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 • Andy Hospodor

SSRC Storage Systems Research Center

Center for Research in Storage Systems

The University of California · Let There Be Light · 1868

Baskin Engineering
UC Santa Cruz

What's the trouble with media for digital preservation?

- Growth in storage media density is slowing
  - Disks still have to be replaced every 5 years
    - Other media may not need to be replaced as often
  - Tapes require a lot of infrastructure
    - Readers (and maybe robots) are expensive
    - Infrastructure is personnel-intensive
- No longer true that replacing media is a big win
  - Replacement media aren't much larger than older ones!
  - We'd like to keep the older ones if they'll continue to work…
- Is disk still the cheapest way to build long-term archives?
  - What about flash memory?

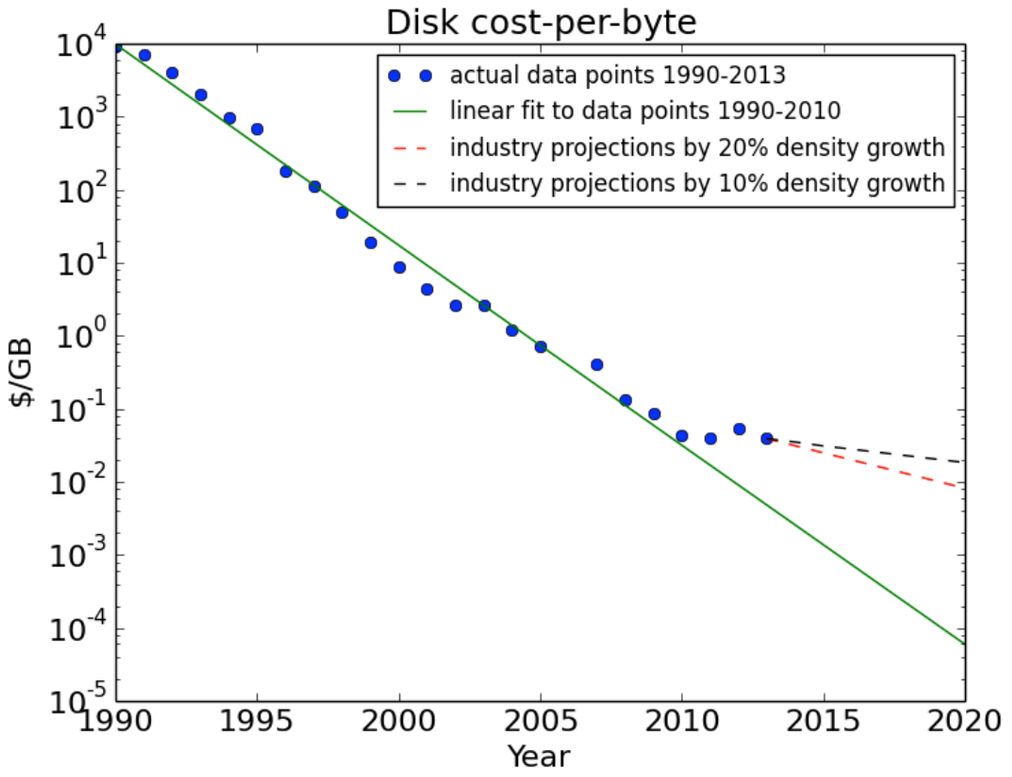

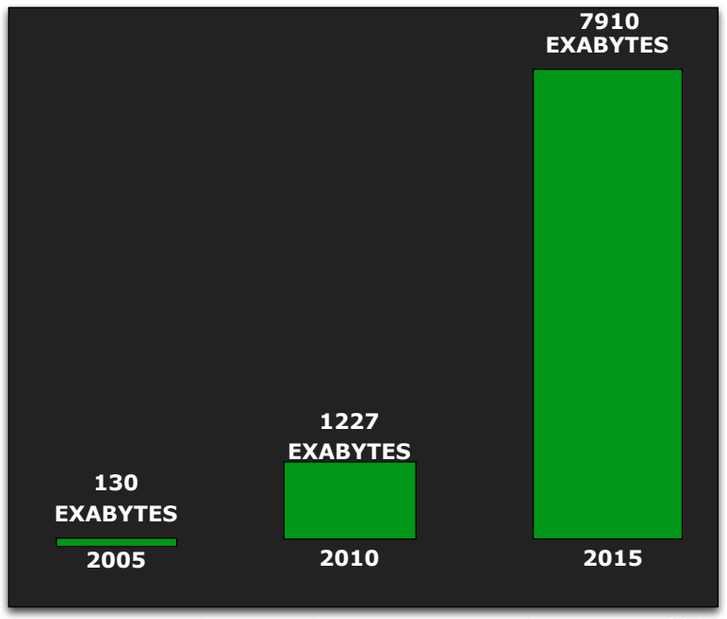

Credit: http://www.emc.com/collateral/about/news/idc-emc-digital-universe-2011-infographic.pdf

CENTER FOR RESEARCH IN STORAGE SYSTEMS

Comparison of storage media

	Tape	Disk	SSD	Optical disk
Bandwidth	Very high	High	High	Moderate
IOPS	Very low	Moderate	High	Very low
Reader?	Yes	No	No	Yes
Reliability	Low-moderate	Moderate	High	Low-moderate
Operational cost	Moderate	Moderate	Low	Moderate
Capital cost	Low	Low	**High**	Low
Longevity	Long*	**Short**	Moderate	Long* (?)
Density	Moderate	High	High	Moderate

How do we model archive cost?

- Need to use a model!
  - Can't actually build the archive (time, money)
  - Compute costs based on future assumptions

- Model inputs (for a given type of device)
  - Media service life: maximum lifetime of a device
  - Operational costs: cost of operating a device (cooling, power, and space)
- Overall model inputs (for an entire archive)
  - Data refresh costs: cost to migrate data from one device to another
  - Planning horizon: length of time over which organization considers costs

- But where does the money come from?

Storage business models

- Rented storage: pay as you go
  - Amazon S3
- Monetized storage: paid by ads
  - Google Mail
- Retrieval model: pay on retrieval
  - Only a small fraction of data is ever retrieved…
- Endowed storage: pay up front
  - Princeton's POSF
- Model uses endowed storage
  - Amount of endowment is a parameter

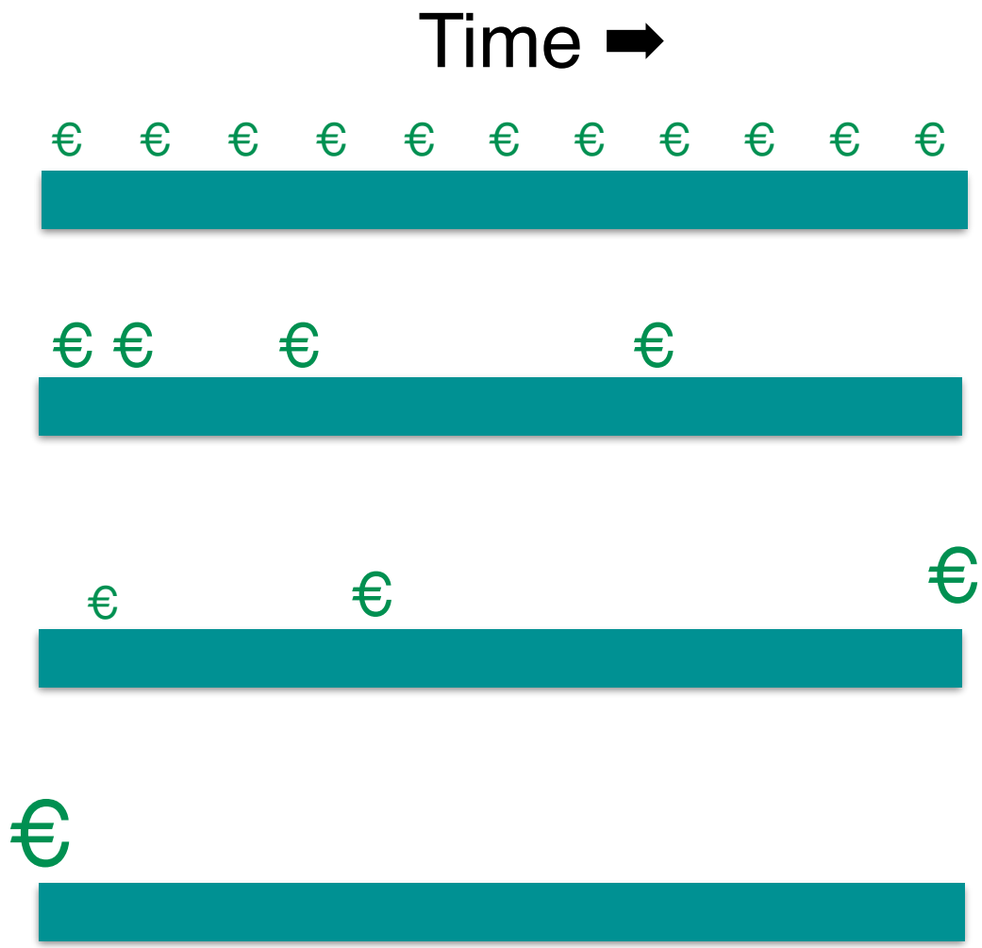

Using the economic model

- Monte Carlo simulations to calculate the endowment
  - Money grows by varying amounts each year, as in the real world
  - Variation in interest rates affects long-term success!

- Replace devices when they do not justify their running costs
  - Consider space & power consumption
  - Look forward for the duration of the planning horizon
  - Pay for new devices using money from the endowment
  - Simulation ends (in failure) if the money runs out

- Costs drop by the storage density growth rate each year
  - New storage is cheaper to purchase and operate
  - Existing storage doesn't get any cheaper!

Simulation assumptions

Default Parameters			
Disk cost	$100	Disk service life	5 years
Flash cost	$500	Flash service life	15 years
Disk operational cost	$60/year	Simulation duration	100 years
Flash operational cost	$20/year	Storage density growth rate	15% / year

- Flash follows same storage density growth curve as disk
  - Both flash and disk are experiencing much slower growth moving forward

Varying service life

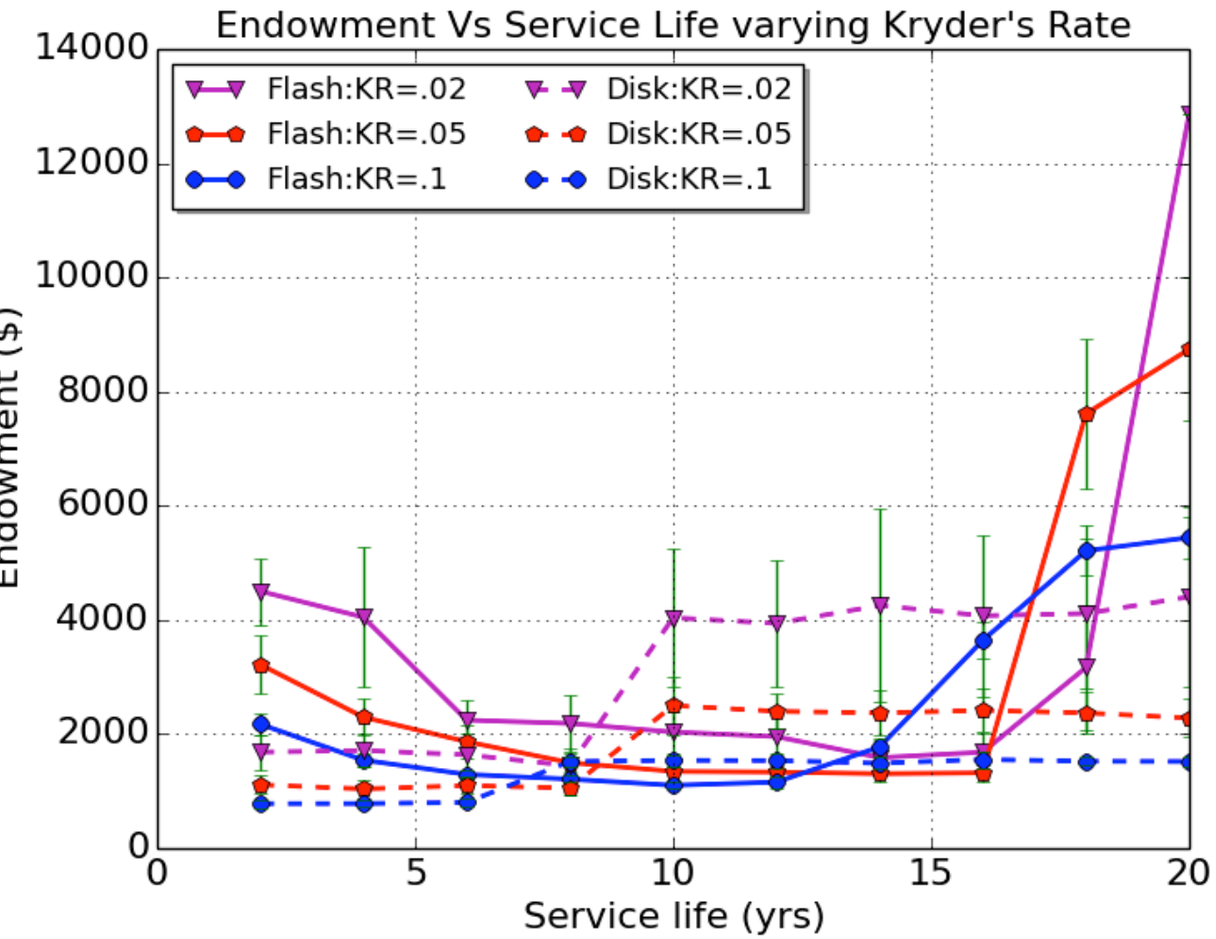

Planning Horizon - 7 yrs

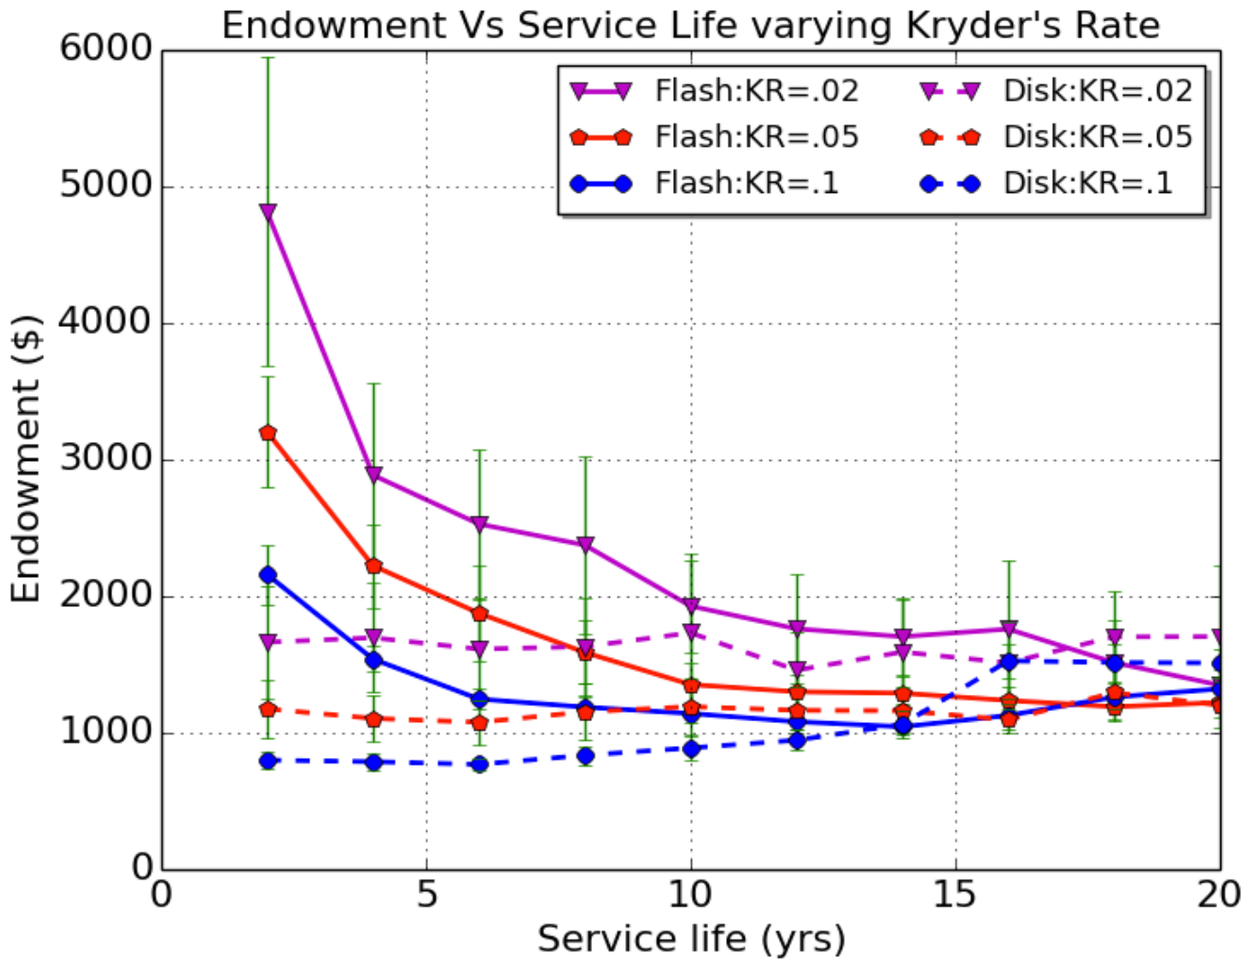

Planning Horizon - 20 yrs

- Old devices need to be kept for long time to keep endowment low
  - Purchasing new devices requires more money from the endowment
- Endowment for flash with service life 10 years is comparable to disk with service life 5 years
  - Need for “Archival Flash”

Varying purchase cost

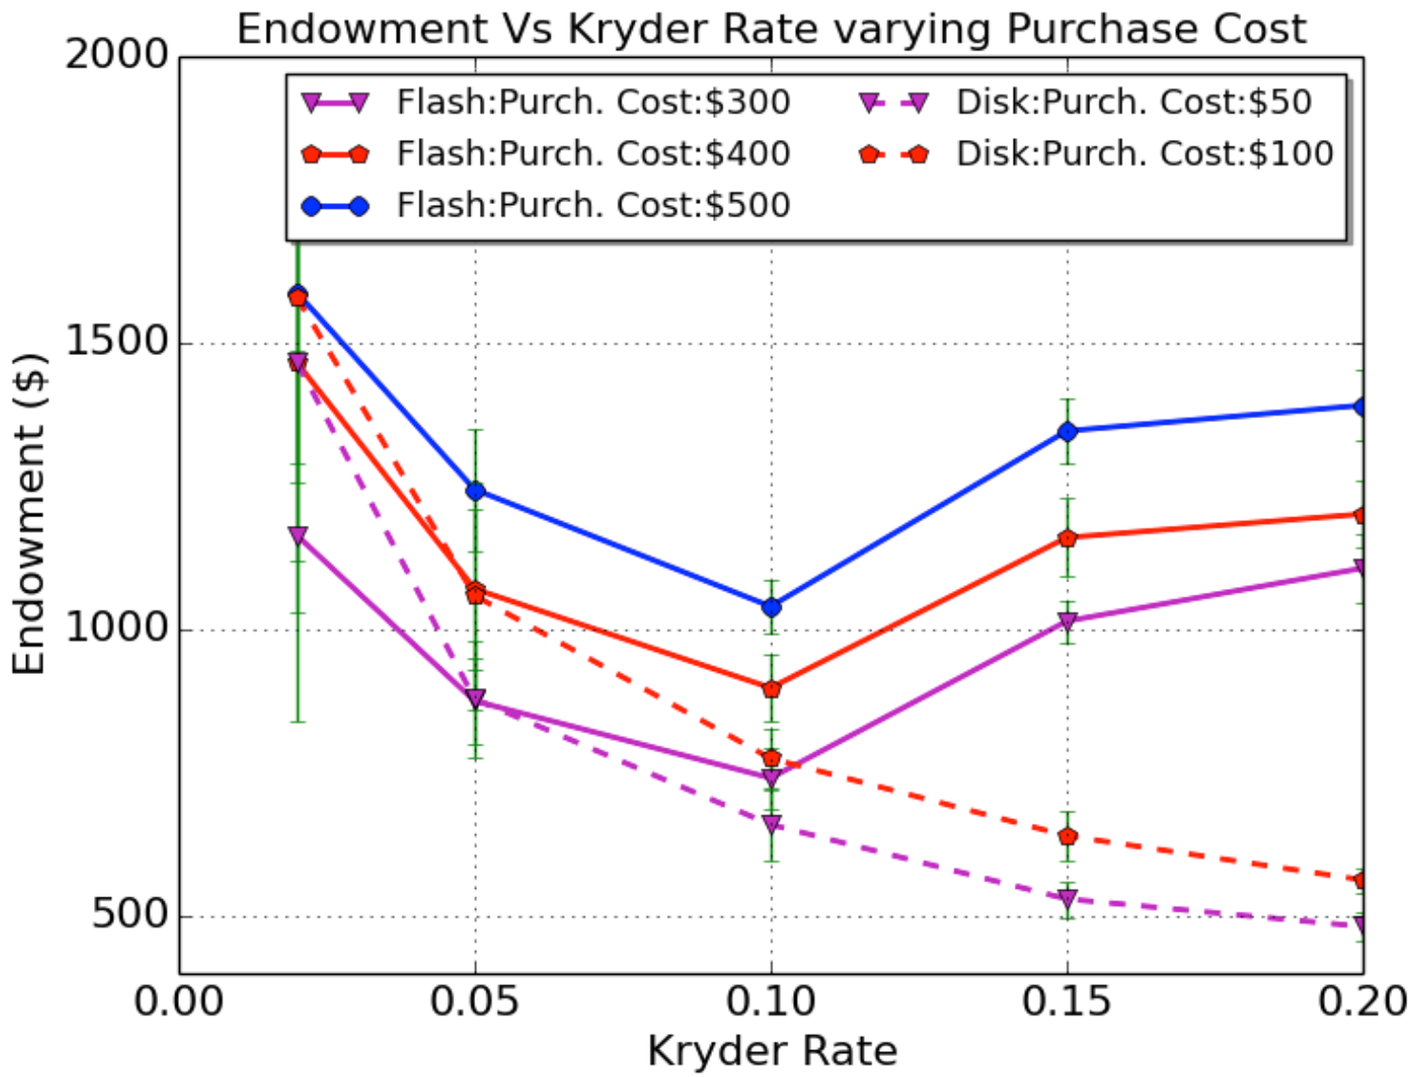

- Low storage growth rate ⟹ initial cost is less important
  - Slight variation in purchase cost doesn't affect long-term preservation costs if the Kryder rate is low

Discussion

- Flash can be cost-competitive with disk
  - Flash can be made to live longer without much extra cost, unlike disk
  - Alternative technologies like memristors and phase-change memory will evolve in the future

- We also found that
  - Operational expenses dominate as storage gets denser
  - Capital expenses are less important as storage growth rates slow, especially for longer-lived devices
  - No time to describe these graphs—ask after the talk

- **Digital preservation community should consider alternative technologies to traditional disk (and tape)**

Ongoing work

- What are the chances of data loss?
- What are the long-term economics of various reliability models?
  - Choose redundancy or longevity?
- What are the best devices to use, given a certain capacity to bandwidth ratio?
- On-demand or on-premise access?

- Building model to accommodate wide range of variation in input parameters
  - Capacity, per-device cost, bandwidth: ***rate*** of change varies with time
  - Annual failure rate (varying with device age)
- Evaluate performance as well as cost
  - Meet given performance / capacity goals at minimum cost
  - How performant is the cheapest archive with a given cost?

Questions or comments ?

Further information at:

http://www.ssrc.ucsc.edu/proj/archive.html

Varying archive lifetime

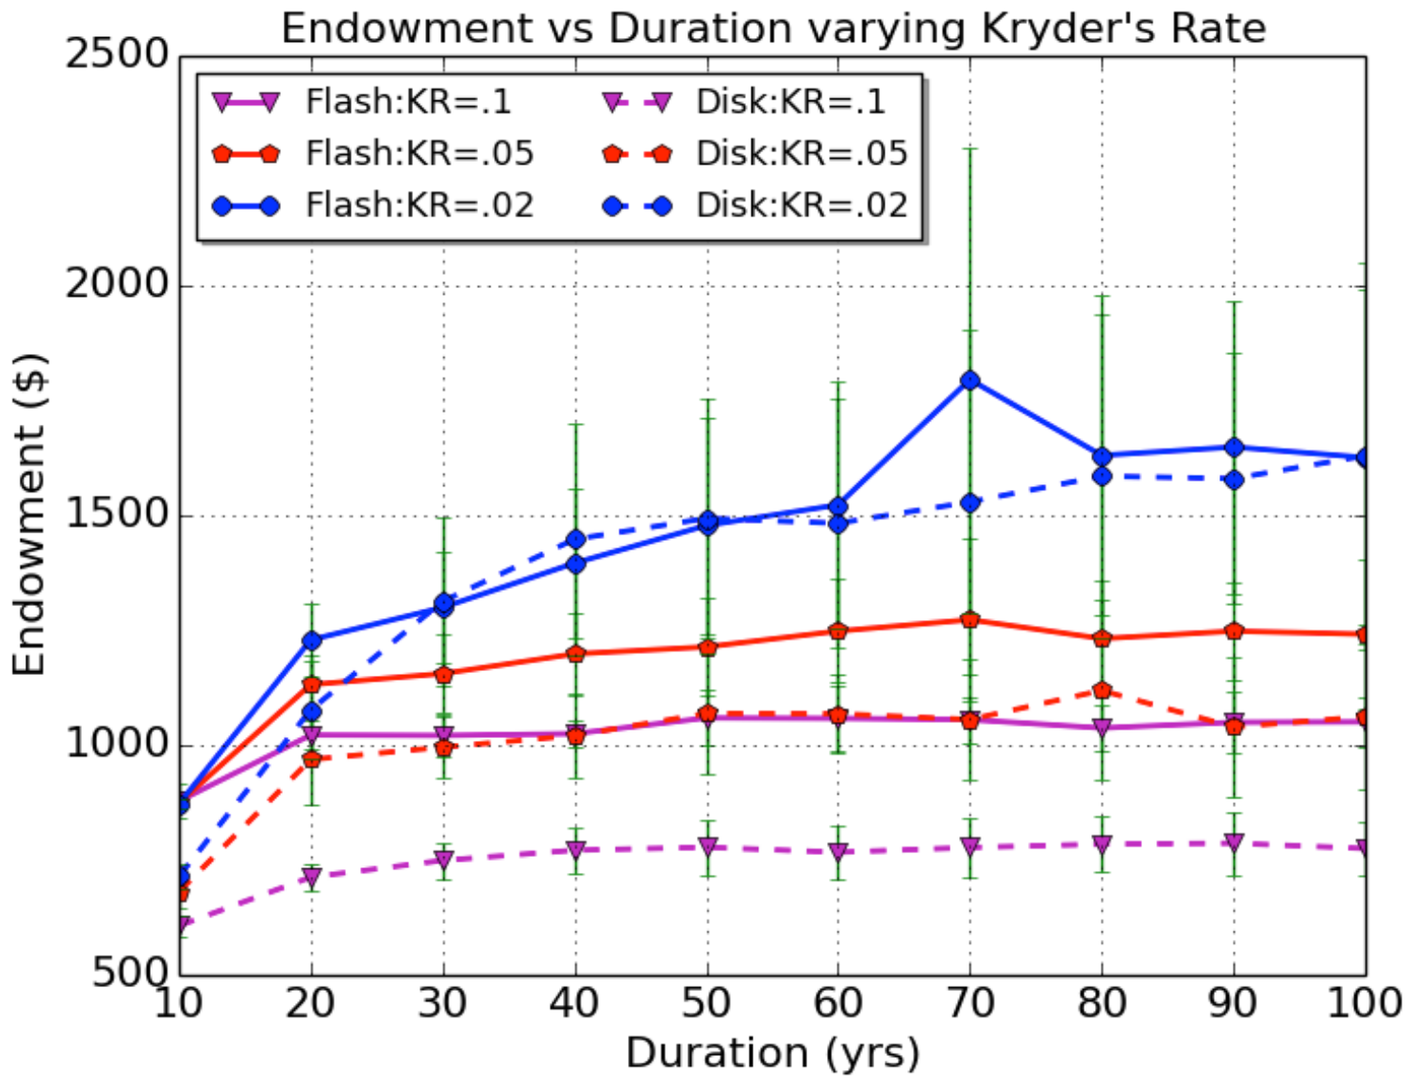

- Declining storage density growth rate negatively impacts long term preservation costs
  - Effect more pronounced for long duration
- Flash is cost-effective for longer duration!

Varying operational costs

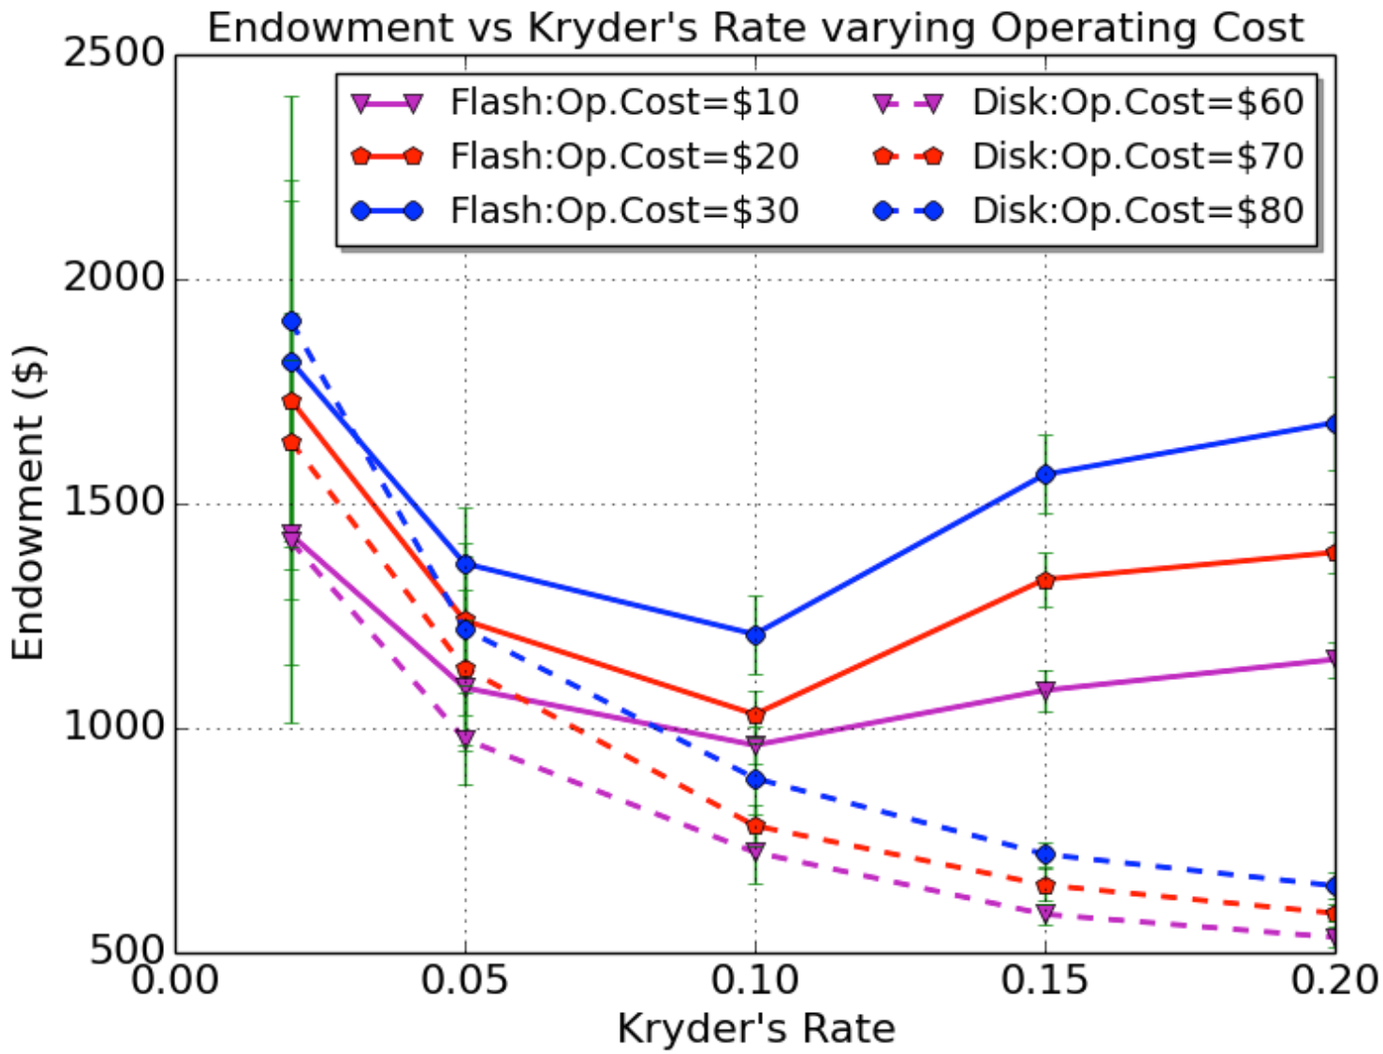

- Reduced operational cost critical as Kryder rate drops
- Technologies with low operational costs can be as good as traditional disk archives

Varying data refresh costs

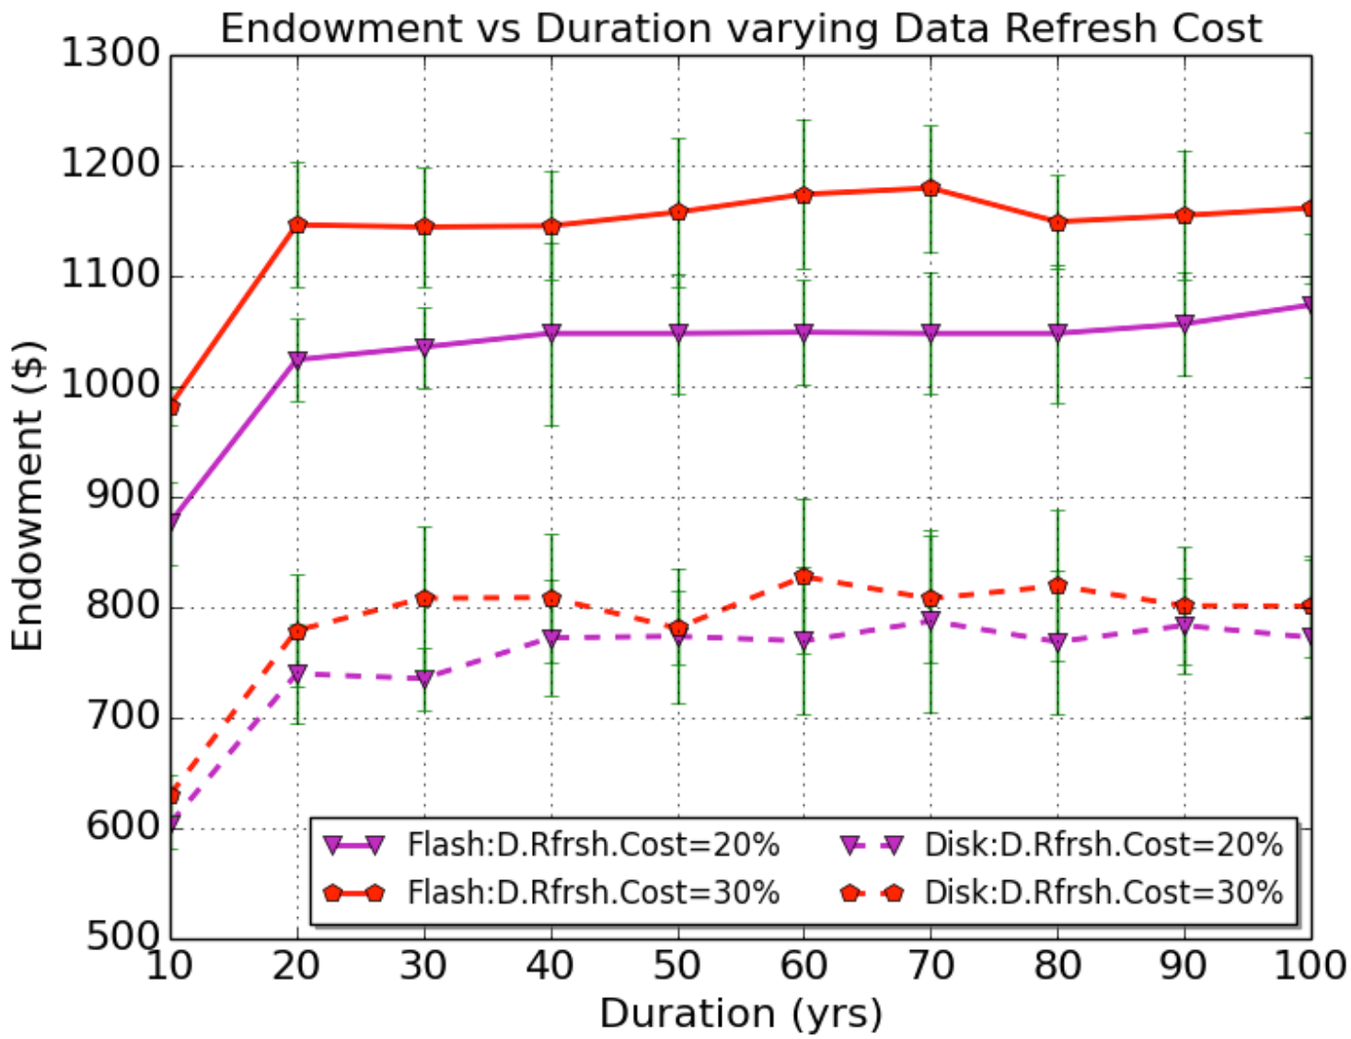

- Data refresh costs are significant in the initial years
  - Frequent device replacement and high operational costs (for both flash and disk)
- Data refresh costs matter little after 20 years
  - Data size is much smaller relative to new device capacity