UPPSALA UNIVERSITET

Adapting The Secretary Hiring Problem for Optimal Hot-Cold Tier Placement under Top-K Workloads

Ben Blamey*, Fredrik Wrede*, Johan Karlsson†,
Andreas Hellander*, and Salman Toor*

DBDM, CCGrid (IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing)
Larnaca, Cyprus, 14-17 May 2019

ben.blamey@it.uu.se

*Department of Information Technology, Division of Scientific Computing, Uppsala University, Sweden
†Discovery Sciences, Innovative Medicines, AstraZeneca, Gothenburg, Sweden

Outline

“Adapting The Secretary Hiring Problem for Optimal Hot-Cold Tier Placement under Top-K Workloads”

1. What’s a Top-K Workload? Why?
2. Hot and Cold Storage Tiers, AWS Costs.
3. The Optimization Problem (and solution). (Case Study 2).
4. The Secretary Hiring Problem (and solution).
5. Evaluation.
6. Related Work.
7. Future Work, HASTE Project.

ben.blamey@it.uu.se

Why a Top-K Workload?

- …because there is too much data to….

  …store

  …compute

  …transport

  …for a human to inspect
- Solution: take the most relevant/high quality/interesting data, from our batch or stream window.

ben.blamey@it.uu.se

What is a Top-K Workload?

```
# Pseudocode

def interestingness(document):
        # TODO: return some interestingness metric for the document.
        # How relevant, interesting, important, etc.
        return 0.0…1.0

# Stream window of N documents:
N_doc_window = InputStream(window_size=N)

top_K = N_doc_window.takeOrdered(K, key=lambda doc: interestingness(doc))
```

ben.blamey@it.uu.se

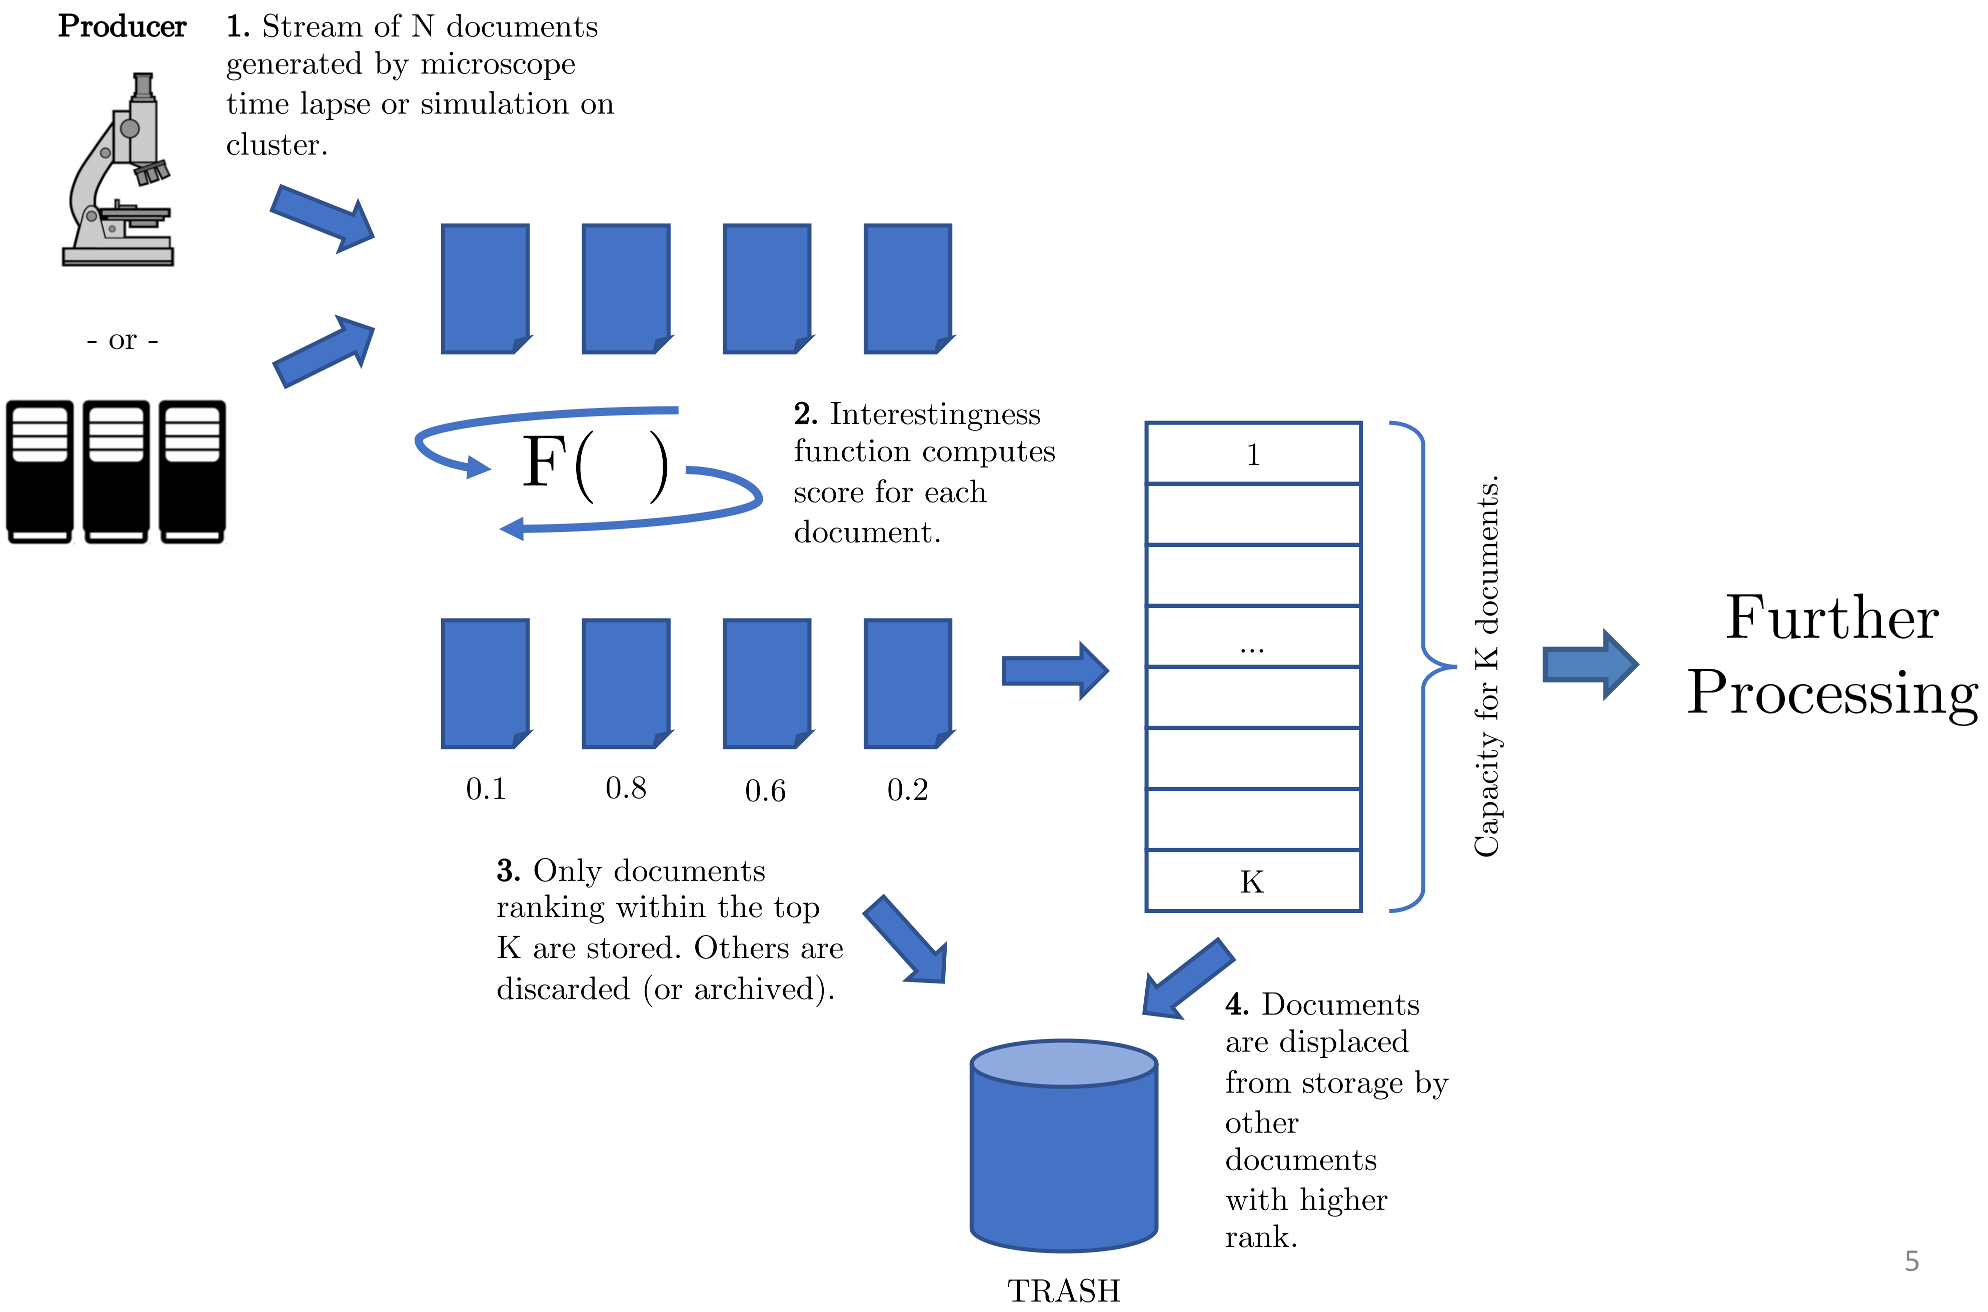

ben.blamey@it.uu.se

Hot and Cold Storage Tiers (AWS) ("Case Study 2")

		Read Cost	Write Cost	Rental Cost	Best For
Hot	AWS Elastic File Storage	0	0	\$0.33 / GB-Month	High IO, short term.
Cold	AWS S3 (Object Storage)	\$0.005 / 1,000 requests	\$0.005 / 1,000 requests	\$0.023 / GB-Month	Long term, Low IO, archive.

ben.blamey@it.uu.se

The Optimization Problem (K=1)

$$P(i\text{th document is best so far}) = \frac{1}{i+1} \qquad (5)$$

$$\mathbb{E}(\#\text{ writes}) = \sum_{i=0}^{N-1} \frac{1}{i+1} \qquad (6)$$

This partial sum of the Harmonic series can be approximated:

$$\mathbb{E}(\#\text{ writes}) = \ln N + 0.57722 \qquad (7)$$

ben.blamey@it.uu.se

The Optimization Problem (K>1)

$$P(i\text{th doc. in top } K \text{ when observed} \,|\, i < K) = 1 \tag{9}$$

$$P(i\text{th doc. in top } K \text{ when observed} \,|\, i \geq K) = \frac{K}{i+1} \tag{10}$$

$$\mathbb{E}(\text{cumulative number of writes} \,|\, i < K) = i \tag{11}$$

$$\mathbb{E}\left(\begin{matrix}\text{cumulative} \\ \text{no. of writes}\end{matrix} \,\middle|\, i \geq K\right) = K + K \cdot \ln(i+1) \tag{12}$$

ben.blamey@it.uu.se

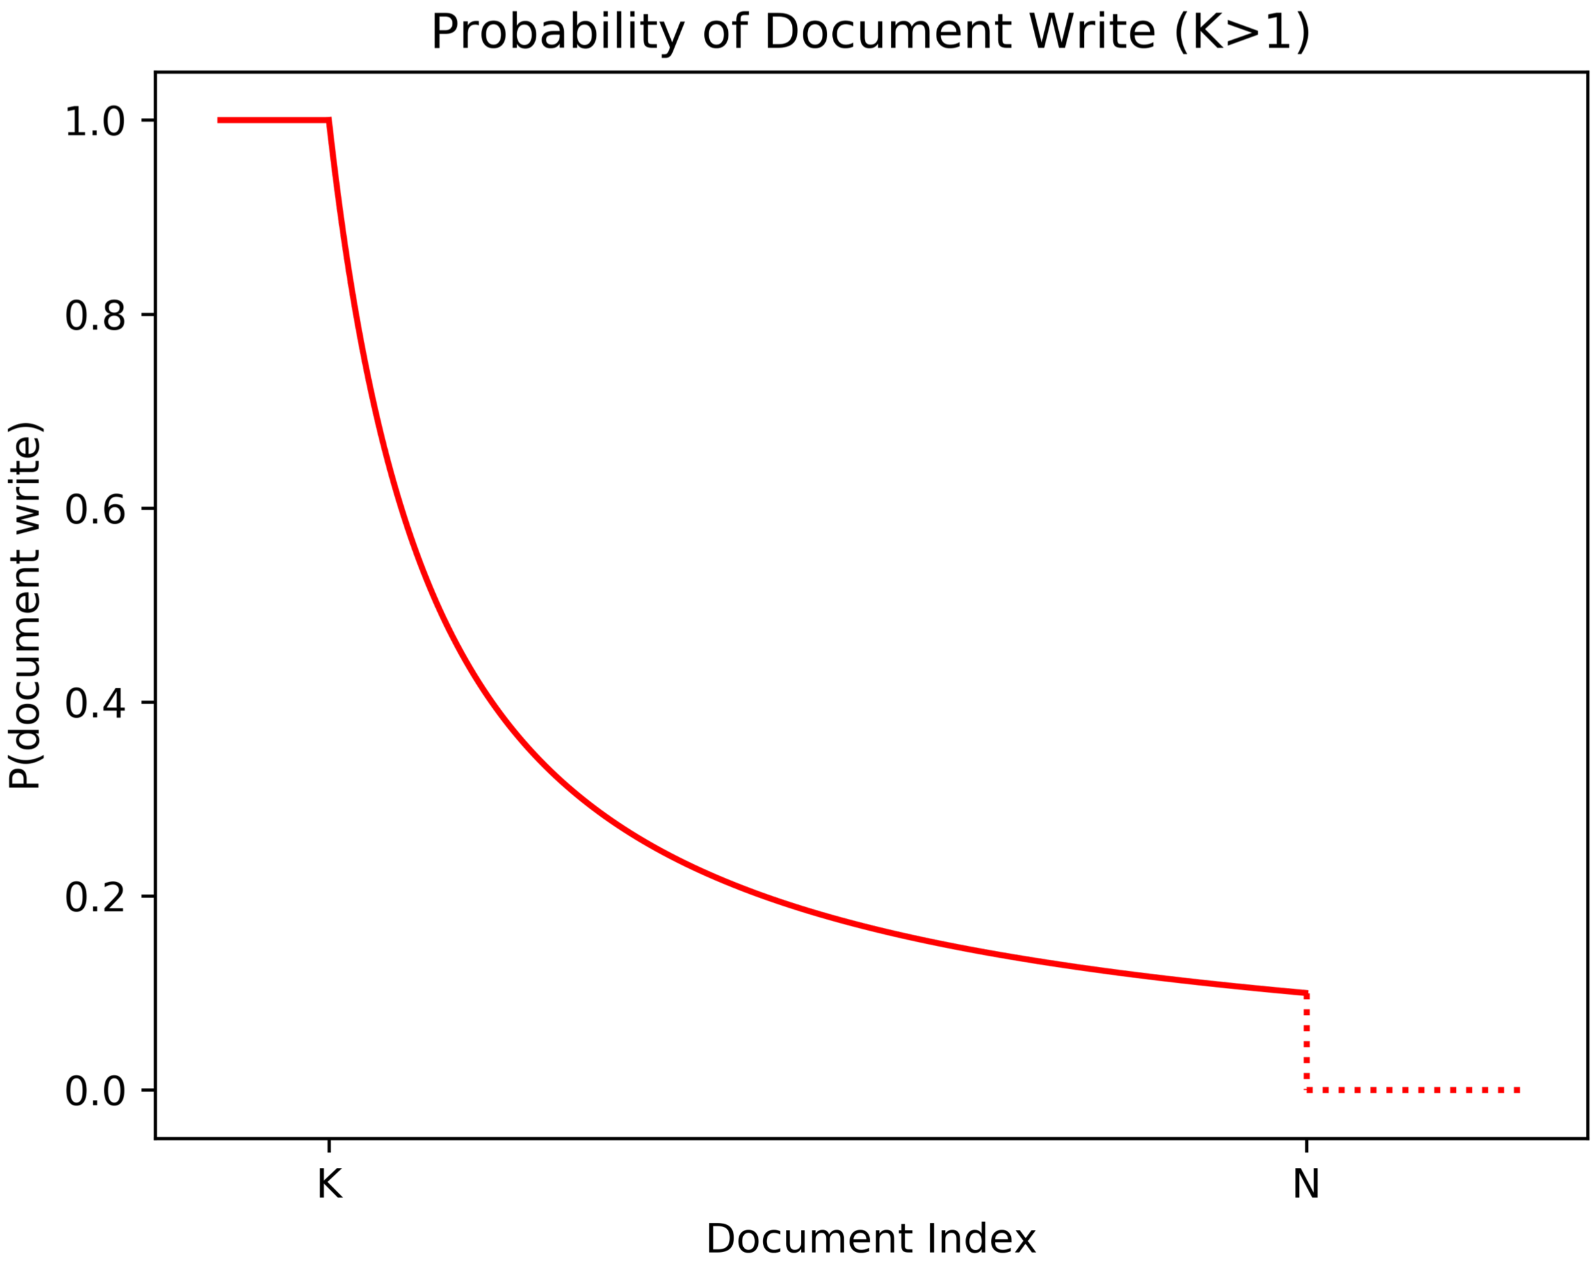

ben.blamey@it.uu.se

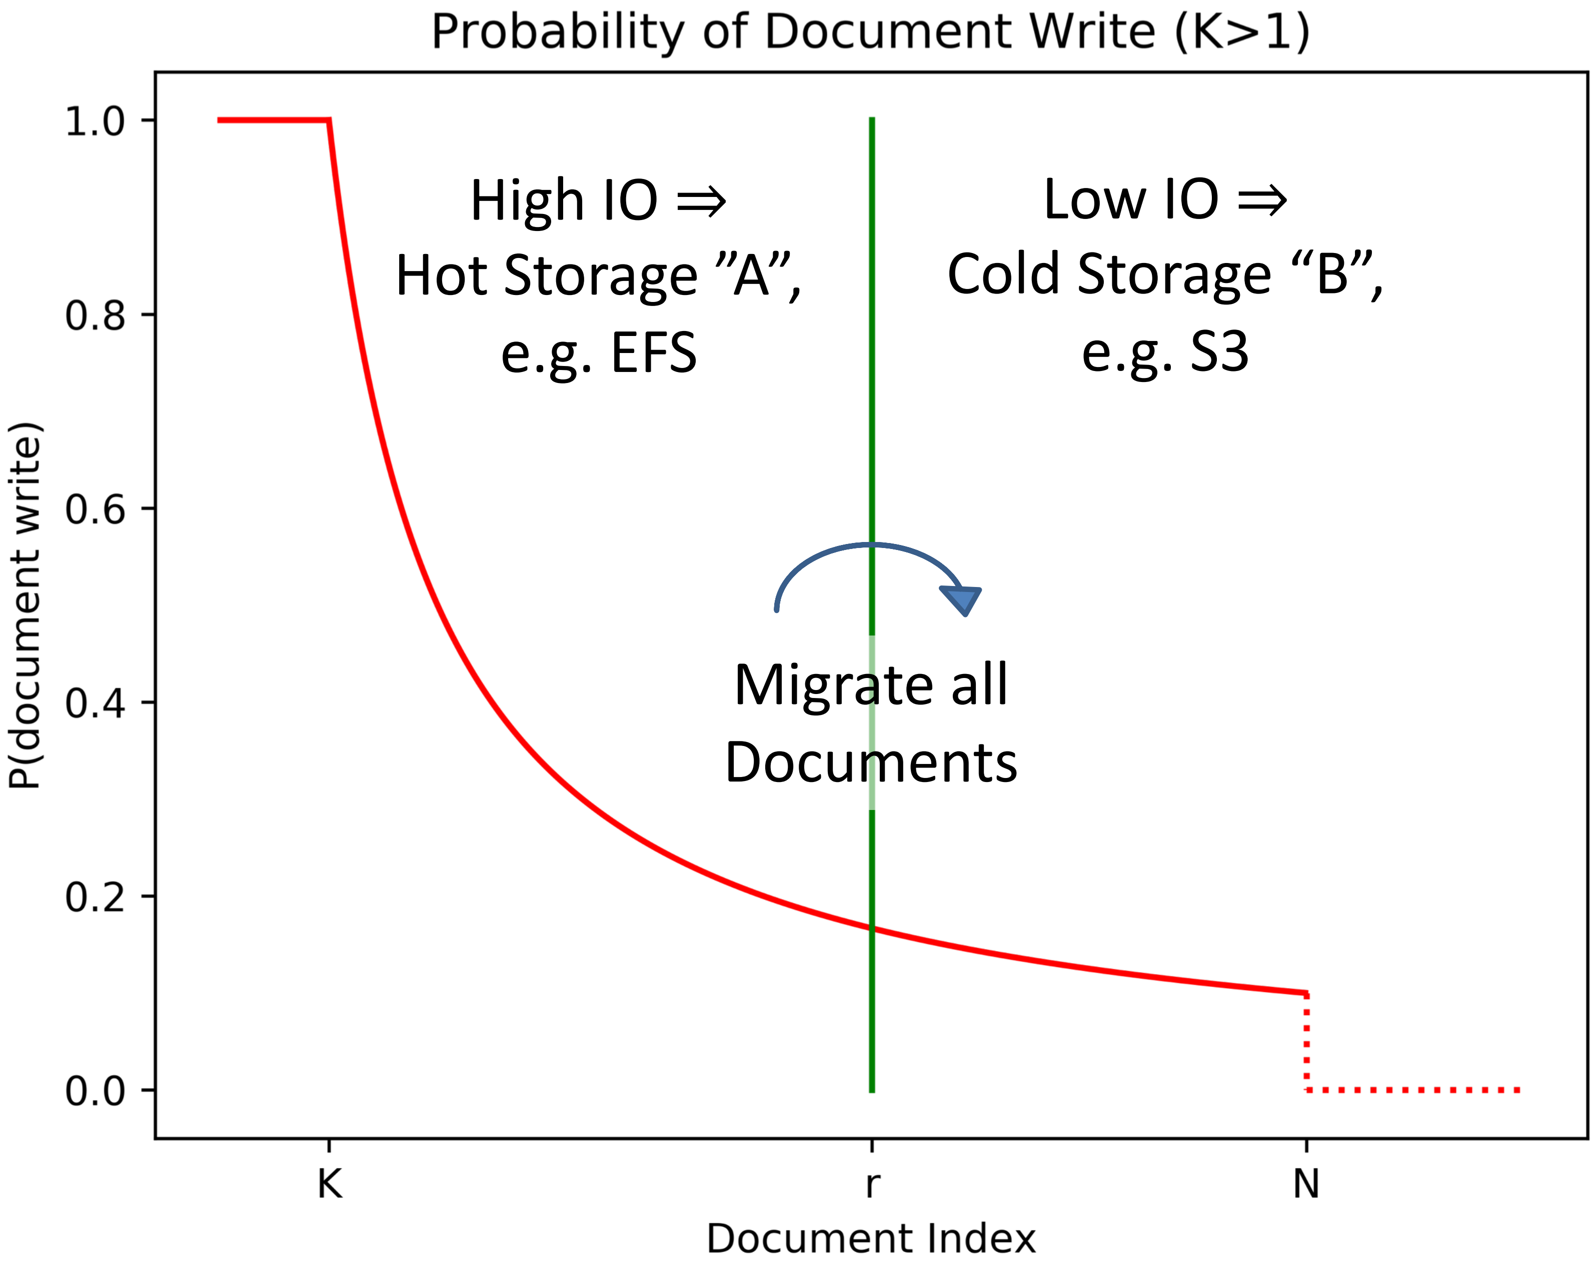

ben.blamey@it.uu.se

What's the optimal value of *r*?

$$\mathbb{E}(\Sigma \text{ cost writes}) = \sum_{i=0}^{N-1} P\binom{i\text{th document is in top}}{K \text{ when observed}} \cdot (\text{cost of writing it})$$

$$\mathbb{E}(\Sigma \text{ cost writes}) = K \cdot cost_{\text{write, A}} + K \cdot (\ln r - \ln K) \cdot cost_{\text{write, A}} + K \cdot (\ln N - \ln r) \cdot cost_{\text{write, B}}$$

$$\mathbb{E}(\Sigma \text{ cost rental}) = K \cdot \left(\frac{r}{N} \cdot cost_{\text{rental, A}} + (1 - \frac{r}{N}) \cdot cost_{\text{rental, B}}\right)$$

$$\mathbb{E}(\Sigma \text{ cost total}) = \mathbb{E}(\Sigma \text{ cost writes}) + \mathbb{E}(\Sigma \text{ cost rental}) + (\text{cost migration}) + (\text{cost reads})$$

$$\frac{d(\mathbb{E}(\Sigma \text{ cost total}))}{dr} = \underbrace{\frac{K}{r}(cost_{\text{write, A}} - cost_{\text{write, B}})}_{\text{writes}} + \underbrace{\frac{K}{N}(cost_{\text{rental, A}} - cost_{\text{rental, B}})}_{\text{rental}} = 0$$

$$\frac{r_{\text{optimal}}}{N} = \frac{cost_{\text{write, A}} - cost_{\text{write, B}}}{cost_{\text{rental, B}} - cost_{\text{rental, A}}}$$

ben.blamey@it.uu.se

What's the optimal value of *r*?

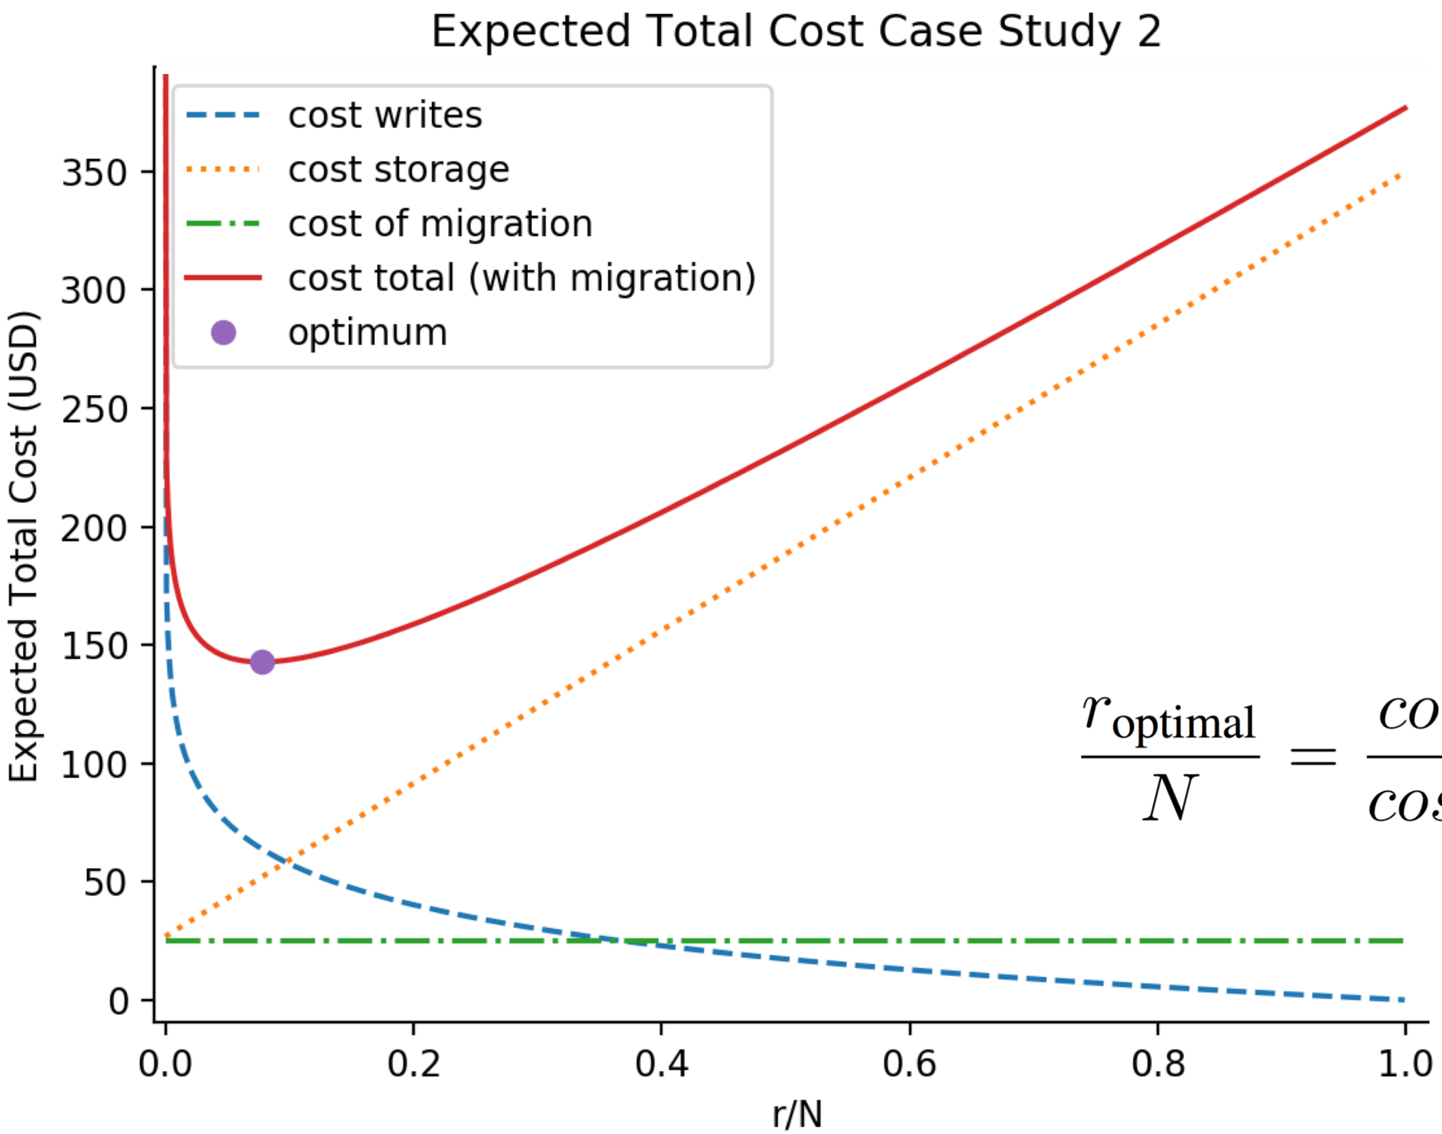

$$\frac{r_{\text{optimal}}}{N} = \frac{cost_{\text{write, A}} - cost_{\text{write, B}}}{cost_{\text{rental, B}} - cost_{\text{rental, A}}}$$

ben.blamey@it.uu.se

Relation to The Secretary Hiring Problem

You want to hire exactly one person for a job. You interview N people in turn, after each interview, you can choose to hire the person (and stop interviewing), or choose not to (and continue interviewing) - decision is irrevocable. The candidates are ranked, but you only see rank against the candidates already observed. How can you maximize your chance of choosing the best overall candidate?

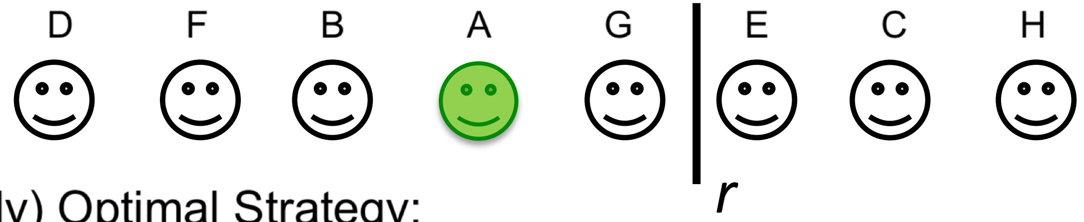

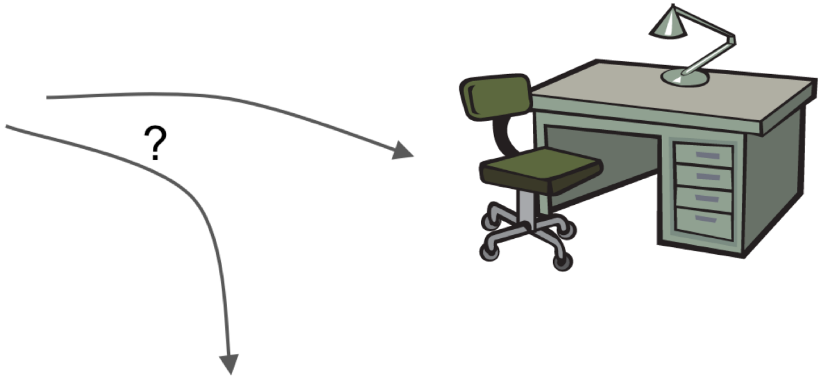

(Provably) Optimal Strategy:
Observe *r* candidates, then, hire the next candidate observed who outranks all of the first *r*.
Otherwise, hire the last remaining candidate.

$$N = \text{Number of documents} \tag{1}$$

$$\frac{r_{\text{optimal}}}{N} = \frac{1}{e} \tag{2}$$

$$P(\text{hiring overall best}) = \frac{1}{e} = 0.367 \tag{3}$$

$$\mathbb{E}(\#\ \text{writes}) = 1 \tag{4}$$

(NB. This is a special case of *bin packing*.
There are many variants, PostDoc hiring problem, k-Secretary hiring problem…)

Evaluation

- Computational Systems Biology: Numerical Simulation of Gene Regulatory Networks.
- Parameter Sweep to find values which give realistic natural behavior. (High-Dimension Parameter Space).
- SVM Classifier, (re-trained after each batch of N).
  - Interestingness Function is Label Entropy.
  - i.e. we’re interested in misclassified documents.
- ~ 14TB over 30 days with 60 vCPUs.

ben.blamey@it.uu.se

Evaluation

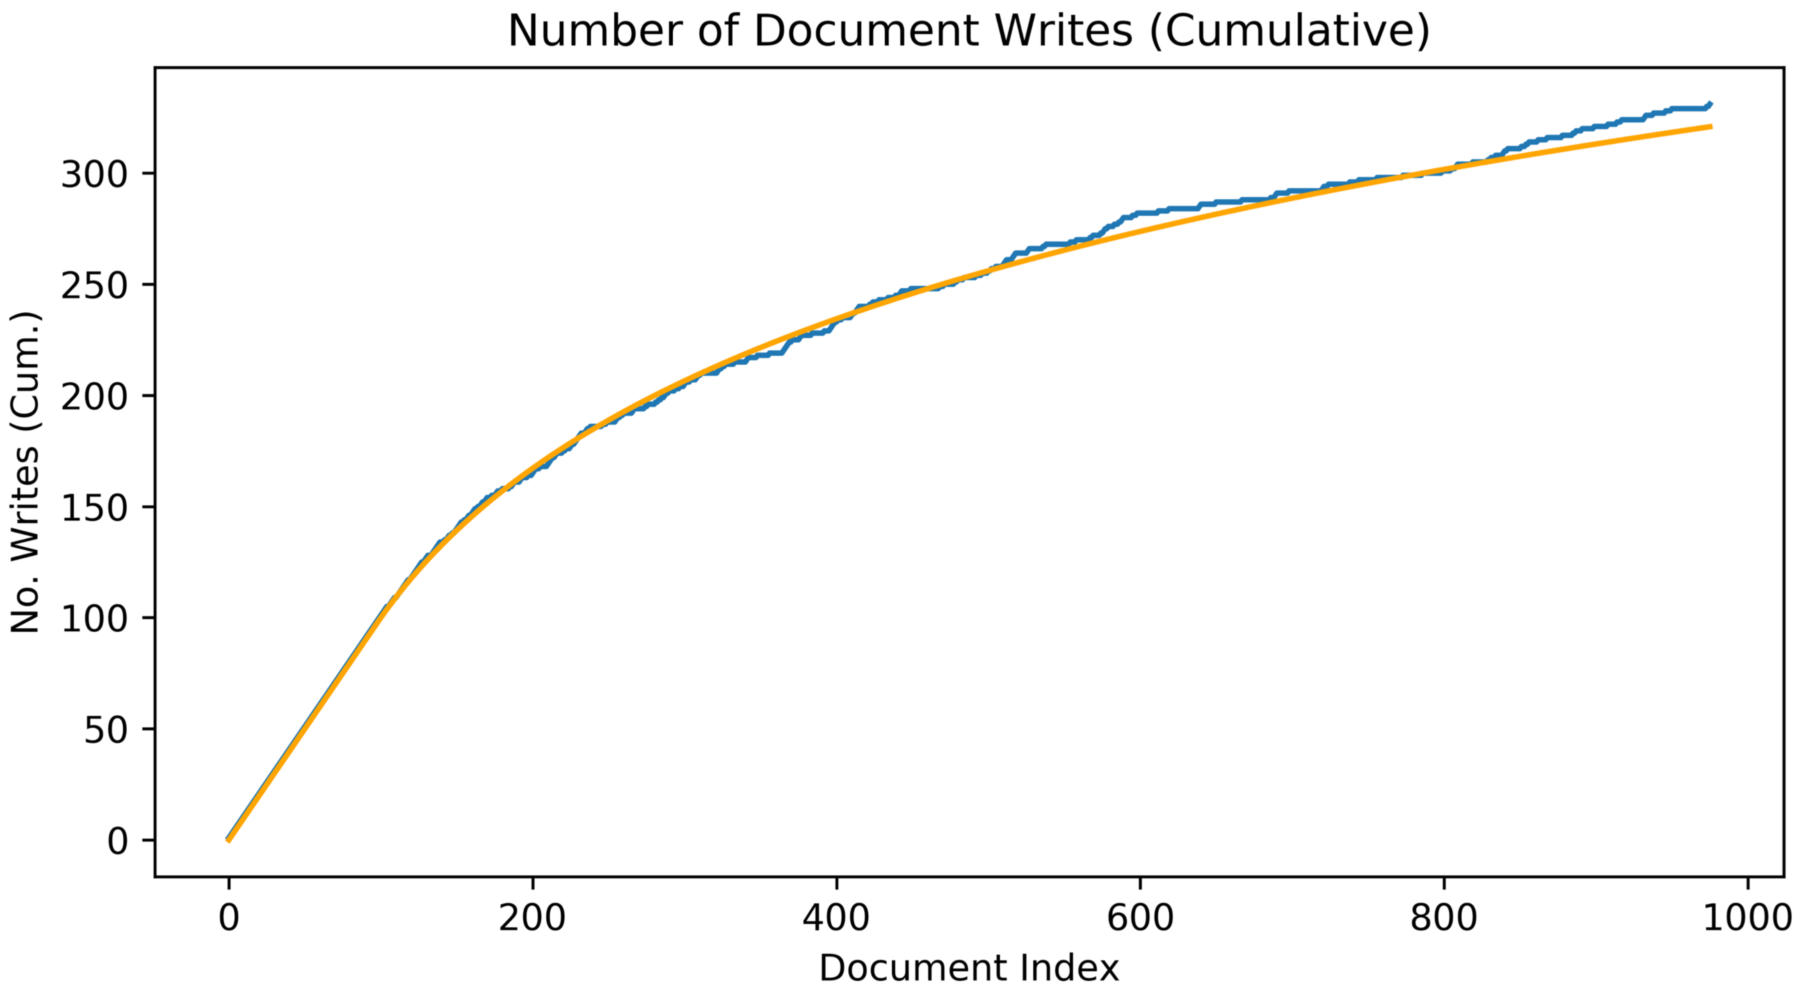

Blue: trace of entropy values. Orange: analytic solution.

ben.blamey@it.uu.se

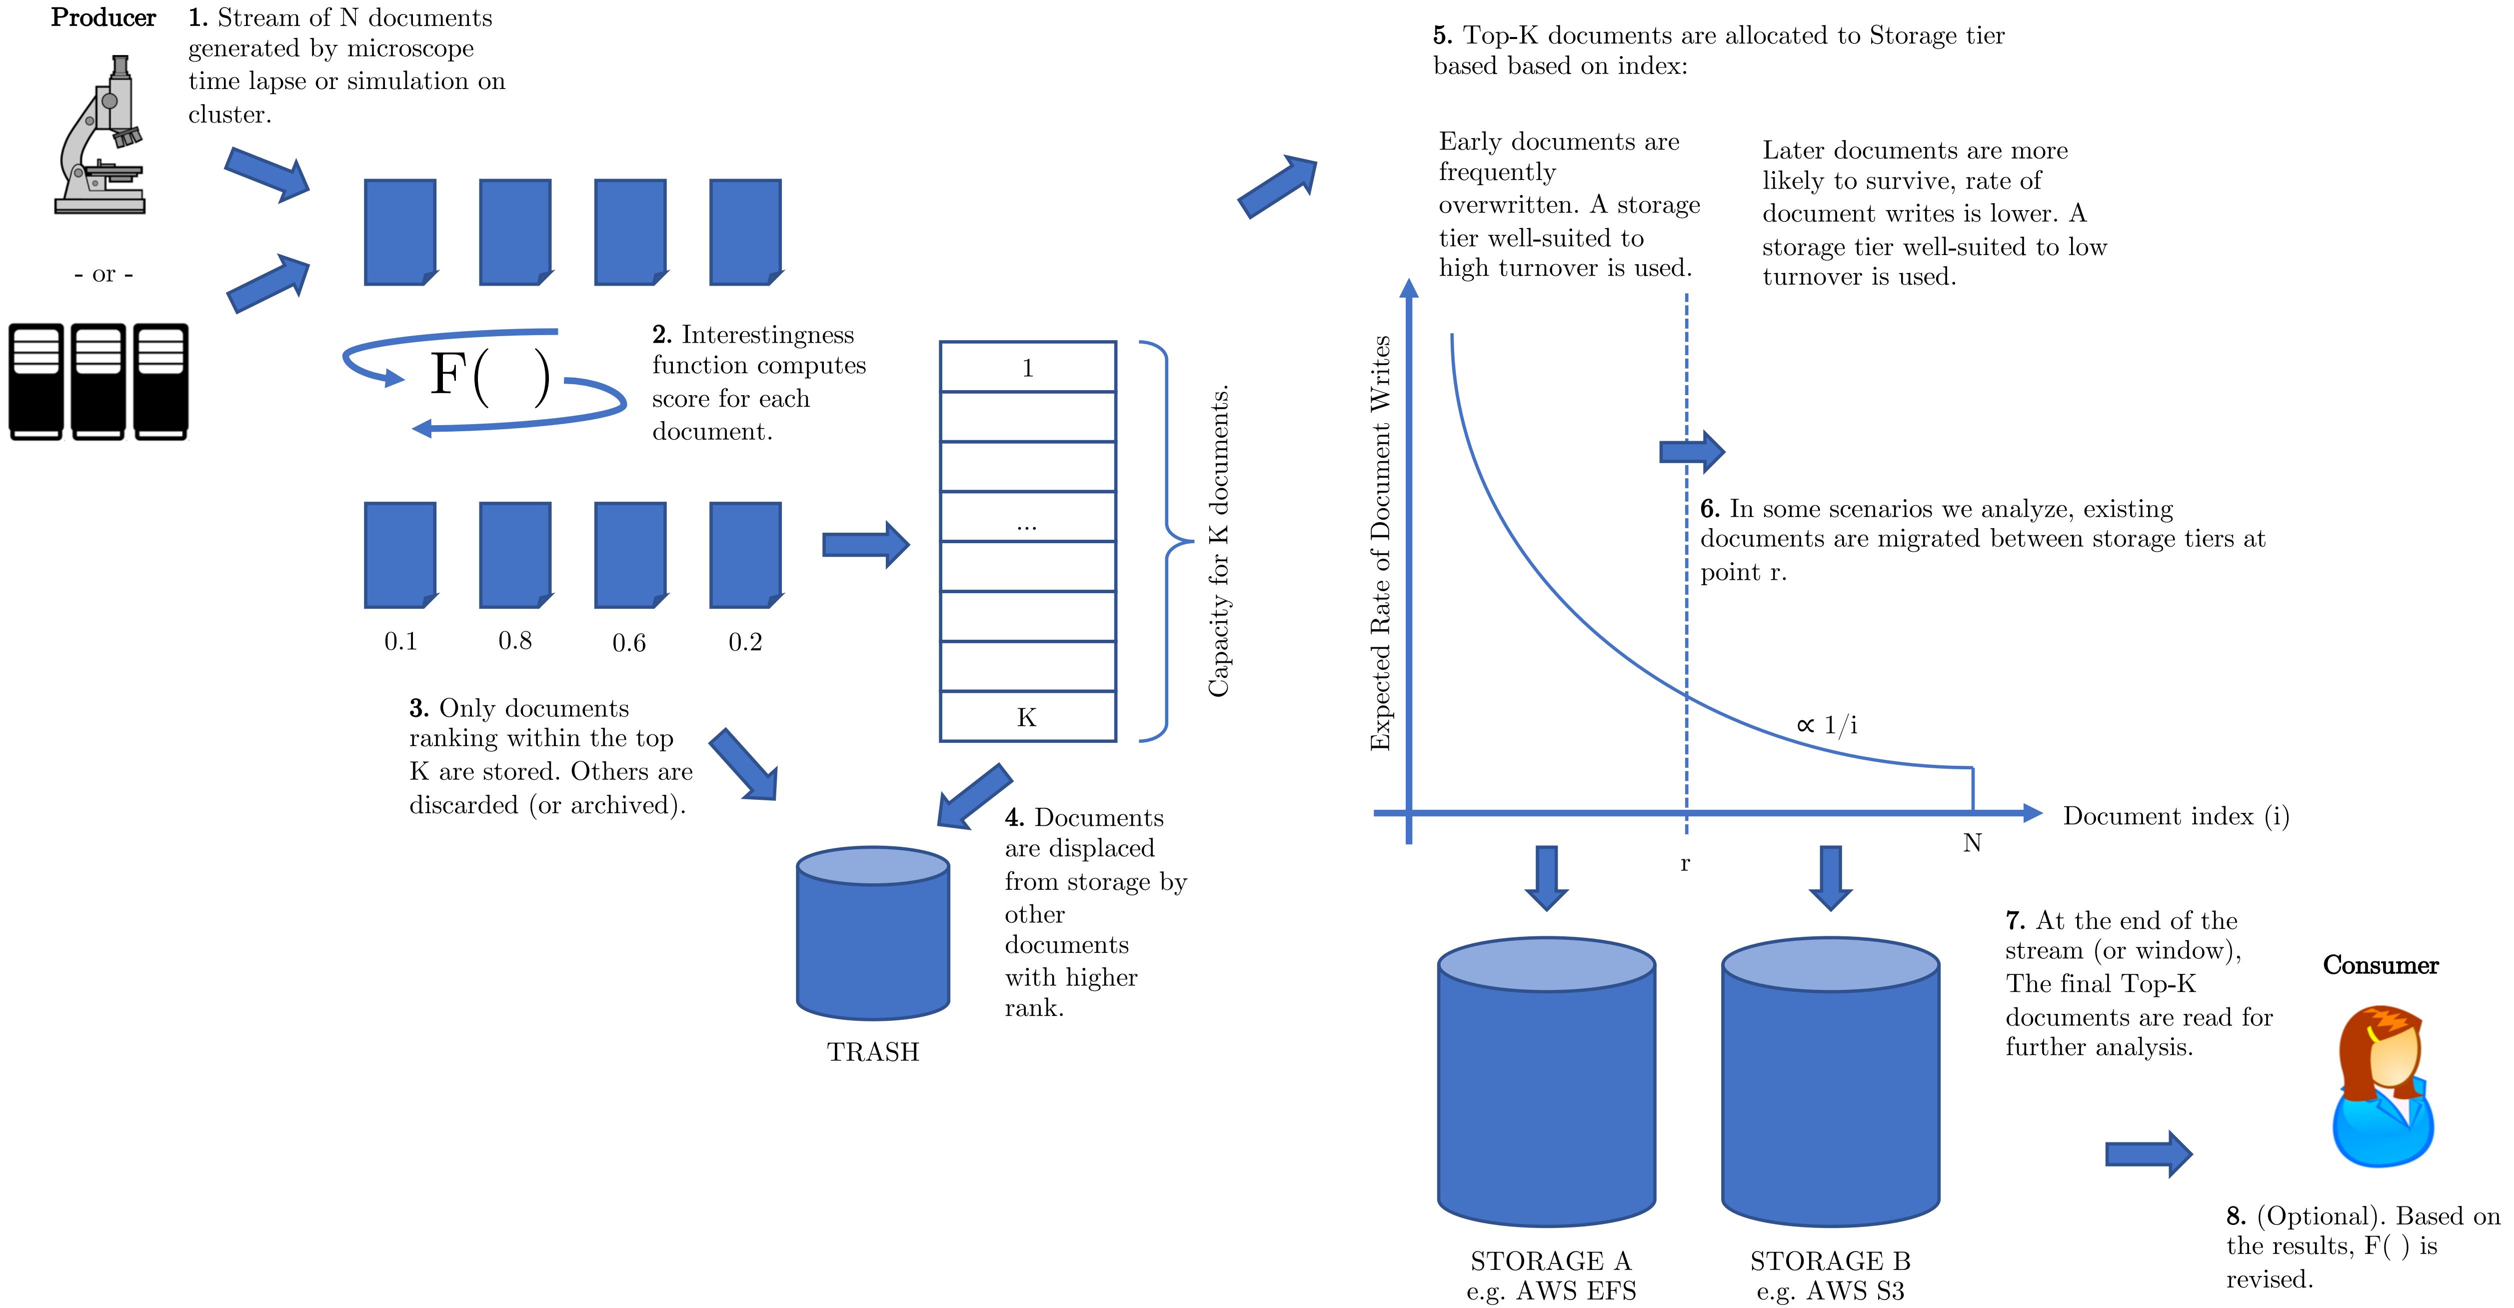
Producer
1. Stream of N documents generated by microscope time lapse or simulation on cluster.
- or -
F( )
2. Interestingness function computes score for each document.
0.1
0.8
0.6
0.2
3. Only documents ranking within the top K are stored. Others are discarded (or archived).
1
...
K
Capacity for K documents.
4. Documents are displaced from storage by other documents with higher rank.
TRASH
5. Top-K documents are allocated to Storage tier based based on index:
Early documents are frequently overwritten. A storage tier well-suited to high turnover is used.
Later documents are more likely to survive, rate of document writes is lower. A storage tier well-suited to low turnover is used.
Expected Rate of Document Writes
6. In some scenarios we analyze, existing documents are migrated between storage tiers at point r.
∝ 1/i
Document index (i)
r
N
STORAGE A
e.g. AWS EFS
STORAGE B
e.g. AWS S3
7. At the end of the stream (or window), The final Top-K documents are read for further analysis.
Consumer
8. (Optional). Based on the results, F( ) is revised.

UPPSALA UNIVERSITET

Case Study 1

- 2 Clouds (AWS + Azure)
- Include Communication Costs
- Applicable to: IoT/BigData Cloud Upload:
  - Motivation: Microscopy Image upload to the cloud

$$\frac{r_{\text{optimal}}}{N} = \frac{cost_{\text{write, A}} - cost_{\text{write, B}}}{cost_{\text{read, B}} - cost_{\text{read, A}}} \tag{18}$$

ben.blamey@it.uu.se

UPPSALA UNIVERSITET

Related Work

- Optimization Problems in Cloud Storage – ‘Pro-active’:
  - Often Linear Optimization
  - Ski Rental Problem (buy vs. rent?)
- Monitoring Application Load – ‘Reactive’:
  - Most Frequently/Recently Accessed ‘Hot’ files.
- **Our Contribution:**
  - **Secretary Hiring Problem ~ Top-K Workload**
  - **Cloud Storage/Transport**

ben.blamey@it.uu.se

Future Work

- Other Interestingness Functions
  - Microscopy Imaging – is image in focus? debris? etc.
- More than 2 storage tiers.
- Integration with the HASTE Core Pipeline
  - *Scientific Computing as a Service*.

ben.blamey@it.uu.se

The HASTE Project, Funding

- HASTE: Hierarchical Analysis of Spatial TEmporal data
  - hierarchical approach to acquisition, analysis, and interpretation of image data.
  - Interestingness measurement, to form intelligent spatial and temporal information hierarchies,
  - distributing data for computation and storage based on image features.
- Funded by (SSF) under award no. BD15- 0008, and the eSSENCE strategic collaboration for eScience (Sweden).
- http://haste.research.it.uu.se/ https://github.com/HASTE-project

ben.blamey@it.uu.se

Questions?

ben.blamey@it.uu.se