Analysis of and Optimization for Write-dominated Hybrid Storage Nodes in Cloud

Shuyang Liu[1], Shucheng Wang[1], **Qiang Cao** [1], Ziyi Lu[1], Hong Jiang [2], Jie Yao[1], Yuanyuan Dong[3] and Puyuan Yang[3]

*Huazhong University of Science and Technology
UT Arlington
Alibaba

Outline

- ✓ Background
- ✓ Trace Analysis
- ✓ Design of SWR
- ✓ Evaluation
- ✓ Conclusion

Hybrid Storage

- Combine SSD and HDD to maximize performance and capacity while minimizing cost
  - SSD: high GB/s(0.5-3), low latency(us), high $/GB(0.5-2.6)
  - HDD: low GB/s(0.2), high latency(ms), low $/GB(0.2-0.45)

- SSD as write buffer (SSD Write Back, **SWB** mode)
  (1) First write incoming data into SSD
  (2) Then flush them into HDD in the background

Pangu

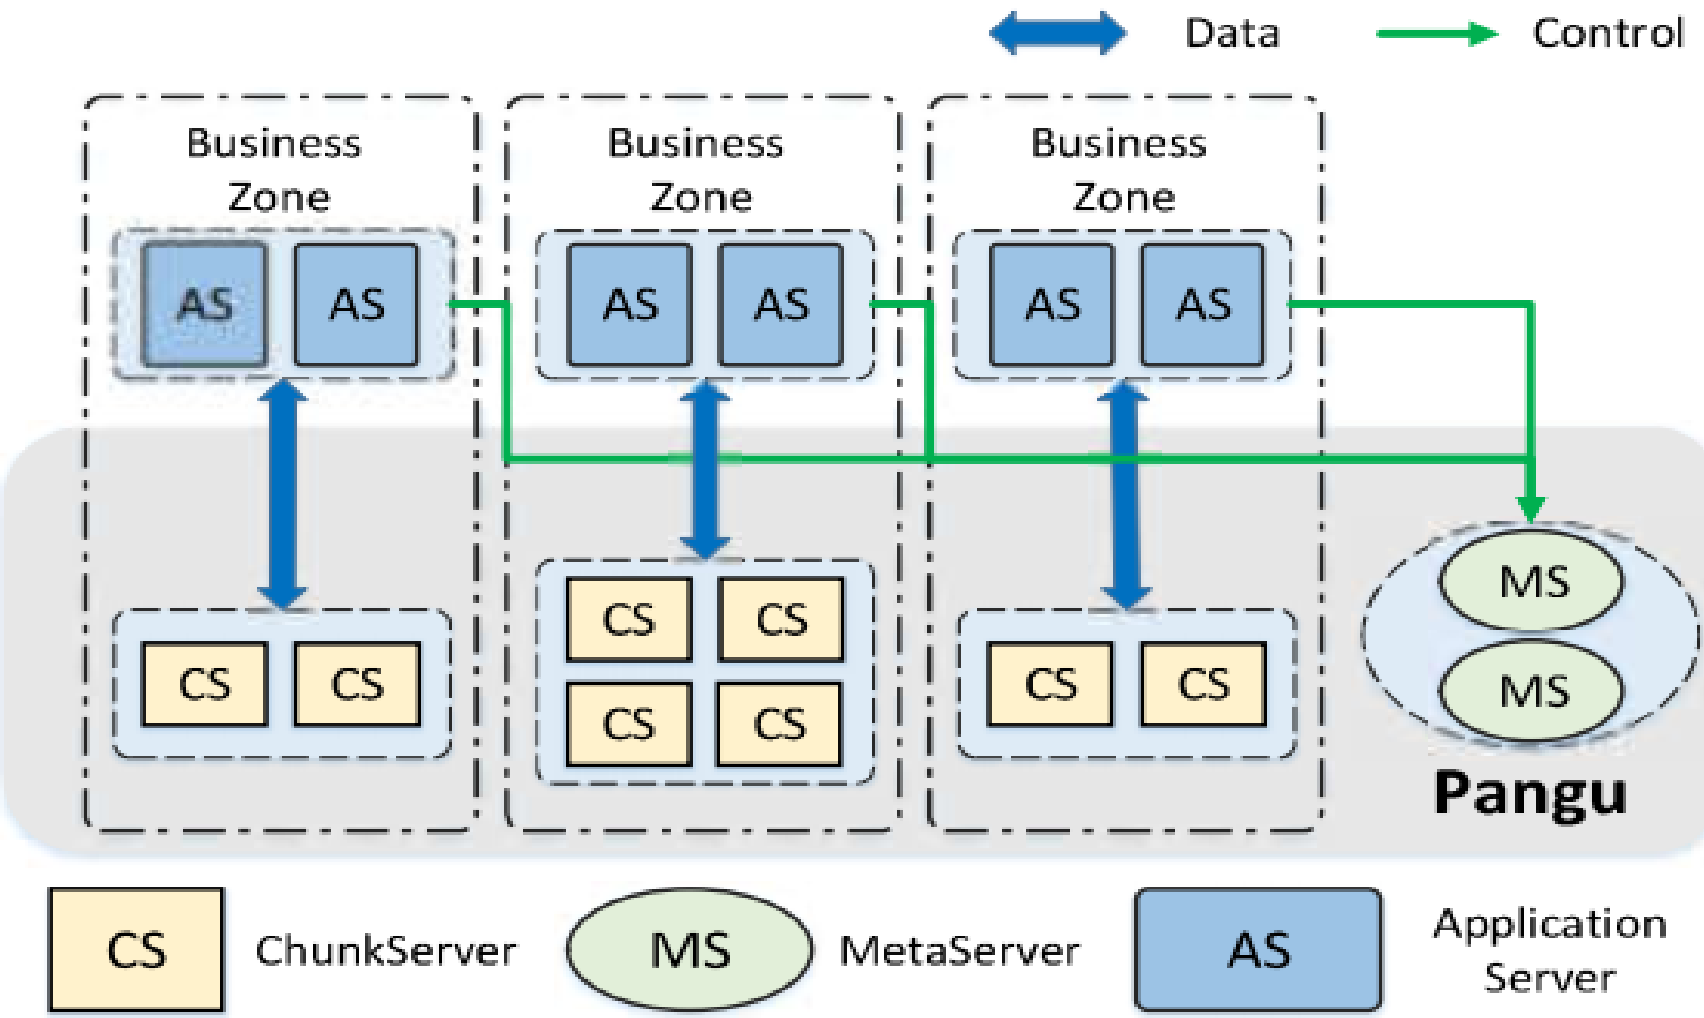
Data
Control
Business Zone
AS
AS
Business Zone
AS
AS
Business Zone
AS
AS
CS
CS
CS
CS
CS
CS
CS
CS
MS
MS
Pangu
CS
ChunkServer
MS
MetaServer
AS
Application Server

Chunk Server

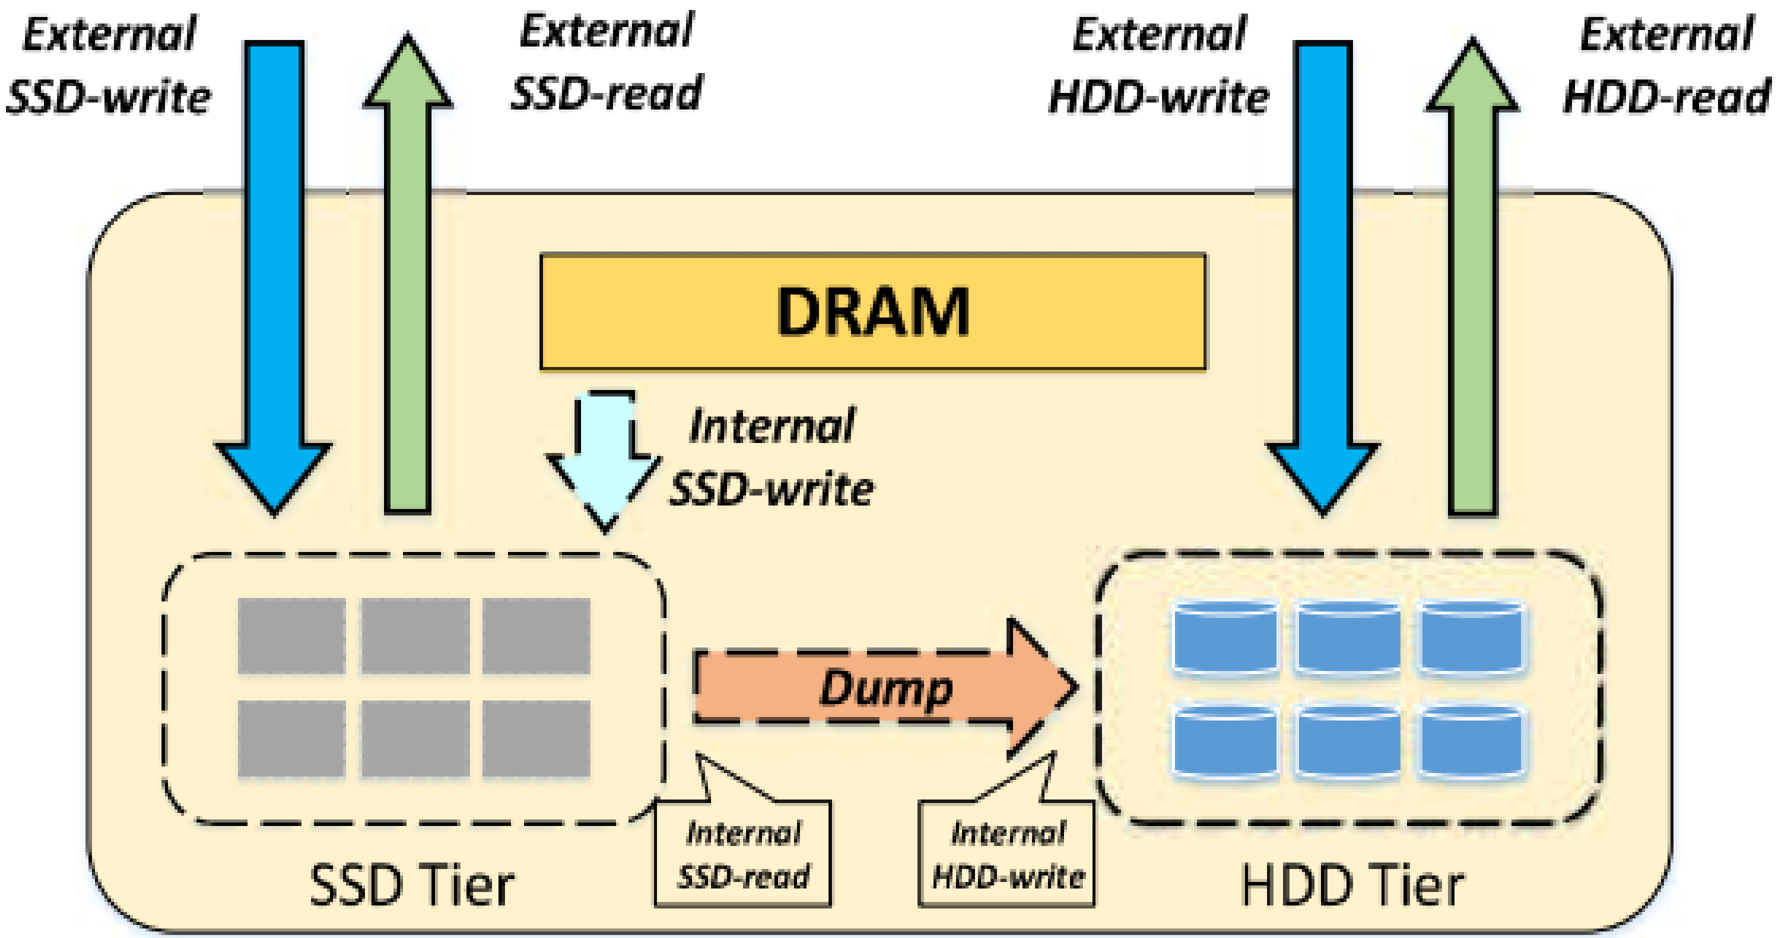

Write-dominated Storage Nodes

- WSNs: ChunkServers in Pangu experience a write-dominant workload behavior.
- Feature:
  - 77%-99% of requests are writes.
  - The amount of data written is much larger than data read.
- Reason:
  - Frontend applications with their own cache layers need rapidly flush all writes into Pangu and reserve their local storage for hot data .
  - Pangu provides a unified persisent platform.

Outline

- ✓Background
- ✓Trace Analysis
- ✓Design of SWR
- ✓Evaluation
- ✓Conclusion

Trace Analysis Summary

Problems according to trace analysis on Pangu production traces

- SSD overuse
- Long-tail write latency
- Low utilization of HDD

Workload Traces

- **Three Business Zones**: A(Cloud Computing), B(Cloud Storage), C(Structured Storage).
- **Nodes**: A1, A2, B, C1, C2
- **Time duration**: 0.5-22hour
- **Number of requests**: 28.5-66.9 millions
- **SSD ratio**: 1 Low(<10%), 2 Mid(10%-33%), 2 High(>33%)
- **Write request ratio**: 77.2%-99.3%
- **Average IO interval**: 62us-2ms
- **Average request size**: 4.1-177 KB

Trace Record: Example

- TimeStamp: 2019-01-24 11:20:36.158678 (us)
- **Operation**: SSDAppend
- ChunkId: 81591493722114_3405_1
- SATADiskId: -1
- **SSDDiskId**: 1
- Offset: 56852480 (byte)
- **Length**: 16384 (byte)
- **Waiting delay**: 76 (us)
- **IO delay**: 213 (us)
- **QueueSize**: 1
- ……

Load Behaviors across Chunkservers

- Load balancing across ChunkServers.
- Load Intensity varying over time

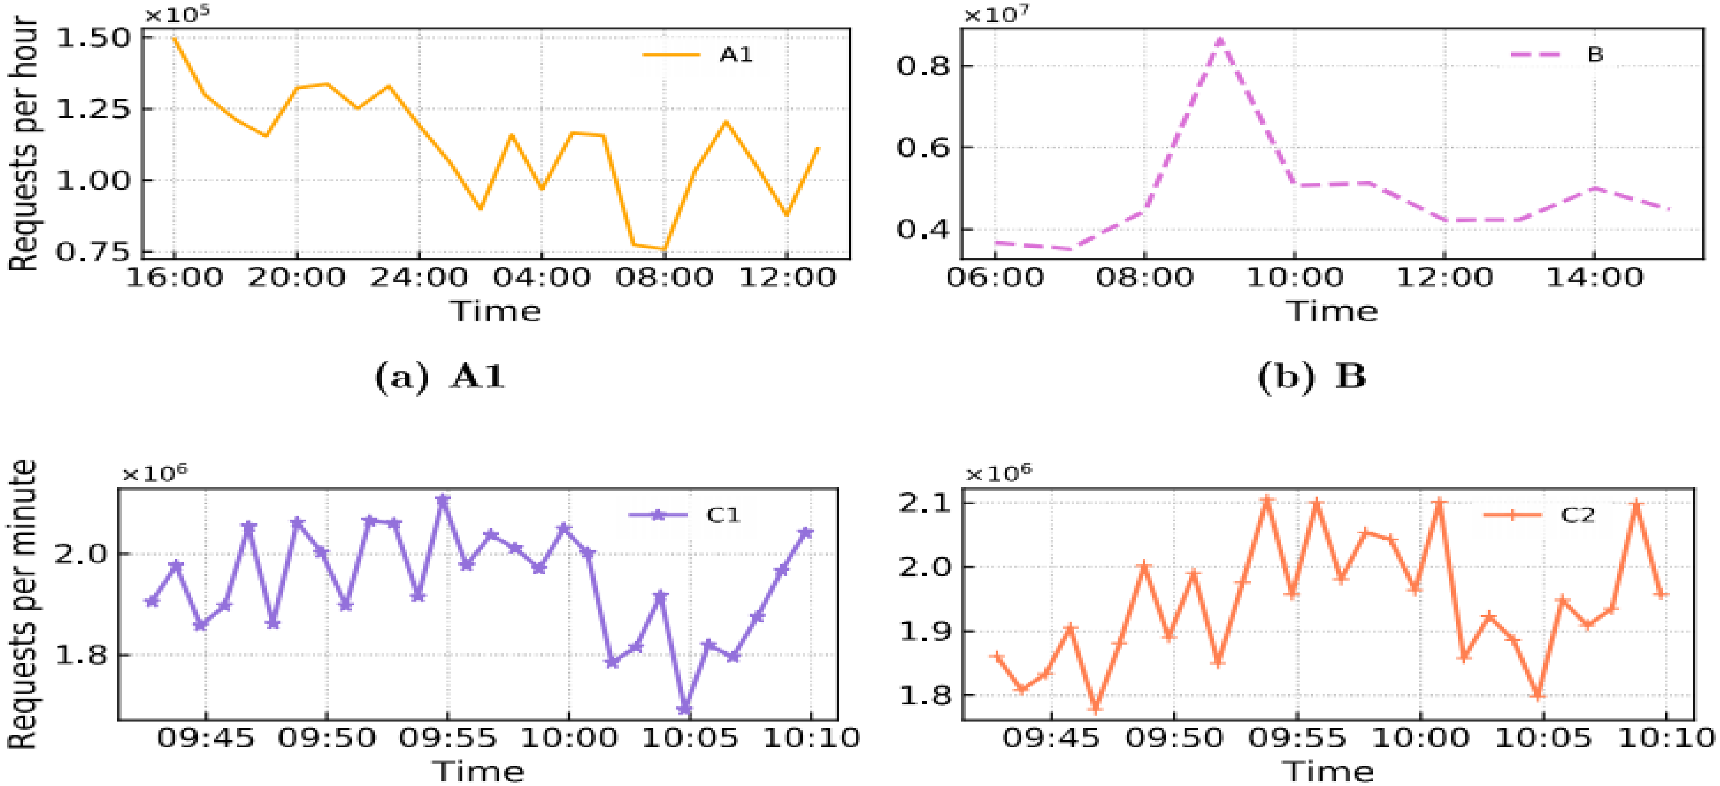

(a) **A1** (b) **B**

(c) **C1** (d) **C2**

Load Behaviors across Disks within Chunkservers

- load balancing across internal disks

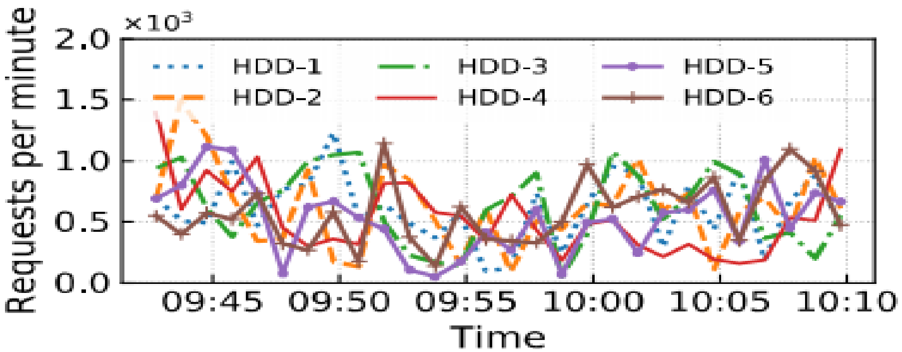

(a) External HDD-write(requests)

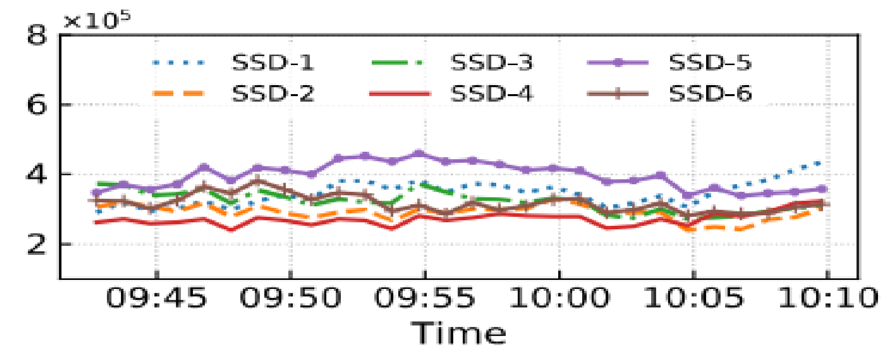

(b) External SSD-write(requests)

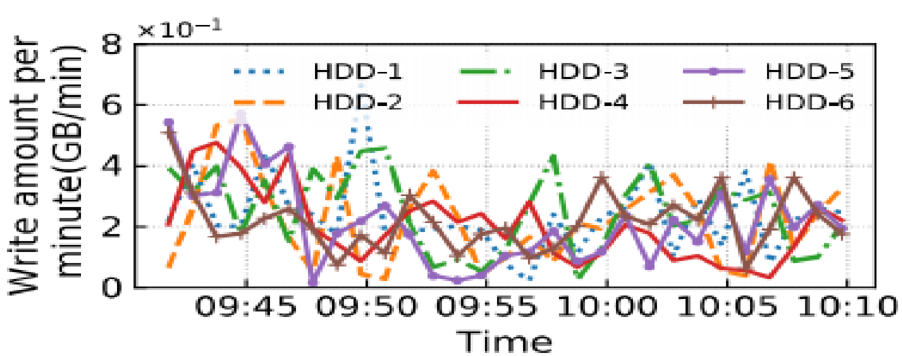

(c) External HDD-write(data)

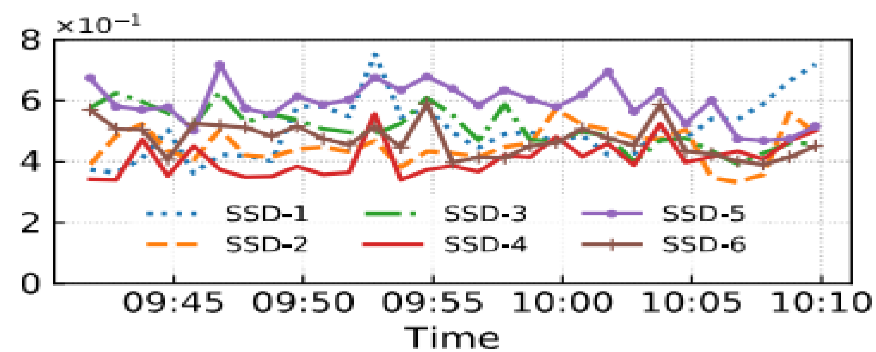

(d) External SSD-write(data)

Operation type and Proportion

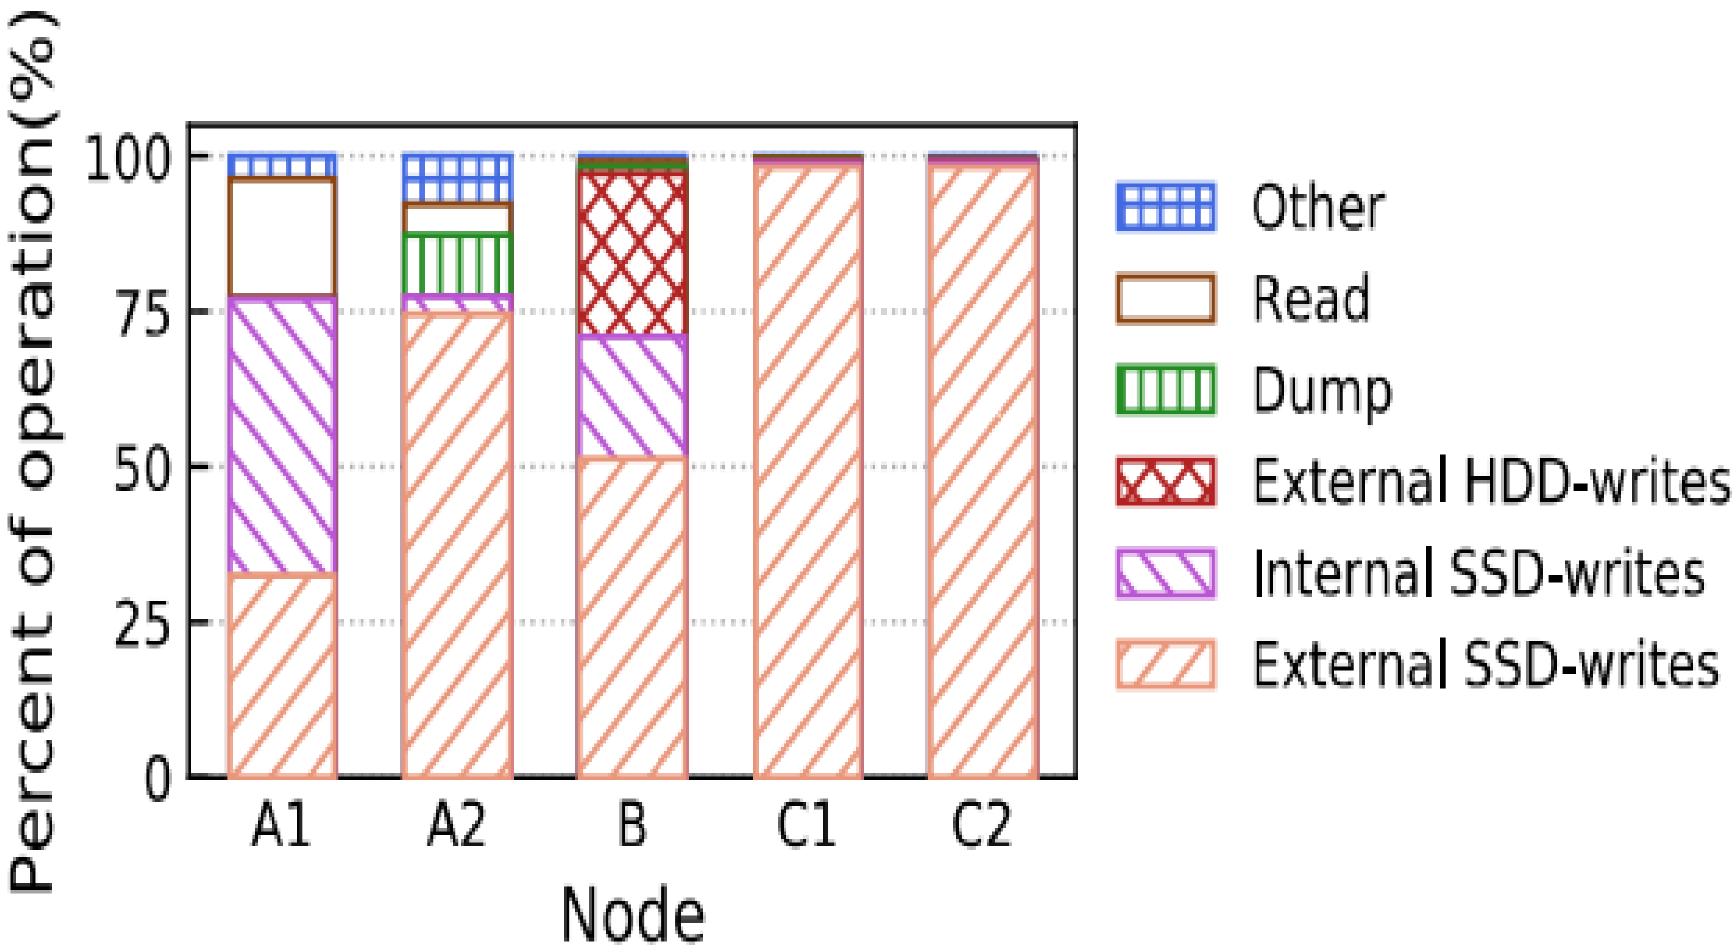

Problem 1: SSD overuse

- The amount of data written to/read from SSD/HDD in 24 hours.

	$A1$	$A2$	B	$C1$	$C2$
SSD-writes(GB)	138	575	3071	820	1027
HDD-writes(GB)	0.1	3.7	555	384	355
SSD-reads(GB)	3.1	1.3	0.2	201	2.4
HDD-reads(GB)	0.3	12.3	0.2	33.6	7.2

- Calculating an SSD's lifespan in B node
  - ☐ 500GB, 300TBW(Terabyte written), 3TB (DWPD)
  - ☐ Lifespan=300TB/3TB/30=3.3month
- SSDs wear out quickly in the write-dominated behavior
- Limit DWPD but increase the number of SSDs

Problem 2: Long Tail Latency

- Long tail latencies appear in different business zones and write operations

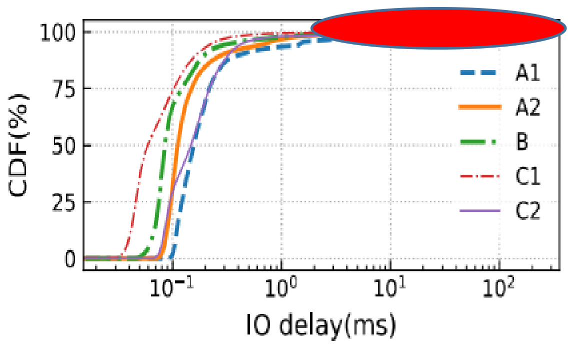

(a) CDF of IO delays of external SSD-writes

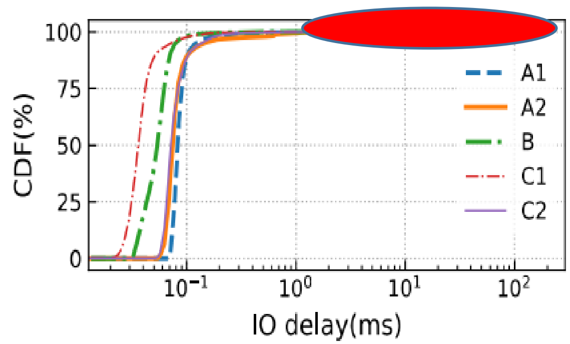

(b) CDF of IO delays of internal SSD-writes

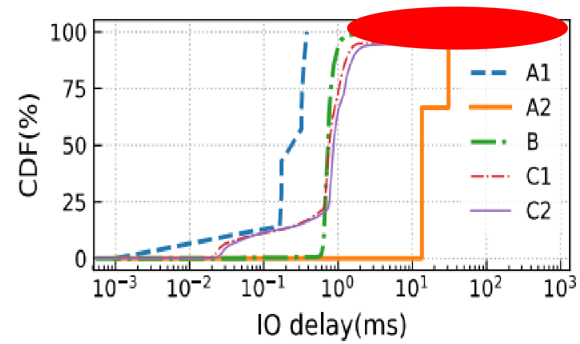

(c) CDF of IO delays of external HDD-writes

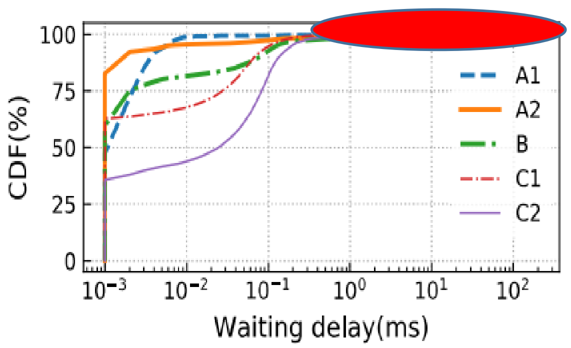

(d) CDF of waiting delays of external SSD-writes

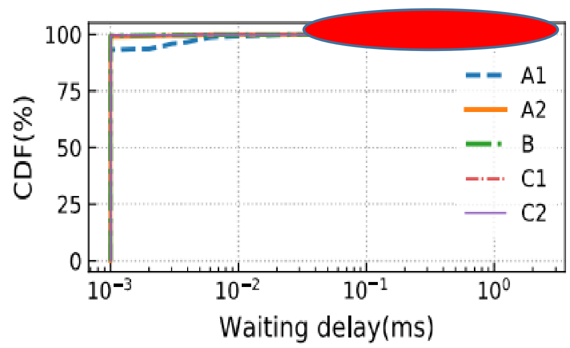

(e) CDF of waiting delays of internal SSD-writes

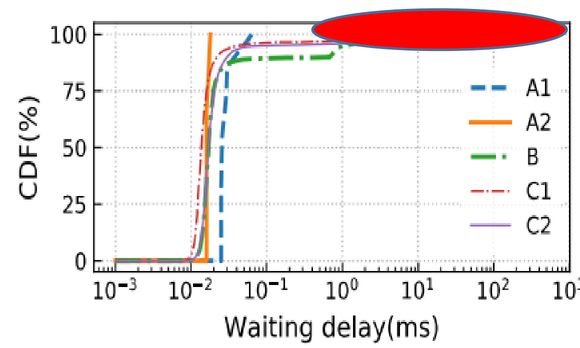

(f) CDF of waiting delays of external HDD-writes

Average/Peak Latency

- External SSD-write: Peak latency is 100-300x larger than average latency.
- Internal SSD-write: Peak latency is 90-2000x larger than average latency.

Node type	External SSD-write	Internal SSD-write	External HDD-write
***A*1**	0.8/113.0	0.09/39.5	0.3/0.4
***A*2**	0.3/94.1	0.1/9.2	19.2/31.0
B	0.1/31.8	0.05/14.4	1.3/415.7
***C*1**	0.1/25.3	0.04/6.7	4.3/613.3
***C*2**	1.0/302.0	0.09/184.1	6.9/774.8

Why is there a long tail delay?

Queue Blockage

- When SSD queue length reaches 2, 90^{th} waiting time is 1000x larger than that without queuing, and average waiting time is 100x.
- Outstanding requests can cause long waiting time.

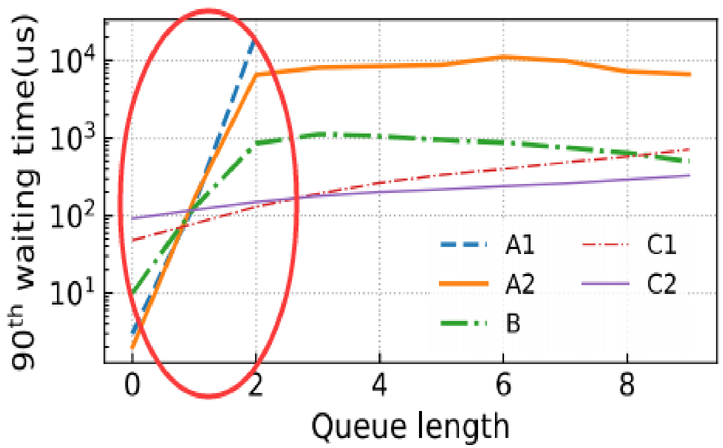

(a) The 90^{th}-percentile waiting delay.

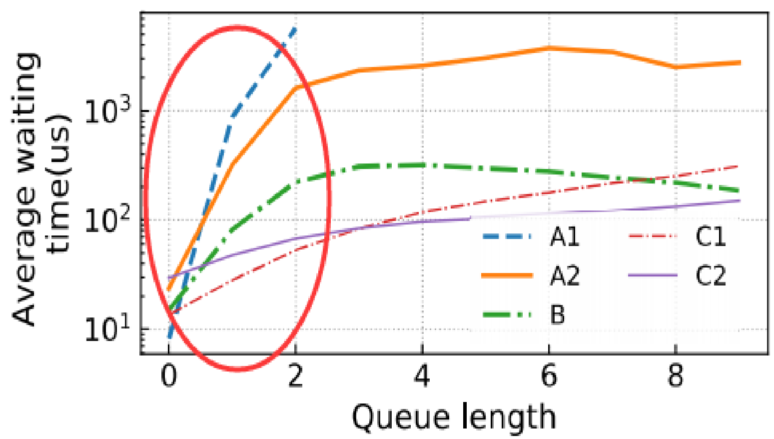

(b) The average waiting delay.

What causes queue blockage?

Blockage Causes

- The reasons behind queue blockage:
  - Large IO
  - Garbage collection

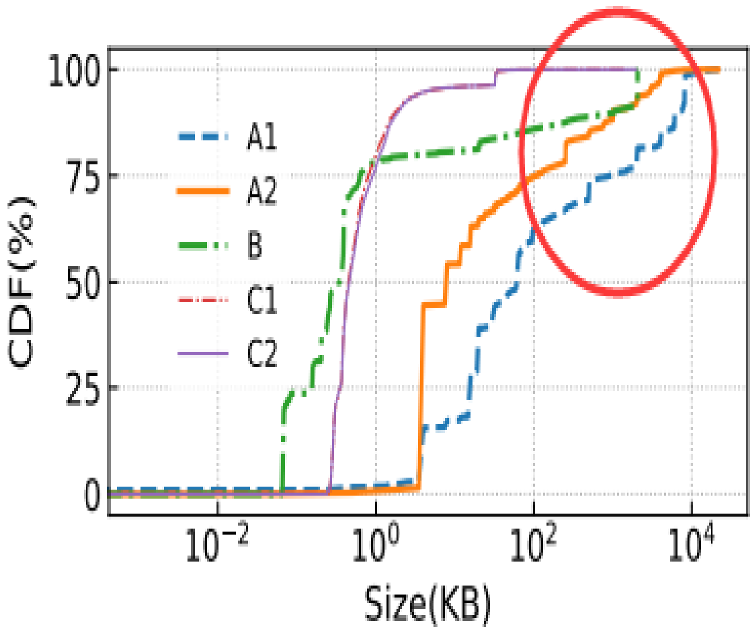

(a) CDF of size of queue-head requests causing queue blockage.

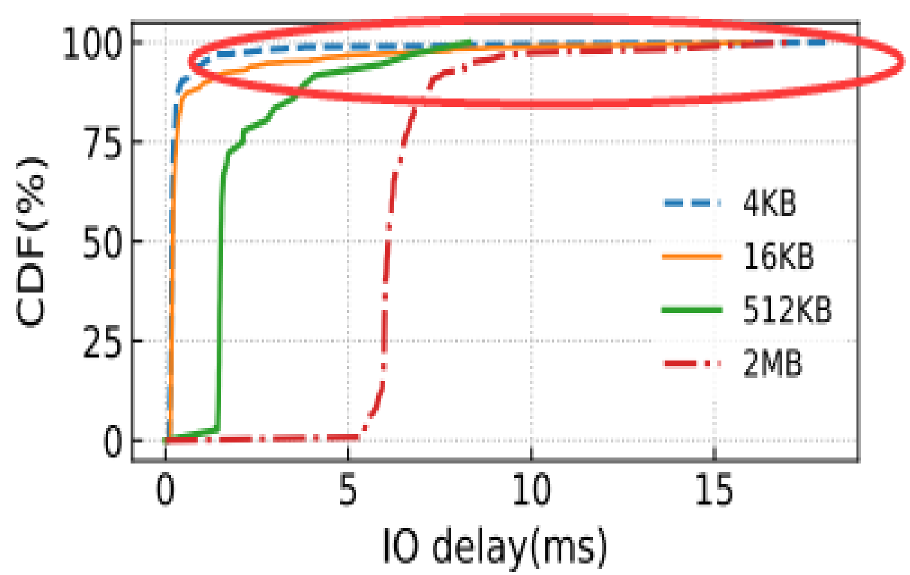

(b) CDF of IO delay of queue-head requests causing queue blockage.

Problem 3: Low Utilization of HDD

	A1	*A2*	*B*	*C1*	*C2*
SSD-writes(GB)	138	575	3071	820	1027
HDD-writes(GB)	0.1	3.7	555	384	355
SSD-reads(GB)	3.1	1.3	0.2	201	2.4
HDD-reads(GB)	0.3	12.3	0.2	33.6	7.2

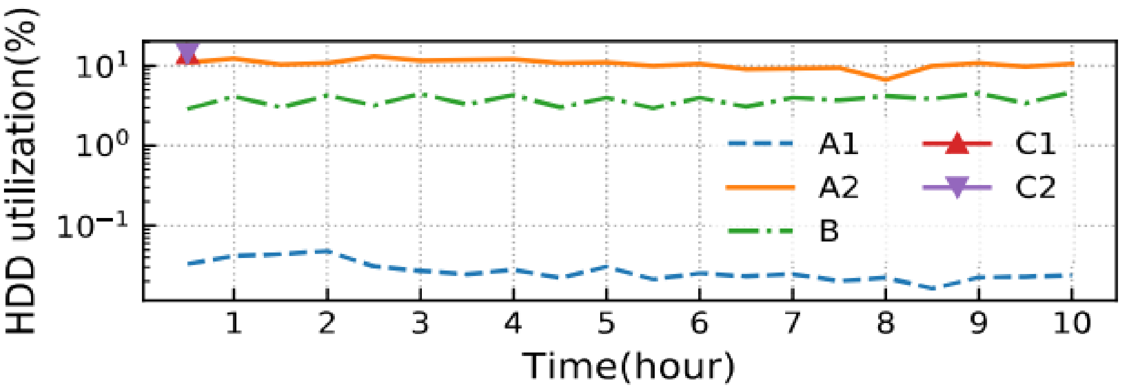

- In *A*1, the amount of data written by SSD-write is 1380x larger than HDD-write.
- The HDD utilization in *A*1 is far less than 0.1% on average, while the maximum is 14.3%.

Outline

- ✓Background
- ✓Trace Analysis
- ✓Design of SWR
- ✓Evaluation
- ✓Conclusion

Architecture Of SWR

- *SSD Write Redirect* (SWR), a runtime IO scheduling mechanism for WSNs.
- Relieve SSD write pressure by leveraging HDDs while ensuring QoS

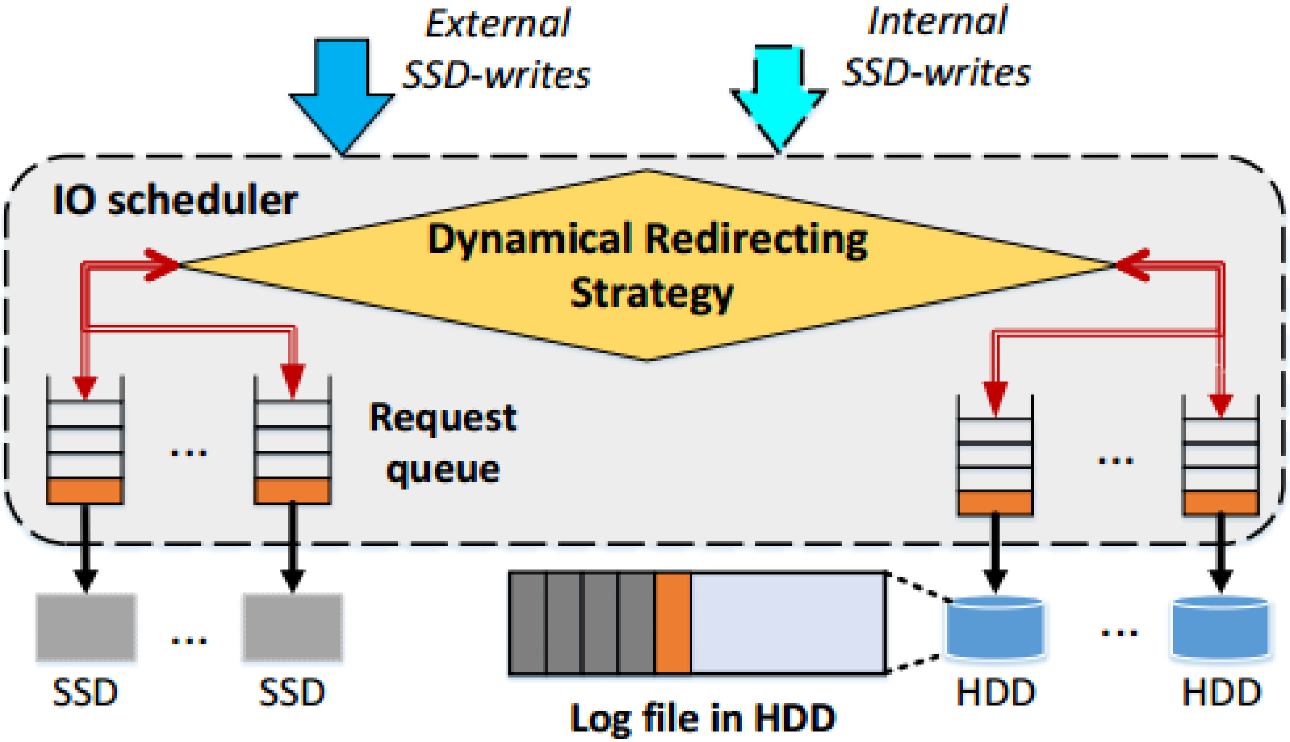

Key Parameters

Idea: redirects large SSD-writes to an idle HDD

(1) ***S***: When a request's size exceeds **S**, it will be redirected.

(2) ***Smax***: Initial value of S.

(3) ***L***: When SSD queue length exceeds L, S will be decreased.

(4) ***p***: SWR gradually decreases the size threshold S with a fixed step value p.

Redirecting Strategy

```
Set S = Smax
for request i in the write queue:
        if OPi == HDD-write:
                put i in HDD queue
        else
                if LSSD(t) > L:
                        S = S – p*Smax
                if LHDD(t) == 0 and Sizei > S:
                        put i in SSD queue
                else
                        put i in HDD queue
```

Logging HDD-Writes

- Using DIRECT_IO to accelerate the data persistence process.

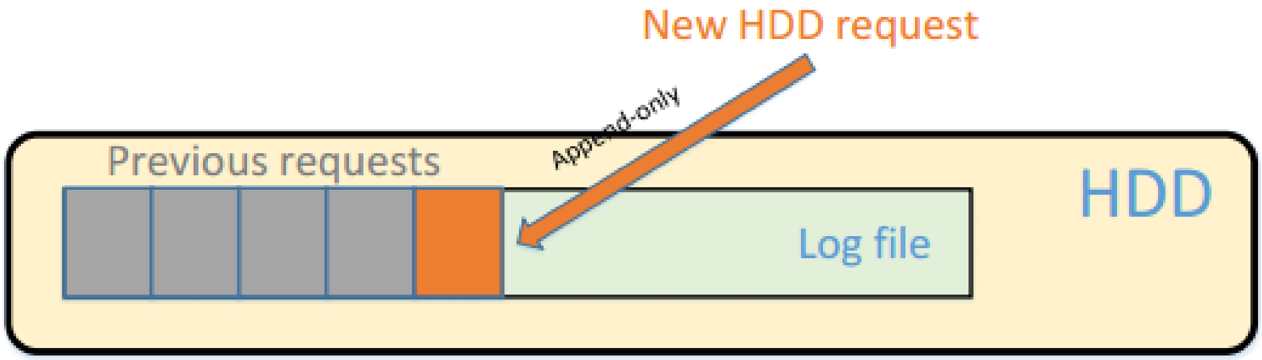

Outline

- ✓Background
- ✓Trace Analysis
- ✓Design of SWR
- ✓Evaluation
- ✓Conclusion

Experiment Setup

- Two types of SSDs:
  - A1, A2: a 256GB Intel 600p SATA with 0.6 GB/s peak writes
  - B, C1, C2: a 256GB Samsung 960 EVO NVMe-SSD with 1.1GB/s peak writes

- HDD: 4TB Seagate ST4000DM005 HDD with 180 MB/s peak write

Trace Replaying on the Test Platform

- Trace: 1 SSD and 1 HDD; 1 hour.
- Average write latency per minute

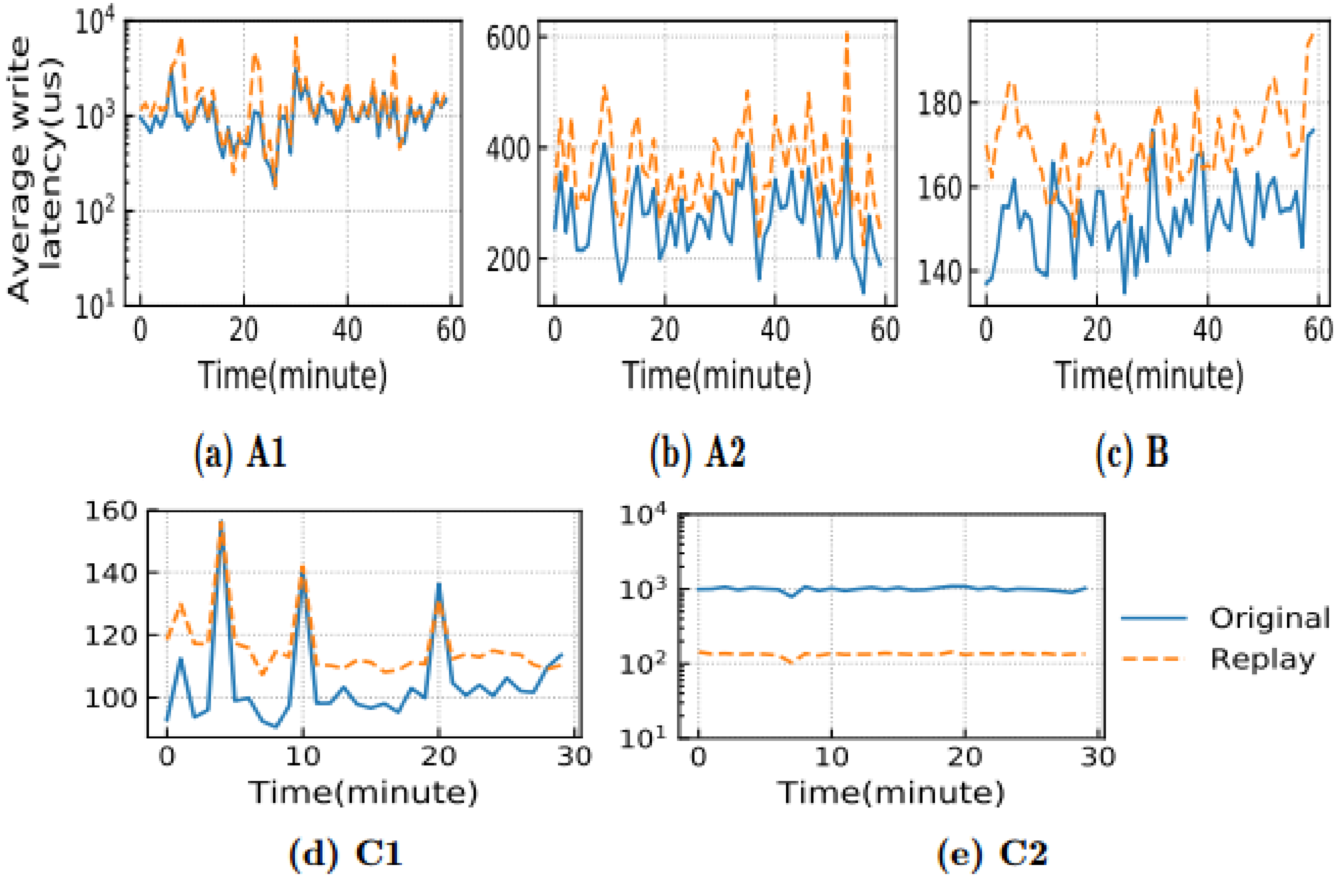

Parameters Selection

- Smax: 99th-percentile block size of SSD-writes
  - The redirected writes should be tiny in number but large in request size.
  - Large IO requests blocking the queue typically account for only 1.1% of all requests.
- L: 6 for *A*1, 5 for *A*2, 30 for *B*, 40 for *C*1 and 57 for *C*2
- p: proportion to S , p = {0, 1/8, 1/4, 1/2,1}

SSD-write Reduction

- SWR effectively reduces the amount data written to SSD, by 70% in *B* and about 45% in the other four nodes.

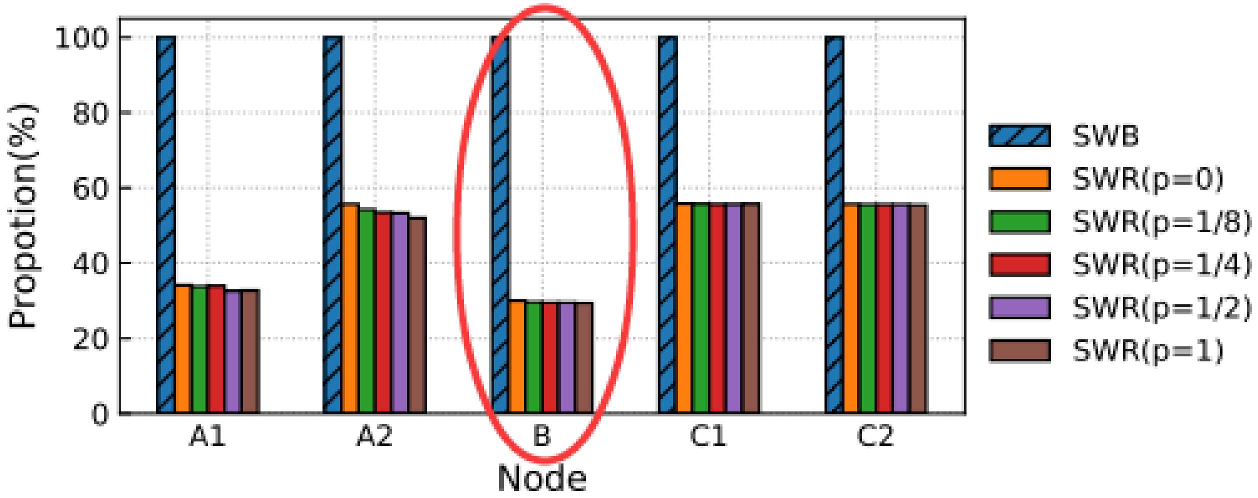

- *p* has no effect on the write reduction.
  - Only effective for the rare burst cases triggering the adjustment of S.

SSD-write Reduction

- By redirecting less than 2% write requests from SSDs to HDDs, SWR is able to reduce 44%-70% of the data written to SSD

	A1	A2	B	C1	C2
SSD data written with SWB(GB)	7.3	24.1	126	11.1	14
SSD data written with SWR(GB)	2.5	13.4	37.8	6.2	7.7
Redirected requests(%)	1.8	1.0	0.7	2.0	2.0

SWR may indirectly increases the SSD lifetime by up to 70%.

Average Write Latency

- SWR reduces average latency by:
  - External SSD-Writes: -10%(B) ~ +13%(A2)
  - Internal SSD-Writes: +52%(A1), +11%(A2), +19%(B) 
  - External HDD-Writes: -95%~-70%(B) 

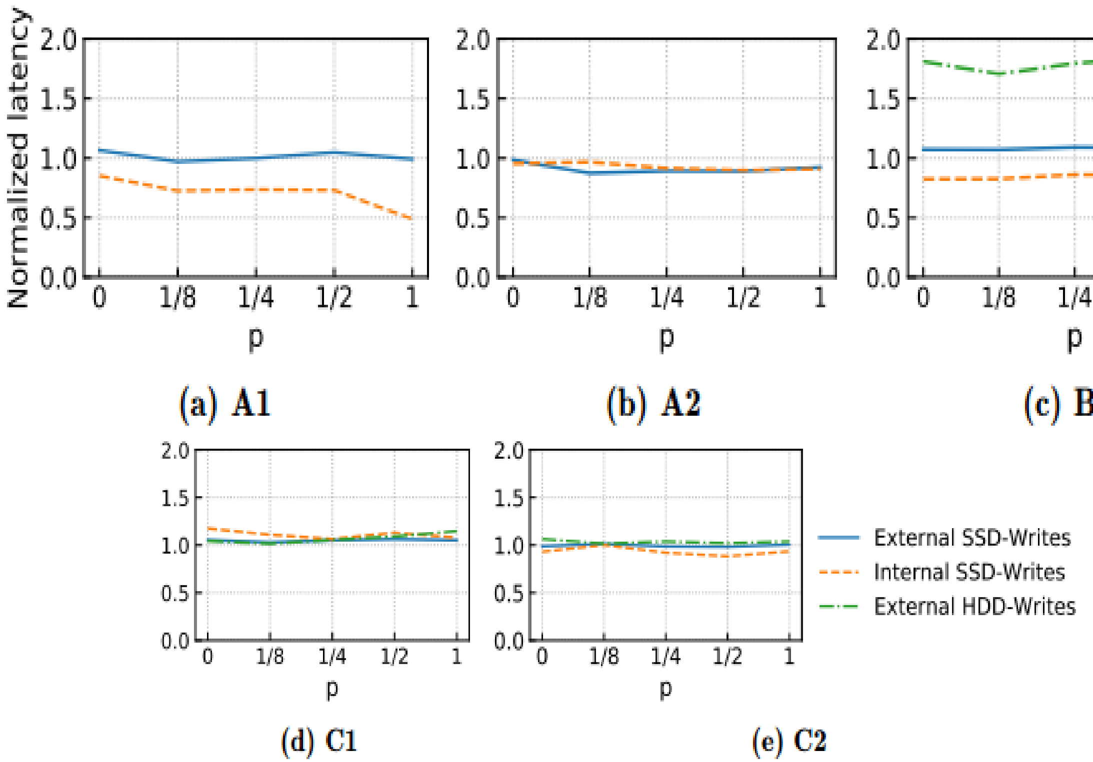

99th Write Latency

- SWR reduces 99th latency by:
  - External SSD-Writes: + 12%(C1)~ +47%(A2)
  - Internal SSD-Writes: + 13%(C2) ~ +79%(A1,B)
  - External HDD-Writes: -169%~-130%(B),-50%~-9%(C1,C2)

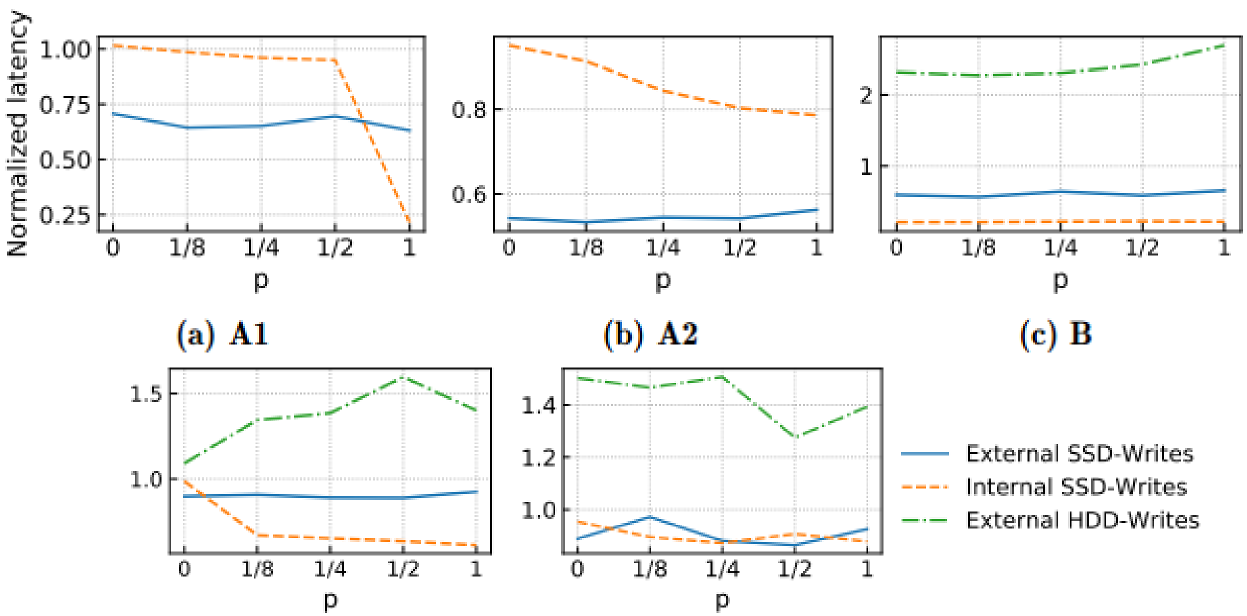

(a) A1 (b) A2 (c) B

(d) C1 (e) C2

HDD Competition

- Reason for an increase in External HDD-Writes average 99th latency:
  - HDD competition between external HDD-writes and redirected SSD-writes
- Can be alleviated by forwarding HDD-writes to the remaining tens of HDDs.
- The avg. and 99th write latency of External HDD-Writes of SWR scheduling upon two HDDs in node *B*.

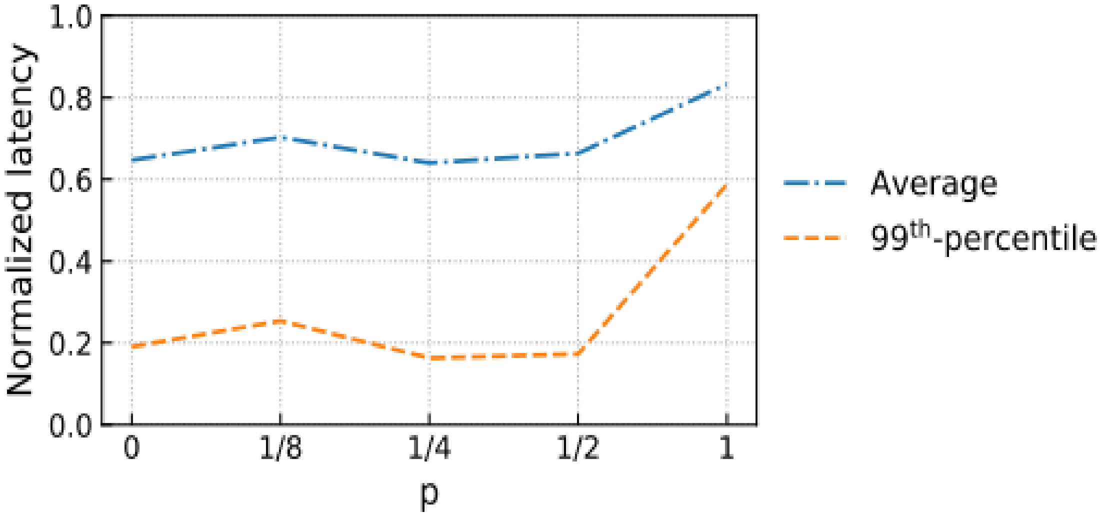

Latencies of Redirected Writes

- In the worst case, the average latency of 0.7% writes in *B* can increase from 0.94 ms with SWB to 7.29 ms with SWR(lower than SLA(50ms at the average))

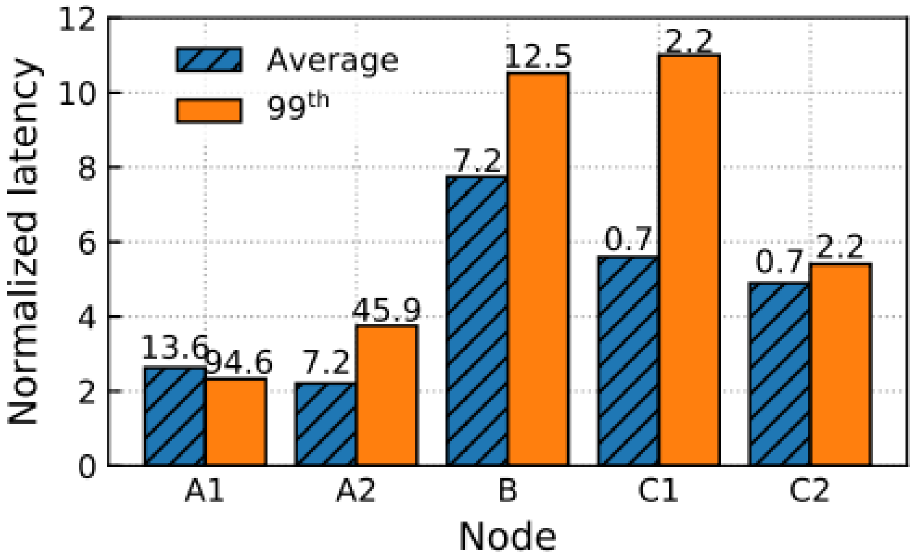

SWR reduces of both data written to SSDs and tail-latency at the expense of a tiny percentage of writes(up to 2%).

Outline

- ✓ Background
- ✓ Trace Analysis
- ✓ Design of SWR
- ✓ Evaluation
- ✓ Conclusion

Conclusion

- Some hybrid storage nodes in Pangu have write-dominated workload behaviors.
- Current request serve mode in such nodes leads to SSD overuse, long-tail latency, and HDD low-utilization.
- Redirecting large SSD write requests to HDDs and dynamically optimize for small and intensive burst requests.

Thank you !

Questions ?