



Effects of Prolonged Media Usage and Long-term Planning on Archival Systems


Preeti Gupta, **Avani Wildani**¹, Ethan L. Miller,
David S.H. Rosenthal², Darrell D.E. Long

MSST 2016

UC Santa Cruz

¹ Emory University

² Stanford Libraries



EMORY
COLLEGE
OF ARTS AND
SCIENCES



Baskin
Engineering
UC SANTA CRUZ



Preserving Data for the Long Haul

- ❖ We make a lot of data, and prioritization is hard
- ❖ Techniques
 - Disk scrubbing
 - Media replacement
 - Fund management
- ❖ Issues
 - Unpredictable data growth
 - Varying duration requirements
 - Technology shifts



Tablet of Shamash



image(c)<http://www.flickr.com/photos/jvdc/>
Sumerian Cuneiform

Who Cares?

- ❖ Libraries / Museums
- ❖ Government Compliance
- ❖ Health Records
- ❖ Scientific Instruments



Key Questions


- ❖ **When do you retire/replace media?**
 - Manufacturers suggest conservative lifetimes
 - Devices improve unevenly and non-linearly
- ❖ **How long should you plan for?**
 - What range are you optimizing over?

Failure Scenarios


- ❖ **Data is lost through device failure**
 - Traditional reliability models
- ❖ **Data is lost through economic failure**
 - Monument models
- ❖ **Data is lost through regulatory failure**
 - Copyright police (future work)

Media Replacement Failure Rates

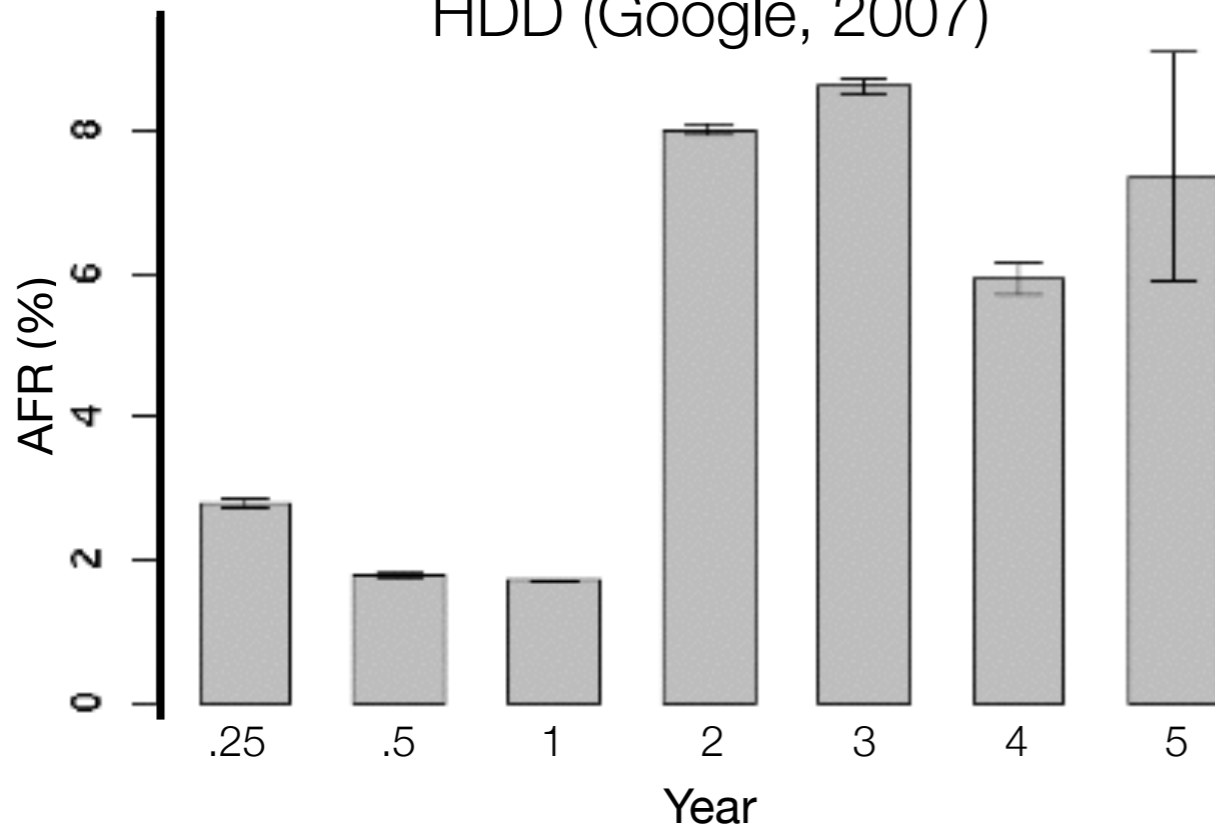
SSD (Intel, 2014)



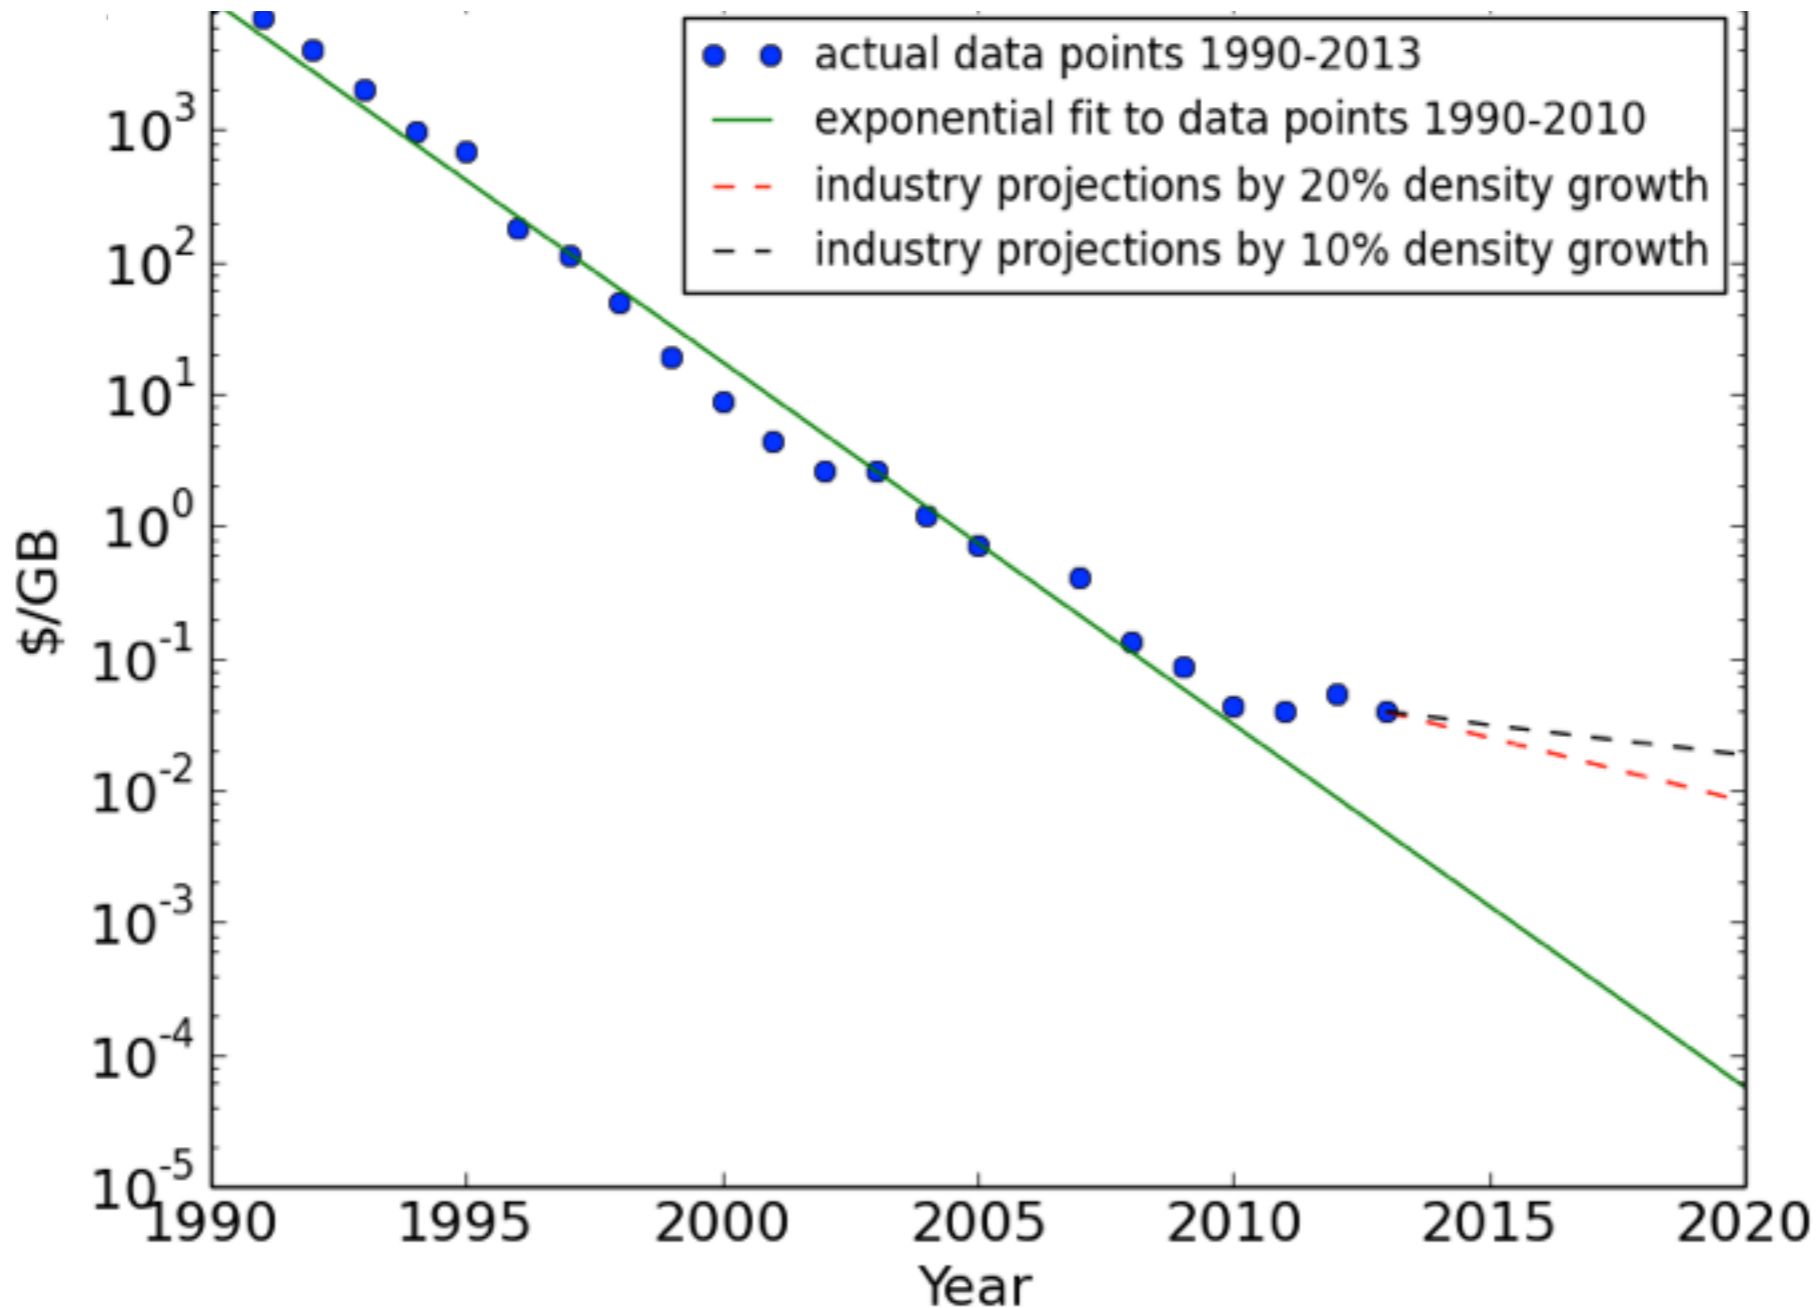
HDD (BackBlaze, 2015)



HDD (Google, 2007)



Media Replacement Cost Improvements



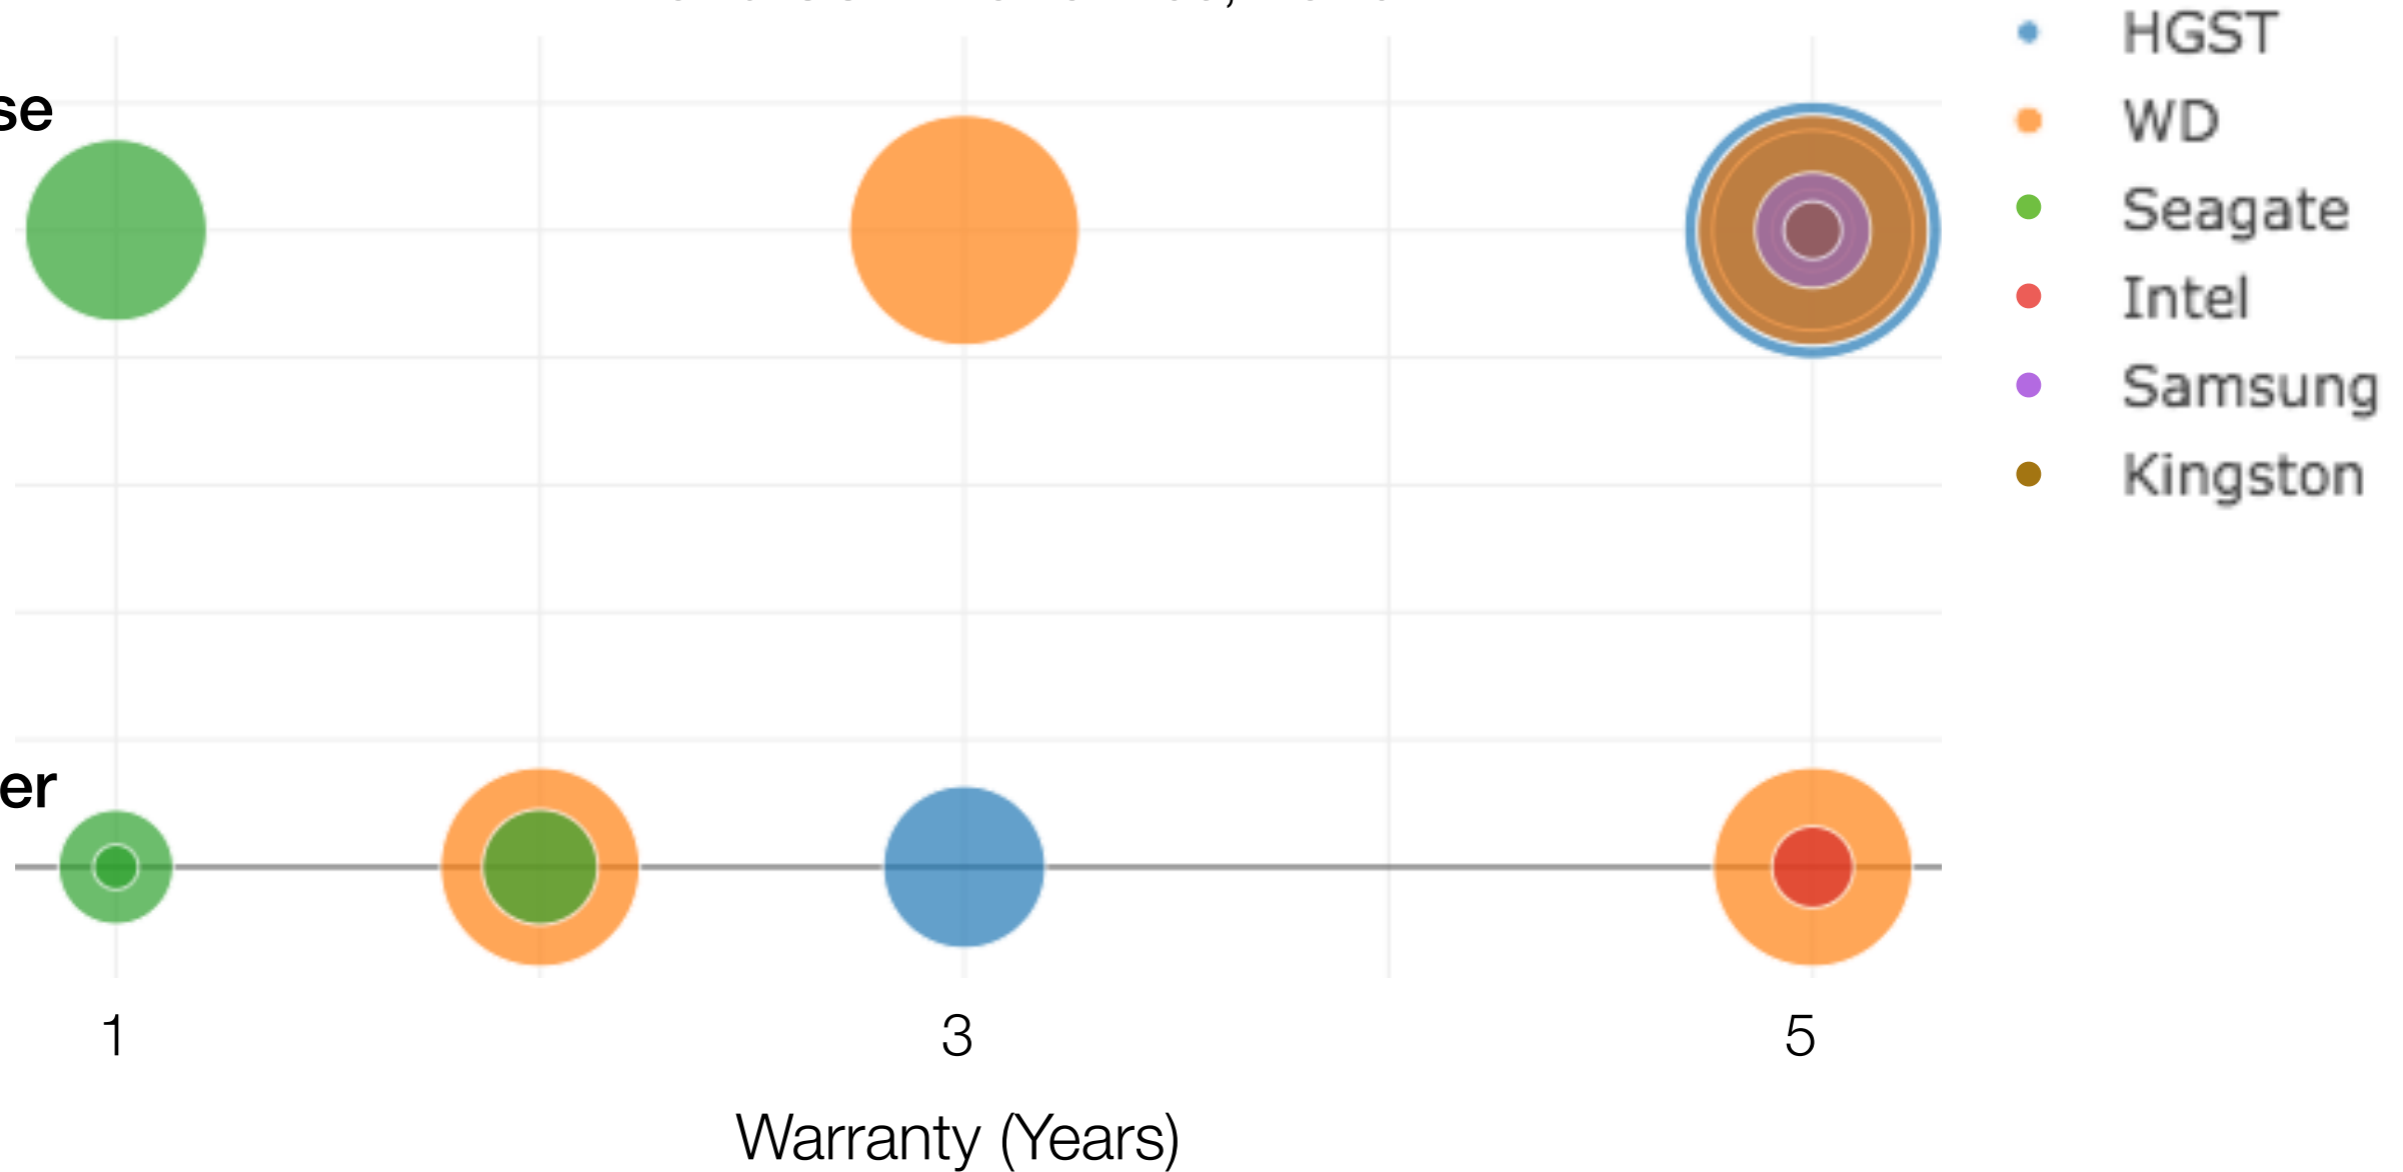
Gupta, P. *et al.* "An economic perspective of disk vs. flash media in archival storage." *Modelling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS), 2014 IEEE 22nd International Symposium on*, 2014.

Media Replacement Manufacturer Warranties

HDD and SSD Warranties, 2016

Enterprise

Consumer




Circle size corresponds to drive size

Media Replacement Core Questions

- ❖ Should media be used past their manufacturer suggested service life or warranty period?
- ❖ What limits typical service life to 3-5 years?
 - HGST Assumptions:
 - Less than 333 power on hours per month.
 - Seeking/Writing/Reading operation is less than 20% of power on hours. (~2hr/day)
 - **For archival usage patterns SSDs and HDDs may last longer**

How long should we plan for?



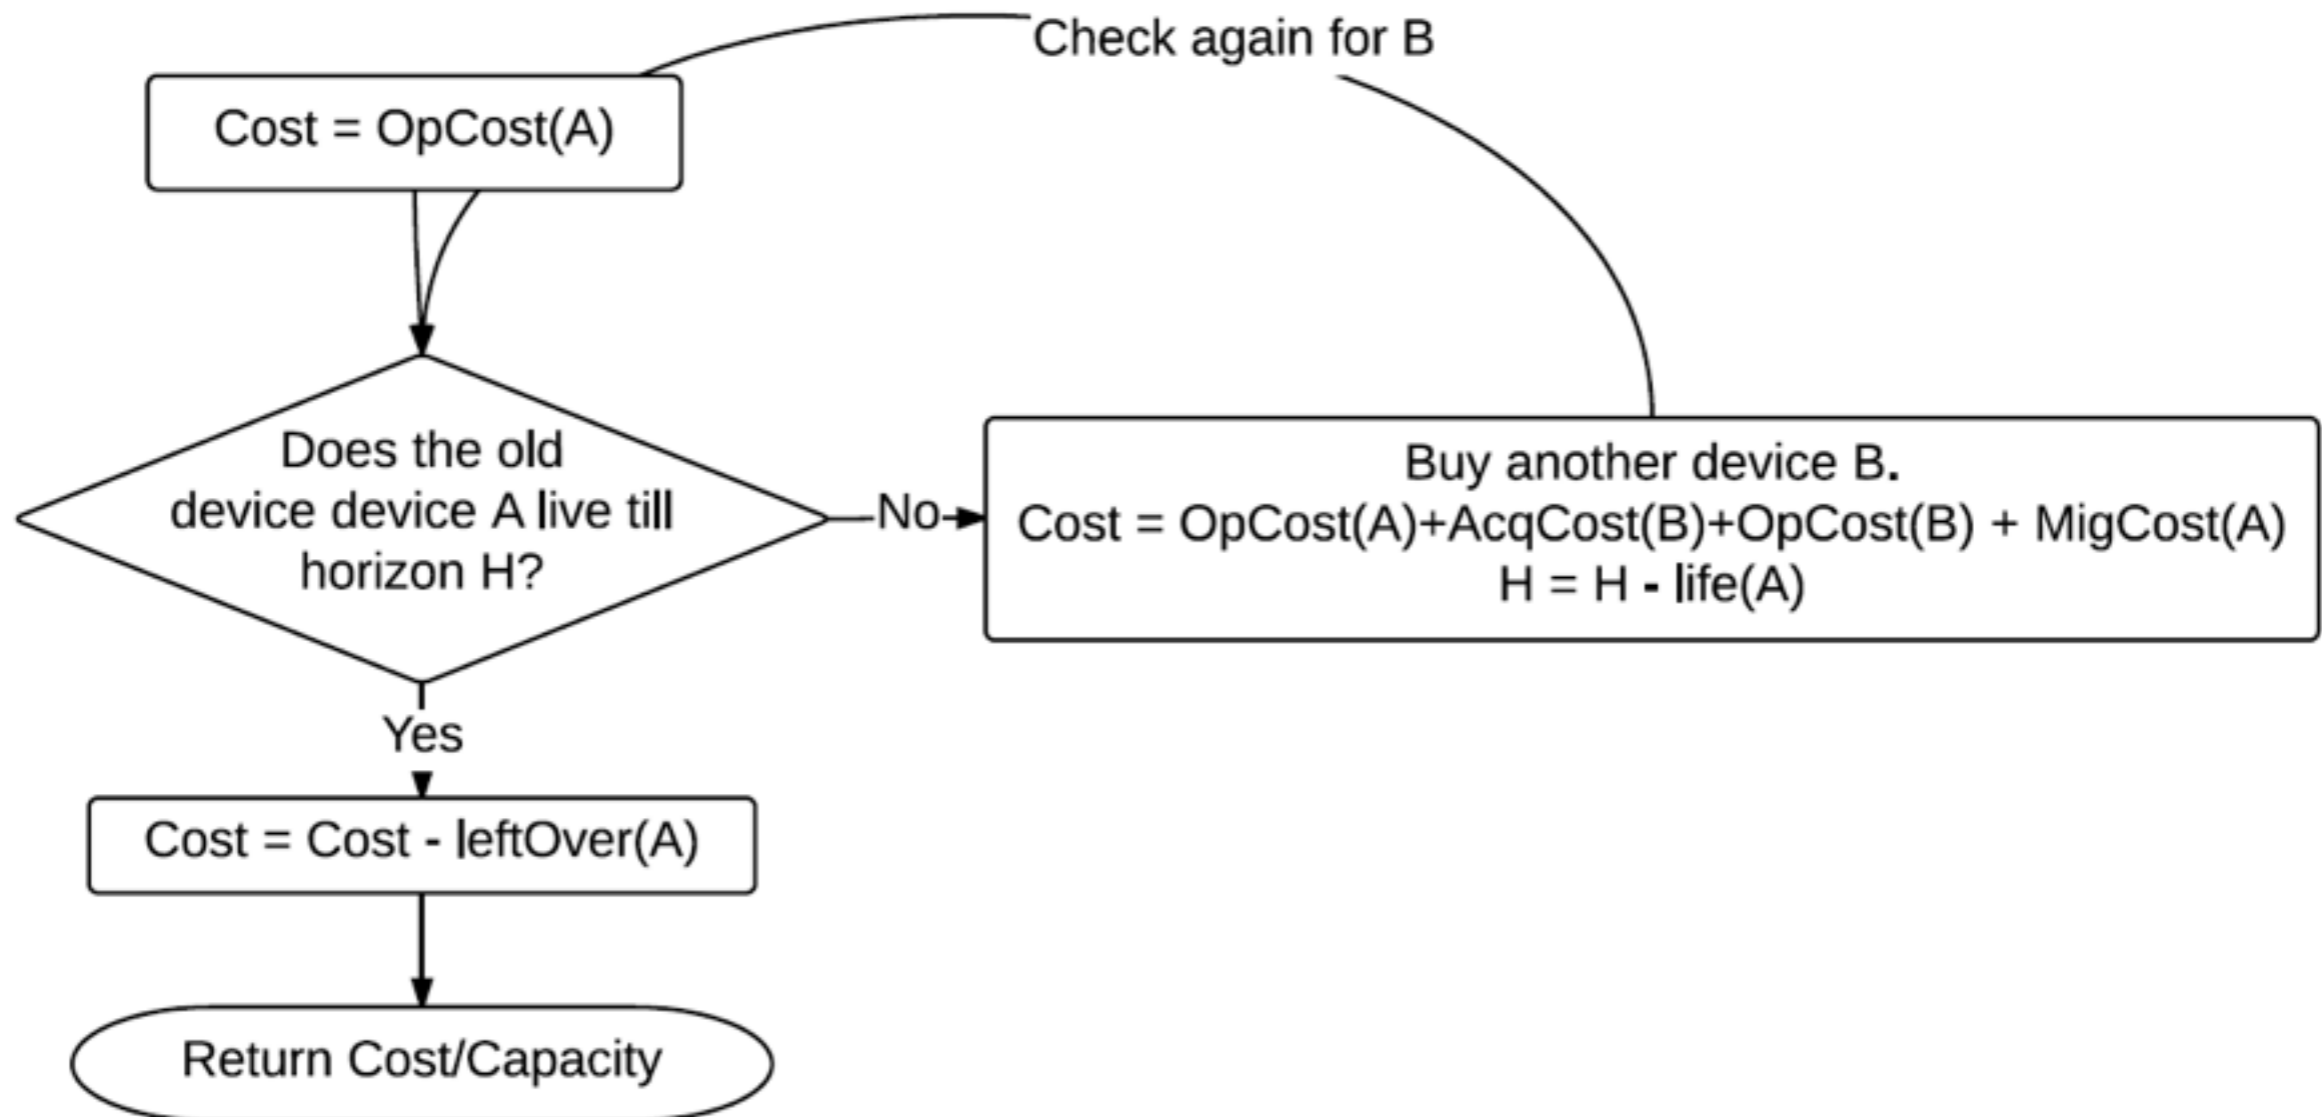
Service life of new media

Left-over life of old media

How long should we plan for?

Short planning horizon: Low risk and easy to plan.

Long planning horizon: High risk, hard to plan, and does not fit with some corporate strategies.



Terms

- **Kryder's rate (K_r):** Annual device density growth rate
- **Planning Horizon:** Length of the time organizations plan for
- **Service Life:** Length of time after which a device is discarded

- **Device Longevity:** Amount of time a device is operational
- **Capital Cost (CAPEX):** Acquisition costs
 - *E.g.* purchase, move-in
- **Operational Cost (OPEX):** Cost of operating a device
 - *E.g.* cooling, power, labor, and space
- **Data Refresh Cost:** Cost of migrating data between devices

- **Device group:** Reliability group of devices purchased together

Calculating Costs

- ❖ Total Cost of Ownership = $CAPEX + (OPEX \times duration) - residual\ value$
 - Residual value: $(s - h) / s \times p$
 - s = service life, h = planning horizon, and p = purchase cost
- ❖ Future Costs are **discounted**
 - Discounted operations: $OPEX / (1 + r)^{duration}$
 - r = discount rate
 - Discount rate accounts for future fluctuations in monetary value

Effect of planning horizon on device replacement

❖ Planning horizon = 2 years

Device	Capacity	CAPEX (\$)	OPEX (\$/yr)	Left-over life
A	4	500	30	4
B	4.4	500	20	5

→ pOld (Continue using A)

- cost = $\$30 \times 2 = \60

- per_cap = $\$60 / 4 \times 2 = \mathbf{\$7.50}$

→ pNew (Replace A by B)

- cost = $\$500 + \$20 \times 2 - \$500 / 5 \times 3 = \240

- per_cap = $\$240 / 4.4 \times 2 = \mathbf{\$0.50}$

Decision: Replace A with B

Effect of planning horizon on device replacement

❖ Planning horizon = 5 years

Device	Capacity	Service Life	CAPEX (\$)	OPEX (\$/yr)	Left-over life
A	4	5	500	30	4
B	4.4	5	500	20	5
C	12	5	500	20	5

→ pOld (Continue using A (use A for 4 years and then use C for one year))

- cost = $(30 \times 4) + 500 + (20 \times 1) - 500 / 5 \times 4 = \240.00

- per_cap = $\$240 / (4 \times 4 + 12 \times 1) = \mathbf{\$8.00}$

→ pNew (Replace A by B (use B for 5 years))

- cost = $500 + 20 \times 5 - 0 = \600.00

- per_cap = $\$600 / 4.4 \times 5 = \mathbf{\$27.27}$

Decision: Continue using A

Economic Model

Inputs

- Data size
- Duration
- Device configurations
- Performance constraints
- Redundancy level
- Planning horizon



Monte Carlo Simulation

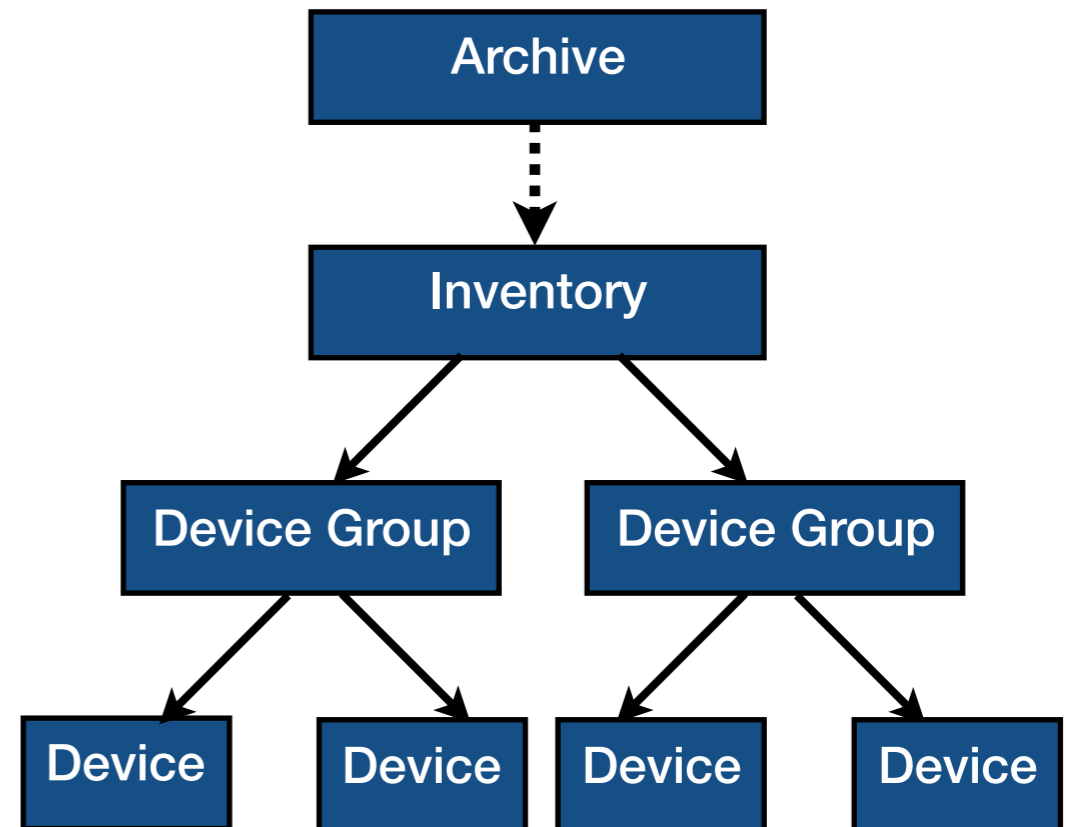
Outputs

- Total cost
- System characteristics

* Each run uses a random interest rate

Assumptions

- ❖ Devices grow in capacity each year
- ❖ Storage demand increases every year
- ❖ Media fails randomly
- ❖ Future expenditures are discounted
 - Discount rate: amount of interest paid as a percentage of a balance at the end of an annual period



A device group is a homogenous group of devices purchased at the same time

- ❖ Initial phase
 - Purchase

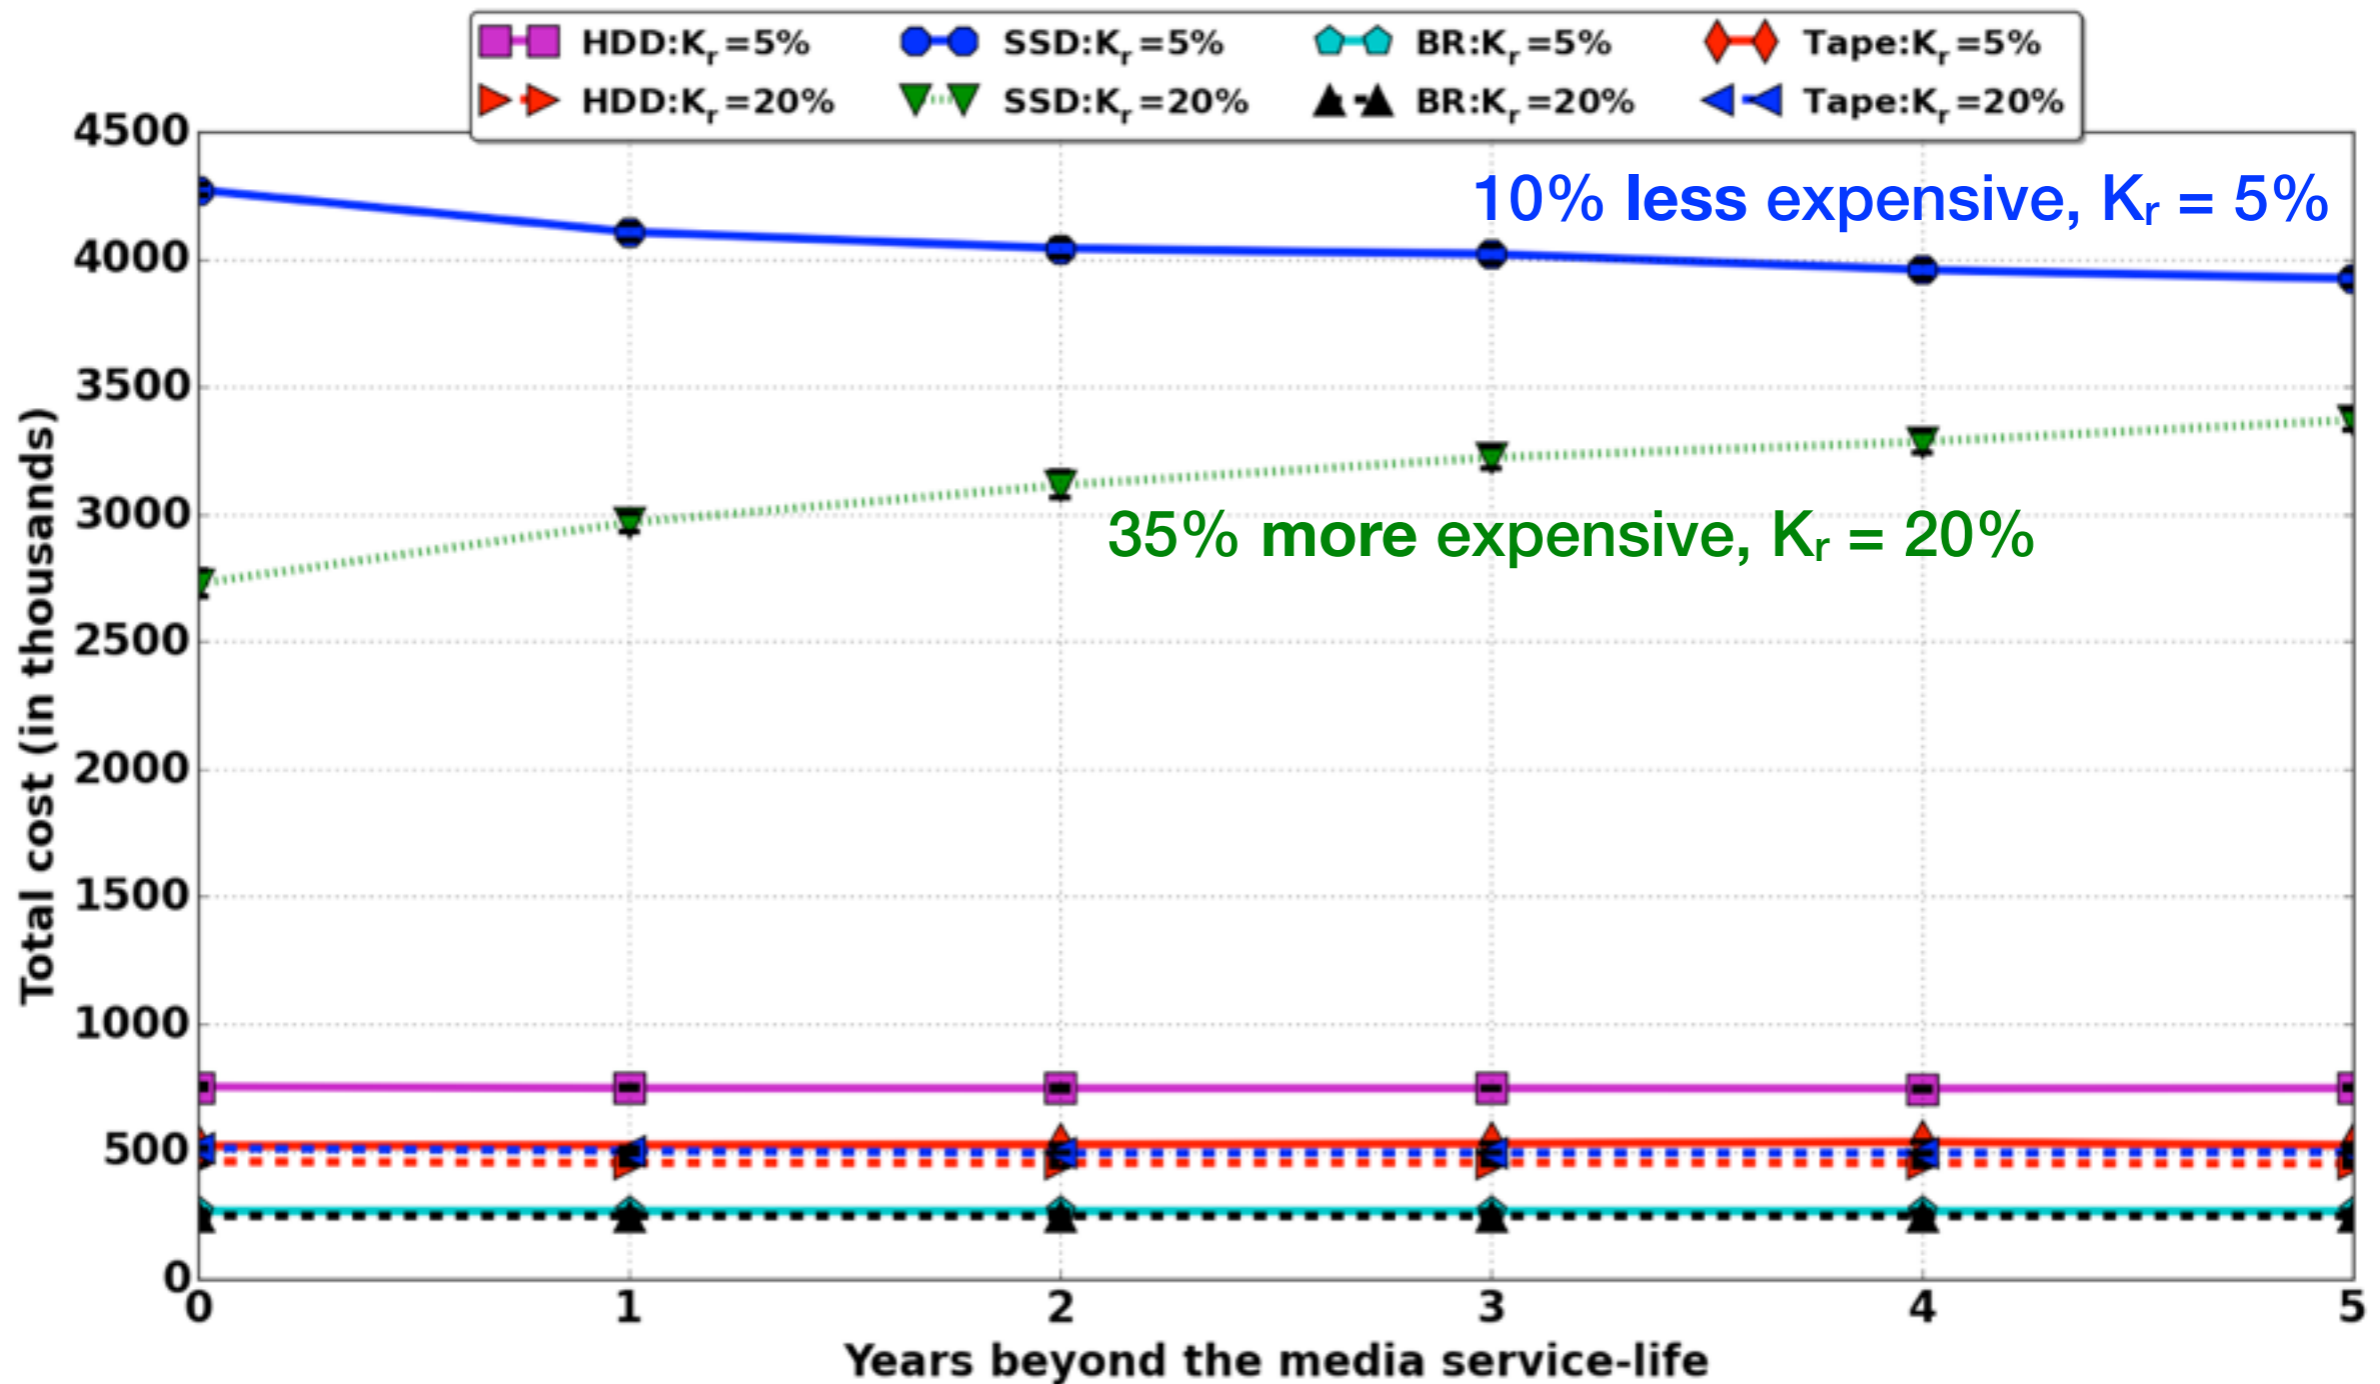
- ❖ Ongoing phase
 - Retire
 - Review
 - Purchase

- ❖ Pay liabilities
- ❖ Calculate total expenditure

Results

Service Life (Cost)

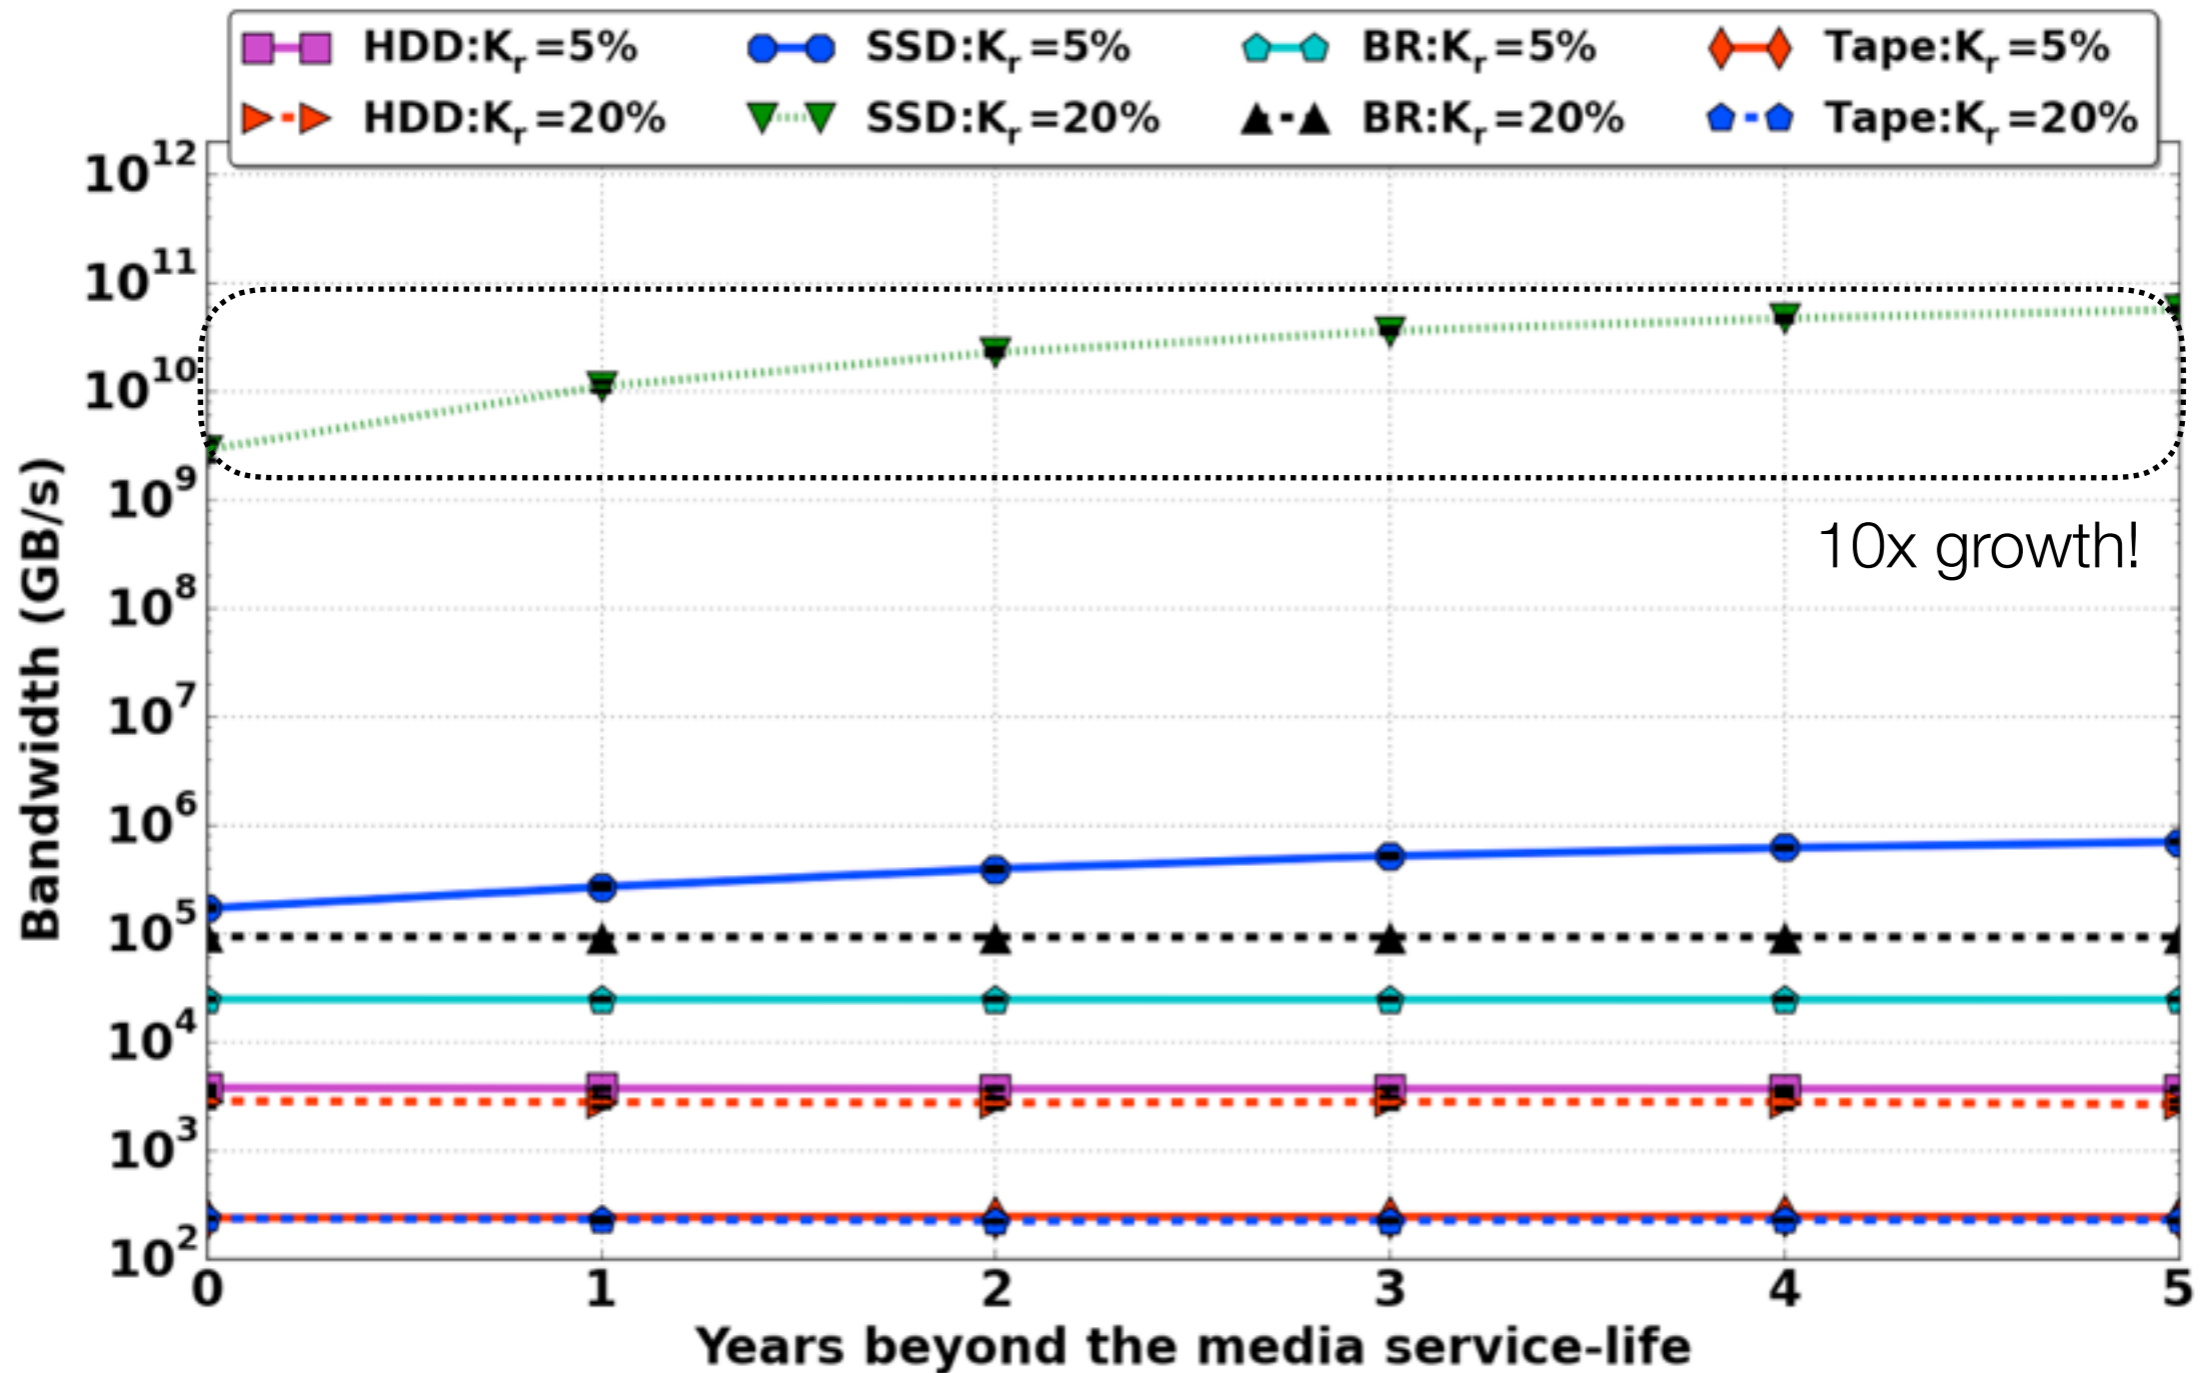
K_r denotes annual media density growth rate.



SSD cost savings are strongly dependent on density growth rate

Results

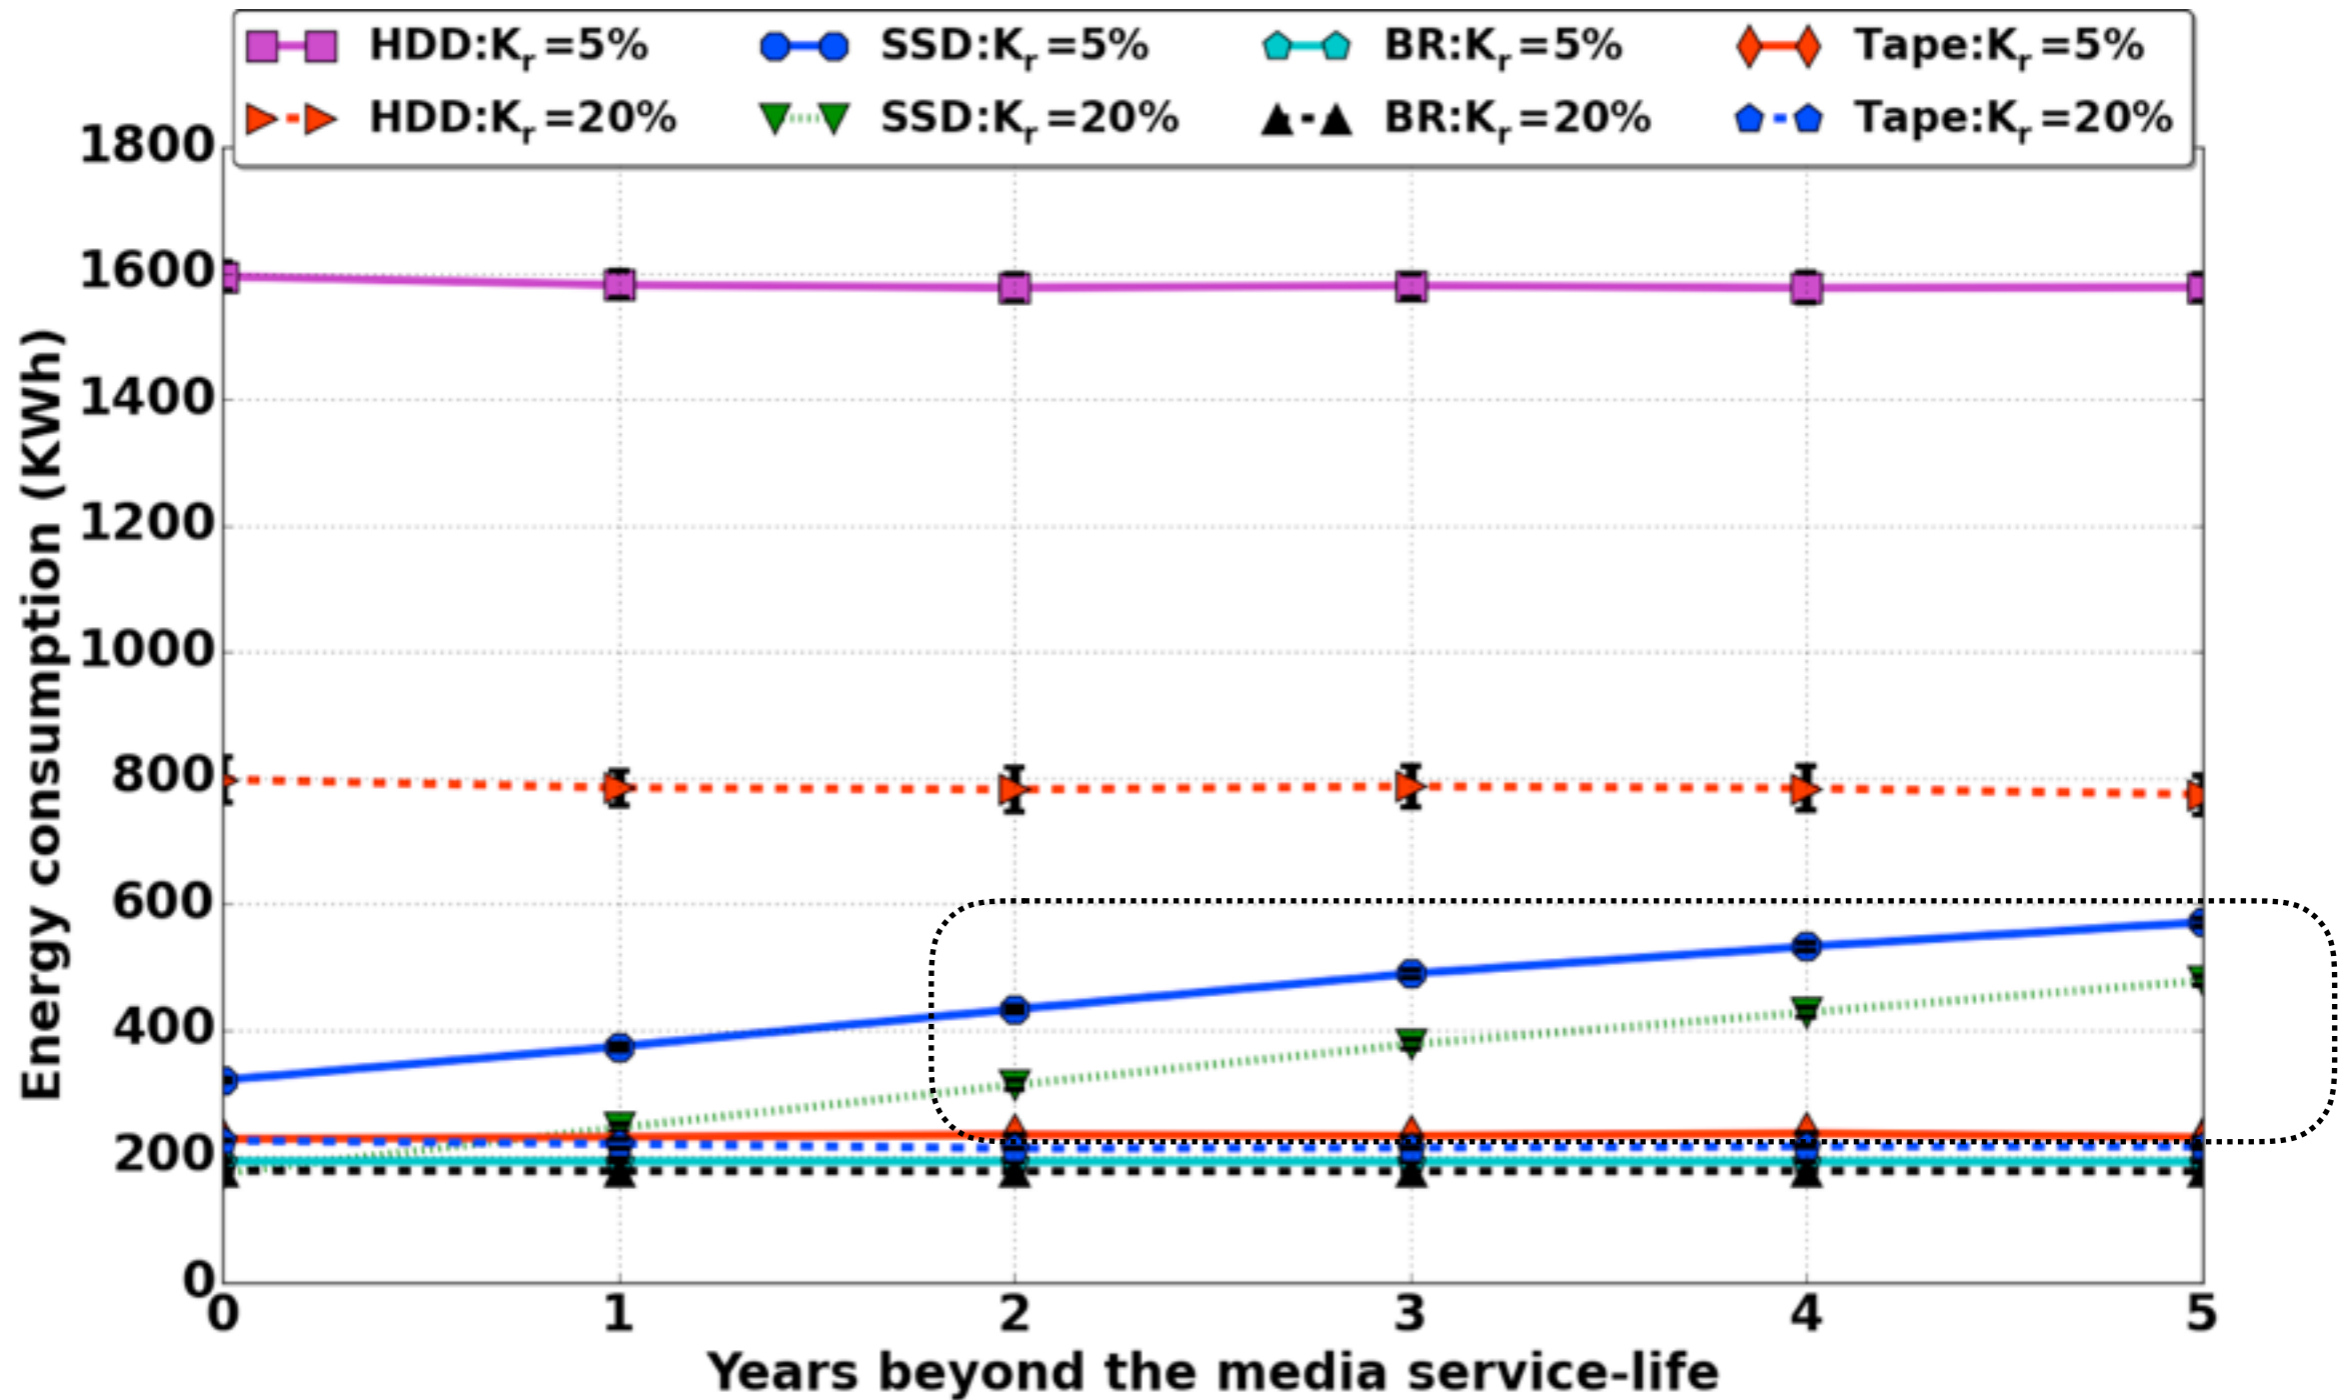
Service Life (Bandwidth)



Increased bandwidth over time due to lower average device density

Results


Service Life (Power)



For SSDs, retaining devices increases power consumption

Results

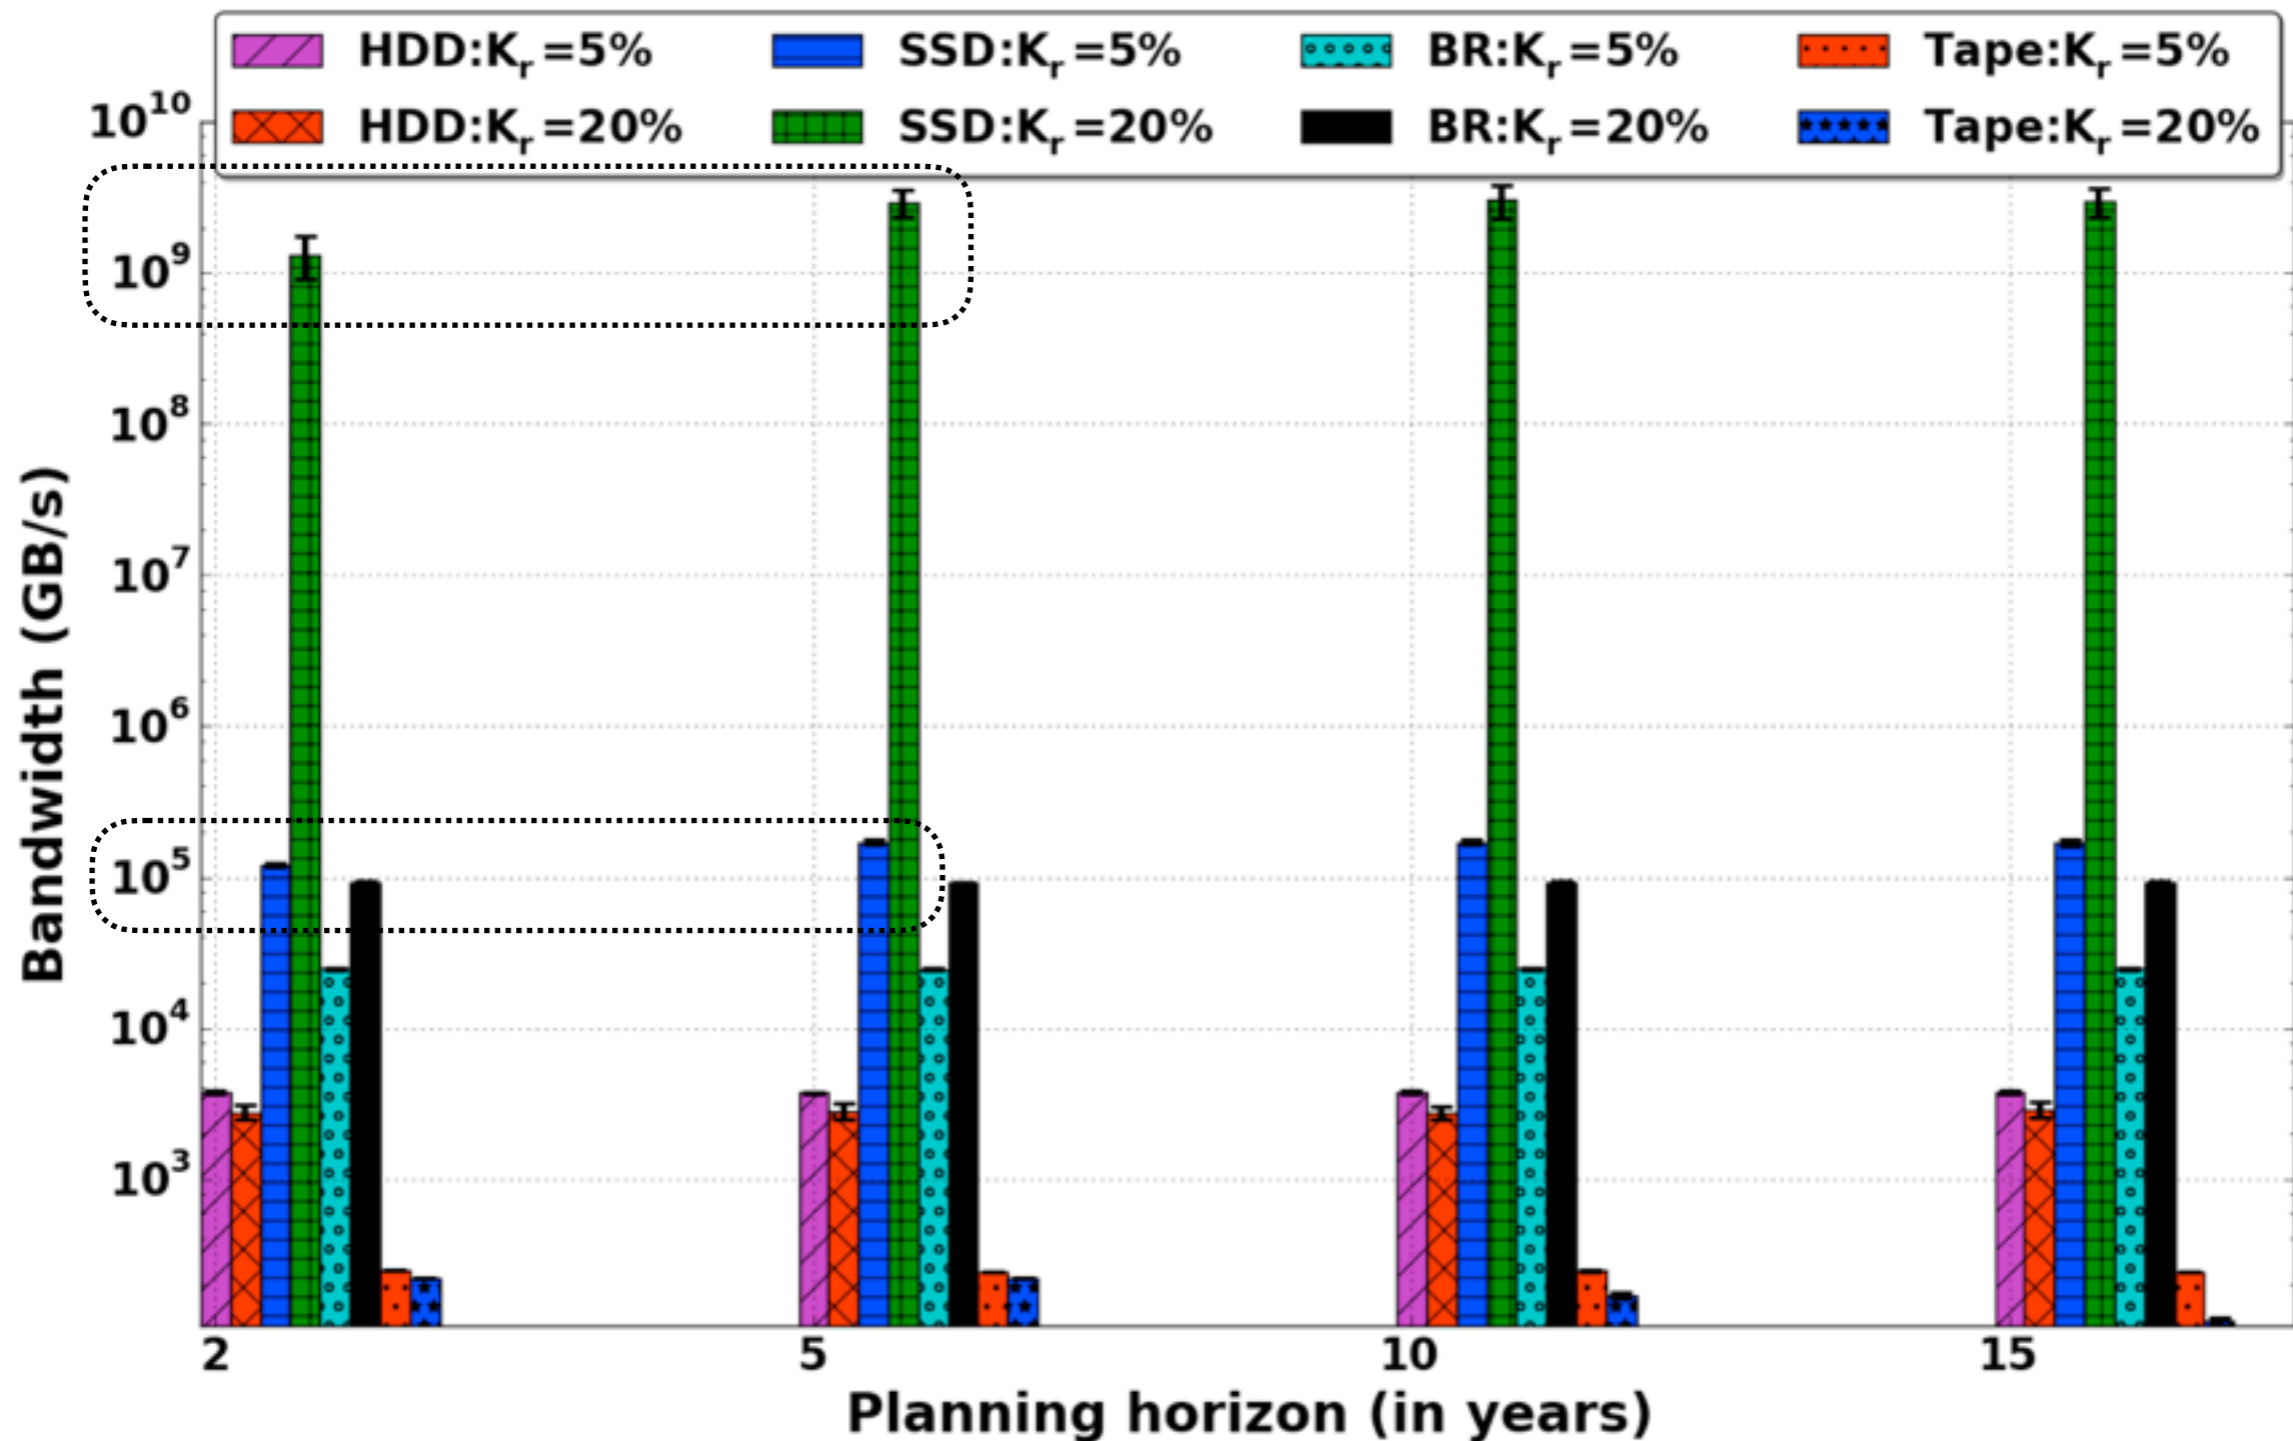
Planning Horizon (Cost)



Planning horizon cost shifts are also strongly dependent on density growth

Results

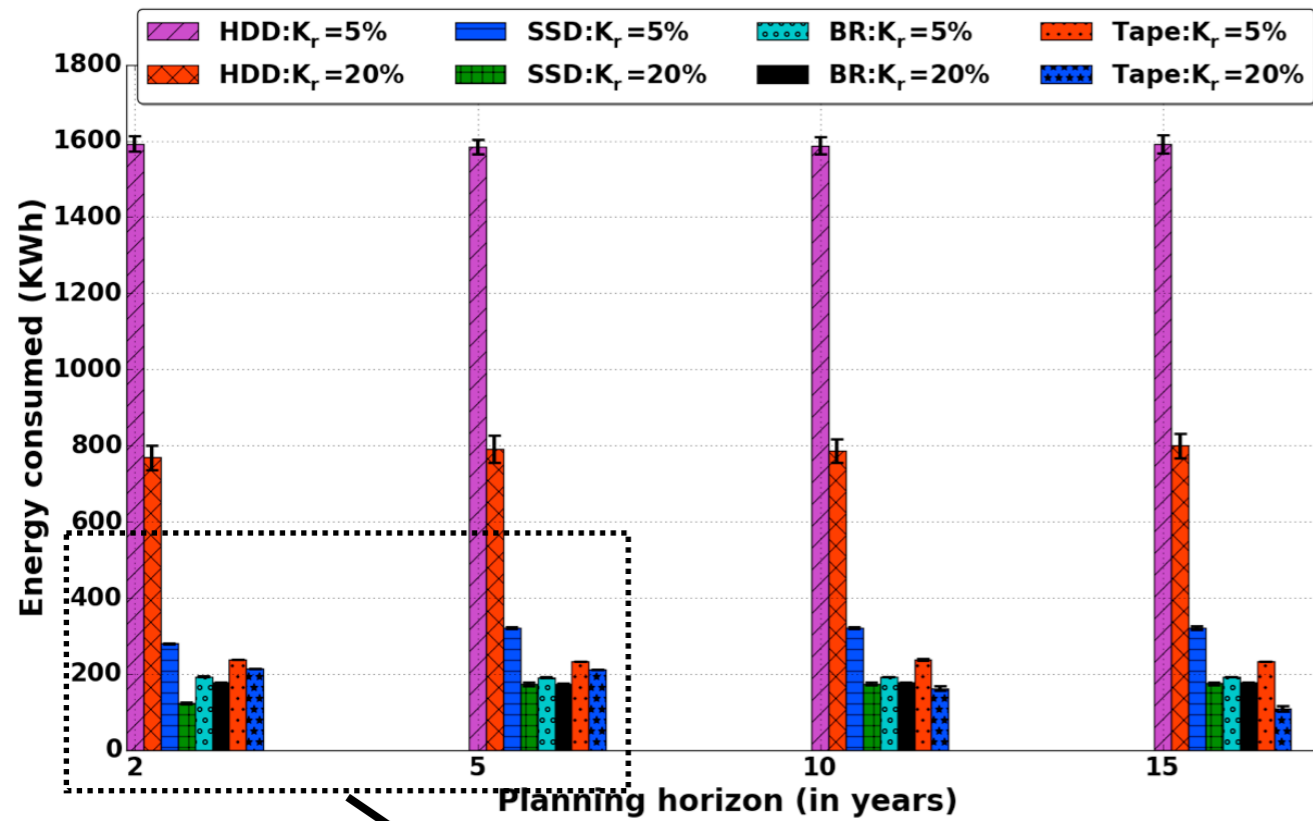
Planning Horizon (Bandwidth)




For SSDs, 5X growth in bandwidth at $K_r = 5\%$

Results

Planning Horizon (Power)



For SSDs, longer planning horizons increase power consumption



Conclusion

- ❖ In times of slow media density growth, media service-life needs to be extended for some device types
- ❖ Optimal planning horizon is at least as long as the media service life
- ❖ Expected planning horizon will have a major impact on the design of future archival storage media

Future Directions

❖ Heterogeneous generations

- Optimize on a subset of axes

❖ Decision analysis

- Methods in the business world to balance cost and planning horizon

❖ Active Archives

	Tape	Disk	SSD	Optical disc
Bandwidth	Very high	High	High	Moderate
IOPS	Very low	Moderate	High	Moderate
Reader?	Yes	No	No	Yes
Reliability	Moderate	Moderate	High	Low moderate
Op. cost	Moderate	Moderate	Low	Moderate
Capital cost	Low	Low	High	Low
Longevity	Long	Short	Moderate	Long
Density	High	High	High	Moderate

Acknowledgements



preetigupta25@soe.ucsc.edu



darrell, elm@soe.ucsc.edu



avani@mathcs.emory.edu



dshr@stanford.edu