

Carver: Finding Important Parameters for Storage System Tuning

18th USENIX Conference on File and Storage Technologies
(FAST'20)

Zhen Cao¹, Geoff Kuenning², and Erez Zadok¹

¹Stony Brook University; ²Harvey Mudd College



Outline

- **Motivation**
- Background
- Design of Carver
- Evaluation
- Related Work
- Conclusions & Future Work

Motivation

- Tuning storage system parameters can provide significant gains
 - ◆ Default settings are often sub-optimal
- Storage system parameter space is huge
- Manual tuning vs. auto-tuning

Motivation (cont.)

- More challenges to solve
 - ◆ “Curse of dimensionality”
 - ◆ Tuning efficiency is critical for storage systems
- Storage parameters come with different “importance”
- **Parameter selection is a critical part of auto-tuning systems**

Carver Overview

- Measurement for storage parameter importance
- Sampling method
 - ◆ Pick a small subset of configurations
- Parameter selection algorithm

Key Contributions

- Quantitative analysis of importance of storage parameters on system performance
- Carver can pick important parameters efficiently
- Evaluated Carver on various datasets
 - ◆ Carver can find a near-optimal subset of important parameters in all cases

Outline

- Motivation
- **Background**
- Design of Carver
- Evaluation
- Related Work
- Conclusions & Future Work

Concepts

- Storage system
 - ◆ File system, underlying storage hardware and any layers between them
- Parameters
 - ◆ Configurable options
 - ◆ *E.g., block size*
- Parameter values
 - ◆ *E.g., 1K, 2K, 4K (Ext4 block size)*
- Configuration
 - ◆ Combination of parameter values
 - ◆ *E.g., [Ext4, 4K, data=ordered]*
- Parameter space
 - ◆ All possible configurations

Challenges

- Large parameter space
 - ◆ Ext4: 59 parameters, 10^{37} configs
 - ◆ Distributed, large-scale
- Evaluation is time-consuming
- Nontransferable tuning results
 - ◆ Hardware, workload, etc.
- Discrete and non-numeric
 - ◆ Linux I/O scheduler: noop, cfq, deadline

Dimensionality Reduction

- Feature extraction **X**
 - ◆ Project high-dimensional data into low-dimensional spaces
 - ◆ Constructed features
 - ◆ Lose physical meaning of original features
- Feature selection
 - ◆ Directly selects a subset of features from the original ones
 - ◆ Relationship between parameters
 - ◆ **Impact of parameters on target variable**
 - Throughput

Outline

- Motivation
- Background
- **Design of Carver**
- Evaluation
- Related Work
- Conclusions & Future Work

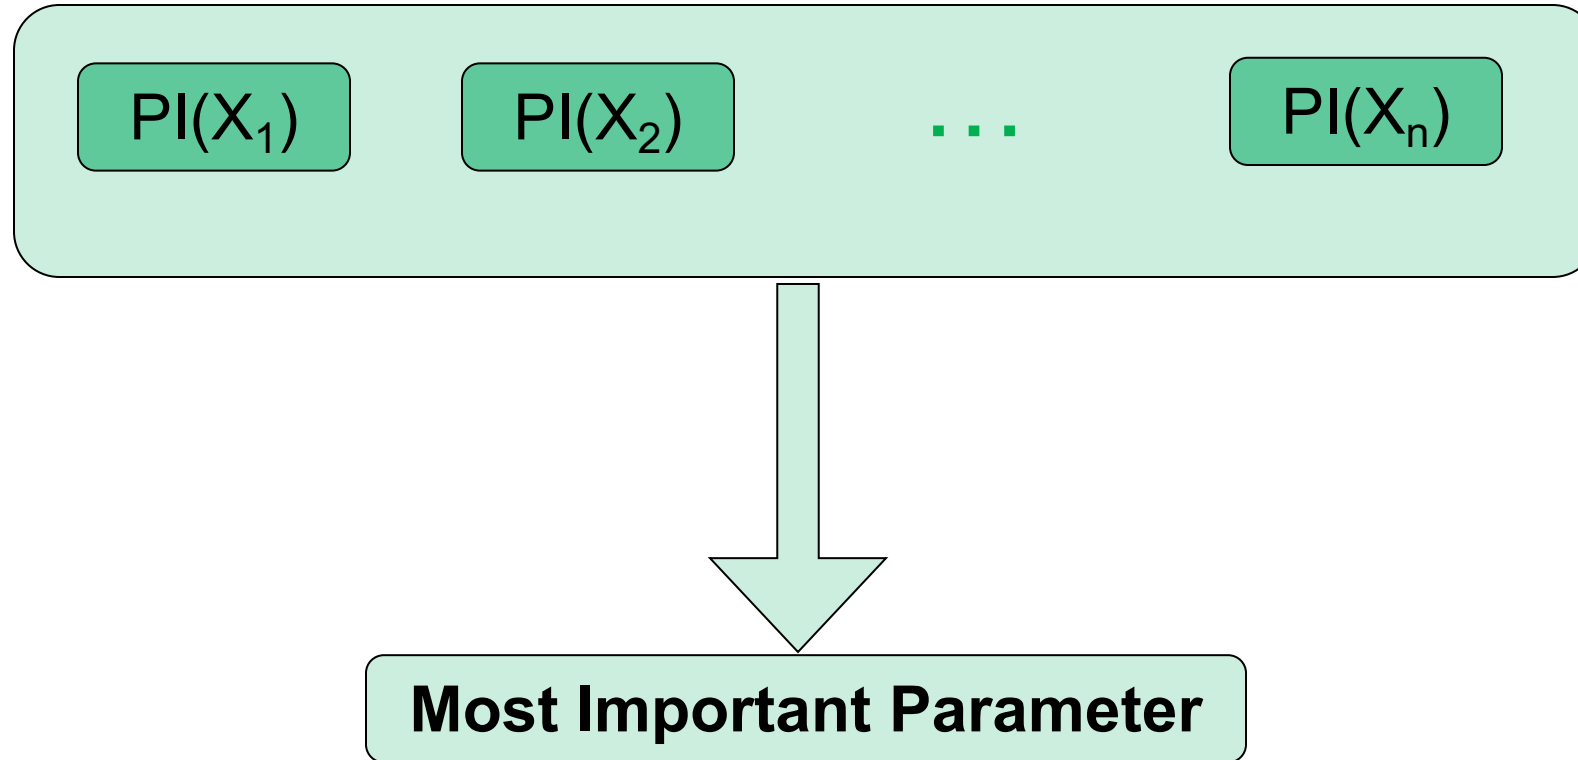
Carver: Measuring Parameter Importance

- Categorical parameters & continuous target variable
- Parameter Importance
 - ◆ Reduction in variance
- Conditional parameter importance
 - ◆ $\text{CPI}(P_2 \mid P_1 = p)$
 - ◆ E.g., given P_1 is fixed to “p”, find importance of P_2

Carver: Sampling

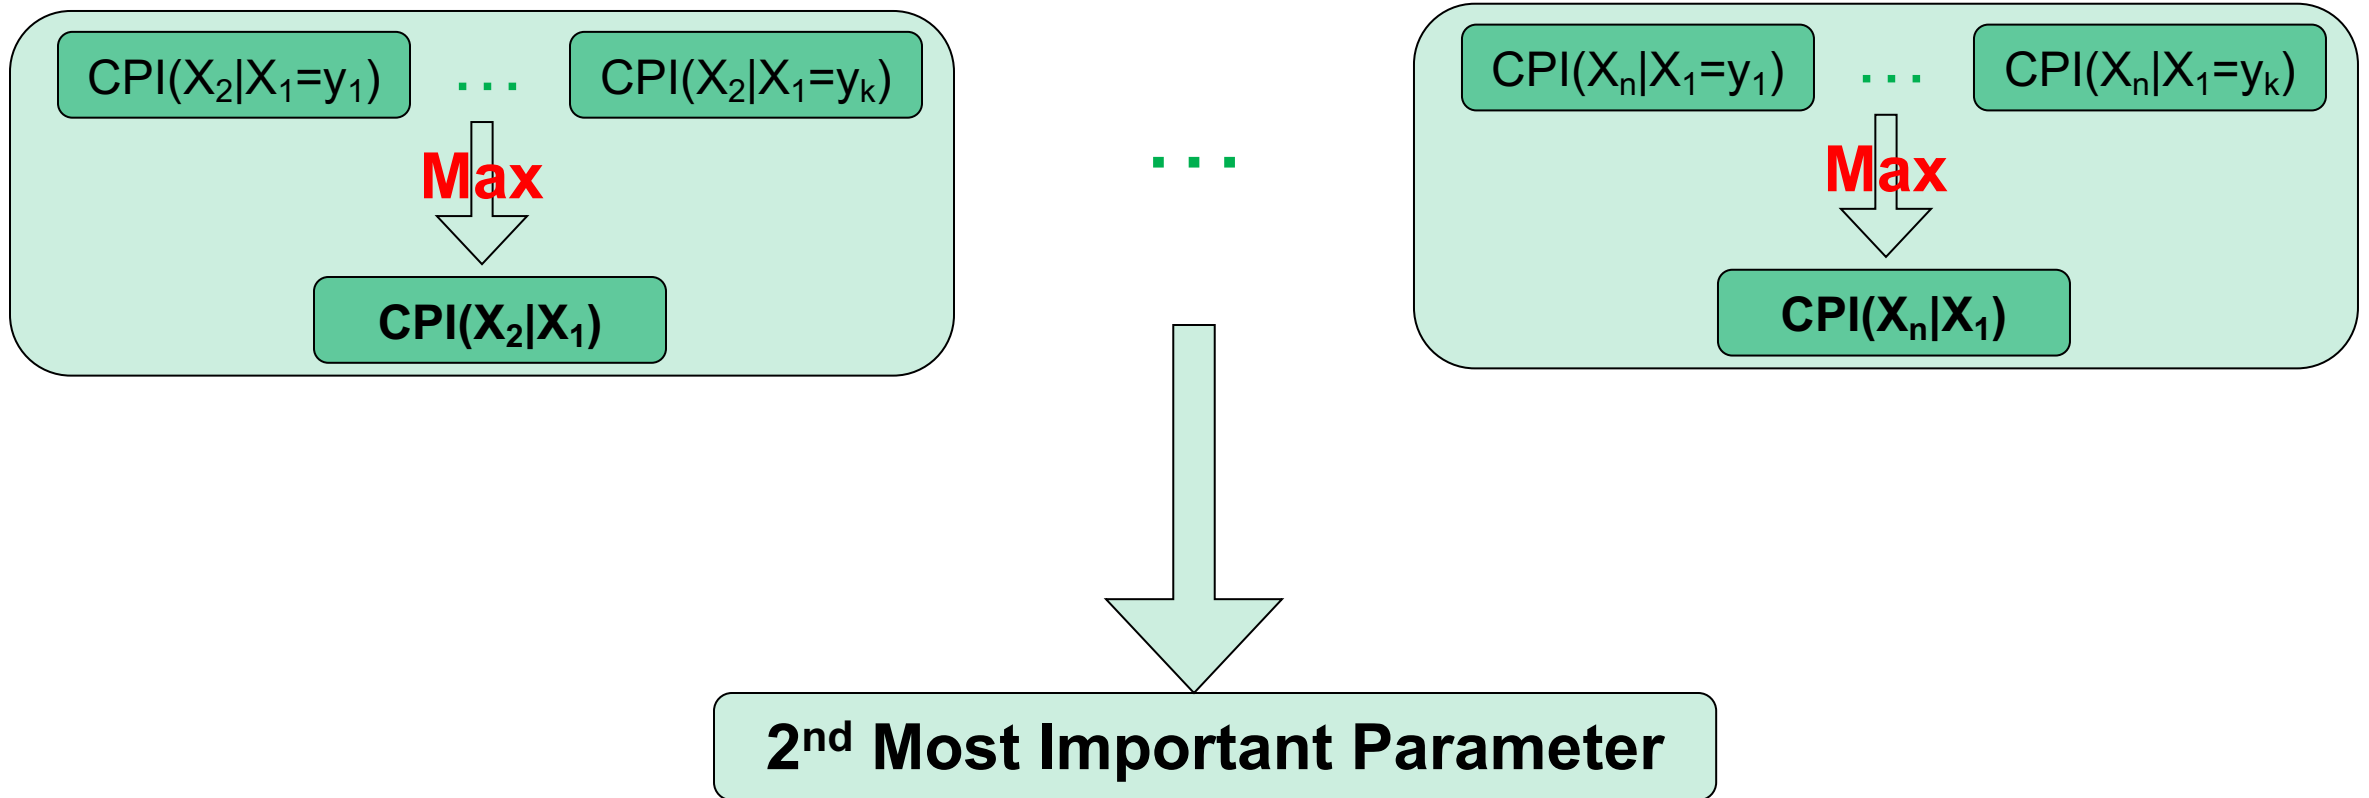
- Limit the number of evaluations
 - ◆ Evaluation of storage configurations is costly
- Sampling
 - ◆ Latin Hypercube Sampling
 - Latin square: one sample in each row and each column
 - Generalization to higher dimensions
 - Proved useful in system analysis

Carver: Parameter Selection (1)



Carver: Parameter Selection (2)

(Assume X_1 is as most important parameter.)



Carver: Parameter Selection (3)

- Repeat the process until stopping criterion is met
- Example stopping criteria:
 - ◆ Certain number of parameters have been selected
 - ◆ Variance within subsets of configurations (fixing values of selected parameters) falls below a certain threshold.

Outline

- Motivation
- Background
- Design of Carver
- **Evaluation**
- Related Work
- Conclusions & Future Work

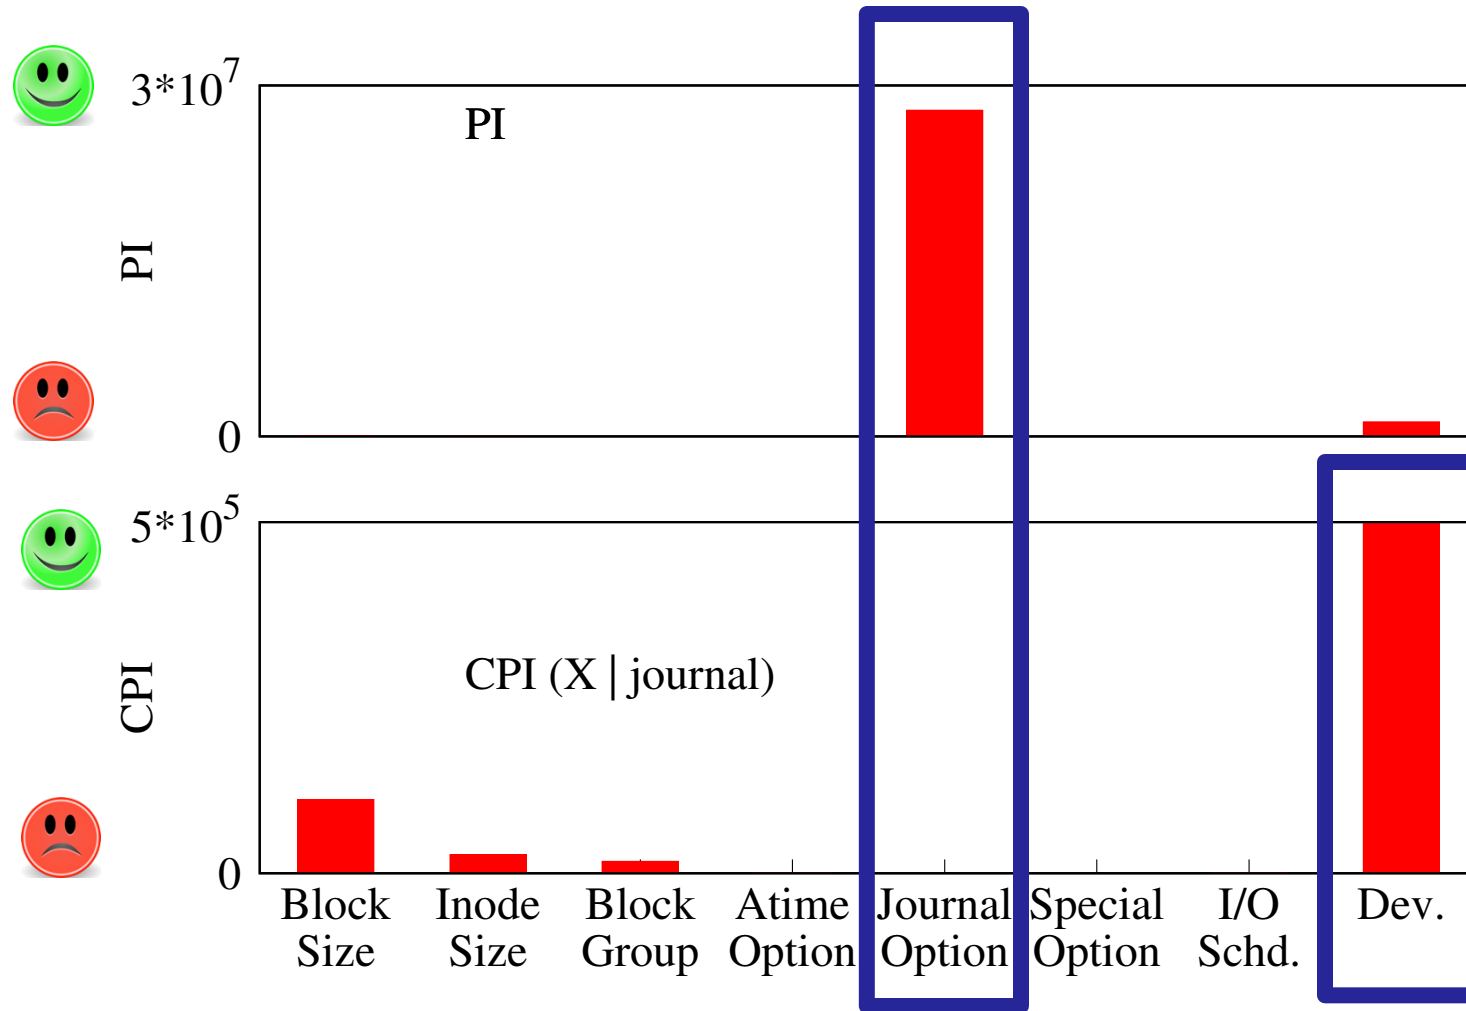
Evaluation: Testbed

- Hardware
 - ◆ Intel Xeon quad-core 2.4GHz CPU, 24G RAM, 4 drives (SAS-HDD 500GB, SAS-HDD 146GB, 1 SATA-HDD, SSD)
- Filebench
 - ◆ Workloads: fileserver, mailserver, webserver, dbserver
- Two Settings:
 - ◆ **S1**: *default working set size*
 - ◆ **S2**: *4G RAM, 10G working set size*

Evaluation: Settings

- Parameters spaces
 - ◆ 7 file system types, 4 workloads
 - ◆ inode size, block size, block group, journal options, mount options, special options, I/O scheduler, etc.
- Methodology
 - ◆ Pre-collect benchmark results for all configurations
 - 3+ runs for each configuration
 - **500,000 data points**
 - ◆ Evaluate Carver on collected datasets

Evaluation: Parameter Importance

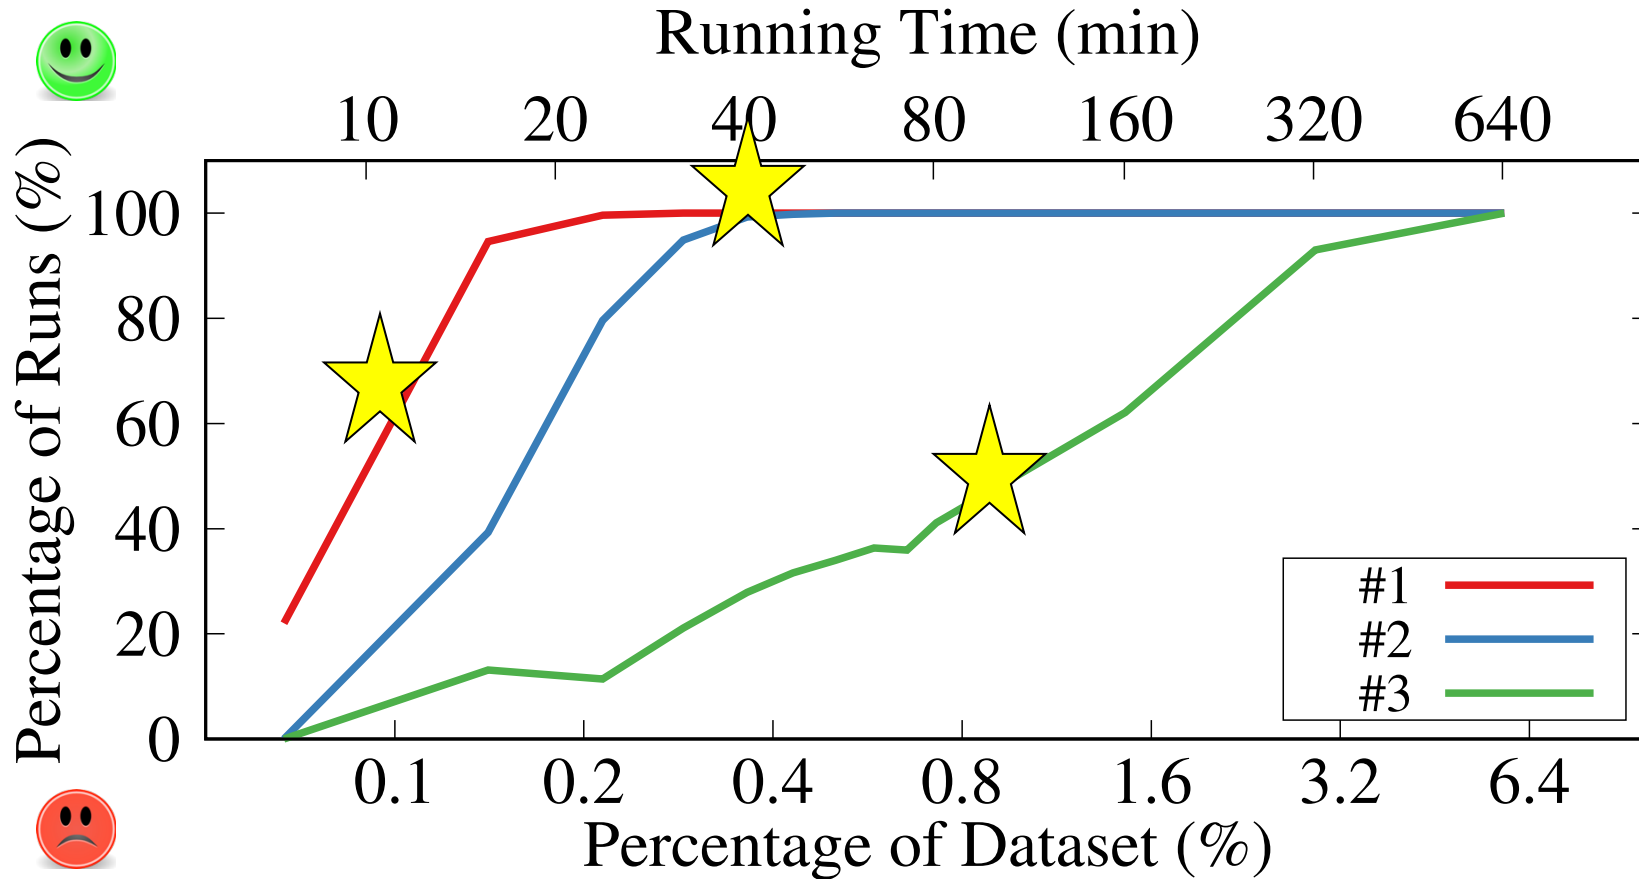


Ext4, S1, Fileserver-default

Parameter Importance Summary

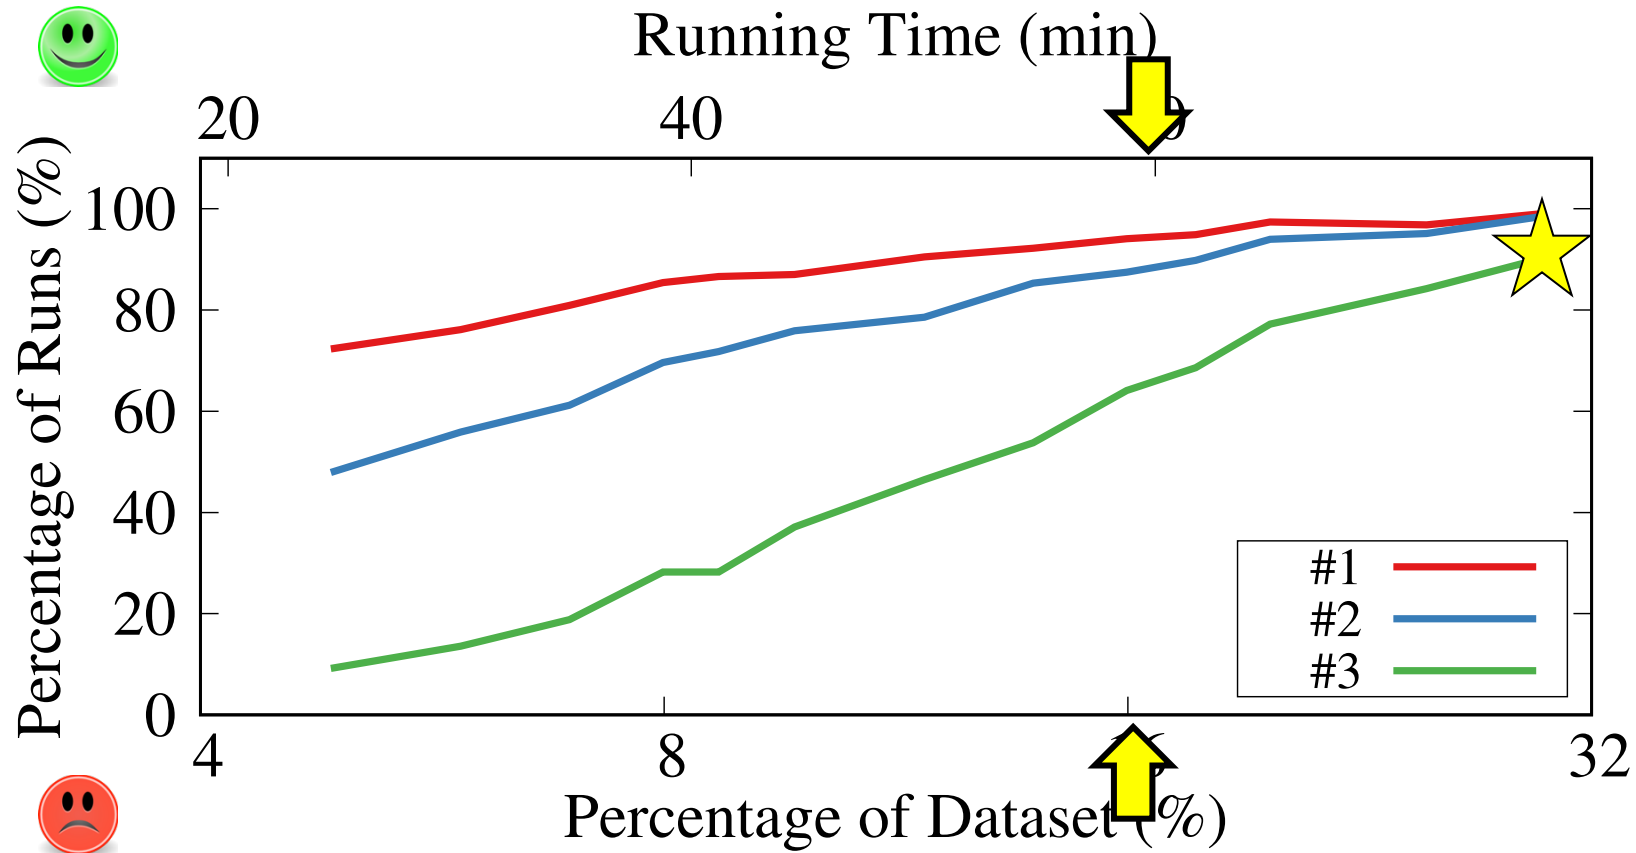
Setting	Workload	F/S	#1	#2	#3	#4
S2	fileserver-10gb	Ext4	Journal Option	I/O Scheduler	Inode Size	-
S2	dbserver-10gb	Ext4	Block Size	Inode Size	I/O Scheduler	Journal Option
S2	webserver-10gb	Ext4	Inode Size	Flex BG	Block Size	Journal Option
S1	fileserver-def	Btrfs	Special Option	Node Size	Device	-
S1	webserver-def	Btrfs	-	-	-	-

Evaluation: Carver



Ext4, S1, Fileserver-default

Evaluation: Carver (Cont.)



Btrfs, S1, Fileserver-default

Outline

- Motivation
- Background
- Design of Carver
- Evaluation
- **Related Work**
- Conclusions & Future Work

Related Work

- Parameter selection for systems
 - ◆ Lasso [Aken et al.]
 - ◆ Adaptive Sampling [Duan et al.]
 - ◆ Probabilistic Reasoning [Sullivan et al.]
- Auto-tuning storage systems
 - ◆ Genetic Algorithms, Deep Q-Networks, etc.
 - ◆ Comparative Study [Cao et al.]

Outline

- Motivation
- Background
- Design of Carver
- Evaluation
- Related Work
- **Conclusions & Future Work**

Conclusions

- Proposed Carver for efficiently selecting storage parameters
- Provided thorough study of storage-parameter importance
- Demonstrated Carver's efficiency by testing it on small fractions of the configuration space

Future Work

- Extend Carver to support other parameter-selection techniques
 - ◆ Group Lasso, ANOVA, etc.
- Evaluate Carver with more optimization objectives and larger storage-parameter spaces
- Extend Carver's parameter selection algorithm into multi-objective optimization
- Integrate Carver with auto-tuning algorithms

Carver: Finding Important Parameters for Storage System Tuning

Zhen Cao, Geoff Kuenning, and Erez Zadok

Thank You Q&A

