A Statistical Evaluation of the Impact of Parameter Selection on Storage System Benchmarks

Nohhyun Park
Weijun Xiao
Kyubaek Choi
David J. Lilja

University of Minnesota

Outline

- Problem with Current Storage Performance Evaluation Techniques.
- Design of Experiments and Proposed Solution.
- Example Case using existing benchmarks.
- Effect of Benchmark Parameters.
- Benchmark Coverage Testing.
- Conclusion and Future work.

Outline

- Problem with Current Storage Performance Evaluation Techniques.
- Design of Experiments and Proposed Solution.
- Example Case using existing benchmarks.
- Effect of Benchmark Parameters.
- Benchmark Coverage Testing.
- Conclusion and Future work.

Motivation

- How do we compare the performance of two different storage systems?

1-3 Benchmarks

1-5 Configurations

P1

P2

It only tells us that system1 is better for the synthetic workloads tested.

- If the P1 > P2, can we say that system1 performs better than system2?

Problem

- Synthetic workload does not reflect all if any real workloads.
  - Burstiness
  - Locality

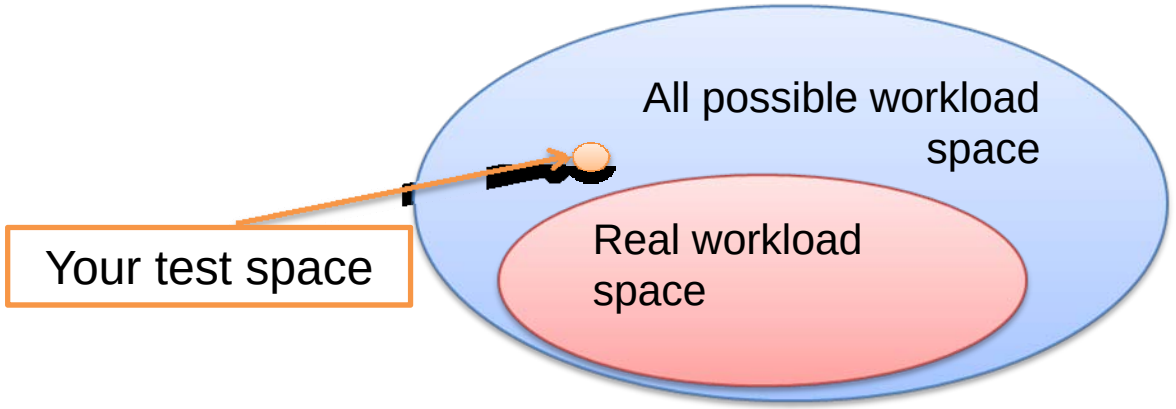

- The fact that system1 performs better than system2 for the few synthetic workloads may have no meaning at all!

More Problems

- We can use the real traces for performance testing but…
- There exists no representative workload.
  - IO workload is more diverse than the applications.
    - Stateful response of storage systems.
    - Shared storage.
- Using real trace limits the applicability of the performance numbers(Coverage)!

Ideal Case

- Generate real workloads using set of parameters.
  - How can we describe workload in a generic manner?
  - Future work.

- Can we generate a benchmark input sequence that will –
  - Increase chance of result being true in a statistical manner?
  - Provide reasonable test time and coverage?

Goal

- To design experiment that
  - Generate input vector for the storage benchmarks that will statistically guarantee correctness of the experiment.

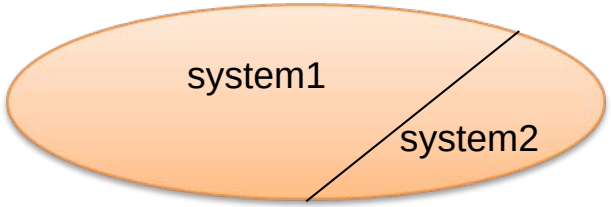

  - Experiment must be able to be carried out in a reasonable manner and provide a sufficient coverage.

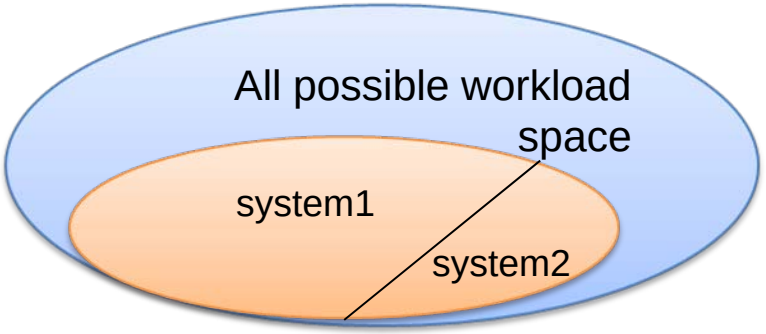

Outline

- Problem with Current Storage Performance Evaluation Techniques.
- Design of Experiments and Proposed Solution.
- Example Case using existing benchmarks.
- Effect of Benchmark Parameters.
- Benchmark Coverage Testing.
- Conclusion and Future work.

Design of Experiment

- Most accurate method would be to test the workload(parameter) space exhaustively.

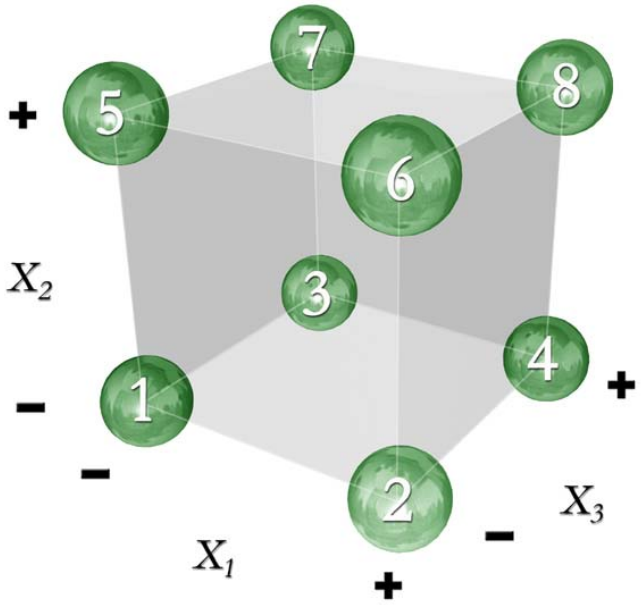

	FIO	Postmark	IOZone
Number of Parameters	20	10	11
Experiment Required	1,048,576	1,024	2,048
Time (10min/exp)	19 years	7 days	14 days

Plackett and Burman Design

- Exploring subset of parameter space such that performance variation can be estimated using only O(N) experiments.

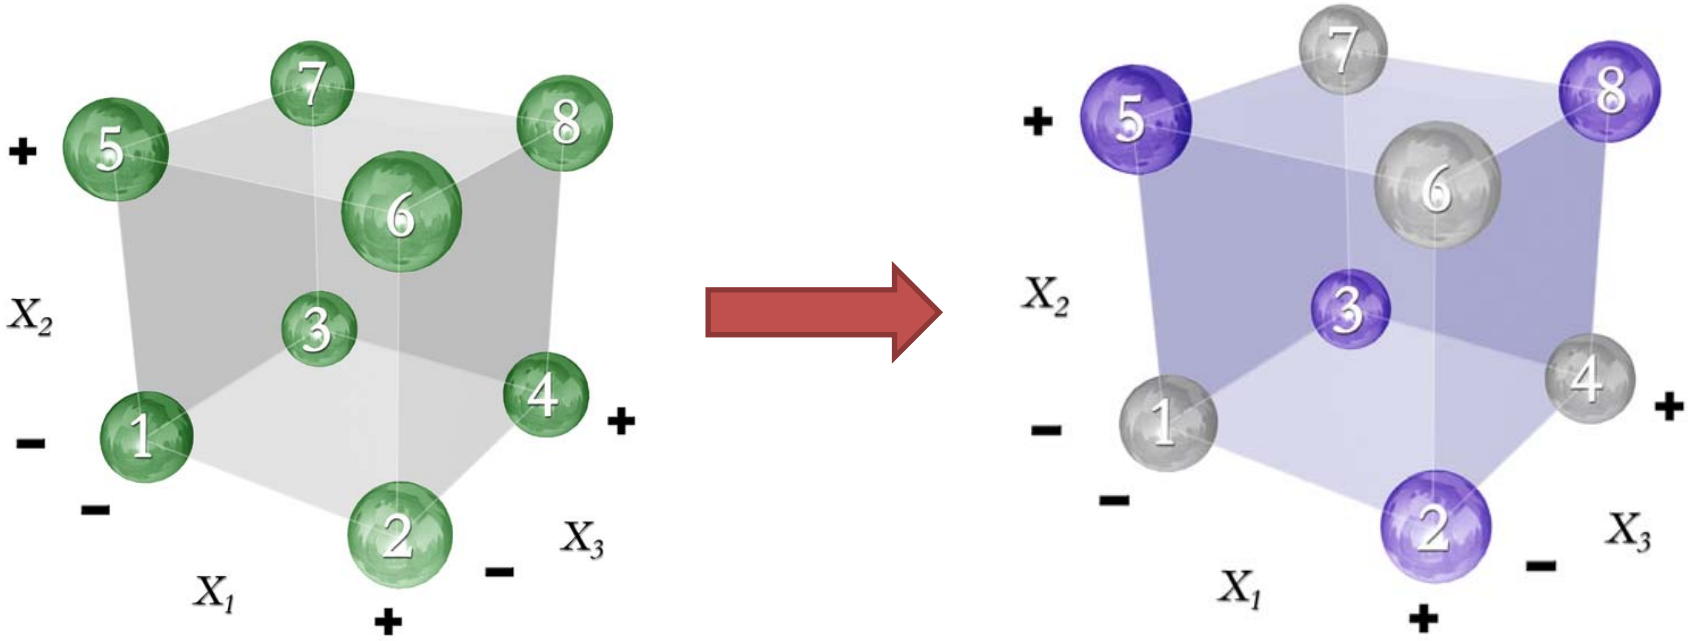

Plackett and Burman Design

- Exploring subset of parameter space such that performance variation can be estimated using only O(N) experiments.

	FIO	Postmark	IOZone
Number of Parameters	20	10	11
Experiments Required	1,048,576	1,024	2,048
Time (10min/exp)	19 years	7 days	14 days
Experiments Required (PB)	24	12	12
Time (10min/exp)	4 hours	2 hours	2 hours

Limitations

- Since we assume the main effects account of 100% of the performance variation;
  - Effect of interactions are cofounded in the main effects.
  - We assume that the interactions are weak enough that it does not affect the rank of the primary effects.

- Since only limits of a parameter is explored;
  - We assume the effect of a parameter is monotonic.

Outline

- Problem with Current Storage Performance Evaluation Techniques.
- Design of Experiments and Proposed Solution.
- Example Case using existing benchmarks.
- Effect of Benchmark Parameters.
- Benchmark Coverage Testing.
- Conclusion and Future work.

Parameter and Level Selection

Name	Level Low	Level High
min_file_size	512B	4KiB
max_file_size	4KiB	16MiB
init_file_count	1000	10000
transaction_count	10000	100000
read_size	512B	32KiB
write_size	512B	32KiB
file_system_buffer	false	true
read_append_ratio	1:9	9:1
create_delete_ratio	5:5	9:1
directory_count	1	1000

Name	Level Low	Level High
operation	read	write
access_pattern	sequential	random
files_used	1	100
min_file_size	512B	1MiB
max_file_size	1MiB	1GiB
min_block_size	512	4KiB
max_block_size	4KiB	64KiB
io_depth	1	100
overwrite	false	true
fsync	false	true
thinktime	0	1000
write_buffer_sync	false	true
file_service	roundrobin	random
thread_count	1	8
threads_similarity	false	true
posix_fadvise	false	true
async_io_engine	false	true
io_engine_queue	false	true
directio	false	true
buffer_alloc	malloc	mmap

Parameter and Level Selection

- Limits of parameters were observed from set of traces available at http://iotta.snia.org/.

- For the categorical parameter, each parameter was tested to identify best/worst performing values.

- If response is non-monotonic the parameter can be split into multiple parameters whose range yields monotonic response.

PB experiement

$$\mathbf{PB}(12) = \begin{bmatrix} 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & -1 \\ -1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 \\ 1 & -1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 \\ -1 & 1 & -1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 & -1 \\ -1 & -1 & 1 & -1 & 1 & 1 & -1 & 1 & 1 & 1 & -1 \\ -1 & -1 & -1 & 1 & -1 & 1 & 1 & -1 & 1 & 1 & 1 \\ 1 & -1 & -1 & -1 & 1 & -1 & 1 & 1 & -1 & 1 & 1 \\ 1 & 1 & -1 & -1 & -1 & 1 & -1 & 1 & 1 & -1 & 1 \\ 1 & 1 & 1 & -1 & -1 & -1 & 1 & -1 & 1 & 1 & -1 \\ -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & -1 & 1 & 1 \\ 1 & -1 & 1 & 1 & 1 & -1 & -1 & -1 & 1 & -1 & 1 \\ -1 & -1 & -1 & -1 & -1 & -1 & -1 & -1 & -1 & -1 & -1 \end{bmatrix}$$

- Each row corresponds to an experiment.
- Each column responds to a benchmark parameter.
- value 1 corresponds to high value and -1 corresponds to low value.

Experiment Flow

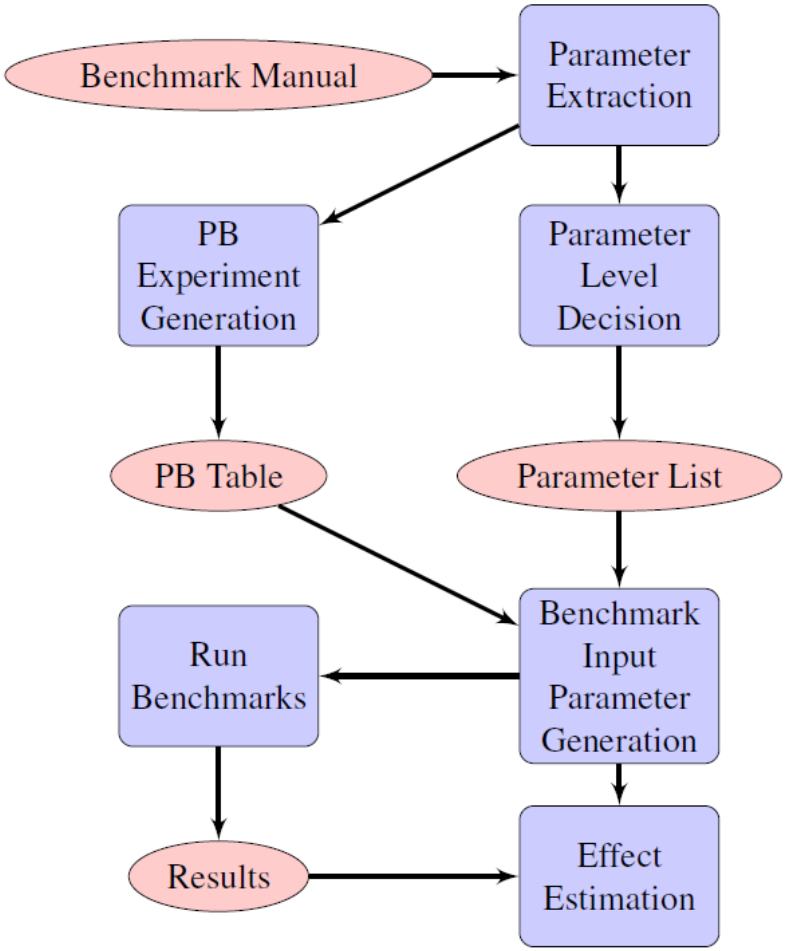

Figure 1. Experiment flow for estimating the parameter effects on performance

- PB method determines what combination of those parameter levels needs to be tested.
- Once the results are generated, we can use simple variance analysis to calculated the effects.

Outline

- Problem with Current Storage Performance Evaluation Techniques.
- Design of Experiments and Proposed Solution.
- Experiment Flow.
- Effect of Benchmark Parameters.
- Benchmark Coverage Testing.
- Conclusion and Future work.

Experimental Setup

Table 1. Description of system under test.	
CPU	2*2cores*2SMT Xeon 3GHz
Memory	10GB
OS	Lucid Server
Kernel	2.6.32
FileSystem	EXT4, noatime
Storage Conf.	RAID5
Controller	PERC3
Disks	SAS Cheetah 15K

- Each Experiment:
  - runtime: 10 minutes
  - repeat: 3 times (averaged)
  - Filesystem formatted and remounted between the runs.
- Dedicated RAID set used.
  - Single LUN
  - Single partition

Postmark Effects

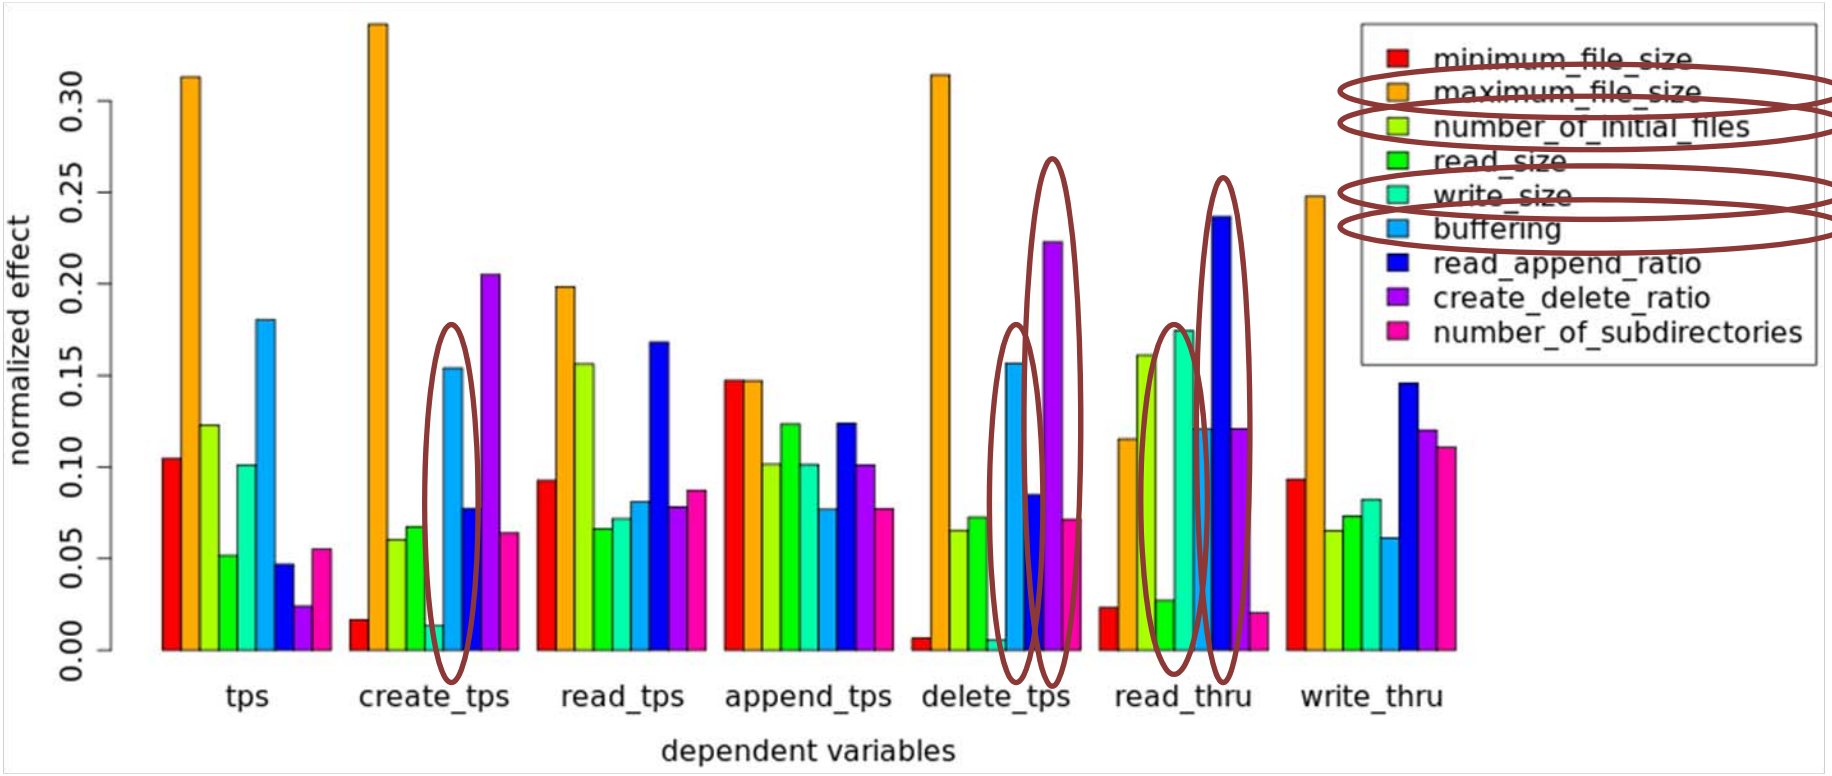

FIO Effects

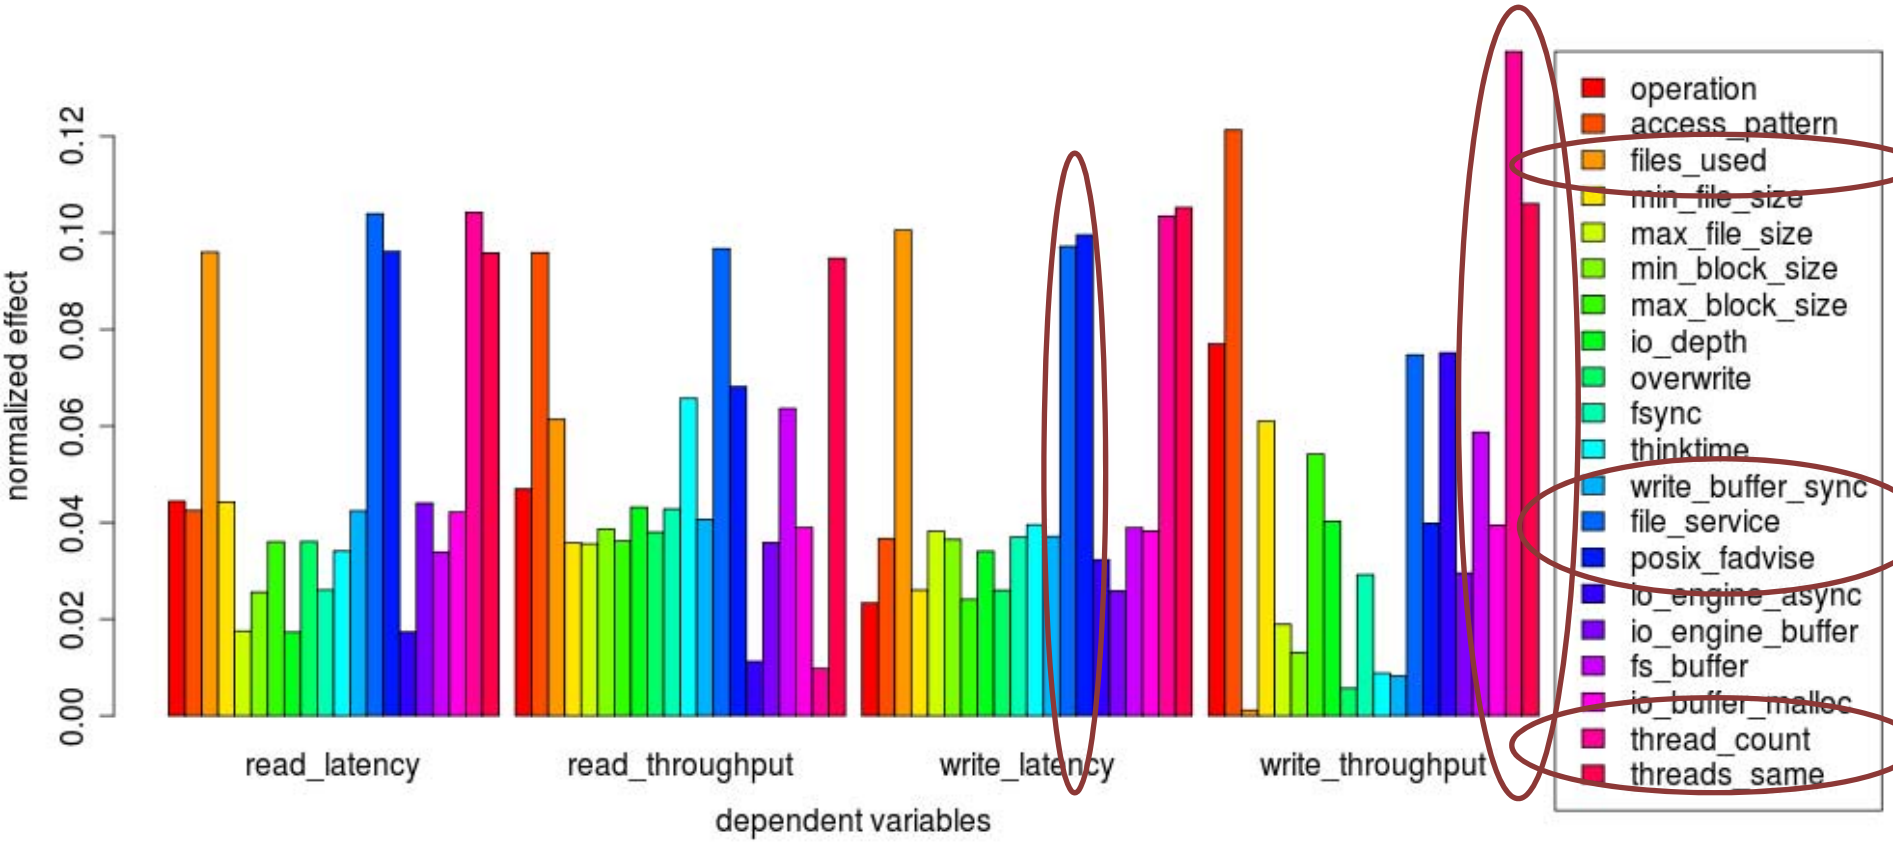
normalized effect
0.00
0.02
0.04
0.06
0.08
0.10
0.12
read_latency
read_throughput
write_latency
write_throughput
dependent variables
operation
access_pattern
files_used
min_file_size
max_file_size
min_block_size
max_block_size
io_depth
overwrite
fsync
thinktime
write_buffer_sync
file_service
posix_fadvise
io_engine_async
io_engine_buffer
fs_buffer
io_buffer_malloc
thread_count
threads_same

IOZone Effects

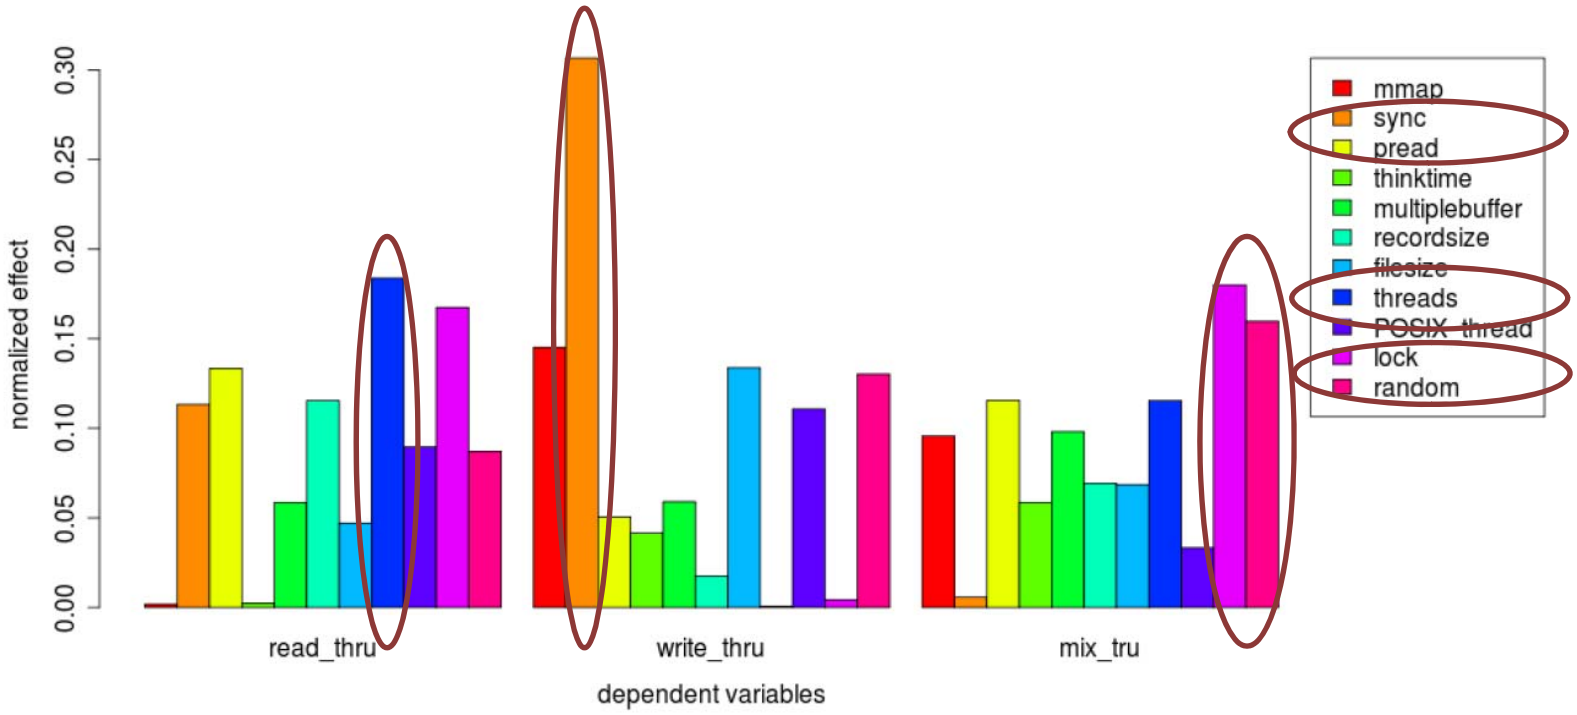

Figure 4. IO Zone parameter effects.

Full Factorial Design

	FIO	Postmark	IOZone
Number of Parameters	20	10	11
Experiments Required	1,048,576	1,024	2,048
Time (10min/exp)	19 years	7 days	14 days
Experiments Required (PB)	24	12	12
Time (10min/exp)	4 hours	2 hours	2 hours
Number of Key Parameters	6	4	4
Experiment Required	64	16	16
Time (10min/exp)	4+10hours	2+2hours	2+2hours

Goal

- To design experiment that
  - Generate input vector for the storage benchmarks that will statistically guarantee correctness of the experiment.

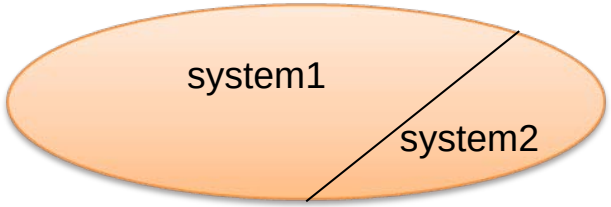

  - Experiment must be able to be carried out in a reasonable manner and provide a sufficient coverage.

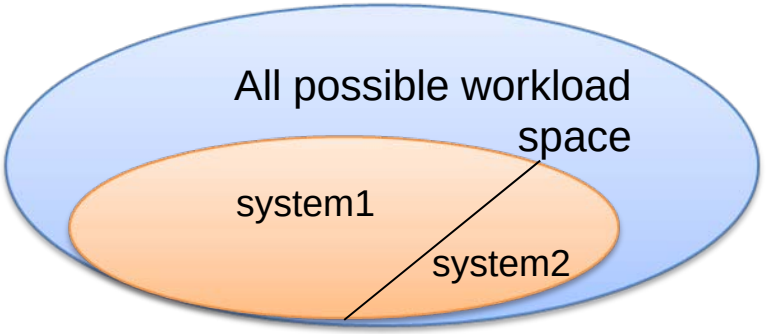

Coverage Testing

- Extract independent components that affects the performance.
- Each independent components actually has no physical meaning.
- Use K-mean to cluster the components based on the Euclidian distance.

Coverage

Overall coverage is not that different!

Read Throughput

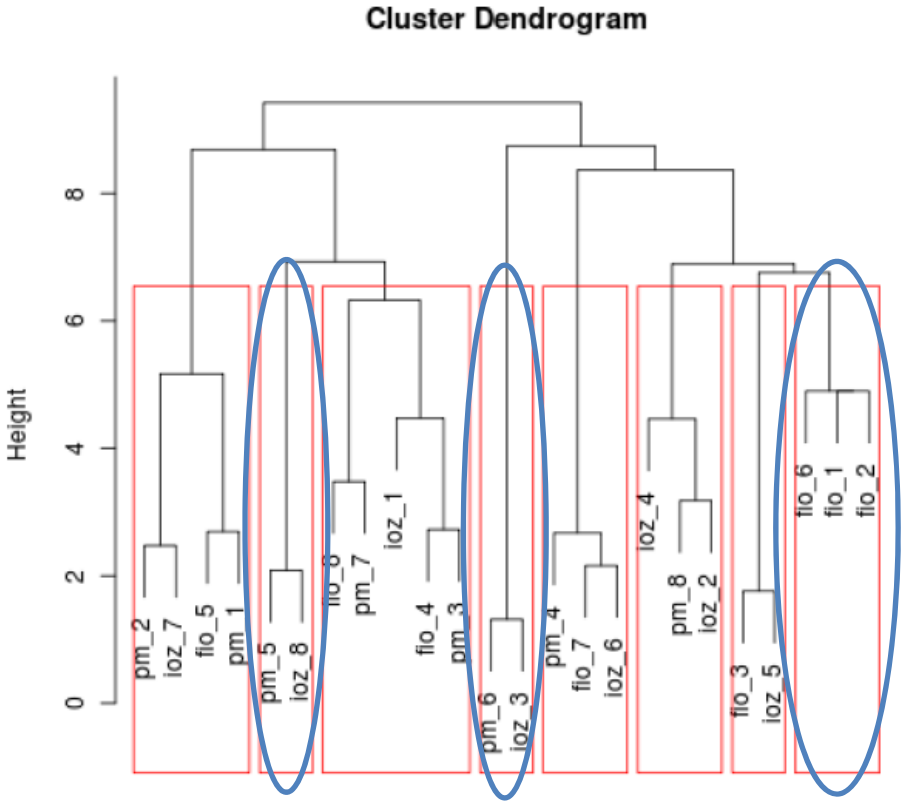

Write Throughput

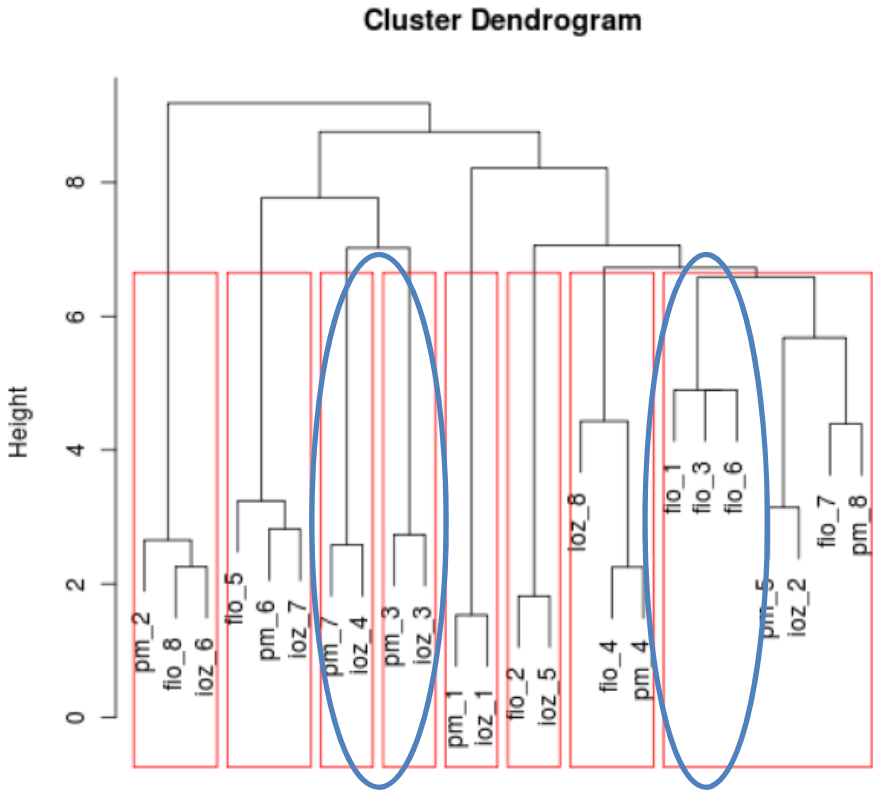

Outline

- Problem with Current Storage Performance Evaluation Techniques.
- Design of Experiments and Proposed Solution.
- Example Case using existing benchmarks.
- Effect of Benchmark Parameters.
- Benchmark Coverage Testing.
- Conclusion and Future work.

Conclusion

- Benchmark parameters can be reasonably evaluated using PB method.
- 2-4 parameters are enough to evaluate ~50% of performance variation.
- Coverage of the benchmark can be compared using ICA.

- It is better to use a single benchmark program varying high ranking parameters than use multiple benchmarks with few ad hoc settings.

On Going Work

- Evaluating parameter space for different storage technologies.
- Evaluate how file systems change the parameter effects.
- Model storage devices in terms of their sensitivity to the workload parameters.

Thank you and Questions?

Special Thanks to:

LSI and Seagate for their generous equipment donation and helpful comments.

UNIVERSITY OF MINNESOTA