The Orchestration of Security, Performance, and Reliability for Stored Data

Haruo Yokota

Tokyo Institute of Technology

Our Recent Research Topics
-Related to Storing Information -

- Well utilizing the ***Primary and Backup*** configuration
  - With keep the reliability for storing information
- ***Adaptive Overlapped Declustering***
  - Balancing both access load and data amount among processing nodes with quick recovery
- ***Replica-assisted Migration***
  - Data migration with keeping QoS
- ***Backup Assisted 1.5 Phase Commit Protocol***
  - Distributed atomic commit protocol
- ***Backup Assisted Revocation***
  - Security management for re-encryption in storage systems

Encryption in a storage system

- Encryption schemes for the security on a network storage [Riedel et al., 2002]
  - Encrypt-on-wire scheme
    - Data is stored in clear, and encrypted when transmitted (e.g., SSL: Secure Socket Layer)
  - Encrypt-on-disk scheme
    - Data is stored in cipher, and transmitted without any encryption process
- Encrypt-on-disk scheme is more efficient than encrypt-on-wire scheme for the performance and security.
  - Storage server does not require as much encryption work with data transfer.
  - Encrypt-on-disk scheme protects data in storage while encrypt-on-wire scheme cannot.

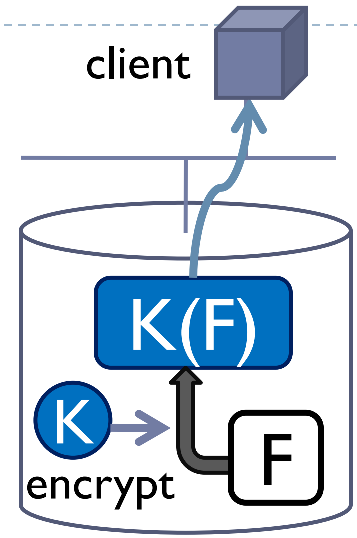

encrypt-on-wire

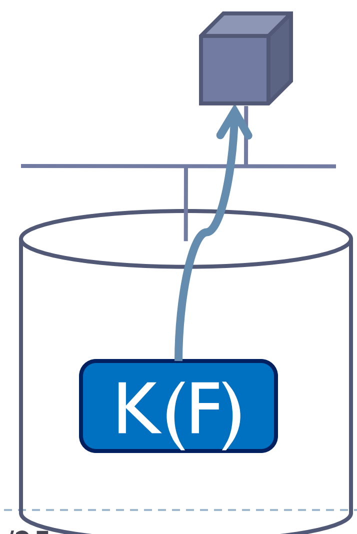

encrypt-on-disk

Revocation methods on encrypt-on-disk (1/2)

- With encrypt-on-disk, shared files must be re-encrypted when revocations occur.
  - There are possibilities of information leakage, if the revoked user holds the cryptographic key and intercepts the files.
- Re-encryption methods [Fu, 1999]
  - Active Revocation:
    - Files are re-encrypted immediately after the revocation.
    - It is enough secure
      - Revoked users are immediately unable to the decrypt data
    - It has a problem of performance
      - Even authorized users cannot access them until re-encryptions are completed.

Revocation methods on encrypt-on-disk (2/2)

- Lazy Revocation:
  - The re-encryptions is delayed until the files are next updated
  - It is more efficient in respect of performance
    - Encryption involved in update process can be combined with the re-encryption required for revocations
    - The re-encryption work for multiple revocations are performed together if the file is not frequently updated
  - It is vulnerable
    - Data stored before update are still encrypted with the old key, which can be accessed by the revoked users.

- There is a trade-off problem between performance and security.

Backup Assisted Revocation (BA-Rev)

- We have proposed BA-Rev to attack the trade-off problem.
  - BA-Rev utilizes the primary-backup configuration.
- Outline
  1. Stores backup data with encrypted by key (K_2) different from that in primary (K_1)
  2. When a revocation occurs, their roles is changed
  3. Old primary data is re-encrypted by another key (K_3) and stored as backup
- Do not need wait for re-encryption
  - Authorized users can access the file immediately after the revocation.
  - Re-encryption processes are performed in background.
    - Revocation does not so frequently occur

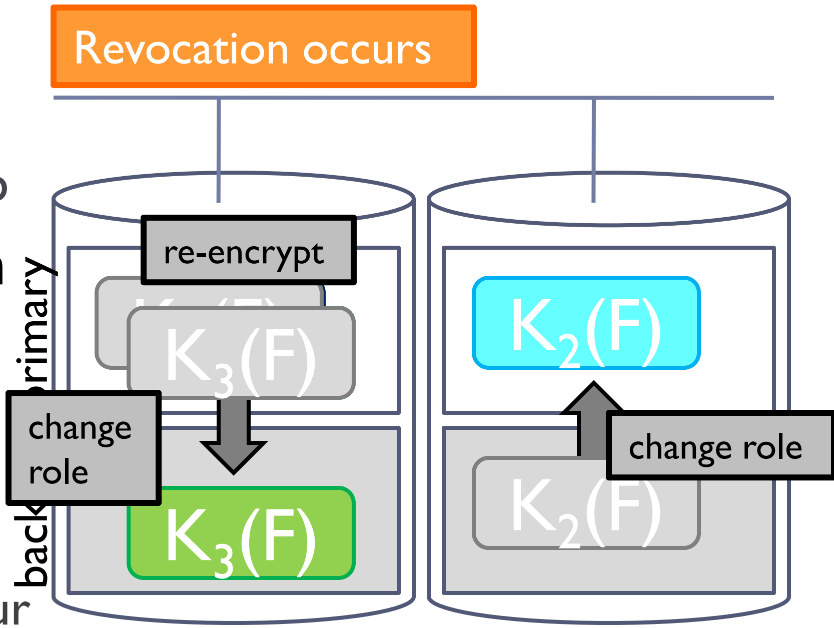

Update performance with BA-Rev

- Naïve BA-Rev is disadvantaged in its update performance.
  - Because the differential data must be re-encrypted for backup, 2 re-encryption processes for an update make response time long.
    - The graph shows average response time of update process when average arrival interval of requests is 400 ms.

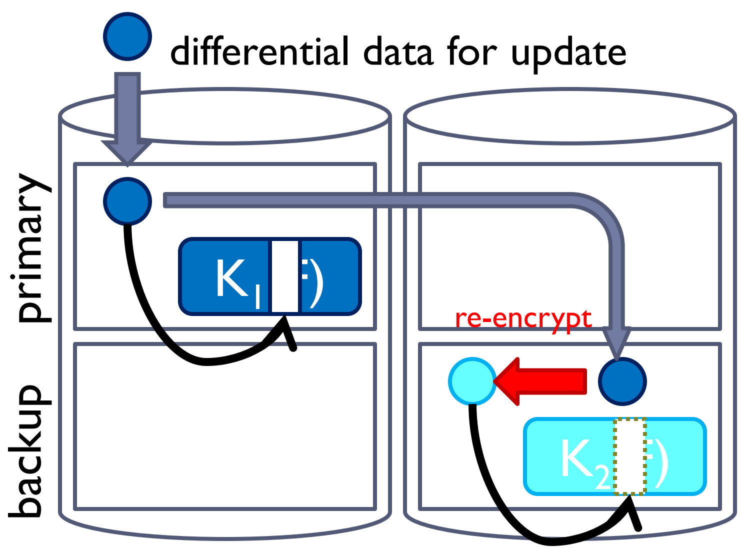

update: BA-Rev

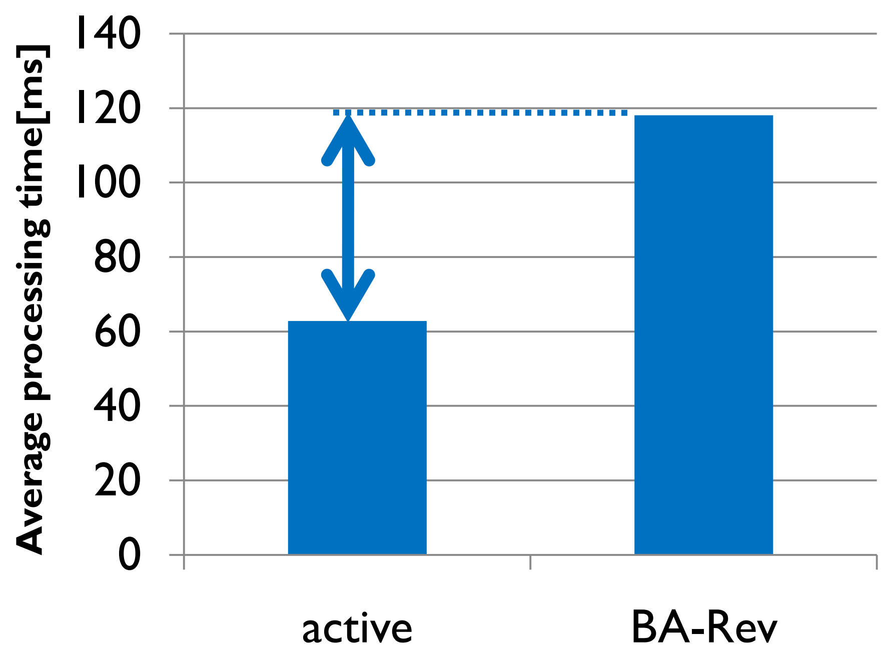

To improve the BA-Rev update performance

- DW(delayed writing) strategy
  - In backup, the process of writing update data is delayed
  - The re-encrypted update data are hold on the memory.
- DRW(delayed re-encrypting and writing) strategy
  - The processes of re-encrypting and writing data are delayed.
- Multiple update processes are aggregated in writing for both strategy

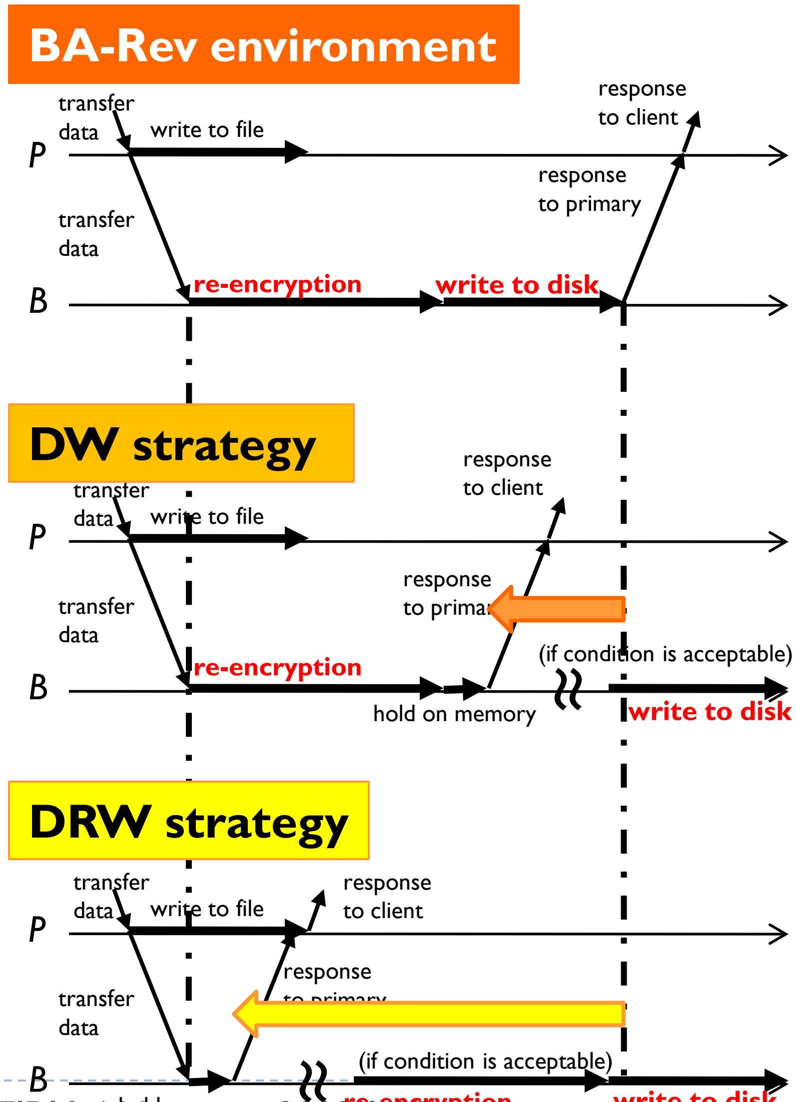

Timing of writing data into disks in DW/DRW

- Timing of writing data into backup disks affects the performance.
  - The timing is decided by conditions on data stored in memory.
- We consider three types of conditions:
  1. Basic condition (DW_{raw} / DRW_{raw})
     - When the amount of unapplied updated data exceeds a threshold, or
     - When a revocation occurs
  2. $DW_{const:n}$ / $DRW_{const:n}$
     - When the above basic condition is met, and for *n* seconds after an update occurs
  3. $DW_{load:n}$ / $DRW_{load:n}$
     - When the basic condition is met, and the number of active threads for constant interval is lower than *n*

Experimental environment

- We use three PCs as clients and other three PCs as storage nodes.
  - Files are stored in 3 storage nodes
  - Each client sends requests of GET and UPDATE to each storage node
    - Files are selected in accordance with Zipf function
    - The access interval is determined by an exponential distribution
  - We measure average response times

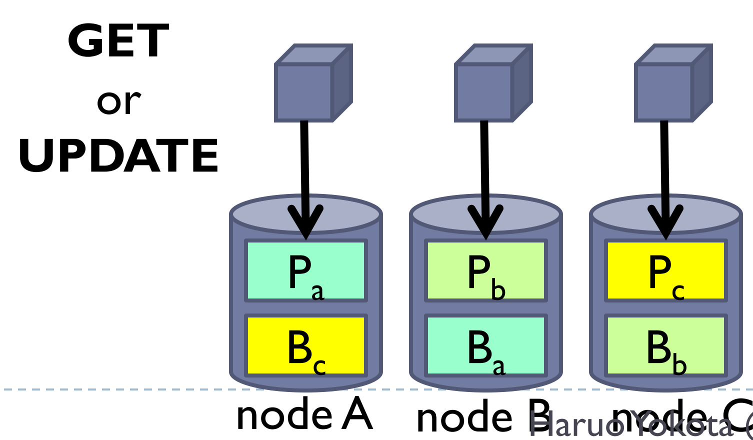

CPU	AMD Athlon XP-M1800+ (1.53GHz)
Memory	PC2100 DDR SDRAM 1GB
HDD	TOSHIBA MK3019GAX (30GB, 5400rpm, 2.5inch)
Network	TCP/IP + 1000BASE-T
OS	Linux 2.4.20
Java VM	Sun J2SE 1.5.0_03 Server VM

Secret key algorithm	AES 128bit
Public key algorithm	RSA 1024bit
Encryption mode	ECB
Padding	PKCS5
Zipf parameter θ	0.7
Number of storage nodes	3
File size	1MB
Number of files	500 /node
Size of updated data	100KB

Response times of usual accesses

- We measured average response times of UPDATE for the average arrival interval from 500ms to 150ms
- There was little benefit in applying the DW strategy to BA-Rev
  - Disk write operations are quickly processed by caching

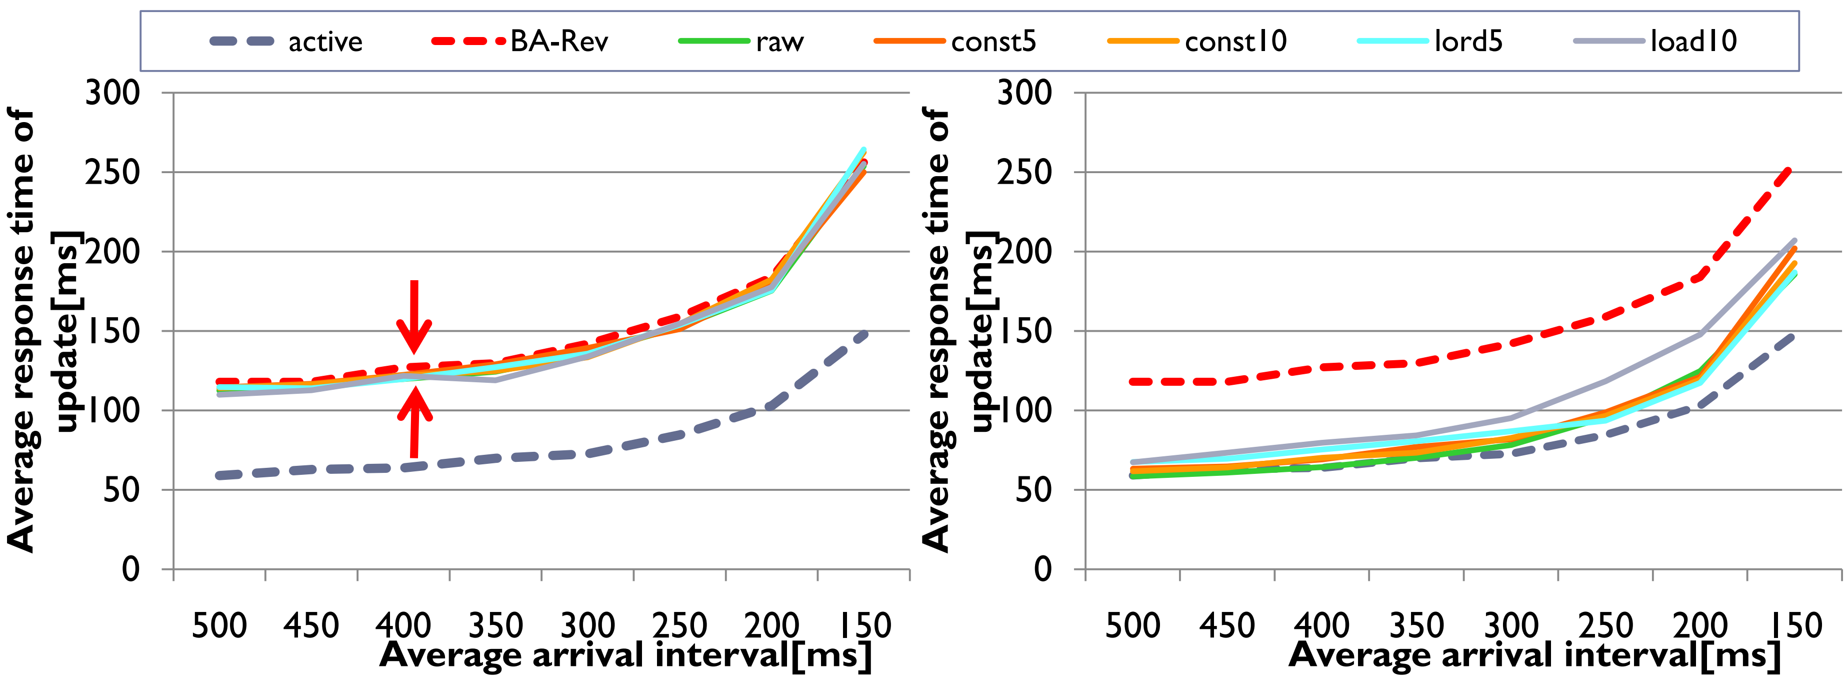

apply DW strategy

apply DRW strategy

Response times of usual accesses

- DRW strategy significantly improved the update performance which is very close to active revocation
  - The delay of re-encryption process have good effect
  - The performance of DRW$_{load:10}$ is inferior at high load
    - Multiple update operations are influenced during the delay

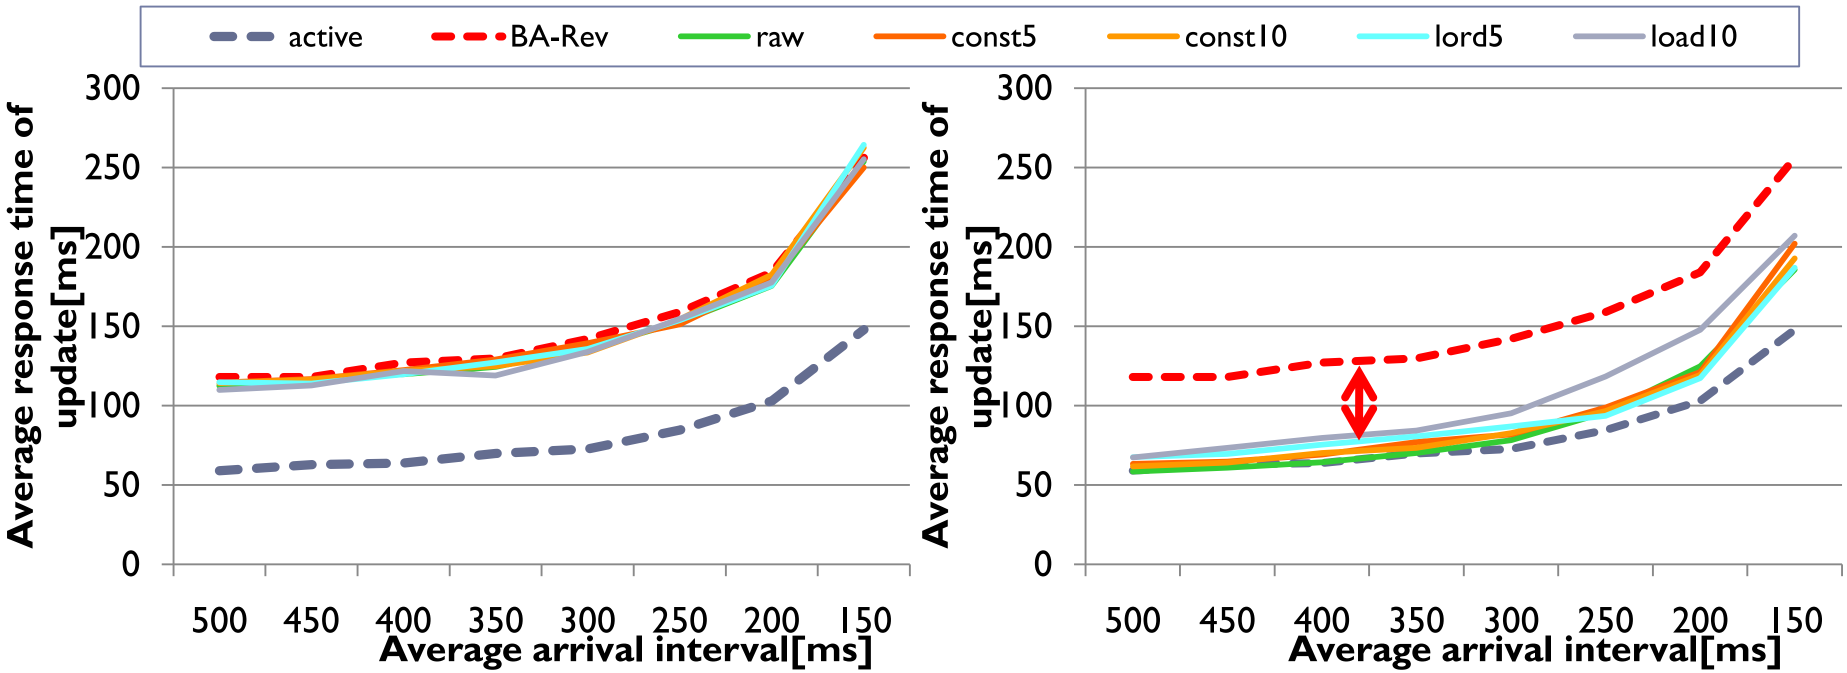

apply DW strategy

apply DRW strategy

Performance under concentrated revocations

- We enforced revocation processes for 50 files stored in node B, and measured response times of 100 accesses
- Response times of BA-Rev with or without DW/DRW are shorter than those of Active revocation
  - node B: Re-encryption processes are executed in background
  - node C: Only the role change of backup data is done

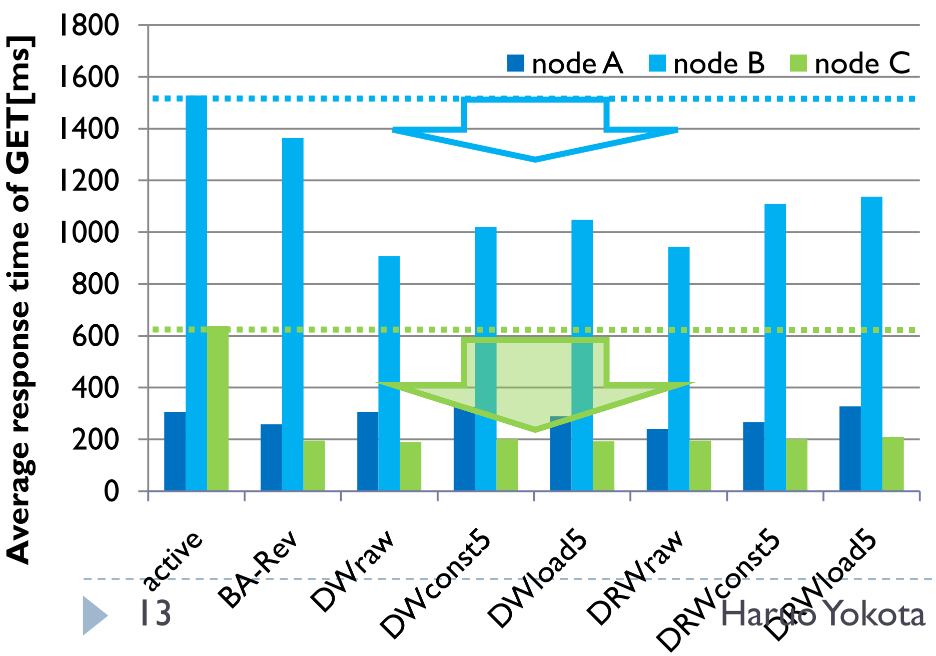

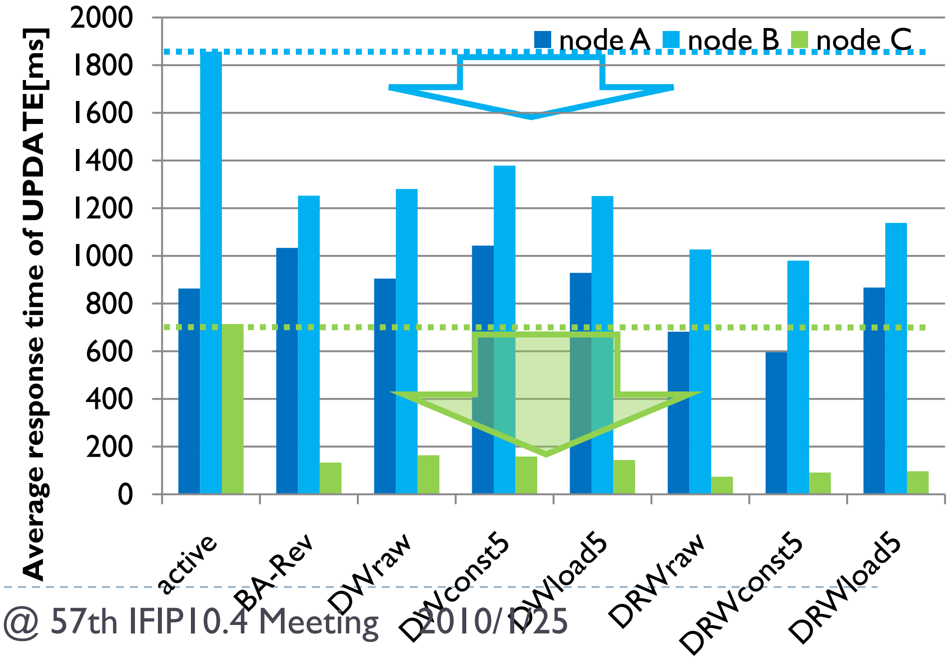

Performance under concentrated revocations

- Comparing DW and DRW with BA-Rev,
  - Two strategies have better GET performance because they reduce the number of writes of updated data in backup
  - DRW, in which re-encryption in backup is delayed, has better UPDATE performance than others

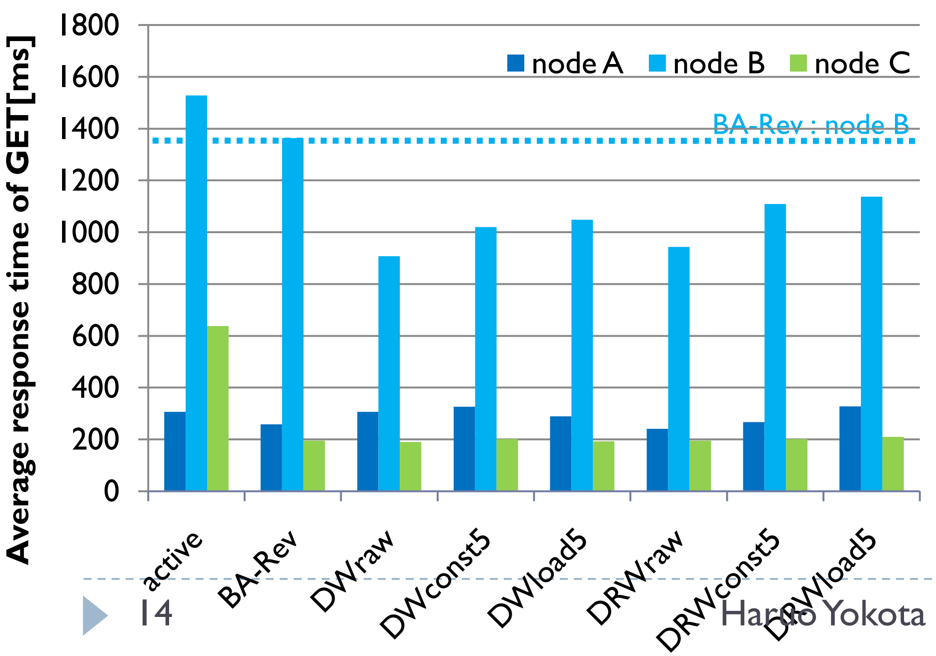

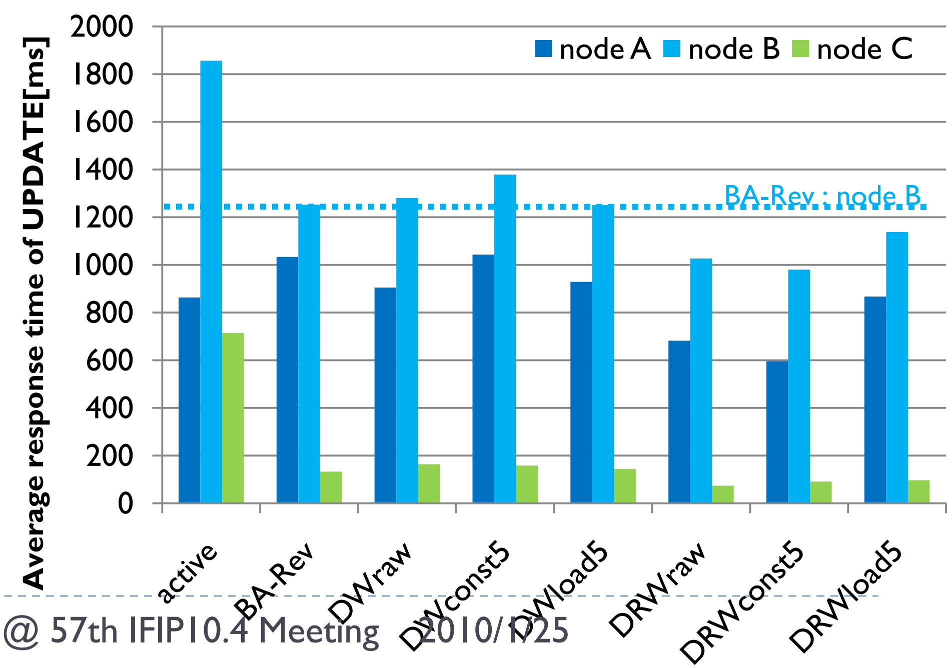

Performance under distributed revocations

- We enforced multiple revocations for 15 files in each node and measured average response times of 100 accesses.
- Each proposed environment has better performance for both GET and UPDATE compared with Active revocation.
- In particular, the UPDATE response time of DRW is the shortest.
  - The update cost is low because of delayed re-encryption and the load affected by revocation processes becomes low.

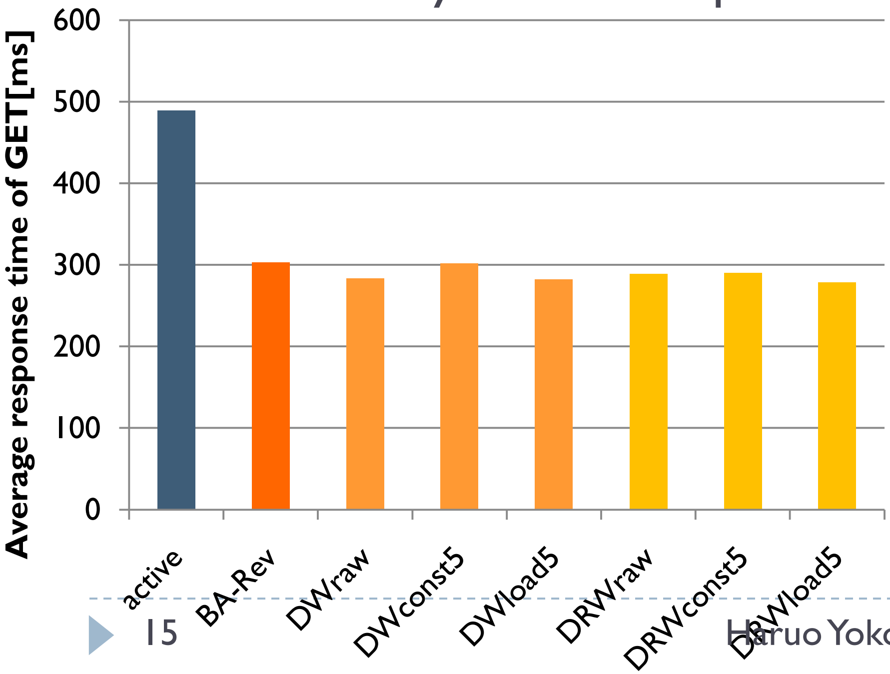

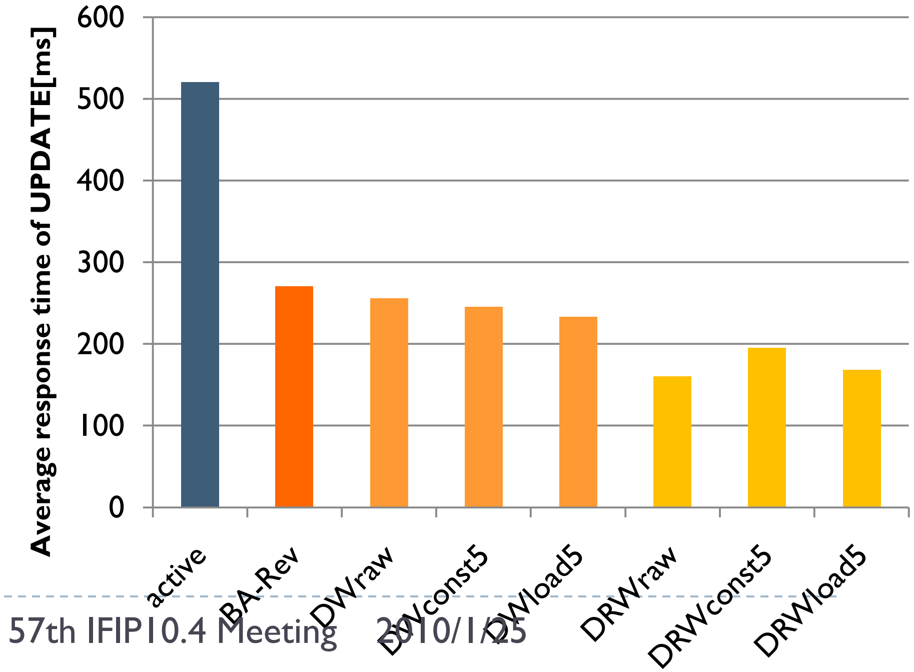

Reliability Estimation with DW/DRW

- To evaluate the reliability of DW/DRW strategy
  - The possibility of data loss may be higher because the updated data form backup are kept in volatile memory.
    1. MTTR (Mean Time To Repair) increases because unapplied updated data must be written to disk at failure restoration.
    2. Data are lost if disk and power supply failures occur at same time .
- We estimate MTTDL (Mean Time To Data Loss) about node failure and power supply failure in order to evaluate the reliability.
  - We assume UPS (Uninterruptible Power Supply) for each node
  - We compare
    - CD (Chained Declustering) : normal environment in which the applying updated data is not delayed such as Active revocation or BA-Rev
    - DW strategy
    - DRW strategy

Reliability Estimation with DW/DRW

- We calculated MTTDL for disk failure and power failure independently.
- The lines of MTTDL for disk failure are almost overlapped.
  - The difference is smaller than one percent, though that in DRW is worst because re-encryption process must be performed at failure restoration.
- MTTDL for power failure is much longer than that for disk failure.
  - The probability of data loss related to power failure is small.

→ The reliability degradation with DW and DRW is very small.

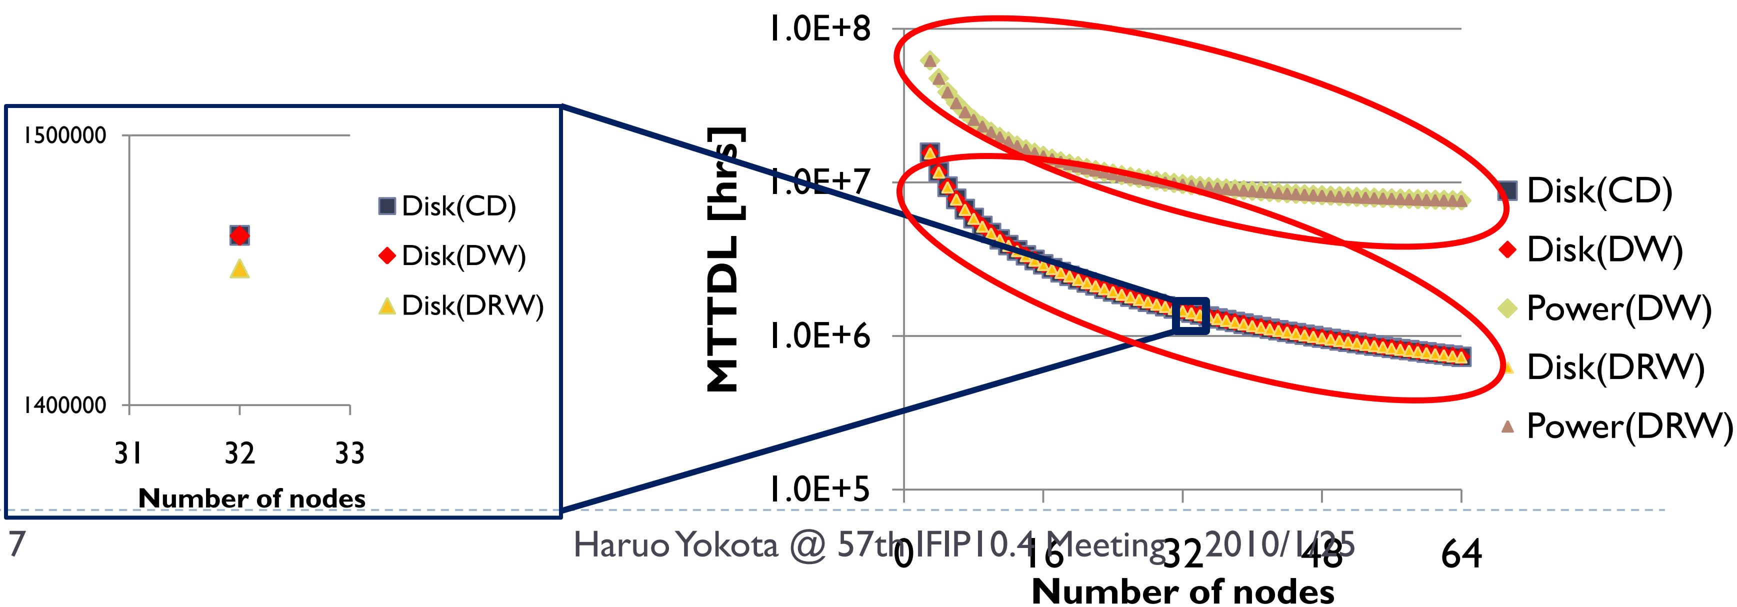

Summary

- The orchestration of security, performance, and reliability for stored data
  - We proposed BA-Rev (Backup Assisted Revocation), an efficient re-encryption method for revocation than active revocation.
  - We apply DW/DRW strategy, in which applications of updated data to backup data are delayed, to improve the update performance.
    - BA-Rev with DRW realize the update performance equivalent to active revocation, and improved revocation performance.
  - We estimate reliability of BA-Rev with DW/DRW and show that the decrease of MTTDL is very small.
- Future work
  - Evaluate the proposed approach in actual environments including different size of files accessed from heterogeneous applications.

Thank you for your attention!