An Economic Perspective of Disk vs. Flash Media in Archival Storage

Preeti Gupta*, Avani Wildani†, Ethan L. Miller*,
Daniel Rosenthal§, Ian F. Adams‡, Christina Strong*, and Andy Hospodor*
*University of California, Santa Cruz †The Salk Institute ‡EVault Inc. §HP

***Abstract*—For three decades, Kryder's law correctly predicted an exponential increase in bit density on disk platters, leading to an exponential drop in cost per gigabyte, and thus to an entrenched expectation that if data could be stored for a few years the incremental cost of storing it forever would be minimal. However, disk now is over 7 times as expensive as Kryder's law would have predicted, and industry projections suggest that in 2020 the gap will reach 200 times, disrupting this expectation.**

Our model shows that archives based upon alternative media are surprisingly cost competitive with archives based upon traditional disk media over the long-term. We propose using Archival Flash for long-term data preservation, with the trade off between longer data retention period and lower write cycles.

I. Introduction

IDC's 2012 Digital Universe study predicted [8] that there will be 40 zetabytes of data on the planet by 2020, which was 14% more than previous forecasts, and approximately 5TB of data per person. Much of this data has long-term value and should be stored permanently. Earlier studies of access to archival data [16] showed infrequent, sparse patterns encouraging the use of long-latency storage such as tape. More recently, archives report increasing demand for data-mining, requiring at least one copy on low-latency media. Here we consider only the costs of the low-latency copy.

Long-term storage system costs comprise regularly replacing media, whose unit capacity increases with bit density at approximately constant cost, and operational costs such as power, cooling, space, and staff, which are approximately constant per unit. Thus per-byte storage costs decrease as bit density increases at the Kryder rate (annual storage density growth rate) [30]. Pre-paid storage services (e.g. Princeton's POSF (Pay Once, Store Forever) [23] and Longaccess [11]) price based on predictions of future storage costs, making the Kryder rate their most important parameter. Compared to pre-2010 projections, Figure 1 shows that per-byte disk cost is now 7 times more expensive, and in 6 years would be around 200 times, more expensive.

Using an enhanced version of the Long-Term model introduced in our prior work [29] to study a fixed sized dataset stored for 100 years we show the effects of varying Kryder rate, device service life, device operational cost, and data refresh costs. Although it is technically feasible to build hard drives with service lives longer than 5 years, a 2009 Seagate study showed that "robust consumer grade drives are not profitable" [6]. At pre-2010 Kryder rates the benefits of moving to newer, denser media outweighed those of longer lasting but more expensive media. We show that at post-2010 Kryder rates technologies with longer service life and lower operational costs become more competitive despite higher capital acquisition costs.

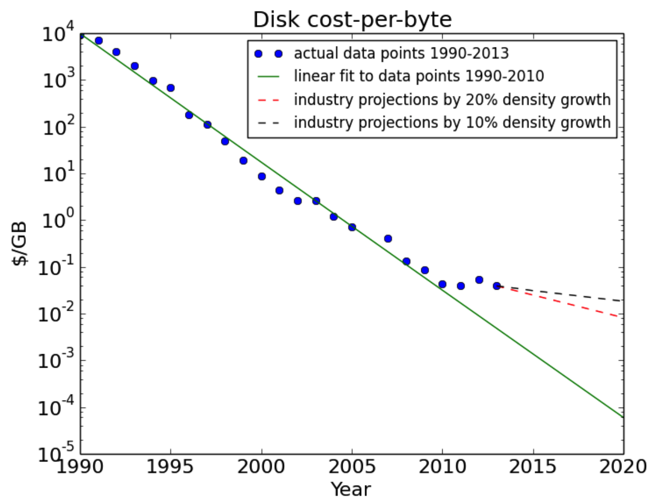

Fig. 1. Cost-per-byte decrease slowed dramatically in 2010 [2].

II. Background

Kryder Rate (K_r) is the annual proportional change in bit density and thus, we assume, decrease in per-byte cost (e.g. $K_r = .2$ means next year's per-byte cost would be 80% of this year's per-byte cost.

An **Endowment** is a capital sum, deposited with the data and invested, believed to be sufficient to pay storage costs for the duration.

Projected Kryder rates for various media have changed in recent years:

Tape: IBM recently demonstrated a major increase in bit density on tape [13], greatly improving tape's per-byte cost advantage in short-term. Even this new technology is less than 30 times less dense than disk; as the disk Kryder rate slows tape bit density will catch up and then run into the same technological limits. However, tape does not provide low-latency access.

Optical Media: Despite their advantages of low cost and long life, optical media such as Blu-ray disks also do not provide low-latency access. Facebook's recent development of Blu-ray cold-storage [14] explicitly tolerates long access latency.

Disk: Disk manufacturing has consolidated as total shipments have declined (by 4.3% in 2013 [3]), leading to increased margins. The transitions to the next disk recording technology (Heat-Assisted Magnetic Recording) and its probable successor (Bit-Patterned Media) turn out to be vastly more difficult and expensive than expected, delaying further bit density improvements and thus decreasing the Kryder rate.

Solid State Device (SSD): Flash and its successor technologies are much more expensive per byte than hard disk but offer much lower power, space, cooling costs, and better access latencies. As flash scales down, its write endurance and data retention worsen. The first is not a problem for archival purposes (write once, read rarely, overwrite rarely), the second is. Flash controllers are optimized for performance, but different trade-offs can be made to improve data retention and reduce write endurance.

Current technologies for both disk (PMR) and solid state (flash) have limited scope for improvement, and their successors are some way from mass production. This means that near-term investment in increasing the supply, and thus decreasing the price, of disk and solid state storage would not be profitable. The lead time on such investment is about 5 years, so we know roughly how much disk and flash could be manufactured over the medium term. It is much less than the projected demand, so significant decreases in price per byte are unlikely.

Most work on the economics of digital preservation lacks consideration of future long-term storage costs in general, and these trends in particular. Some work ignores storage cost. CMDP [24] focused on early activities, preservation planning, and ingest. A Blue Ribbon Task Force on Sustainable Digital Preservation was funded by the National Science Foundation and other organizations. Their final report does not consider storage costs [7].

Other work examined storage costs but not cost trends. Chapman [20] compared historic storage costs for analog items in the Harvard Depository with those for digital objects. LIFE [31], Prestoprime [17], KRDS [19] , and the California Digital Library developed TCP [15]developed tools for estimating preservation costs based on cost history without a model of future storage costs.

Rosenthal *et al.* [29] showed the effect of K_r slowing down on long-term preservation cost. Our work is an enhancement: we provide comprehensive experiments with flash and disk using various parameters and discuss their effects on long-term preservation cost. Adams *et al.* [16] suggested that an appropriate system architecture could make flash's total long-term storage cost-competitive with disk.

III. Business Models For Storage

In this section we present the main economic scenarios for long-term storage.

Rented Storage (e.g. Amazon S3**):** Storage is rented. Rent is based on storage occupancy and retrieval. Rents have decreased over time, but much slower than K_r [28]. The risk in this model is that slight increase in rent or fluctuations in the customer's money supply could lead to permanent data loss.

Monetized Storage (e.g. Google Mail): Storage is free, supported by advertising. As archived data is accessed rarely [16] or primarily by machines there will be little advertising revenue. The risk is that the customer is the advertiser, not the data owner, who has little or no leverage over the storage supplier.

Endowed Storage (e.g POSF [23], Longaccess [11]): In this model the data is deposited in the storage service together with sufficient funds to cover the entire lifetime. Costs are paid from the endowment as they occur; the residue is invested at prevailing interest rates. It faces two opposing risks. One is that the projection of future costs is optimistic and data runs out of money. The other is that it is pessimistic, leading to an endowment so large that no one would pay.

IV. Monte-Carlo Model

Our methodology is based on the endowment model because it allows for true apples-to-apples comparison between different approaches requiring different overlays at different times. To compare them, expenditure must be converted to its Net Present Value (NPV) and summed to obtain the endowment needed to fund storage in that case. The standard technique to do so is Discounted Cash Flow (DCF). It computes the amount which, invested at a chosen constant interest rate (discount) would amount to the expenditure at the time it occurs.

The constant interest rate averages out the effects of period of very high and very low interest rates. Over long periods DCF is thus incorrect. Correct computation of NPV requires Monte-Carlo simulations [22], averaging over repeated runs that choose randomly from a model of varying interest rates [29] to determine the most probable outcome. Our model [29] uses an interest rate model based on the history of US Treasuries [1].

For simplicity, the model used here follows a constant size archive migrating between successive media as they are replaced either because their service life is over or they are no longer cost-effective against more modern media. We ignore the granularity of storage media.

Terms:

Media represents a device in our model. **Media Service Life** is the time after which a device is replaced.
Planning Horizon is the length of time organizations plan for. For example, if a company is planning for the next ten years, its planning horizon is ten years. In the context of our model, when taking decisions on media replacement, costs and benefits beyond the planning horizon are ignored.
Total cost of ownership (TCO) represents the Purchase cost plus the Data Refresh cost plus the Operating cost over the planning horizon or device service life, whichever is shorter.
Data Refresh Costs represent the cost of migrating the data from old device to new device.
Operating Costs models all costs of operating the device except the purchase and data refresh costs.
Duration is the period over which we store the data.

TABLE I. PARAMETER VALUES USED UNLESS OTHERWISE SPECIFIED

Disk Purchase Cost	\$100	Disk Operating cost	\$66
Flash Purchase Cost	\$500	Flash Operating cost	\$20
Disk Service Life	5 years	Flash Service Life	15 years
Duration	100 years	Planning Horizon	20 years

Model Framework:

Disk and flash should be expected to have different K_r values. These studies investigate possible future paths for their costs by treating K_r as an independent variable. The operation of the model is shown in algorithm 1. Default parameters, shown in Table I, are chosen based on current market trends and the expectations of long-term preservation systems. For archival purposes, disk spin down helps on these aspects but it increases the disk failure rate [10]. We chose disk operational costs based on the data from San Diego Computer Center [26]. Flash operational cost is set much lower than disk to reflect lower power, cooling, and space requirements.

V. Simulation Results

Algorithm 1 Monte Carlo Model

Set initial endowment, media density growth rate, purchase cost, operational cost, data refresh cost, planning horizon, and duration.
For each year in the duration:

- Set interest rates.
- Adjust purchase costs, operating costs, and the data refresh costs by K_r.
- Retire expired old media.
- Calculate cost to keep the old media running.
- If old media is not cost-effective (compare amortized upfront cost plus operational costs) in comparison with new media, based on the lesser of the remaining old device's service life or the planning horizon, prematurely retire old media.
- Purchase new media.
- Spend the yearly cost (media replacement plus operating cost plus migration cost) from the starting endowment.
- Earn interest on the remaining endowment.
- Use this as initial endowment for the next year and repeat the process.
- End of duration, if endowment does not run out, reduce the initial endowment and re-run the simulation.

K_r has been predicted to be around .2 [9] and it is predicted that this rate will decrease rapidly in future, due to reasons discussed in Section II. We chose values of K_r ranging from $.2 - .02$ to simulate the range of potential values as storage becomes denser. The error bars represent the standard deviation of the endowment. The lower K_r the more sensitive the endowment is to the variations in the interest rate, so the error bars are bigger. It also shows the uncertainty in the endowment prediction since we do not know what K_r will be in the future. Also, the historical interest rates are taken from the period of $1990 - 2010$, which was a period of wildly varying interest rates. We do not have archival flash yet but we are proposing that such flash can be built, discussed further in Section VI. SanDisk's recent announcement of an SSD with a 10 year long warranty [12] supports this hypothesis.

A. Varying Device Service Life

As shown in Figures 2(a),2(b), flash is expensive in the face of high K_r. For both disk and flash, as K_r goes down, costs increase. And increasing disk service life did not affect endowment much even in the face of low K_r because major contributing factor for disk is its high operating costs.

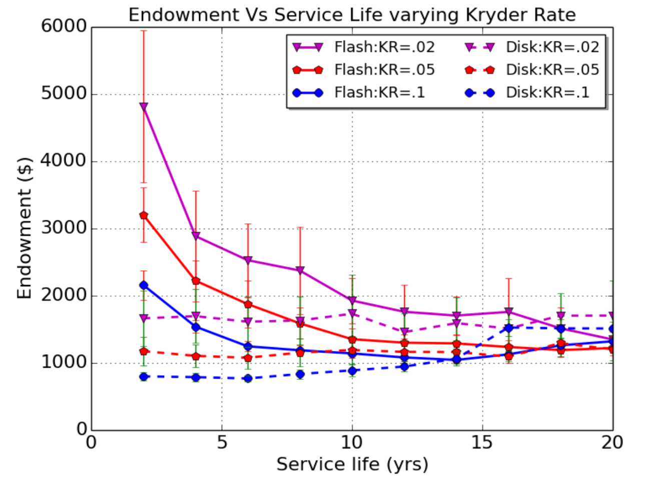

(a) 20 year planning horizon

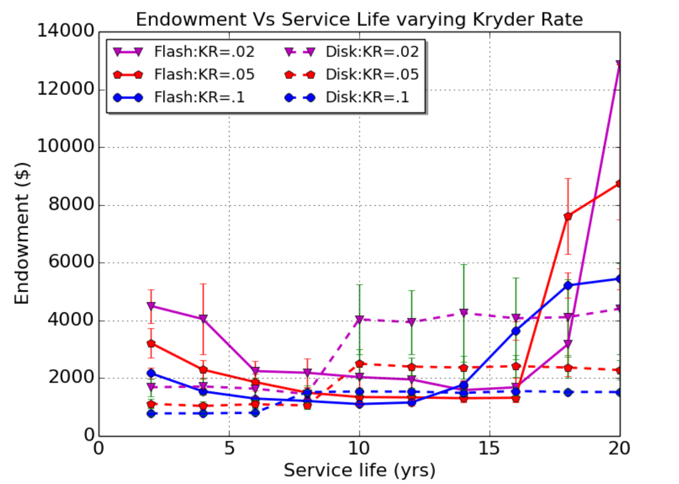

(b) 7 year planning horizon

Fig. 2. As K_r slows down we need longer lived devices to keep the long-term storage costs low.

For a longer horizon, longer than device life time, as shown in Figure 2(a), long-lived flash significantly reduced the costs than short-lived flash. Though the endowment is low for both long-lived flash and long-lived disk, because of reasons discussed in Section I, it is not cost-effective to sell long-lived disks [21]. Because flash's high upfront costs are amortized over long period of time, endowment for flash goes down when run for long.

Figure 2(a) used a planning horizon of 20 years. However, few organizations are able to plan that far ahead. Figure 2(b) shows the effect of short-term planning by imposing a 7-year horizon, beyond which benefits are ignored. For disk, as the life exceeds the planning horizon costs rise because disks are replaced prematurely. For flash, device life needs to exceed the planning horizon by a significant amount because the benefits of low operational cost are greater. As K_r slows down organizations that can plan well ahead would realize benefits from using long-lived flash drives for archival storage.

B. Varying Operating Cost

The operating costs, paid on regular basis, are a major contributor to the costs of long-term preservation. For high values of K_r preservation costs were higher for long-lived flash than those for disk over long-term because of flash's premature replacement in the face of high K_r, as shown in Figure 3; disk was replaced earlier too, however, because of its low upfront costs and high K_r it was cost-effective.

For low K_r, flash was used for its lifetime while disks were still replaced every 5 years, at the end of their service life. For low K_r, flash is comparable to disk only systems. As K_r increase from .1 on wards, long-term costs for flash increase because of premature replacement; the replacement time decreases as K_r increases. Early replacement of flash is caused by cost-effectiveness of new flash devices (purchase cost plus data refresh costs amortized over their service life and operational costs over their service life) against old flash devices (purchase cost plus data refresh costs amortized over their service life plus operational costs for the remaining service life). At the time of device replacement purchase cost and data refresh costs are paid from endowment. Earlier the flash is replaced, the expensive it is in long-term.

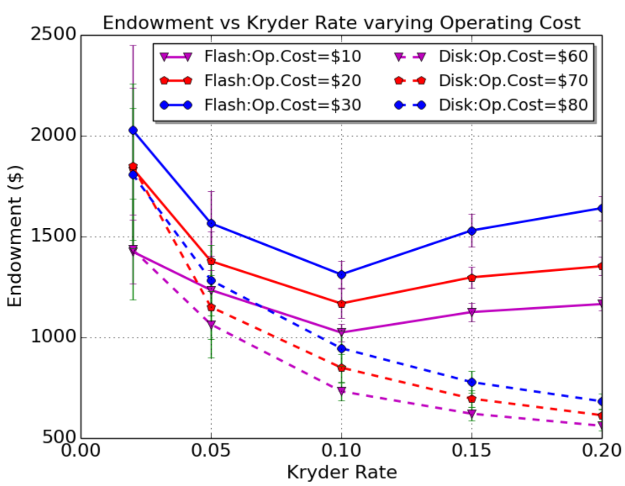

Fig. 3. Flash is cost-competitive to disk in the face of low K_r irrespective of varying operational costs.

C. Varying Purchase Costs

Figure 4 shows the endowment increases steeply as K_r decreases, for both flash and disk. The long-term cost of storing the data with flash for 100 years with $K_r = .2$ is more than with $K_r = .1$ because it was getting replaced earlier than its lifetime. For disks the long-term cost increases continuously as K_r slows down because they are getting replaced in 5 years even in the face of low K_r.

However, for flash with a service life 15 years, costs are comparable to those of disk having service life 5 years in the face of low K_r. Even if disk cost-per-byte comes down, flash is giving cost competition to disk for long-term preservation because of reduction in device density combined with longer lifetime.

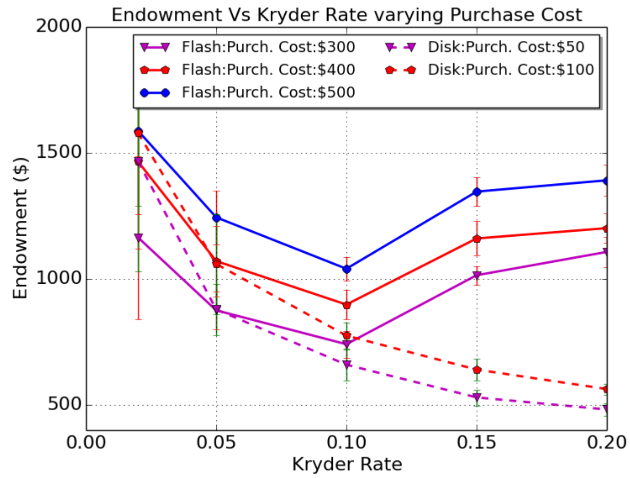

Fig. 4. For all purchase costs, long-term preservation costs for the flash were as good as disk in the face of low media density growth rates.

We experimented with disk purchase costs \$50, \$100 and flash purchase costs \$300, \$400, and \$500 for K_r varying from $0.2 - .02$ to reflect real world. This experiment confirms that a decreasing Kryder's rate has great influence on long-term preservation cost. For varying purchase costs, flash is cost-competitive to disk in long-term preservation systems in an environment of low media density growth rates.

D. Varying Duration

As shown in the Figure 5, K_r slown down affects endowment a lot for long duration. For $K_r = .1$, the graph is almost flat for durations longer than 20 years for both flash and disk because devices are 10% denser every year, making newer devices more cost-effective. This confirms the assumption we all have that if Kryder's law continues, the storage portion of the long-term preservation cost will become negligible.

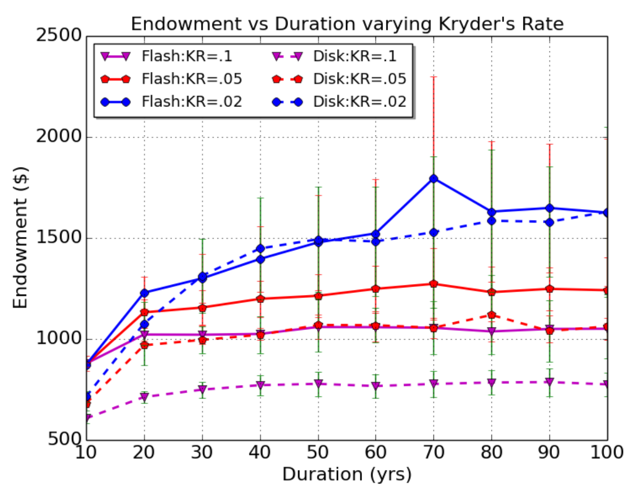

Fig. 5. Flash costs more in initial years, however, it is as cost-effective as disk for long-term.

Doing the same calculation for $K_r = .02$ has a tremendous effect on the endowment as shown in the graph. The curve is not flat anymore which suggests storage costs will not be negligible anymore in long-term. For $K_r = .02$, long-term costs with flash are comparable to long-term costs with disks because the flash devices had a longer life time and lower operating costs than disk. For $K_r = .05$ and $K_r = .02$, the error bars are big, due to the effect of varying interest rates on the endowment calculation in the face of low K_r, but the trend of growth is evident.

E. Varying Data Refresh cost

Labor is the major contributor to data refresh costs, which requires someone to change disks and do the data migration. As devices get denser, more data fits in fewer devices, which brings down the data refresh cost per byte.

Figure 6(b) shows the increase in endowment as data refresh gets expensive. Endowment increased for duration up to 20 years for both flash and disk because of high device replacement and initial operating costs, but in the long-term we do not see much increase because data is already much smaller than the storage capacity. For high K_r (.1 here), flash gets replaced pre-maturely because newer devices became more cost-effective, that increase the long-term preservation cost. Overlapping error bars in Figure 6(a) show for one particular media data refresh costs do not matter much in long-term. Also, flash which was expensive in the face of high K_r is cost-competitive to disk for all data refresh costs.

VI. DISCUSSION

The slow down of K_r will have a negative impact on the long-term storage costs, as confirmed by our experiments. As of 2014, disk prices are predicted to to level off over time [27]. This will hit long-term costs hard at the time of disk replacement. To increase the disk density further, approaches like adding more platters or changing the recording technology to SMR [18], HAMR [25], BMP [4], and MAMR [32] are not likely to return the cost-per-byte down to historical rate.

SSDs store data using electrical charges which leak out due to low insulation which is kept low to achieve high read/write IOPS. Insulation can be made stronger for archival storage which will keep the data intact for longer time and it does not add much to the cost [21]. However, increasing flash life-time further would also need better electronics in the controller. Controllers are optimized for high performance of SSD, because that is what SSD is primarily used for as of today. However, they can be changed for archival purposes by compromising on write endurance and increasing data retention. Wear leveling, strong insulation, and using fewer P/E cycles with flash can increase the data retention for minimal additional cost. Therefore, we believe it is plausible for SSDs with a longer life to enter the archival market, without increasing the manufacturing cost a lot. SanDisk recently announced an SSD with a 10 year long warranty [12]. These SSD are expected to cost \$599 for 960GB.

SSDs are relatively recent and just beginning to penetrate the consumer market [3]. SSDs also have long-term benefits with small form factor, low power, durability, and longer data retention period. As the demand for cheaper high capacity SSDs grow there will be more effort to bring down the cost.

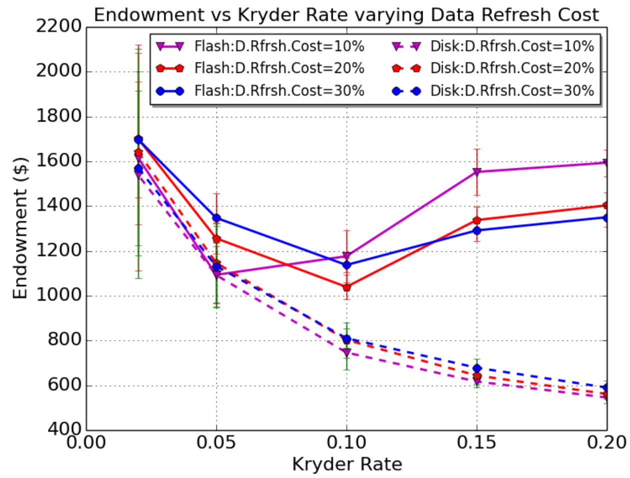

(a) Varying Kryder rate

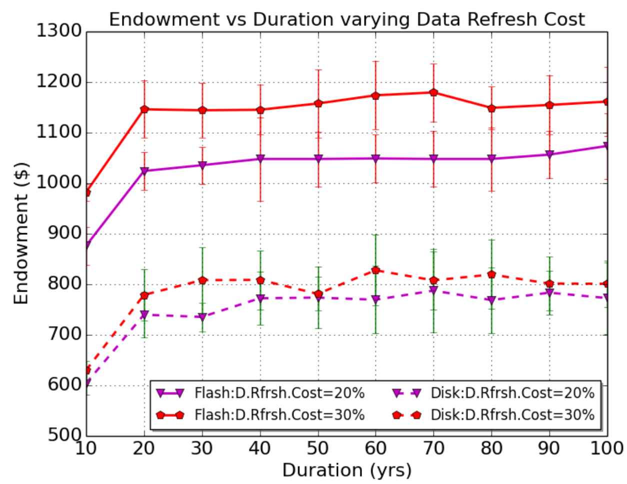

(b) Varying duration ($K_r = .1$)

Fig. 6. Endowment increases with increasing data refresh costs for both flash and disk.

Theoretically, however, SSD can fit more capacity in less space because it does not need any mechanical components to read/write. An example of this is the Flash MicroSD card, announced recently, that stores 128GB [5]. We showed that long-lived flash (15 years life-time) is as effective as disk with carefully chosen long planning horizon. However, if flash can be designed to live longer than 15 years,or have even lower operational costs, it wins over disk. Therefore, we believe that archival flash or alternative technologies, with long life times and low operational costs, will be more cost-effective than traditional disk archives for long-term preservation. Disk and flash may not be existing for 100years. We tried to show that any technology that can live longer, and is cheaper to operate needs to be used for long-term preservation.

VII. FUTURE WORK

A long-term preservation system must balance the trade-off between redundancy and longevity of the data when money is running low. Running the devices longer than their service lives compromises reliability, while reducing the number of devices to bring down the cost compromises redundancy. This question is important as we may run out of money due to technology changes or unpredicted fluctuations in the money supply.

Also, we would like to look into optimizing flash controllers for archival purposes and study its effect on endowment. We also plan to consider other reliability models for long-term storage systems. Also, to build a long-term preservation system it is important to know whether to invest in building an in-house solution or store data in the cloud. The cloud is popular today because of the convenience, performance, and reliability it gives to the consumer. However, its suitability for long-term preservation has not been analyzed yet. Given the decline in storage density growth, we would like to compare cloud storage for long-term preservation of the data with in-house solutions.

VIII. Conclusion

HDD areal density growth have been slowing down, most recently in the range of $20-25\%$ annually which shows the decline of K_r. Most current storage service providers assume continuation of Kryder's law, which indicates storage media costs will be irrelevant over time. Our results show that as K_r slows down there will be an economic crisis in the long-term preservation of data.

We showed SSDs (having 15 years life-time) are cost-competitive to traditional disk archives for long-term preservation, because of their low operating costs. However, we believe that flash or other alternative technologies, having life time longer than 15 years and even lower operational costs, will be highly cost-effective for long-term preservation. Given these observations we propose variant of flash storage optimized for archival purposes for long-term data preservation.

Acknowledgments

Authors would like to thank our colleagues in Storage Systems Research Center at University of California, Santa Cruz for their valuable feedback on the ideas in this paper. A special thanks goes to Dr. David S.H. Rosenthal for the support and guidance he provided during the course of this work. This research was based on work supported in part by the DOE under Award Number DE-FC02-10ER26017/DE-SC0005417 and NSF under award IIP-1266400 and industrial members of the Center for Research in Storage Systems.

References

[1] Daily Treasury Yield Curve Rates. http://www.treasury.gov/resource-center/data-chart-center/interest-rates/Pages/TextView.aspx?data=yield.

[2] Disk Drive Prices (1955-2013). http://www.jcmit.com/diskprice.htm.

[3] Invest in New Technologies or Divest in Market Share. http://www.tomcoughlin.com/Techpapers/Invest%20in%20New%20Technologies%20or%20Divest%20in%20Market%20Share,%20090210.pdf.

[4] Patterned Magnetic Media. https://www1.hgst.com/hdd/research/storage/pm/.

[5] SANDISK INTRODUCES WORLDS HIGHEST CAPACITY microSDXC MEMORY CARD AT 128GB. http://www.sandisk.com/about-sandisk/press-room/press-releases/2014/sandisk-introduces-worlds-highest-capacity-microsdxc-memory-card-at-128gb/.

[6] Archive Drive Study. 2009. http://www.digitalpreservation.gov/meetings/documents/othermeetings/5-4_Anderson-seagate-v3_archive_study.pdf.

[7] Blue Ribbon Task Force on Sustainable Digital Preservation and Access. April 2010. http://brtf.sdsc.edu/biblio/BRTF_Final_Report.pdf.

[8] Data to grow more quickly says IDC's Digital Universe study. Dec 2012. http://www.computerweekly.com/news/2240174381/Data-to-grow-more-quickly-says-IDCs-Digital-Universe-study.

[9] The Impact of Areal Density and Millions of Square Inches (MSI) of Produced Memory on Petabyte Shipments for TAPE, NAND Flash, and HDD Storage Class. 2013. http://storage-conference.org/2013/Presentations/Fontana.pdf.

[10] Backblaze stats show most/least reliable hard drives: Hitachi leads the pack with lowest annual failure rate. 2014. http://www.tuaw.com/2014/01/21/backblaze-stats-show-most-least-reliable-hard-drives-hitachi-le/.

[11] LongAccess Wants To Cold-Store Your Digital Life For 30 Years. 2014. https://www.longaccess.com/pricing/.

[12] SanDisk announces Extreme Pro enthusiast SSDs with crazy long 10-year warranties. 2014. http://www.pcworld.com/article/2357767/sandisk-announces-extreme-pro-ssd-series-backs-drives-with-10-year-warranties.html#tk.fb_pc.

[13] WHOMP! There it is: IBM demos 154TB tape. May 2014. http://www.theregister.co.uk/2014/05/19/ibm_demos_154tb_tape/.

[14] Why Facebook thinks Blu-ray discs are perfect for the data center. Jan 2014. http://arstechnica.com/information-technology/2014/01/why-facebook-thinks-blu-ray-discs-are-perfect-for-the-data-center/.

[15] S. Abrams, P. Cruse, J. Kunze, and M. Mundrane. Total Cost of Preservation (TCP): Cost Modeling for Sustainable Services, 2012.

[16] I. Adams, E. L. Miller, and D. S. Rosenthal. Using Storage Class Memory for Archives with DAWN, a Durable Array of Wimpy Nodes. Technical report, Technical Report UCSC-SSRC-11-07, University of California, Santa Cruz, 2011.

[17] M. Addis and M. Jacyno. Tools for modelling and simulating migration based preservation. Project report, December 2010.

[18] A. Amer, J. Holliday, D. D. Long, E. L. Miller, J. Paris, and T. Schwarz. Data management and layout for shingled magnetic recording. *Magnetics, IEEE Transactions on*, 47(10):3691–3697, 2011.

[19] N. Beagrie, M. Duke, M. Patel, L. Lyon, C. Hardman, D. Kalra, B. Lavoie, and M. Woollard. The krds benefit analysis toolkit: Development and application. *International Journal of Digital Curation*, 7(2):64–67, 2012.

[20] S. Chapman. Counting the costs of digital preservation: is repository storage affordable? *Journal of digital information*, 4(2), 2006.

[21] M. Cornwell. private communication, 2014.

[22] D. J. Farmer and J. Geanakoplos. Hyperbolic discounting is rational: Valuing the far future with uncertain discount rates. 2009.

[23] S. J. Goldstein and M. Ratliff. DataSpace: A Funding and Operational Model for Long-Term Preservation and Sharing of Research Data. 2010.

[24] U. B. Kejser, A. B. Nielsen, and A. Thirifays. Cost model for digital preservation: Cost of digital migration. *International Journal of Digital Curation*, 6(1):255–267, 2011.

[25] M. H. Kryder, E. C. Gage, T. W. McDaniel, W. A. Challener, R. E. Rottmayer, G. Ju, Y.-T. Hsia, and M. F. Erden. Heat assisted magnetic recording. *Proceedings of the IEEE*, 96(11):1810–1835, 2008.

[26] R. L. Moore, J. D'Aoust, R. H. McDonald, and D. Minor. Disk and tape storage cost models. *Archiving 2007*, 2007.

[27] Robin Harris. Hard drive prices (and innovation) decline. April 2013. http://www.zcoughlin_reportdnet.com/hard-drive-prices-and-innovation-decline-7000013399/.

[28] D. S. Rosenthal. Talk at IDCC2013. Januray 2013. http://blog.dshr.org/2013/01/talk-at-idcc2013.html.

[29] D. S. Rosenthal, D. C. Rosenthal, E. L. Miller, I. F. Adams, M. W. Storer, and E. Zadok. The economics of long-term digital storage. *Memory of the World in the Digital Age, Vancouver, BC*, 2012.

[30] C. Walter. Kryder's law. *Scientific American*, page 293, July 2005. http://www.scientificamerican.com/article.cfm?id=kryders-law.

[31] P. Wheatley and B. Hole. LIFE3: Predicting Long Term Digital Preservation Costs. 2009.

[32] J.-G. Zhu, X. Zhu, and Y. Tang. Microwave assisted magnetic recording. *Magnetics, IEEE Transactions on*, 44(1):125–131, 2008.