disruptively change the whole system.

6.2. Applications may require re-restructuring

Application restructuring was demanded by magnetic drives to overcome the performance gap between their sequential and random access. Unfortunately, such performance-patches deteriorate the performance of SSDs. It is expected that these performance-patches will also hurt the performance of emerging storage technologies. Thus, application *re*-restructuring may be required. Moreover, different *re*-restructuring may be demanded by different emerging storage technologies that will result into complex application designs.

6.3. Data-reduction techniques will become irrelevant

Data-reduction techniques, *Compression* and *De-duplication*, reduce the storage demand by optimizing capacity and reducing data footprint. However, they also add computational and storage overhead, and energy costs, besides hindering I/O performance. Considering the extremely huge-storage densities offered by emerging storage technologies, the relevance of data-reduction techniques in big data centres at storage level will become questionable, and hence their applicability will need to be revisited.

6.4. Non-standardization will hinder interoperability

Emerging technologies do not follow standardization for many reasons. Unfortunately, this affects their interoperability. HDS, in particular, has no substantial standards endorsed by the International Standards Organization (ISO). This lack of standardization can work against holographic storage by complicating interoperability between media and drives.

6.5. File system design will require rethinking

File systems were designed long ago with due considerations to many parameters including nature and working of magnetic disk drives, and were augmented to deal with specific needs of the applications [78–80]. In fact, with the advent of SSDs, existing file systems were either substantially-modified or new file systems were designed and developed to better suit the requirements of SSDs. Therefore, it is to be expected that file systems will need to be redesigned from scratch to meet specific requirements of emerging storage technologies. Specifically, to exploit the full potential of extremely huge storage densities and throughputs offered by the emerging storage technologies, file system design will require rethinking and innovation.

6.6. New security, privacy & forensic implications

Security, privacy and forensic implications of current storage technologies are well known, and the measures to combat them are also well studied [81–83]. However, the emerging storage technologies are inherently different from the current storage technologies. Therefore, they demand new, innovative and different set of approaches to combat their security, privacy and forensic implications. As an example, DNA is susceptible to DNA mutations, and mechanisms to reveal whether the mutations were induced naturally or were result of some breach must be devised and enforced.

6.7. Media specific problems need to be addressed

Since these emerging technologies are absolutely unconventional, big data centres will have to cater to different and unique demands of these technologies. As an example, unrecorded holographic media behaves more like the unexposed photographic film. Exposing the media prematurely to light can render it useless, while as no exposure for 3 years may turn it into unrecordable media.

Undoubtedly, adoption of emerging storage technologies has wide-domain, heterogeneous and severe implications. Since current storage technologies are not competent to overcome the data-capacity gap, the gap is inevitable and widening, and the implications are far more severe. As an example, by 2020, 6 ZBs of critical data and 31 ZBs of non-critical data will be lost due to storage shortage. Value associated with this huge volume of lost-data significantly supersedes the value of the implications resulting from the adoption of the emerging storage technologies. However, less-disruptive technology-transition in big data centres needs attention and requires research efforts. Since the implications are wide domain; standardization, evaluation, monitoring and adoption of emerging storage technologies must be institutionalized in both academics and industry.

7. Prospects & challenges

All the emerging storage technologies discussed and evaluated in this paper have received many significant advances in the recent past. Consequently, they offer extremely huge storage densities and suitable throughputs, significantly alleviate energy consumption costs besides incurring minimal operational, maintenance and infrastructure costs, can sustain harsh atmospheric conditions like high temperatures, and can maintain data reliability and readability for much longer time frames. All these features are critical to big data centres for overcoming the data-capacity gap. However, since they vary greatly in advances received and challenges faced, their prospects need to be evaluated and challenges faced highlighted.

ODS renders high storage capacity of hundreds of PBs per disc with an overall throughput of the order of tens of TBps for big data centres. ODS also promises to maintain data reliability and readability under harsh atmospheric conditions for hundreds of years while incurring low operational, maintenance and infrastructure costs due to high energy efficiency and small data centre footprint. These features make ODS an ideal storage technology for big data centres to overcome the data-capacity gap. However, due to absence of rewritable-storage support in ODS, its scope may remain limited to real-time applications that frequently read huge volumes of data with low latency, and to those applications that archive immutable and large volumes of data for longterm and access that data with low latency. Consequently, to further widen its scope, research efforts towards development of rewritable ODS and alleviation in cost of ownership are required.

In contrast, DDS tenders an extremely huge and rewritable storage capacity of around 214 PB per gram of DNA besides offering strong sustainability against severe environmental conditions for thousands of years while incurring minimal operational, maintenance and infrastructure costs due to extremely small data centre footprint. The offerings are highly lucrative for big data centres to overcome the data-capacity gap. However, DDS significantly suffers from low throughput and high cost of sequencing and synthesizing. Even though further advances in DDS to alleviate cost of reading/writing are promising, improvements in throughput seem to be challenging. These pitfalls may limit the application of DDS in big data centres to longterm archival of huge volumes of data that is less frequently accessed with high latency. Nevertheless, future advances in throughput will certainly extend its scope, and research efforts in this direction are needed.

Finally, HDS exhibits storage capacity of few TBs per disc with $\frac{1}{2}$GBps of throughput per unit, which is comparatively better than current storage technologies. Moreover, the technology promises further advances in storage density and throughput rendering it comparable to ODS and DDS. HDS also has a lifetime of 50 years, and incurs low maintenance and operational costs due to being energy efficient, and low infrastructure costs due to small data centre footprint. These attributes also render HDS a better technology than current storage technologies and one comparable to ODS and DDS. In fact, comparing HDS on cost/TB with ODS and DDS reveals the economic strength of HDS. More importantly, all the advances in HDS sufficiently suffice the storage requirements of big data centres to overcome the data-capacity gap. However, HDS has yet to offer rewritable storage which may restrict its domain to a set of applications that is same as that of ODS. Nevertheless, research efforts in storage density, throughput and rewritable media are required to extend its domain as the technology has the potential to progress further.

8. Other emerging storage technologies

In addition to the emerging storage technologies discussed and evaluated in this paper, there are other storage technologies that have the potential to fulfil all the storage requirements of big data centres, but are still in their infancy. These include *Nanocrystal memory*, *Scanning probe memory*, *Carbon nanotube memory*, and *Organic memory*.

8.1. Nanocrystal memory

Nanocrystal memory stores information in nanocrystals by injecting charges. The technology offers low latency, good scalability, and superior reliability. These features allow it to be a strong candidate for future big data centres. The technology uses a simplified fabrication process and a smaller cell area that significantly reduces its cost [84]. Unfortunately, the technology is confronted by some significant challenges that need to be addressed [85]. As an example, nanocrystal shape and size affect its memory properties. With the growth in storage density, the cell size is expected to shrink. This makes it challenging to synthesize same sized and shaped nanocrystals, and to assemble then into a well-ordered matrix. Also, charge stored in nanocrystals tend to leak more easily in high density nanocrystals. Other challenges include lowering operating voltages, demonstrating sufficient data retention periods, and so on.

8.2. Scanning probe memory

Scanning probe memory uses a nanoprobe to record and read bits in storage medium. The technology has demonstrated its ability to store data at super-high densities and with super-fast writing and reading speed while consuming very low energy [86]. This allows the technology to be considered for future big data centres. In recent years, significant progress in the technologies underlying scanned probe memories have been made. These advances have resulted into three main approaches [7]: (a) thermo-mechanical write/read into polymer media, (b) electrical write and read in ferro-electric media, and (c) thermo-electric write/electrical read in phase-change media. While the first two approaches are evolving, the last approach is the most advanced and promising one [87]. However, the technology faces many technical hurdles that need to be overcome.

8.3. Carbon nanotube memory

Carbon nanotubes have some unique and intrinsic electrical and mechanical properties, and these properties have promoted them to be used for a wide variety of applications including digital storage [88]. Carbon nanotube based storage devices support high speed, non-volatility, and small size, rendering the technology suitable for future big data centres. However, Carbon nanotube based memories suffer from poor retention period, fabrication difficulties, high production cost, and so on. Advances made by Natero, Inc. [89] has led to carbon nanotube based memory that can be manufactured in a CMOS fabrication facility. The progress in the technology has resulted into read/write performance superior to existing Flash memory, significant alleviation in power consumption during read and write, and has rendered it scalable to <5 nm. Also, the write-endurance and data-retention demonstrated has been better than that of Flash.

8.4. Organic memory

Organic material based devices exhibit a peculiar characteristic of switching, and hence are promising candidates for future information technologies. Organic memory is based on the deposition of low cost polymer films of active organic material [90], and has been proposed as the basis for organic non-volatile memory technology for future big data centres. In recent past, the technology has received significant attention due to its simple device structures, low fabrication costs, printability, and flexibility [91]. Organic memory can be classified into two: (a) capacitive type, and (b) resistive type. Among the two, organic resistive memory has been the center of focus recently as it exhibits fast switching speed, good scalability, and low power consumption [7]. However, the short retention time and the small endurance cycles severely limit its practical use.

9. Conclusion

Data-capacity gap in big data storage is eminent and inevitable, and can be bridged with storage technologies with greater storage densities, higher throughputs, and longer lifetime. In this paper, we first discussed the working principles of three emerging storage technologies, i.e., ODS, DDS and HDS, that possess these features. We then evaluated the advances received by them in storage density, throughput and lifetime, and quantitatively compared them with the advances in current storage technologies, i.e., HDD, SSD and LTO. We also investigated the implications of the adoption of these emerging storage technologies, evaluated their prospects and highlighted the challenges.

The study suggests that advances in storage density of ODS and DDS has crossed Pbpsi scale areal density, and thus can easily overcome the data-capacity gap in big data centres. However, due to absence of rewritable storage in ODS, its scope is currently limited to longterm archival of immutable data that is accessed frequently in large volumes and with low latency. Similarly, due to low throughput offered by DDS, its scope is currently limited to longterm archival of mutable data that is accessed less frequently and with high latency. Nevertheless, both technologies offer longer lifetimes. For to be considered for all workloads of a big data centre, ODS and DDS must offer rewritable storage and higher throughput, respectively. In addition, to be economically feasible, ODS and DDS must alleviate the cost of ownership and the cost of reading/writing, respectively.

Unfortunately, storage density of HDS is within Tbpsi scale like that of HDD and SSD, and throughput offered is comparable to SSD, while as lifetime is better than all the current storage technologies. Nevertheless, HDS has potential to overcome data-capacity gap, and research efforts towards improving storage density and throughput of HDS are required to compete with ODS and DDS. Also, HDS has yet to offer rewritable media so as to extend its scope to wide variety of applications of big data centres. Interestingly, HDS is better than ODS and DDS on cost/TB basis, supersedes current storage technologies on account of storage density and lifetime, has throughput comparable to that of SSD, and promises to advance further. These characteristics allows HDS to act as an immediate and intermediate solution for data-capacity gap until ODS and DDS technologies further advance to take over. Also, a heterogeneous and multi-tiered storage model that comprises of ODS, DDS, HDS and SSD can be envisioned wherein each tier serves a specific workload.

Finally, adopting emerging storage technologies in any configuration will have wide range of implications on big data centres. However, these implications are less severe than the implications resulting out of the loss of critical data due to data-capacity gap. Nevertheless, research efforts towards smooth and less-disruptive technology-transition in big data centres are required.

Please cite this article in press as: W.A. Bhat, Bridging data-capacity gap in big data storage, Future Generation Computer Systems (2018), https://doi.org/10.1016/j.future.2017.12.066.

References

[1] W.A. Bhat, S.M.K. Quadri, Big data promises value: is hardware technology taken onboard? Ind. Manage. Data Syst. 115 (9) (2015) 1577–1595.

[2] P. Parwekar, From Internet of Things towards cloud of things, in: 2011 2nd International Conference on Computer and Communication Technology, ICCCT, 2011, pp. 329–333. http://dx.doi.org/10.1109/ICCCT.2011.6075156.

[3] H. Bauer, M. Patel, J. Veira, (2014) The Internet of Things: Sizing up the opportunity, McKinsey.

[4] E.E. Schadt, M.D. Linderman, J. Sorenson, L. Lee, G.P. Nolan, Computational solutions to large-scale data management and analysis, Nature Rev. Genet. 11 (9) (2010) 647–657.

[5] V. Turner, J.F. Gantz, D. Reinsel, S. Minton, The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things, IDC, Framingham (MA), 2014.

[6] M. Hilbert, P. López, The worlds technological capacity to store, communicate, and compute information, Science 332 (6025) (2011) 60–65.

[7] L. Wang, S. Gai, The next generation mass storage devices–Physical principles and current status, Contemp. Phys. 55 (2) (2014) 75–93.

[8] D. Xu, Multi-Dimensional Optical Storage, Springer, 2016.

[9] D.A. Parthenopoulos, P.M. Rentxepis, Three-dimensional optical storage memory, Science 245 (4920) (1989) 843.

[10] D. Day, M. Gu, A. Smallridge, Use of two-photon excitation for erasable–rewritable three-dimensional bit optical data storage in a photorefractive polymer, Opt. Lett. 24 (14) (1999) 948–950.

[11] Y. Kawata, H. Ishitobi, S. Kawata, Use of two-photon absorption in a photorefractive crystal for three-dimensional optical memory, Opt. Lett. 23 (10) (1998) 756–758.

[12] D. Day, M. Gu, Effects of refractive-index mismatch on three-dimensional optical data-storage density in a two-photon bleaching polymer, Appl. Opt. 37 (26) (1998) 6299–6304.

[13] X. Li, Y. Cao, M. Gu, Superresolution-focal-volume induced 3.0 Tbytes/disk capacity by focusing a radially polarized beam, Opt. Lett. 36 (13) (2011) 2510–2512.

[14] M. Gu, X. Li, Y. Cao, Optical storage arrays: a perspective for future big data storage, Light: Sci. Appl. 3 (5) (2014) e177.

[15] V.M. Shalaev, Optical negative-index metamaterials, Nature Photonics 1 (1) (2007) 41–48.

[16] J. Chen, Y. Wang, B. Jia, T. Geng, X. Li, L. Feng, W. Qian, B. Liang, X. Zhang, M. Gu, et al., Observation of the inverse doppler effect in negative-index materials at optical frequencies, Nature Photonics 5 (4) (2011) 239–245.

[17] J. Serbin, M. Gu, Experimental evidence for superprism effects in three-dimensional polymer photonic crystals, Adv. Mater. 18 (2) (2006) 221–224.

[18] E. Chow, S. Lin, S. Johnson, P. Villeneuve, J. Joannopoulos, J.R. Wendt, G.A. Vawter, W. Zubrzycki, H. Hou, A. Alleman, Three-dimensional control of light in a two-dimensional photonic crystal slab, Nature 407 (6807) (2000) 983–986.

[19] V.R. Almeida, C.A. Barrios, R.R. Panepucci, M. Lipson, All-optical control of light on a silicon chip, Nature 431 (7012) (2004) 1081–1084.

[20] S. Noda, M. Fujita, T. Asano, Spontaneous-emission control by photonic crystals and nanocavities, Nature Photonics 1 (8) (2007) 449–458.

[21] J. Li, B. Jia, G. Zhou, C. Bullen, J. Serbin, M. Gu, Spectral redistribution in spontaneous emission from quantum-dot-infiltrated 3d woodpile photonic crystals for telecommunications, Adv. Mater. 19 (20) (2007) 3276–3280.

[22] E. Rittweger, K.Y. Han, S.E. Irvine, C. Eggeling, S.W. Hell, STED microscopy reveals crystal colour centres with nanometric resolution, Nature Photonics 3 (3) (2009) 144–147.

[23] M.J. Rust, M. Bates, X. Zhuang, Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM), Nature Methods 3 (10) (2006) 793–796.

[24] M. Gu, Y. Cao, S. Castelletto, B. Kouskousis, X. Li, Super-resolving single nitrogen vacancy centers within single nanodiamonds using a localization microscope, Opt. Express 21 (15) (2013) 17639–17646.

[25] Z. Gan, Y. Cao, R.A. Evans, M. Gu, Three-dimensional deep sub-diffraction optical beam lithography with 9 nm feature size, Nature Commun. 4 (2013).

[26] L. Li, R.R. Gattass, E. Gershgoren, H. Hwang, J.T. Fourkas, Achieving $\lambda/20$ resolution by one-color initiation and deactivation of polymerization, Science 324 (5929) (2009) 910–913.

[27] S. Kosuri, G.M. Church, Large-scale de novo DNA synthesis: technologies and applications, Nature Methods 11 (5) (2014) 499–507.

[28] P.T. Gilham, H.G. Khorana, Slowercasetudies on polynucleotides. I. A new and general method for the chemical synthesis of the C5 -C3 internucleotidic linkage. Syntheses of deoxyribo-dinucleotides1, J. Am. Chem. Soc. 80 (23) (1958) 6212–6222.

[29] J. Bornholt, R. Lopez, D.M. Carmean, L. Ceze, G. Seelig, K. Strauss, A DNA-based archival storage system, Oper. Syst. Rev. 50 (2) (2016) 637–649.

[30] L. Dhar, K. Curtis, T. Fäcke, Holographic data storage: Coming of age, Nature Photonics 2 (7) (2008) 403–405.

[31] L. Anderson, Holographic optical memory for bulk data storage (Holographic laser type memory with bulk storage capacity and short readout time), Bell Lab. Rec. 46 (1968) 318–325.

[32] K. Kubota, Y. Ono, M. Kondo, S. Sugama, N. Nishida, M. Sakaguchi, Holographic disk with high data transfer rate: its application to an audio response memory, Appl. Opt. 19 (6) (1980) 944–951.

[33] H. Ruan, Recent advances in holographic data storage, Front. Optoelectron. 7 (4) (2014) 450–466.

[34] B. Nikoobakht, M.A. El-Sayed, Preparation and growth mechanism of gold nanorods (NRs) using seed-mediated growth method, Chem. Mater. 15 (10) (2003) 1957–1962.

[35] P. Zijlstra, J.W. Chon, M. Gu, Five-dimensional optical recording mediated by surface plasmons in gold nanorods, Nature 459 (7245) (2009) 410–413.

[36] X. Li, J.W. Chon, S. Wu, R.A. Evans, M. Gu, Rewritable polarization-encoded multilayer data storage in 2, 5-dimethyl-4-(p-nitrophenylazo) anisole doped polymer, Opt. Lett. 32 (3) (2007) 277–279.

[37] X. Li, J.W. Chon, R.A. Evans, M. Gu, Quantum-rod dispersed photopolymers for multi-dimensional photonic applications, Opt. Express 17 (4) (2009) 2954–2961.

[38] D. Ganic, D. Day, M. Gu, Multi-level optical data storage in a photobleaching polymer using two-photon excitation under continuous wave illumination, Opt. Lasers Eng. 38 (6) (2002) 433–437.

[39] N. ad Bozinovic, Terabit-Scale orbital angular momentum mode division multiplexing, Science 1237861 (1545) (2013) 340.

[40] S.W. Hell, J. Wichmann, Breaking the diffraction resolution limit by stimulated emission: stimulated-emission-depletion fluorescence microscopy, Opt. Lett. 19 (11) (1994) 780–782.

[41] P.C. Wong, K.-k. Wong, H. Foote, Organic data memory using the DNA approach, Commun. ACM 46 (1) (2003) 95–98.

[42] M. Arita, Y. Ohashi, Secret signatures inside genomic DNA, Biotechnol. Prog. 20 (5) (2004) 1605–1607.

[43] G.M. Skinner, K. Visscher, M. Mansuripur, Biocompatible writing of data into DNA, J. Bionanosci. 1 (1) (2007) 17–21.

[44] M. Yamamoto, S. Kashiwamura, A. Ohuchi, M. Furukawa, Large-scale DNA memory based on the nested PCR, Natural Comput. 7 (3) (2008) 335–346.

[45] M. Ailenberg, O.D. Rotstein, An improved Huffman coding method for archiving text, images, and music characters in DNA, Biotechniques 47 (3) (2009) 747.

[46] G.M. Church, Y. Gao, S. Kosuri, Next-generation digital information storage in DNA, Science 337 (6102) (2012) 1628.

[47] D. Haughton, F. Balado, Repetition coding as an effective error correction code for information encoded in DNA, in: Bioinformatics and Bioengineering, BIBE, 2011 IEEE 11th International Conference on, IEEE, 2011, pp. 253–260.

[48] H.M. Kiah, G.J. Puleo, O. Milenkovic, Codes for DNA storage channels, 2014, ArXiv Preprint ArXiv:1410.8837.

[49] N. Goldman, P. Bertone, S. Chen, C. Dessimoz, E.M. LeProust, B. Sipos, E. Birney, Towards practical, high-capacity, low-maintenance information storage in synthesized DNA, Nature 494 (7435) (2013) 77–80.

[50] R.N. Grass, R. Heckel, M. Puddu, D. Paunescu, W.J. Stark, Robust chemical preservation of digital information on DNA in silica with error-correcting codes, Angew. Chem., Int. Ed. 54 (8) (2015) 2552–2555.

[51] Y. Erlich, D. Zielinski, DNA fountain enables a robust and efficient storage architecture, Science 355 (6328) (2017) 950–954.

[52] J. Bonnet, P. Subsoontorn, D. Endy, Rewritable digital data storage in live cells via engineered control of recombination directionality, Proc. Natl. Acad. Sci. 109 (23) (2012) 8884–8889.

[53] S. Yazdi, Y. Yuan, J. Ma, H. Zhao, O. Milenkovic, A rewritable, random-access DNA-based storage system, 2015, ArXiv Preprint ArXiv:1505.02199.

[54] E. Yong, Synthetic double-helix faithfully stores Shakespeare's sonnets, Nature (2013).

[55] K. Curtis, L. Dhar, A. Hill, W. Wilson, M. Ayres, Holographic Data Storage: From Theory to Practical Systems, John Wiley & Sons, 2011.

[56] H. Yamatsu, M. Ezura, N. Kihara, Study on multiplexing methods for volume holographic memory, in: International Symposium on Optical Memory and Optical Data Storage, Optical Society of America, 2005, p. ThE1.

[57] H. Horimai, X. Tan, Collinear technology for a holographic versatile disk, Appl. Opt. 45 (5) (2006) 910–914.

[58] K. Anderson, K. Curtis, Polytopic multiplexing, Opt. Lett. 29 (12) (2004) 1402–1404.

[59] K.-i. Shimada, T. Ishii, T. Ide, S. Hughes, A. Hoskins, K. Curtis, High density recording using monocular architecture for 500 GB consumer system, in: Optical Data Storage, Optical Society of America, 2009, p. 75050Q.

[60] K. Anderson, M. Ayres, F. Askham, B. Sissom, Holographic data storage: science fiction or science fact? in: SPIE Optical Engineering+ Applications, International Society for Optics and Photonics, 2014, p. 920102.

[61] M.R. Ayres, K.E. Anderson, F.R. Askham, B.J. Sissom, Dynamic aperture holographic multiplexing, Google Patents, US Patent App. 13/875,071, 2013.

[62] T. Ishii, M. Hosaka, T. Hoshizawa, M. Yamaguchi, S. Koga, A. Tanaka, Terabyte holographic recording with monocular architecture, in: Consumer Electronics, ICCE, 2012 IEEE International Conference on, IEEE, 2012, pp. 427–428.

[63] M. Ayres, Coherent techniques for terabyte holographic data storage, in: Optical Data Storage Topical Meeting, ODS 2010, 2010.

[64] K. Curtis, L. Dhar, A. Hill, W. Wilson, M. Ayres, Holographic Data Storage, Wiley Online Library, 2010.
[65] L. Dhar, A. Hale, H.E. Katz, M.L. Schilling, M.G. Schnoes, F.C. Schilling, Recording media that exhibit high dynamic range for digital holographic data storage, Opt. Lett. 24 (7) (1999) 487–489.
[66] K. Anderson, M. Ayres, B. Sissom, F. Askham, Holographic data storage: rebirthing a commercialization effort, Proc. Soc. Photo-Instrum. Eng. 9006 (2014) 90060C.
[67] D.S. Rosenthal, D.C. Rosenthal, E.L. Miller, I.F. Adams, M.W. Storer, E. Zadok, The economics of long-term digital storage, in: Memory of the World in the Digital Age, Vancouver, BC, 2012.
[68] W.A. Bhat, S.M.K. Quadri, hFAT: A high performance hybrid FAT32 file system, in: T. Clary (Ed.), Horizons in Computer Science Research, Nova Science Publishers, NY, 2014, pp. 285–299.
[69] W.A. Bhat, S.M.K. Quadri, Performance augmentation of a FAT filesystem by a hybrid storage system, in: Advanced Computing, Networking and Informatics, Vol. 2, Springer, 2014, pp. 489–498.
[70] L. Sacconi, E. Froner, R. Antolini, M. Taghizadeh, A. Choudhury, F. Pavone, Multiphoton multifocal microscopy exploiting a diffractive optical element, Opt. Lett. 28 (20) (2003) 1918–1920.
[71] J.-i. Kato, N. Takeyasu, Y. Adachi, H.-B. Sun, S. Kawata, Multiple-spot parallel processing for laser micronanofabrication, Appl. Phys. Lett. 86 (4) (2005) 044102.
[72] A. Jesacher, M.J. Booth, Parallel direct laser writing in three dimensions with spatially dependent aberration correction, Opt. Express 18 (20) (2010) 21090–21099.
[73] H. Lin, B. Jia, M. Gu, Dynamic generation of debye diffraction-limited multifocal arrays for direct laser printing nanofabrication, Opt. Lett. 36 (3) (2011) 406–408.
[74] H. Ren, H. Lin, X. Li, M. Gu, Three-dimensional parallel recording with a Debye diffraction-limited and aberration-free volumetric multifocal array, Opt. Lett. 39 (6) (2014) 1621–1624.
[75] https://www.technologyreview.com/s/607880/microsoft-has-a-plan-to-add-dna-data-storage-to-its-cloud/.(Accessed: 13 June 2017).
[76] D. Limbachiya, M.K. Gupta, A natural data storage: A review on sending information from now to then via nature, V(212) (2015) 1–17. ArXiv:1505.04890v1.
[77] M.E. Allentoft, M. Collins, D. Harker, J. Haile, C.L. Oskam, M.L. Hale, P.F. Campos, J.A. Samaniego, M.T.P. Gilbert, E. Willerslev, et al., The half-life of dna in bone: measuring decay kinetics in 158 dated fossils, Proc. R. Soc. B (2012) rspb20121745.
[78] W.A. Bhat, S.M.K. Quadri, restFS: Secure data deletion using reliable & efficient stackable file system, in: Applied Machine Intelligence and Informatics, SAMI, 2012 IEEE 10th International Symposium on, IEEE, 2012, pp. 457–462.
[79] W.A. Bhat, S.M.K. Quadri, suvfs: A virtual file system in userspace that supports large files, in: Big Data, 2013 IEEE International Conference on, IEEE, 2013, pp. 7–11.
[80] W.A. Bhat, Achieving efficient purging in transparent per-file secure wiping extensions, in: Handbook of Research on Security Considerations in Cloud Computing, IGI Global, 2015, pp. 345–357.
[81] W.A. Bhat, S.M.K. Quadri, After-deletion data recovery: myths and solutions, Comput. Fraud Secur. 2012 (4) (2012) 17–20.
[82] W.A. Bhat, S.M.K. Quadri, POSTER: Dr. Watson provides data for post-breach analysis, in: Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security, ACM, 2013, pp. 1445–1448.
[83] W.A. Bhat, S.M.K. Quadri, Understanding and mitigating security issues in Sun NFS, Netw. Secur. 2013 (1) (2013) 15–18.
[84] J. De Blauwe, Nanocrystal nonvolatile memory devices, IEEE Trans. Nanotechnol. 99 (1) (2002) 72–77.
[85] T.-C. Chang, F.-Y. Jian, S.-C. Chen, Y.-T. Tsai, Developments in nanocrystal memory, Mater. Today 14 (12) (2011) 608–615.
[86] C.D. Wright, M.M. Aziz, P. Shah, L. Wang, Scanning probe memories–Technology and applications, Curr. Appl. Phys. 11 (2) (2011) e104–e109.
[87] L. Wang, C.H. Yang, J. Wen, S. Di Gong, Y.X. Peng, Overview of probe-based storage technologies, Nanoscale Res. Lett. 11 (1) (2016) 342.
[88] X. Liu, Z. Ji, M. Liu, L. Shang, D. Li, Y. Dai, Advancements in organic nonvolatile memory devices, Chin. Sci. Bull. 56 (30) (2011) 3178–3190.
[89] S. Kianian, G. Rosendale, M. Manning, D. Hamilton, X.H. Huang, K. Robinson, Y.W. Kim, T. Rueckes, A 3D stackable carbon nanotube-based nonvolatile memory (NRAM), in: Solid-State Device Research Conference, ESSDERC, 2010 Proceedings of the European, IEEE, 2010, pp. 404–407.
[90] D. Prime, S. Paul, Overview of organic memory devices, Phil. Trans. R. Soc. A 367 (1905) (2009) 4141–4157.
[91] T.W. Kim, Y. Yang, F. Li, W.L. Kwan, Electrical memory devices based on inorganic/organic nanocomposites, NPG Asia Mater. 4 (6) (2012) e18.

Wasim Ahmad Bhat is an Assistant Professor in the Department of Computer Science at the University of Kashmir. He received a Ph.D. in Computer Science from the University of Kashmir. His current research interests include big data storage infrastructure, high-density storage technologies, and future-green storage technologies.

Please cite this article in press as: W.A. Bhat, Bridging data-capacity gap in big data storage, Future Generation Computer Systems (2018), https://doi.org/10.1016/j.future.2017.12.066.