Performance Models of Access Latency in Cloud Storage Systems

Qiqi Shuai, Victor O.K. Li and Yixuan Zhu
Department of Electrical and Electronic Engineering
The University of Hong Kong
June15, 2014

Outline

- **Background and Problems**
- **System Model**
- **Numerical Results and Analysis**
  - Access latency performance comparison between replica strategy and MDS codes.
  - Access latency performance comparison of different erasure codes.
- **Conclusion and Future Work**

Outline

- **Background and Problems**
- System Model
- Numerical Results and Analysis
  - Access latency performance comparison between replica strategy and MDS codes.
  - Access latency performance comparison of different erasure codes.
- Conclusion and Future Work

Massive Distributed Storage Systems in the Cloud

- Failure is norm rather than exception.
- Frequent Repairs of failed nodes.
- Ways to tolerate failure:

  simple replication and erasure coding

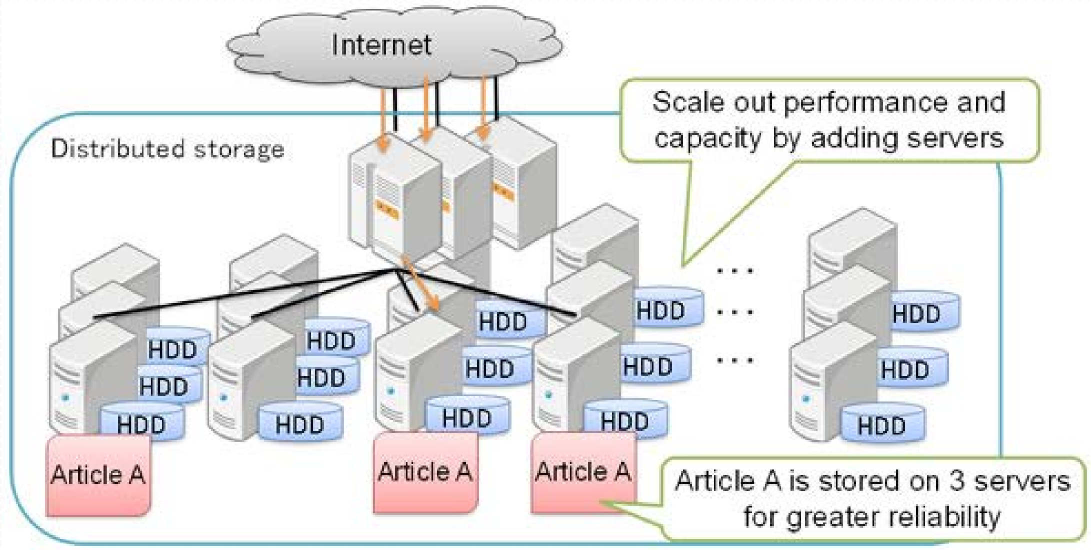

Access Latency

- Indicates the availability of storage systems.
- Greatly impacts user experience.
- For example: Google found that users performed fewer searches because of a 400-millisecond additional delay.

Existing Performance Models

- Analysis from the complexity of coding.
- Comparison of access latency between MDS codes and replica strategy with queueing-theoretic analysis.
- Fork-join queues for parallel processing.
- Redundant requests.

Existing Models Unrealistic

- Assume that each request to storage systems need to access at least k nodes.
- Ignore the usual repair process and degraded reads in cloud storage systems.

Existing Models Unrealistic

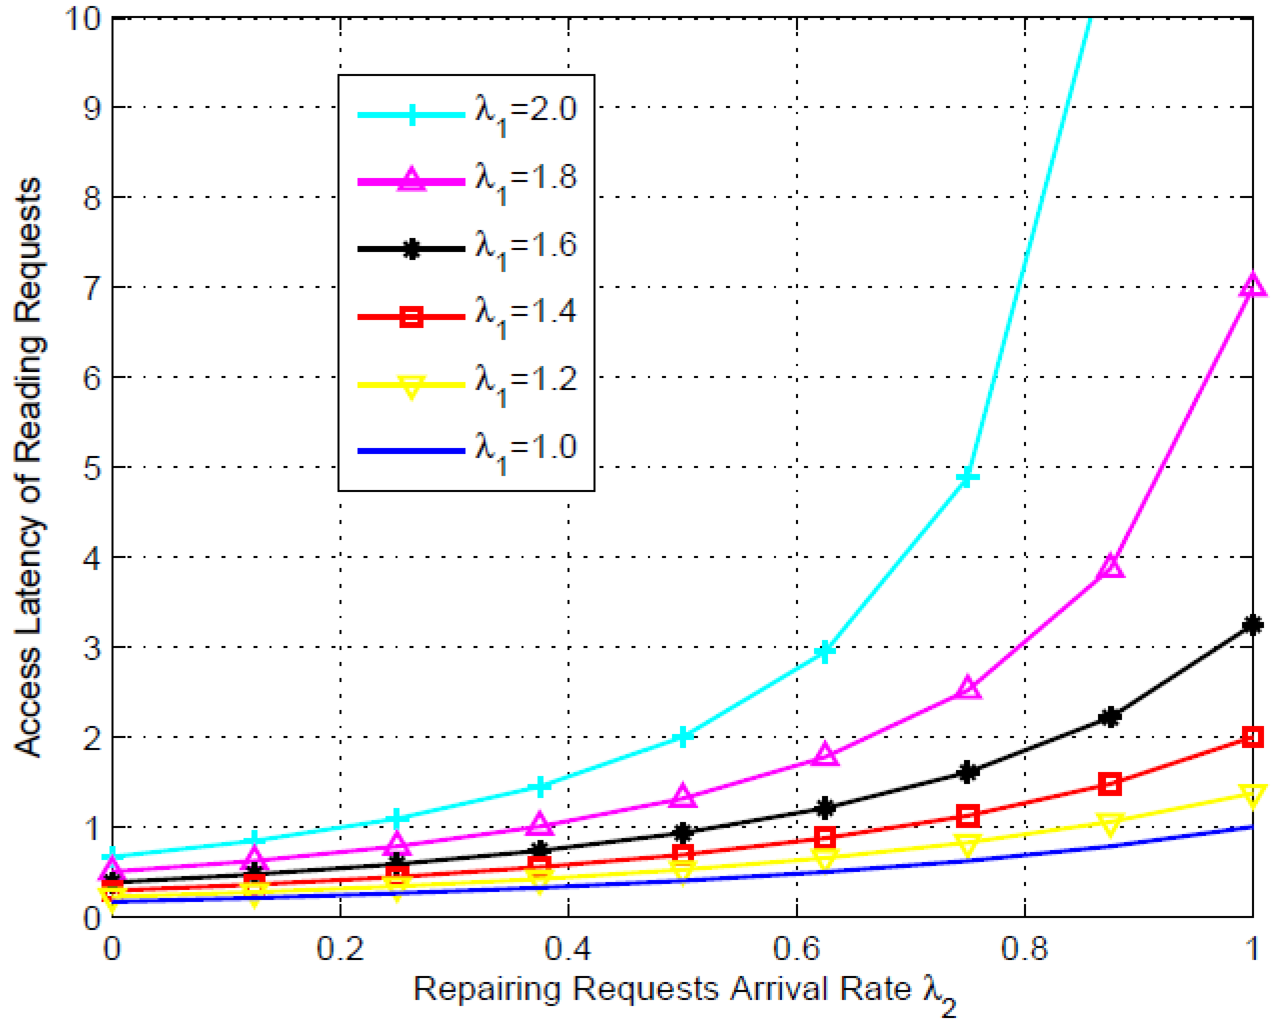

Existing Models Unrealistic

- Only compare the access latency of MDS erasure codes with the replication strategy and cannot compare the latencies of different erasure codes.

Outline

- Background and Problems
- **System Model**
- Numerical Results and Analysis
  - Access latency performance comparison between replica strategy and MDS codes.
  - Access latency performance comparison of different erasure codes.
- Conclusion and Future Work

Key Problems

- Cloud storage systems are very complicated and contain many storage nodes.
- It is almost impossible to keep track of all these storage nodes' concurrent actions on different requests simultaneously.

Our Solutions

- The system is assumed to be homogeneous as in almost all previous work.
- The access latency performance of the whole system can be estimated from an arbitrary data storage node.

Different System Requests

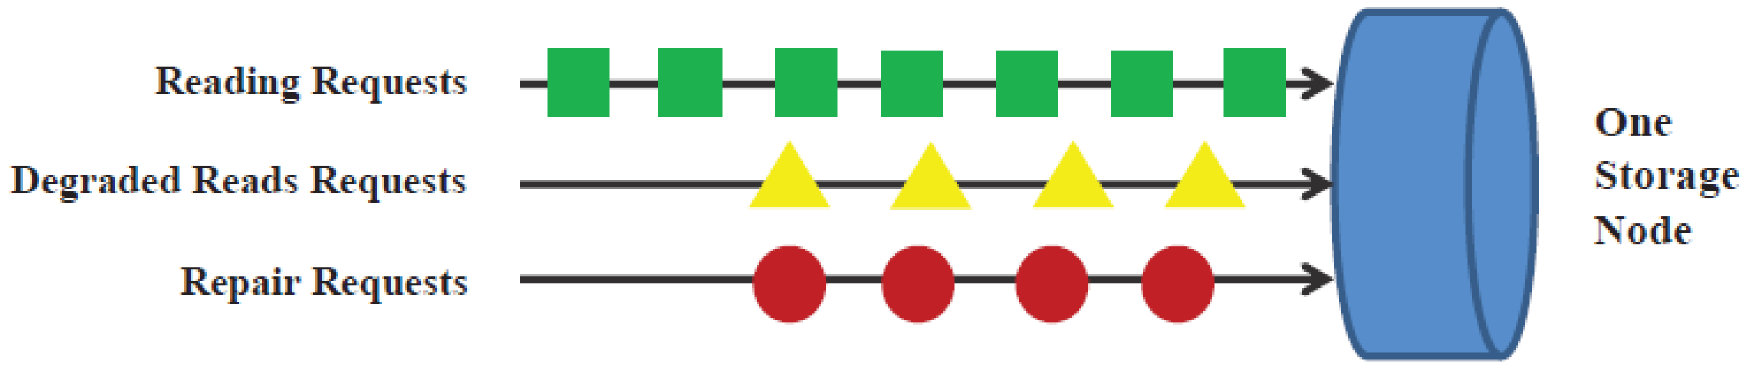

Fig. 2. Different requests to one storage node in cloud storage system.

Different System Requests

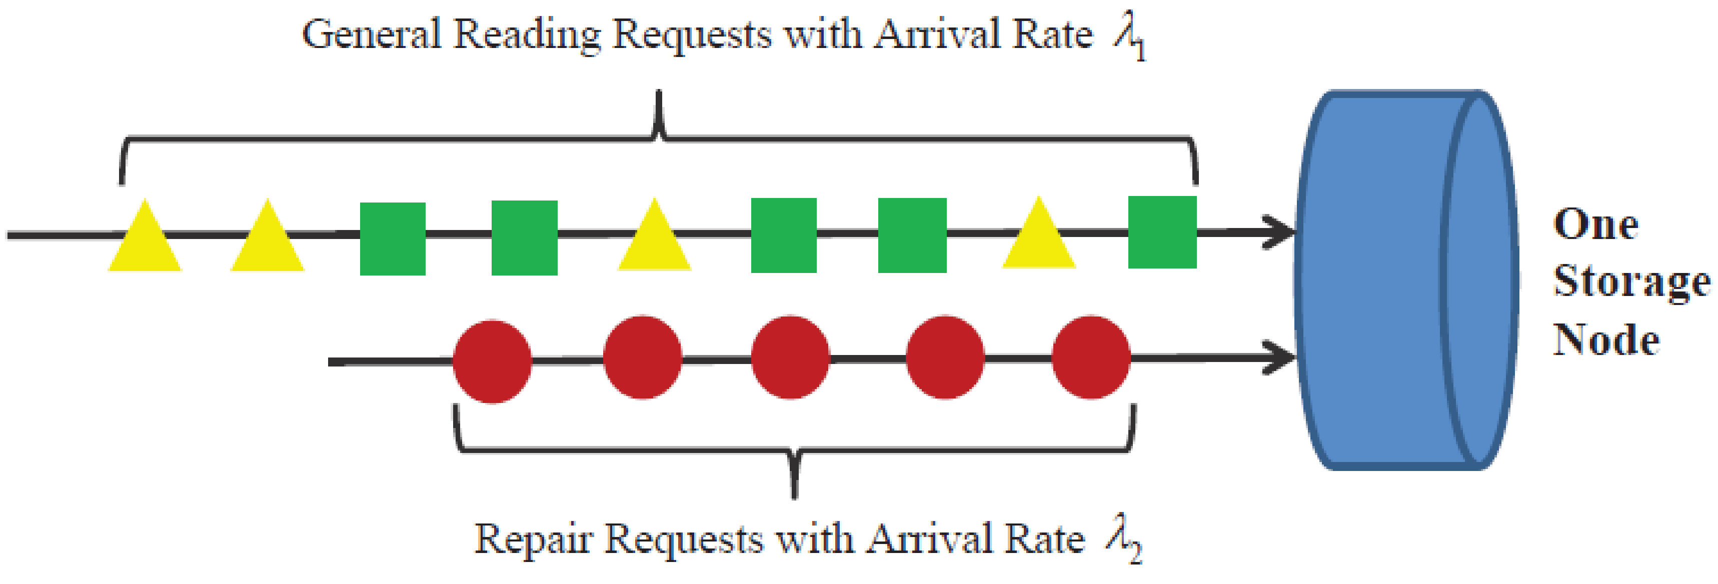

Fig. 3. General reading requests and repair requests to one storage node.

HOL Priority Queueing System

- We model the general reading and repair requests as a head-of-the-line (HOL) priority queueing system.
- Since we try our best to guarantee no data loss in cloud storage systems, repair requests should have higher priority than general reading requests.

HOL Priority Queueing System

- Assume Poisson arrivals and exponential service times for both general reading and repair requests. We can get the access latency of general reading requests:

$$w = \frac{\left(\lambda_1 \mu_2^2 + \lambda_2 \mu_1^2 + \frac{\mu_1^2 \lambda_2^2}{\mu_2 - \lambda_2} + \frac{\lambda_1 \lambda_2 \mu_2^2}{\mu_2 - \lambda_2}\right)}{\mu_1^2 \mu_2^2 - \lambda_1 \mu_1 \mu_2^2 - \lambda_2 \mu_1^2 \mu_2}$$

where λ_1 and λ_2 are the arrival rates of general reading requests and repair requests, μ_1 and μ_2 are the service rates of general reading requests and repair requests, respectively.

Outline

- Background and Problems
- System Model
- **Numerical Results and Analysis**
  - Access latency performance comparison between replica strategy and MDS codes.
  - Access latency performance comparison of different erasure codes.
- Conclusion and Future Work

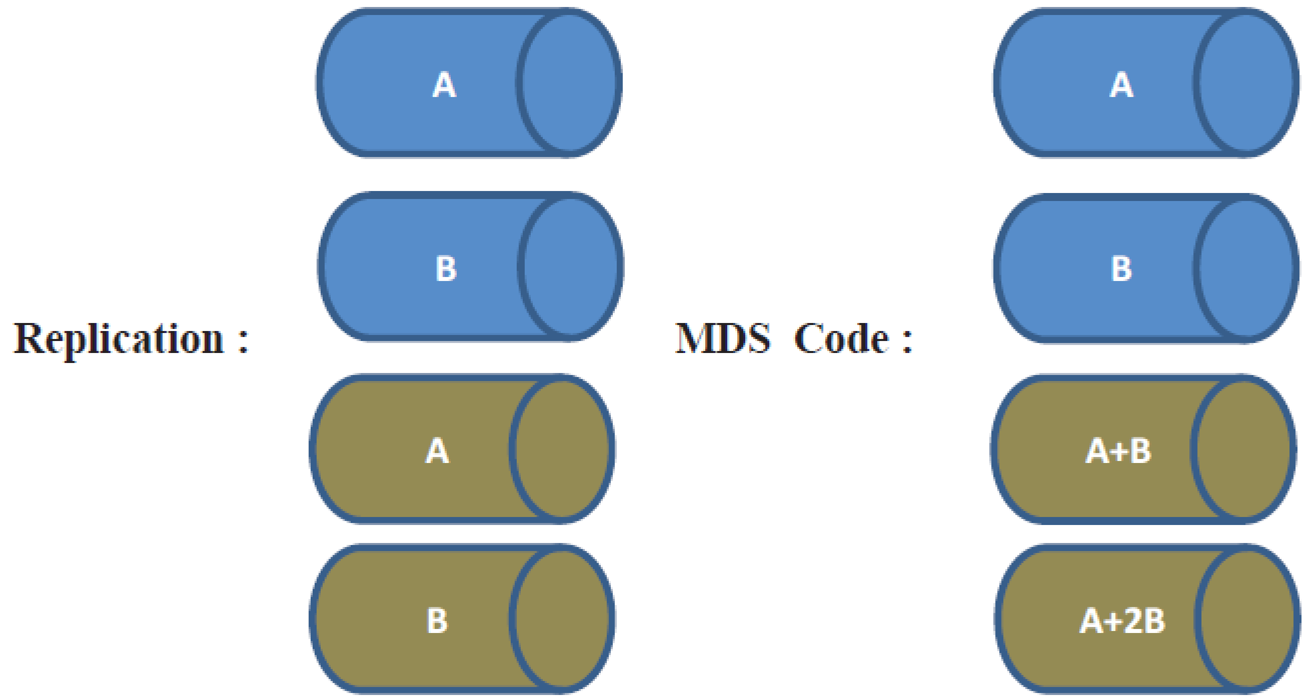

Fig. 4. 2-replica strategy storage and (4,2) MDS code storage.

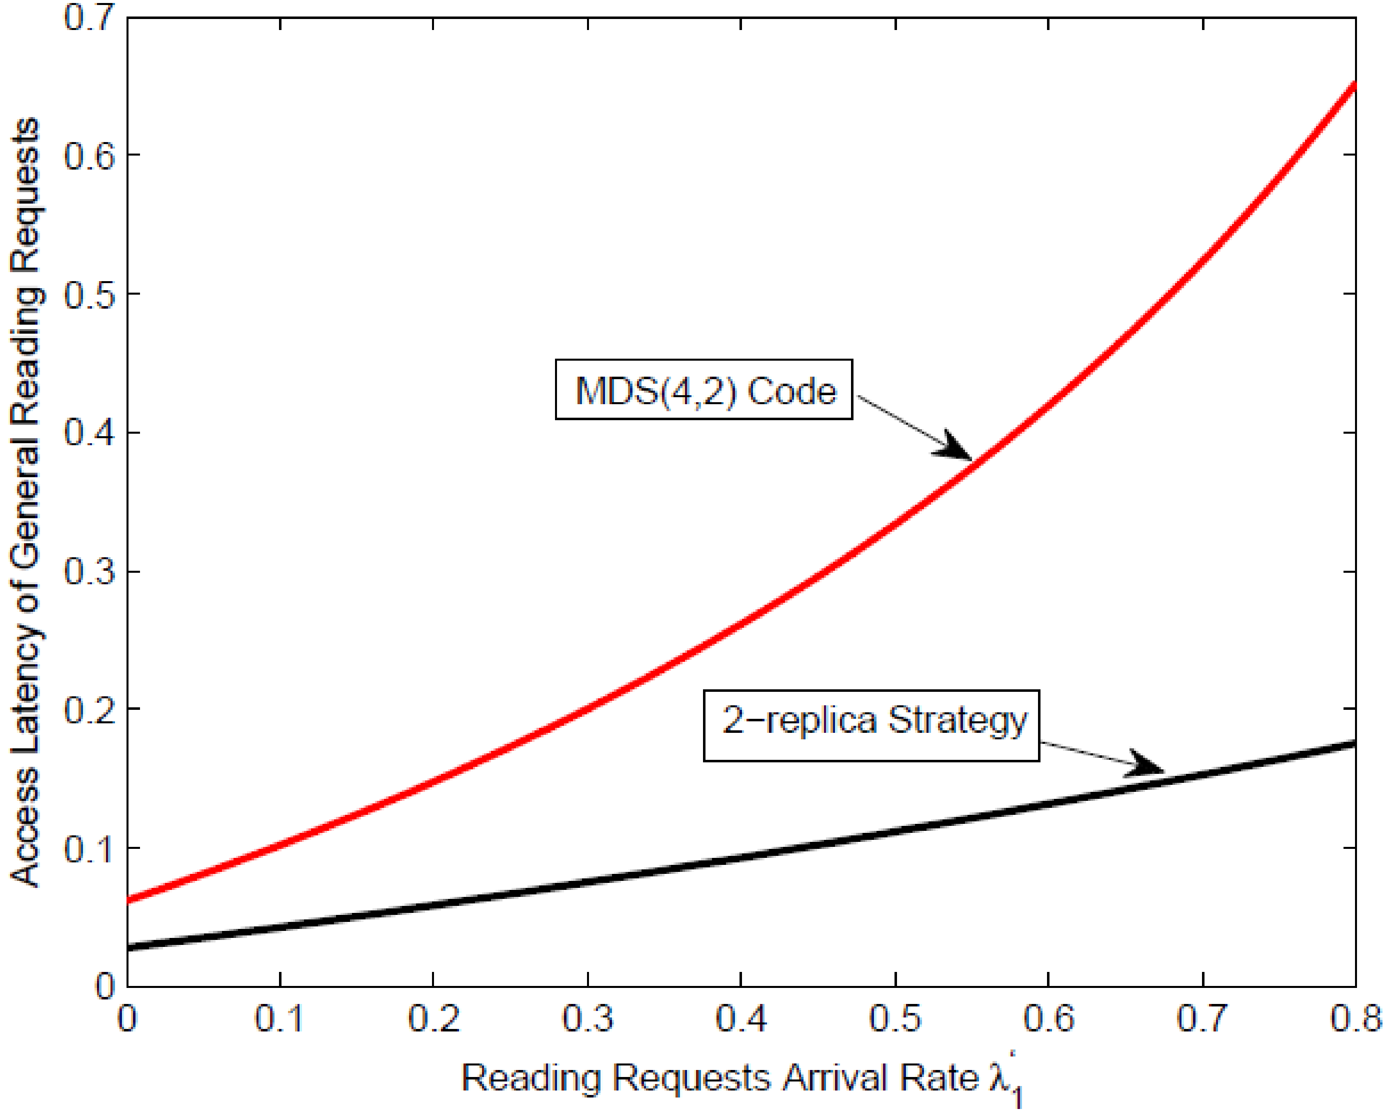

Fig. 5. Access latency comparison between 2-replica and MDS(4,2).

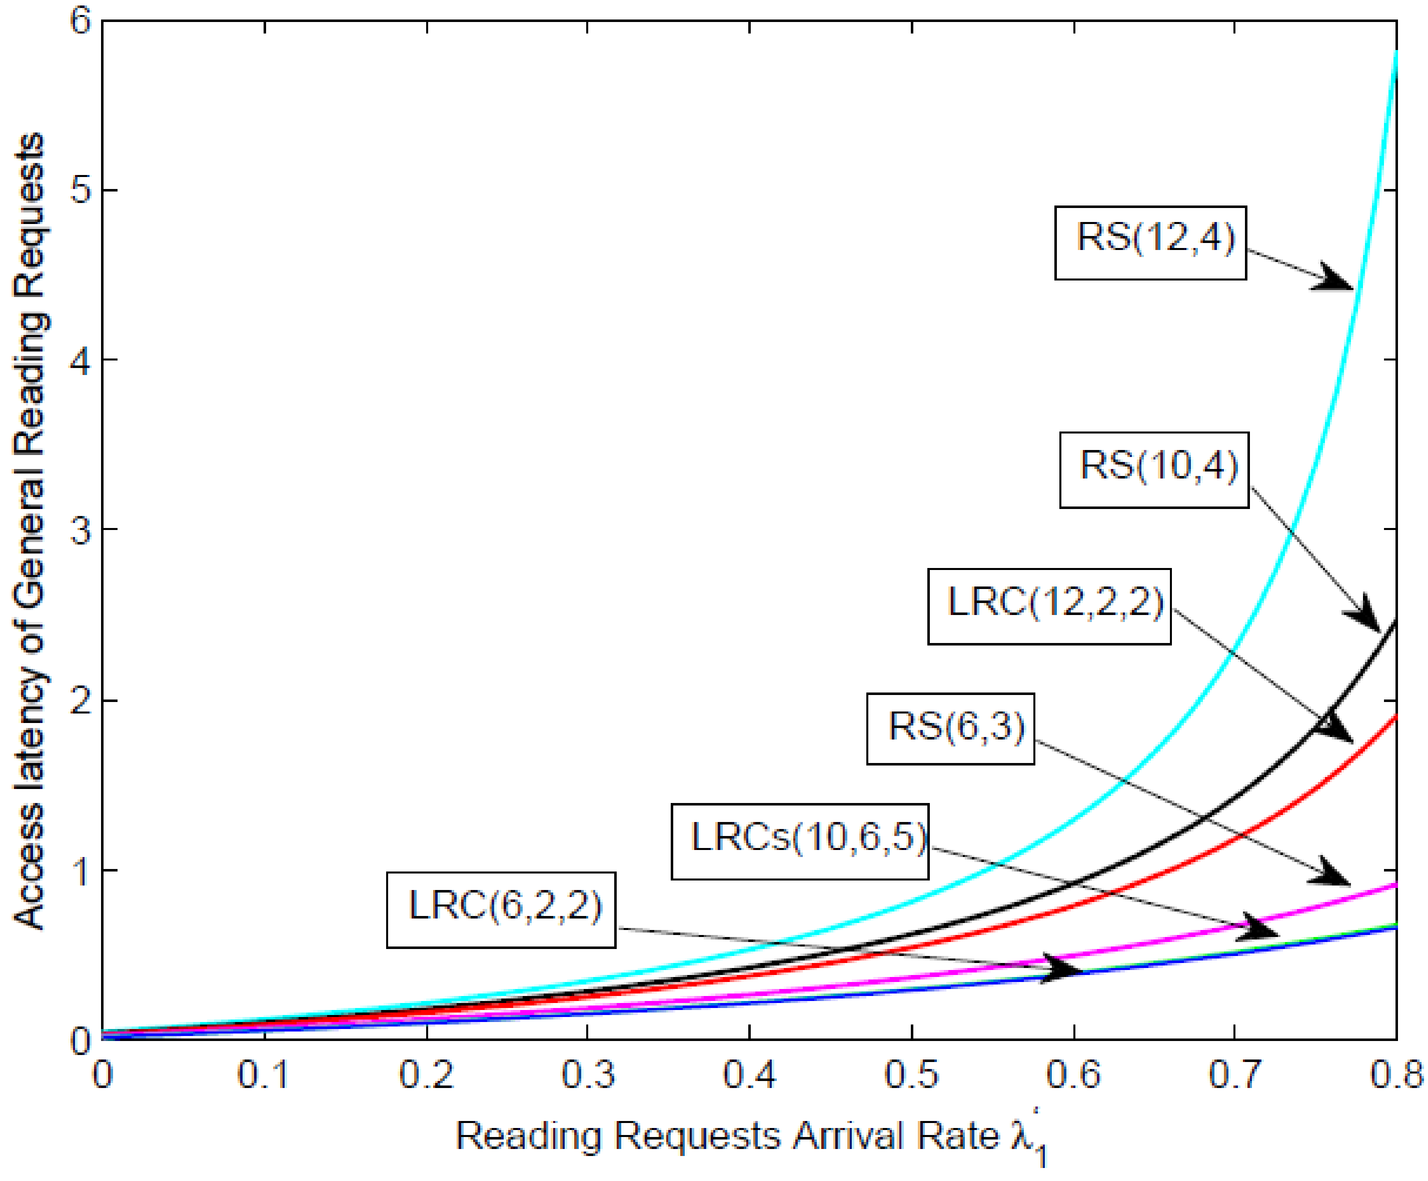

Fig. 6. Access latency comparison of different erasure codes.

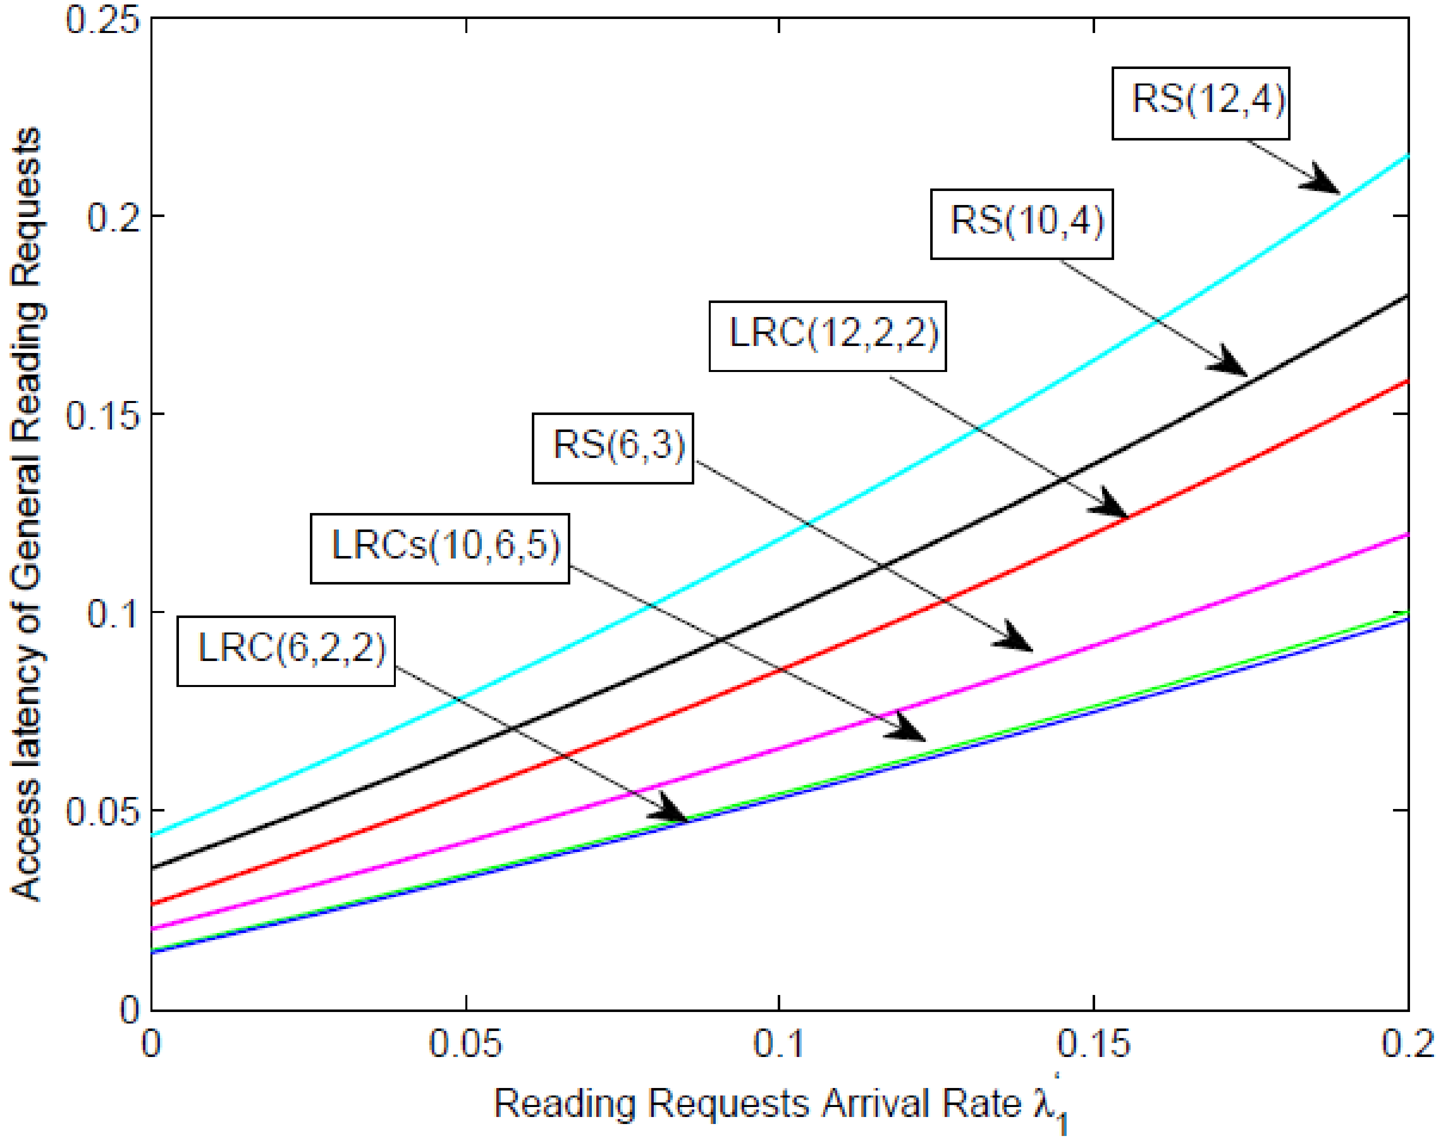

Fig. 7. Access latency comparison of different erasure codes.

Outline

- Background and Problems
- System Model
- Numerical Results and Analysis
  - Access latency performance comparison between replica strategy and MDS codes.
  - Access latency performance comparison of different erasure codes.
- **Conclusion and Future Work**

Conclusion

- In this paper, we propose a more realistic model to measure access latency.
- To the best of our knowledge, we are the first to provide a general model to measure and compare access latency of different erasure codes.

Future Work

- To confirm that our access latency performance models can better reflect the reality in cloud storage systems, we plan to build platforms in Windows Azure Storage (WAS) systems to further verify our models.

- Since computational costs of different codes also impact access latency, it will be interesting to study such impact in the future.

Q & A
Thank you!