

Content-aware Trace Collection and I/O Deduplication for Smartphones

Bo Mao¹, Suzhen Wu¹, Hong Jiang², Xiao Chen¹, Weijian Yang¹

{maobo, suzhen}@xmu.edu.cn, hong.jiang@uta.edu

¹ Xiamen University (<http://astl.xmu.edu.cn/>)

² University of Texas at Arlington

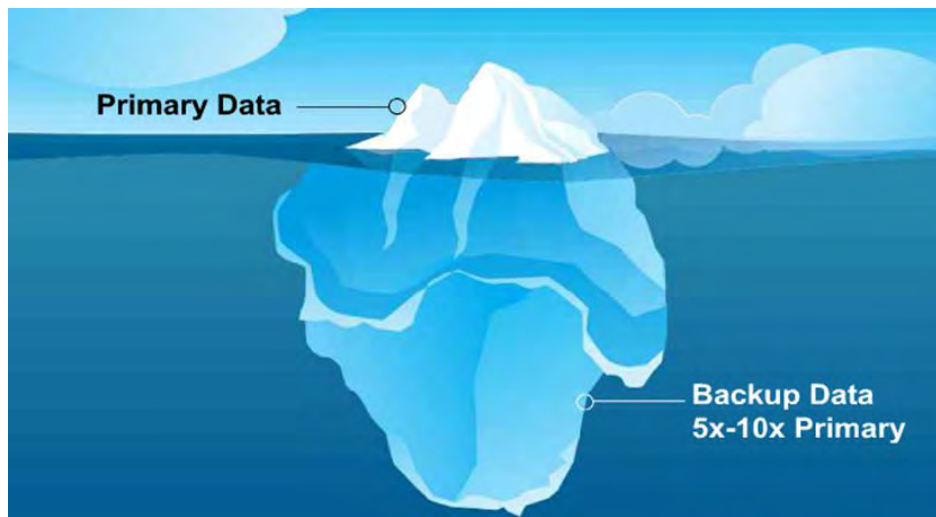


Outline


- Introduction and challenges
 - Trace collection and observations
 - System overview and design
 - Performance evaluations
 - Conclusion
-

Why data deduplication?

- Backup media changed from Tape to HDDs



Capacity

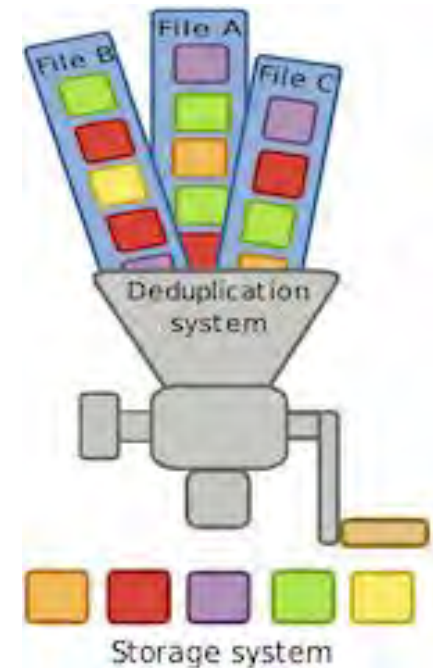


Cost


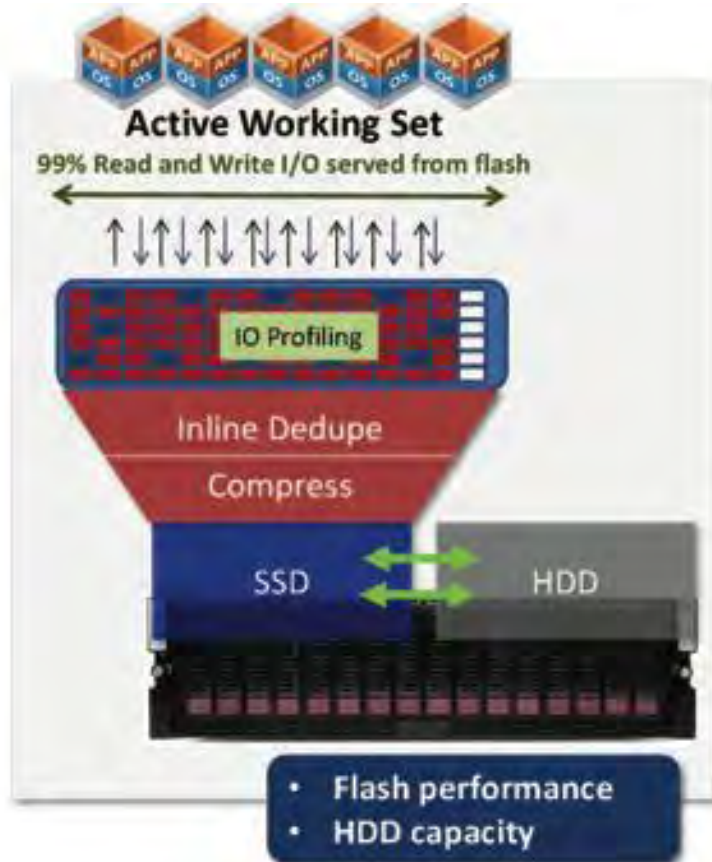
- In primary storage systems, it is also important to shrink the data volume.

Data deduplication

- Data deduplication is widely deployed in secondary storage systems to:
 - ✓ Reduce backup time
 - ✓ Improve storage space efficiency
 - ✓ Improve network bandwidth
 - ✓
- In primary storage systems:
 - ✓ VM-based storage systems (Linux KSM)
 - ✓ Flash storage products (Nimble storage, Tintri, Pure Storage ...)
 - ✓




Deduplication + Flash



- CA-FTL (FAST'11)
- CA-SSD (FAST'11)

Deduplication has become an important feature for flash-based storage!

Flash within Smartphones



Make Android Speed Faster



- Flash (eMMC or UFS) in Smartphones:
 - ✓ Performance tends to degrade after repeated usages.
 - ✓ Limited life cycles affect the Smartphones' reliability.
 - ✓ The cost of upgrading flash capacity is high.

How about applying data deduplication on flash storage within Smartphones?

Workflow of data deduplication



Fixed chunk, CDC, FastCDC...

Write data →

Data Chunking
(Method, Size
et. al)



StoreGPU, Shredder ...

Hash Computing
(GPU, Multicore
et. al)



DDFS, ChunkStash, SiLo ...

Hash Index
(Disk-Index,
Structure et. al)



Index and Metadata
Update (Update,
Distribution et. al)

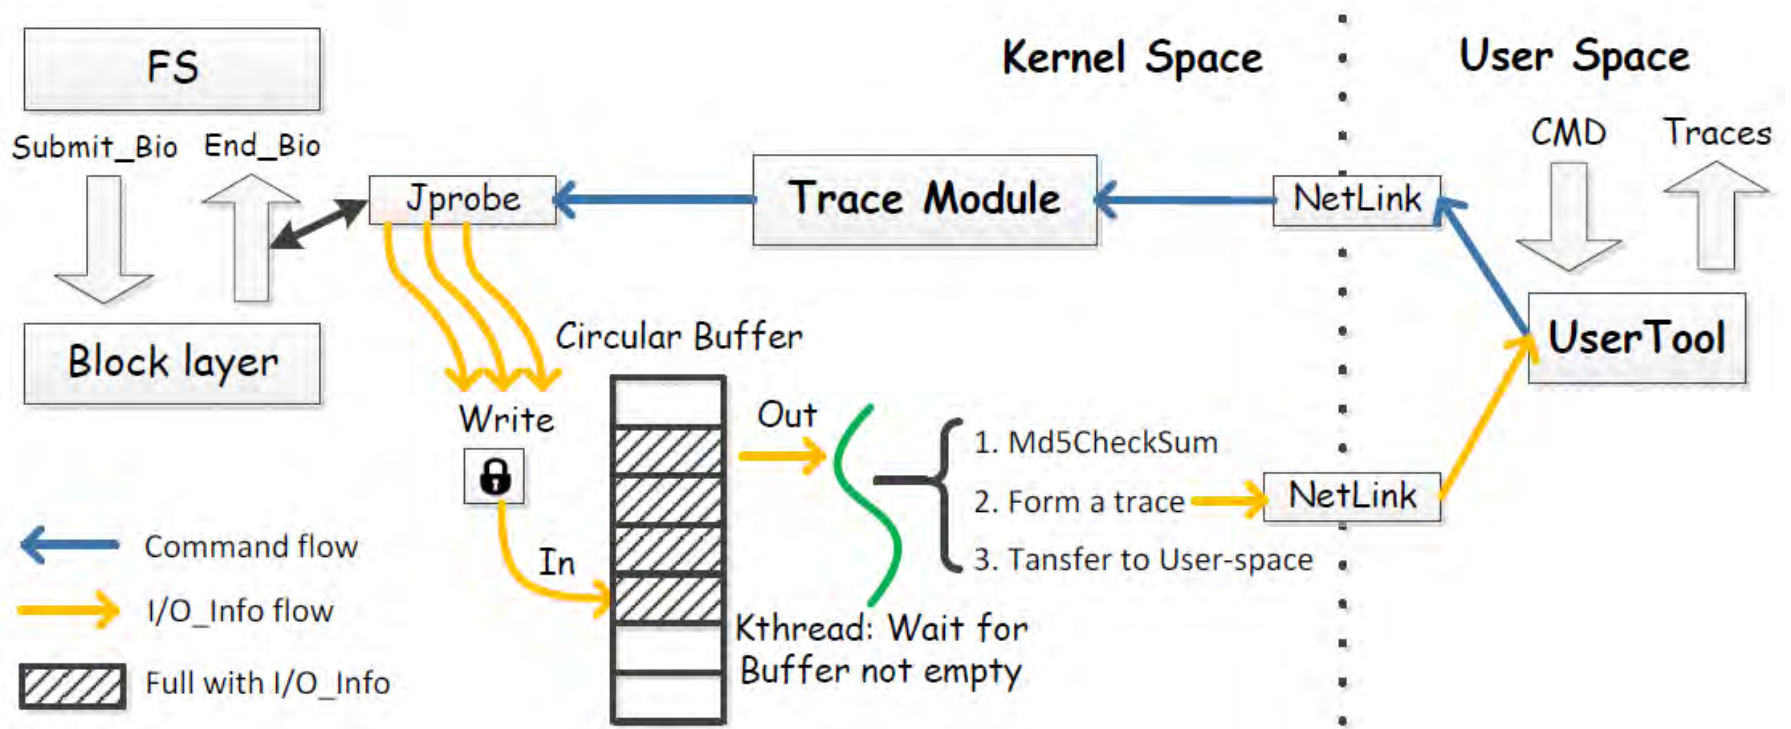
Devices

CPU and Memory Overhead!

Challenges

- Resources in Smartphones:
 - ✓ CPU utilizations affect power consumption.
 - ✓ Limited memory capacity.
 - ✓ Mobile APP usages.
 - Is data deduplication feasible and how to?
 - ✓ How to investigate the redundancy within Smartphones?
 - ✓ How much data redundancy in mobile APPs?
 - ✓ How to design a lightweight data deduplication engine?
-

Content-aware trace collection



Obs 1: Redundancy characteristics*

<i>APPs</i>	<i>Qiu</i>	<i>58City</i>	<i>Baidu</i>	<i>Game2048</i>	<i>Meitu</i>	<i>Moji</i>
<i>Qiu</i>	40.5%	0.2%	1.0%	0.1%	0.3%	0.0%
<i>58City</i>	0.3%	32.2%	3.3%	1.3%	3.6%	3.2%
<i>Baidu</i>	0.4%	6.5%	39.5%	0.2%	2.4%	5.5%
<i>Game2048</i>	0.3%	0.7%	1.0%	31.7%	1.0%	2.0%
<i>Meitu</i>	0.3%	3.1%	2.2%	1.3%	31.3%	2.2%
<i>Moji</i>	0.1%	2.7%	4.7%	2.8%	2.3%	32.4%

- ✓ Moderate to high data redundancy exists within mobile APPs.
- ✓ Amount of data redundancy shared between any two different mobile APPs is minimal#.

*Detailed results and analysis for all the 15 mobile APPs can be found in our paper.

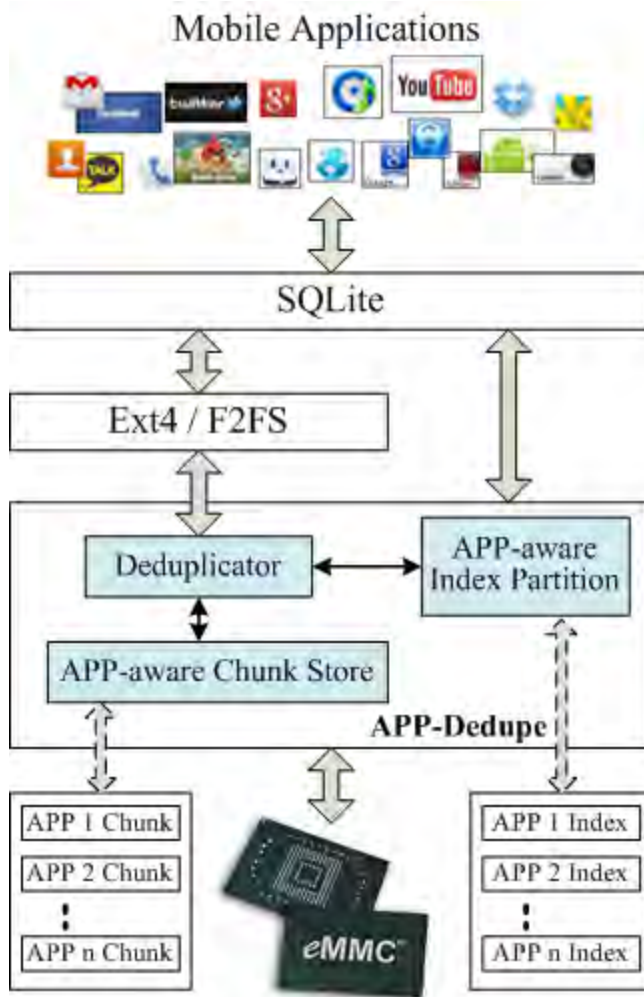
#Y. Fu, H. Jiang, N. Xiao, L. Tian, F. Liu, AA-Dedupe: An Application-Aware Source Deduplication Approach for Cloud Backup Services in the Personal Computing Environment, in Proceedings of IEEE Cluster 2011, Austin, Texas, Sept. 2011.

Obs 2: Lower IOPS

<i>APPs</i>	<i>Redundancy</i>	<i>IOPS</i>	<i>Size (MB)</i>	<i>Write Ratio</i>
<i>Qiu</i>	40.5%	1.6	55.0	87.9%
<i>58City</i>	32.2%	2.6	153.3	69.4%
<i>Baidu Tieba</i>	39.5%	4.3	213.5	88.6%
<i>Game2048</i>	31.7%	1.7	63.2	89.2%
<i>Meitu</i>	31.3%	5.6	131.6	41.9%
<i>Moji Weather</i>	32.4%	2.6	82.1	63.4%
<i>Opera Browser</i>	33.1%	3.0	72.7	57.5%
<i>Fruit Cool</i>	36.6%	6.0	159.9	47.1%
<i>Sohu News</i>	32.2%	2.0	138.2	84.2%
<i>Tencent</i>	34.1%	1.5	82.6	95.4%
<i>Pea Pod</i>	35.4%	2.1	237.0	67.8%
<i>Wechat</i>	30.0%	4.6	345.5	90.4%
<i>Weibo</i>	29.0%	1.4	97.0	88.5%
<i>Xiami Music</i>	27.3%	7.4	122.9	39.4%
<i>Youdao Dict</i>	31.9%	9.2	136.5	35.2%

The I/O intensity is low for most APPs (IOPS)*

System overview

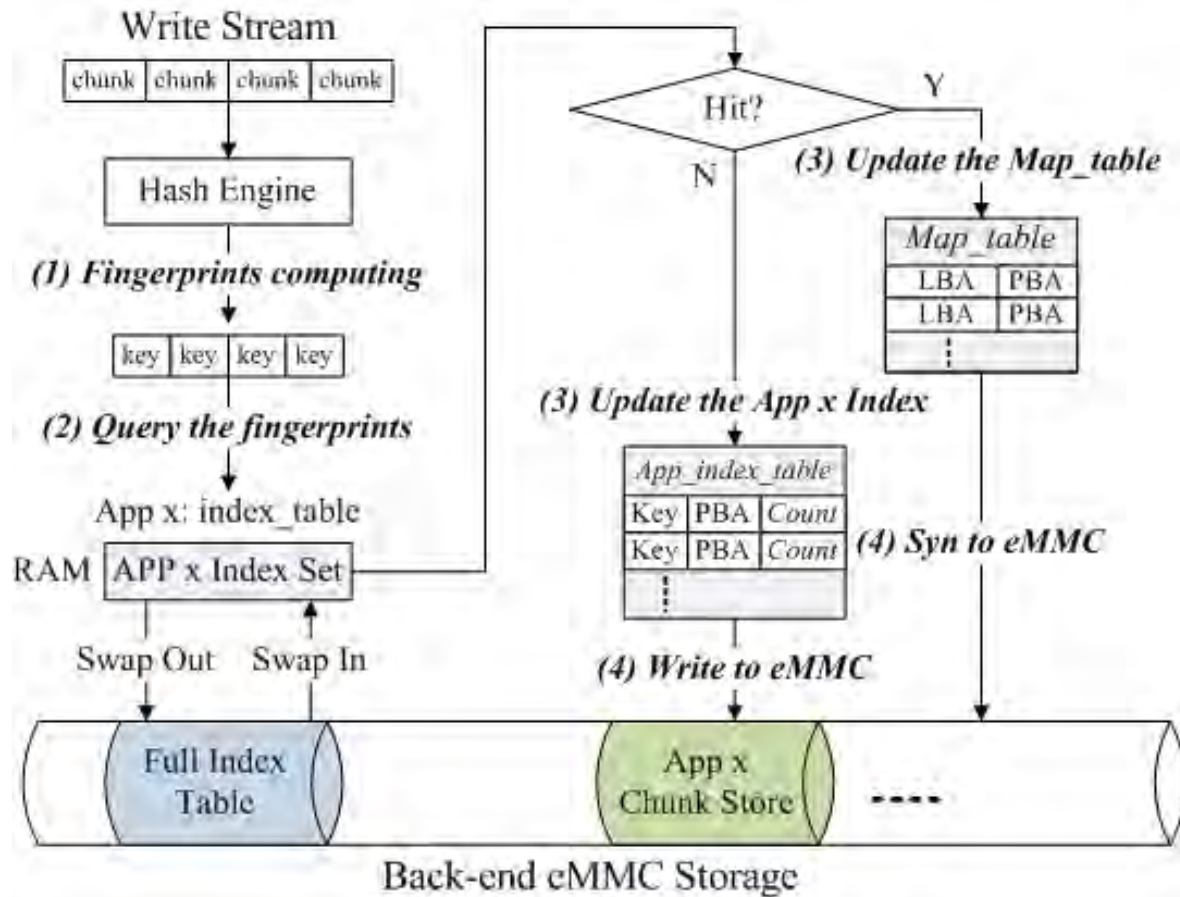


- Independent of upper file systems
- Low overhead design choice:
 - ✓ MD5 hash computing
 - ✓ Fixed chunking (4KB)
- Two optimizations:
 - ✓ Index partition
 - ✓ Chunk store

Design and Optimizations

- APP-aware Index Partition (AIP):
 - ➔ *Memory overhead associated with big hash index table.*
 - ✓ Grouping the hash index according to the APPs.
 - ✓ Swap In/Out between memory and Flash.
 - APP-aware Chunk Store (ACS):
 - ➔ *Data fragmentation associated with data deduplication.*
 - ✓ Storing the data chunks according to the APPs (LBAs).
 - ✓ Concentrating the read accesses to a single container.
-

Write workflow in APP-Dedupe

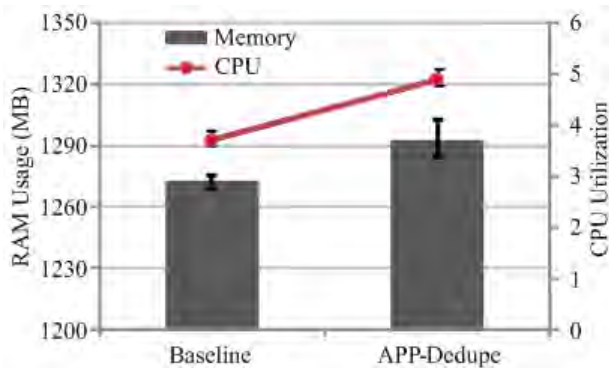


Experimental setup

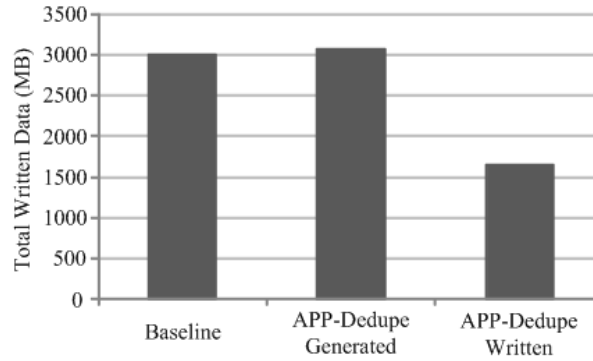
- Google Nexus 5 Smartphone:
 - ➔ *Real system study.*
 - ✓ Qualcomm MSM8974 Quadcore 2.3 GHz, 2 GByte DRAM, 16 GByte eMMC storage.
 - ✓ Android 5.0.1 with Linux Kernel 3.4.
 - ✓ Benchmarks: Monkey tool and A1 SD Bench.
- SSD-based DiskSim simulator:
 - ➔ *Simulation study.*
 - ✓ Replay the traces collected from real system.
 - ✓ Evaluate response time and GC count within flash.

Results and analysis

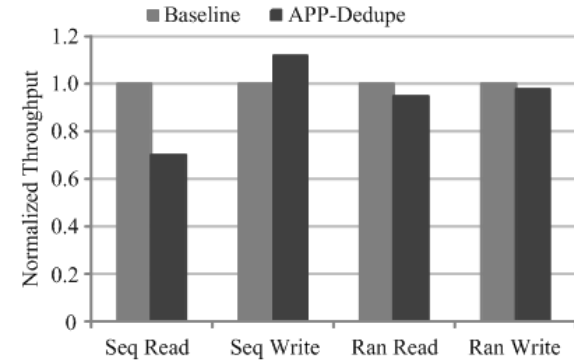
(1) Memory and CPU usages



(2) Total written data

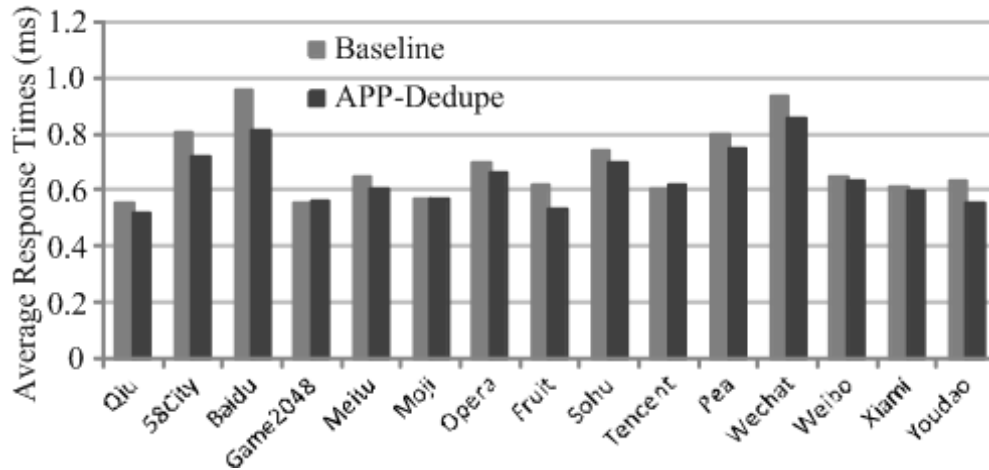


(3) Throughput

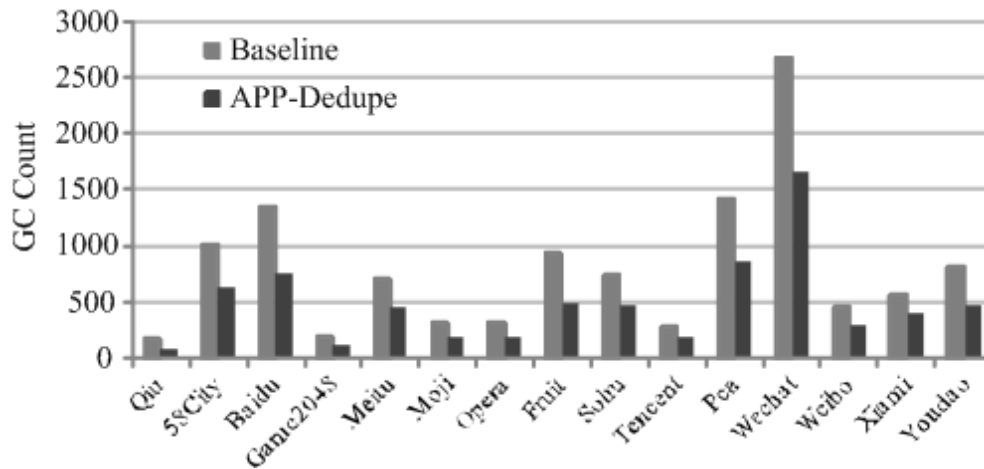


- APP-Dedupe incurs very little memory and CPU overhead, by less than 3%.
- APP-Dedupe reduces the amount of write data to the back-end eMMC storage by an average of 45.2%.
- System throughput performance is complicated.

Results and analysis



By up to 15.4% with an average of 6.2%



By up to 58.7% with an average of 41.5%

Conclusion

- Performance of the storage subsystem in Smartphones plays an important role in the application performance.
 - We investigate the data redundancy characteristics within Smartphones and propose APP-Dedupe that detects and eliminates the I/O redundancy by exploiting the mobile applications' redundancy characteristics.
 - APP-Dedupe reduces the GC overhead by an average of 41.5%, reduces the response times by up to 15.4% and saves the storage capacity by an average of 45.2%.
-

Content-aware Trace Collection and I/O Deduplication for Smartphones

Bo Mao^{1*}, Suzhen Wu¹, Hong Jiang², Xiao Chen¹, Weijian Yang¹

¹ Xiamen University (<http://astl.xmu.edu.cn/>)

² University of Texas at Arlington

* Please feel free to contact me: maobo@xmu.edu.cn for any questions!

Q & A ? Thanks!

