

File Systems Comparison: EXT4, ZFS, NTFS

Aribhit Mishra, Mickey Barboi and Mingi Kim

{aribhitm, barboi, mingi}@cs.wisc.edu

Motivation

- Compare three of the most-widely used modern file systems : EXT4, ZFS and NTFS
- Lack of existing work that compares all three file systems
- Different features and implementations would mean varying performance over different loads

File system Comparison

Three different file systems provides different structures and features.

	EXT4	ZFS	NTFS
Max Filename Length	255 bytes	255 characters	255 characters
Max File Size	16 GiB to 16 TiB	16 EiB	16 EiB
Max Volume Size	1 EiB	256 ZiB	16 EiB
Checksum	✗	✓	✗
Block Journaling	✓	✓	✗
File Changelog	✗	✗	✓
Internal Snapshotting	✗	✓	Partial
Data Deduplication	✗	✓	✓
Allocate-on-flush	✓	✓	✗
Variabe Block Size	✗	✓	✗

Design

1. Macrobenchmark

- Serving images through Apache Server
 - Serve image from static directory to n clients
 - Clients issue sequential requests with no intercommunication amongst clients
 - Images removed from drive, and recopied between tests
- Serving HTTP requests with Go and PostgreSQL
 - Client issues HTTP requests trying with a JSON payload
 - Server creates a new record in database for each request
- Compiling Go
 - Download Go source code and compile it

2. Microbenchmark

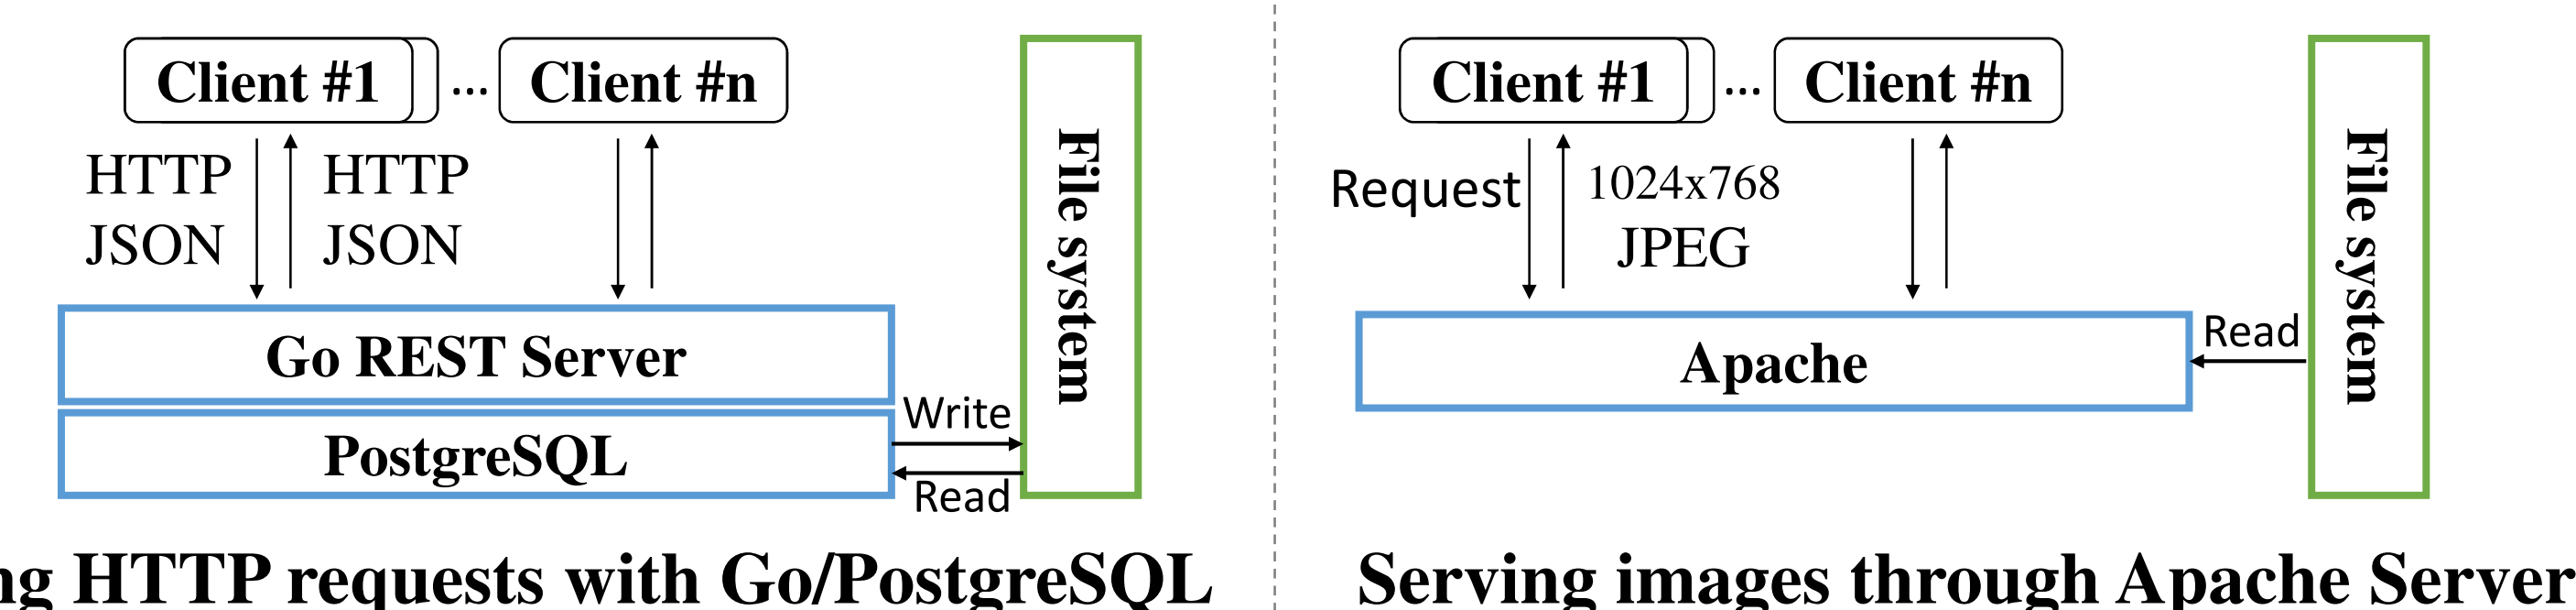
- Single block size read/write operation test
- Large file I/O test
- Small files bulk I/O test

Test Environment

- Intel i7 6700k (8 threads at 4.2GHz, 8MB cache)
- 32GB DDR4 RAM
- System Drive : Samsung SM951 NVMe with 512GB of capacity
- Testing Drive : Western Digital 500GB HDD (7200 max RPM)
- Operating Systems:
 - NTFS run on Windows 10
 - EXT4 and ZFS run on Ubuntu 16.0.4

Benchmark Program Design at a Glance

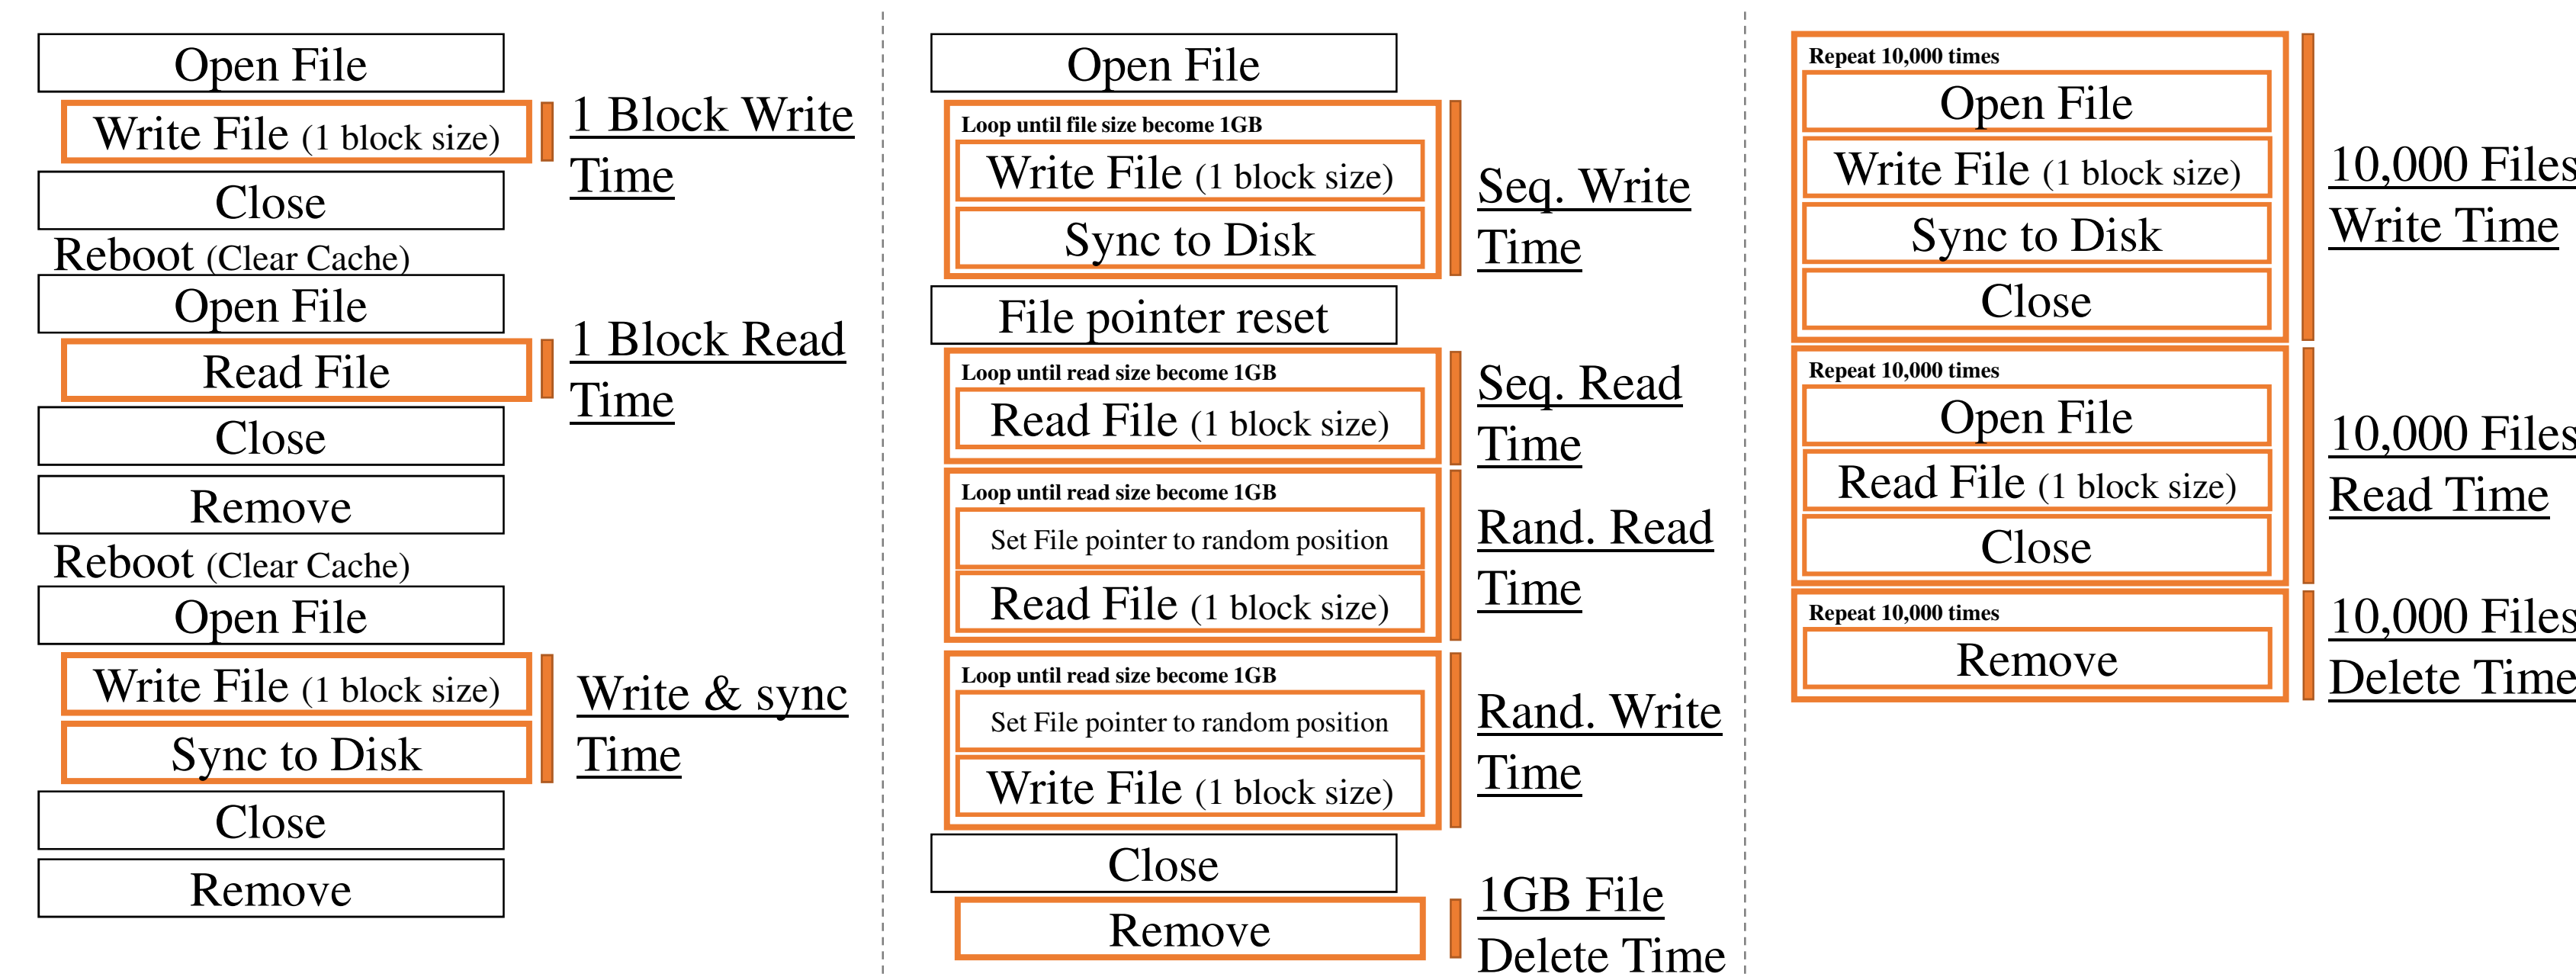
1. Macrobenchmark



Serving HTTP requests with Go/PostgreSQL

Serving images through Apache Server

2. Microbenchmark



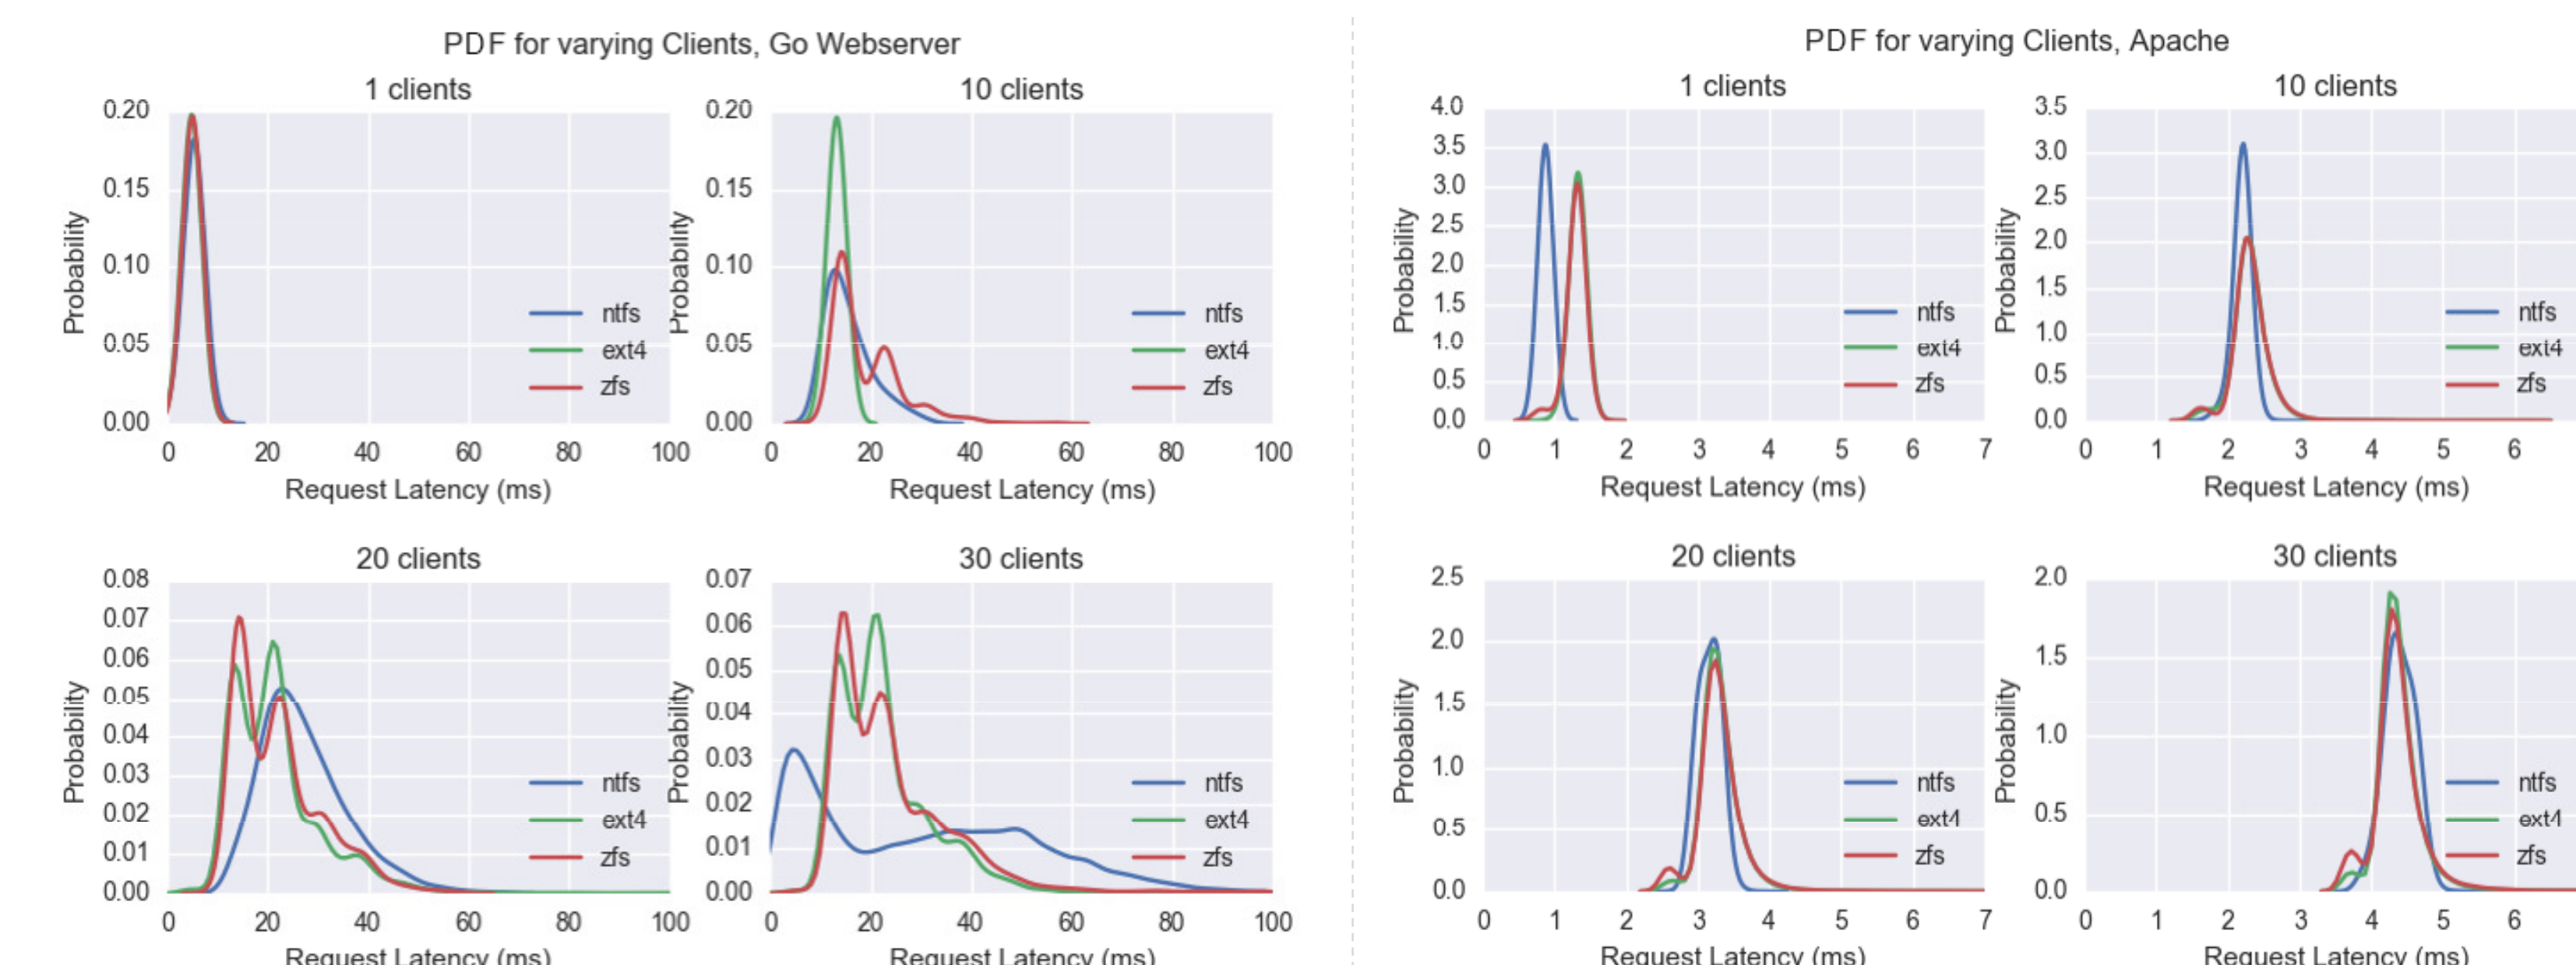
Single block size read/write operation test

Large file I/O test

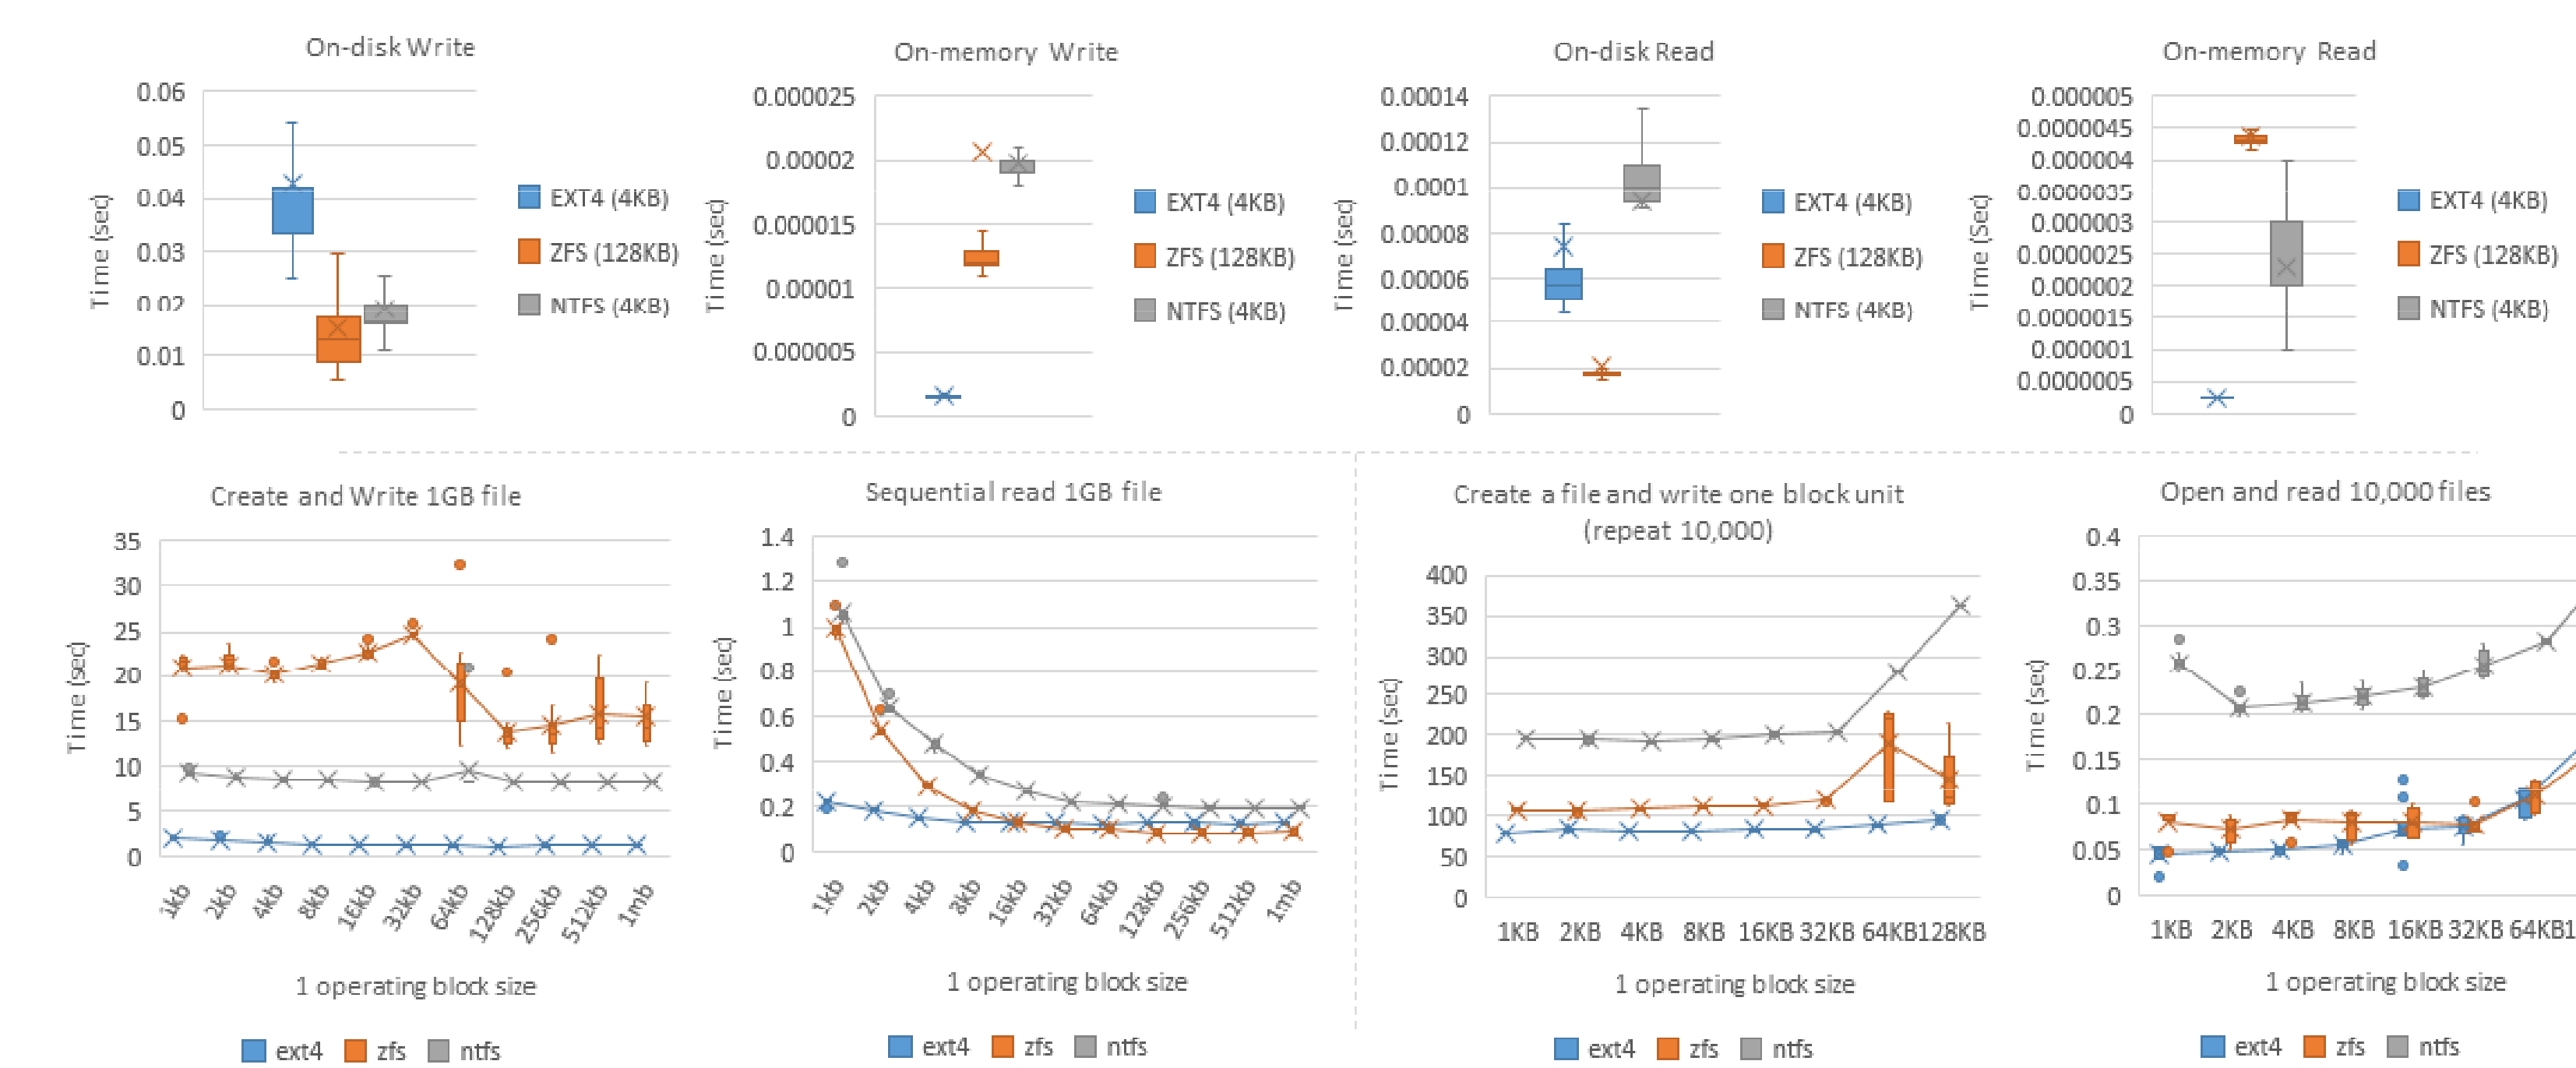
Small files bulk I/O test

Results

1. Macrobenchmark Results



2. Microbenchmark Results



Implementation

1. Macrobenchmark

- General Design
 - n clients: [1, 10, 20, 30]
 - Clients issue next request as soon as current returns
- Go REST HTTP Server
 - HTTP POST with 52B JSON payload
 - Clients do not share data
- Server
 - Commits payload to database
 - Reads it and returns to client
- Apache Webserver
 - Requesting 1024x768 JPEG, 247KB
 - Clients request unique images

2. Microbenchmark

- File I/O System calls
 - Used different system call libraries to see the performance in native operating system environments of each file system

	Open/Create	Read	Write	Move File Pointer
Linux	<i>open()</i>	<i>read()</i>	<i>write()</i>	<i>lseek()</i>
Windows	<i>CreateFile()</i>	<i>ReadFile()</i>	<i>WriteFile()</i>	<i>SetFilePointer()</i>
	Sync to Disk	Close	Remove	
Linux	<i>fsync()</i>	<i>close()</i>	<i>unlink()</i>	
Windows	<i>FlushFileBuffers()</i>	<i>CloseHandle()</i>	<i>DeleteFile()</i>	

2. Time measurement

	Get Time Function	Time Precision
Linux	<i>clock_gettime()</i> with <i>CLOCK_REALTIME</i>	25 nanosecond
Windows	<i>GetSystemTimePreciseAsFileTime()</i>	100 nanosecond

Conclusion

- EXT4 slightly outperforms ZFS with NTFS lagging behind
- EXT4 shows consistent performance across all metrics, while ZFS and NTFS have wider dispersion
- ZFS outperforms when dealing with large files with at least 32KB access size in microbenchmark
- Differences in Operating Systems and Application could have affected results
 - Apache, PostgreSQL and Go are native to Linux
 - Native home of ZFS is Solaris, not Linux

Source Code

- Macrobenchmark
 - <https://github.com/damouse/fs-bench-suite>
- Microbenchmark
 - Single RW
 - Linux : https://github.com/MingiK1m/cs736_rwbm
 - Windows : https://github.com/MingiK1m/cs736_rwbm_windows
 - Large/Small file(s) benchmark
 - Linux : https://github.com/MingiK1m/cs736_fsbm
 - Windows : https://github.com/MingiK1m/cs736_fsbm_windows