# A Practical, Real-Time Auto-Tuning Framework for Storage Systems

A Dissertation Proposal Presented

by

**Zhen Cao**

to

The Graduate School

in Partial Fulfillment of the

Requirements

for the Degree of

**Doctor of Philosophy**

in

Computer Science

Stony Brook University

**Technical Report FSL-18-01**

**April 2018**

## Abstract

A Practical, Real-Time Auto-Tuning Framework for Storage Systems

by

Zhen Cao

Doctor of Philosophy Candidate
in
Computer Science

Stony Brook University
2018

Storage systems come with a large number of configurable parameters that control their behavior. Tuning such parameters can provide significant gains in performance, but is challenging because of huge parameter spaces and complex, non-linear system behavior. Auto-tuning with black-box optimization have shown some promising results in recent years, thanks to its obliviousness to systems' internals.

However, previous work all applied only one or few optimization methods, and did not systematically evaluate them. Therefore, in this thesis proposal, we first apply and then perform comparative analysis of multiple black-box optimization techniques on storage systems from various aspects such as their ability to find near-optimal configurations, convergence time, and instantaneous system throughput during auto-tuning, etc. We also provide insights into the efficacy of these automated black-box optimization methods from a system's perspective.

During our auto-tuning experiments, we noticed that sometimes multiple runs of the same workload—in a carefully controlled environment—produced widely different performance results. So next, we undertook a study to characterize the amount of variability in modern storage systems. We analyzed these variations and found that there was no single root cause: it often changed with the workload, hardware, or software configuration in the storage system. In several of those cases we were able to fix the cause of variation and reduce it to acceptable levels.

Despite some promising early results, we believe several critical features are still missing from traditional black-box optimization methods. Therefore, we propose to investigate and design a more intelligent and practical framework for auto-tuning storage systems in real-time. We define stopping and restarting criteria to stop auto-tuning when "good enough" configurations are found and restart it in response to environment changes. We add a workload modeler to characterize the running workload. Initialization methods will be studied as well, which showed significant impact on the overall efficacy of auto-tuning in our preliminary results. Our framework includes a weighted penalty function, to account for costly configuration changes. We also plan to investigate how Machine Learning (ML) techniques can help on various aspects of our auto-tuning framework (e.g., identify unimportant parameters and eliminate them from the search space).

It is our thesis that real-time auto-tuning storage systems is important, promising, and feasible with a carefully designed framework to include missing yet critical features. This can improve systems' performance efficiency, and save energy and human resources in the long term.

# Contents

**List of Figures** v

**List of Tables** vii

**Acknowledgments** ix

**1 Introduction** 1

**2 Background** 4
- 2.1 Problem Statement . . . 4
- 2.2 Black-box Optimization . . . 7
- 2.3 Machine Learning . . . 8
- 2.4 Unified Framework . . . 9

**3 Related Work** 11
- 3.1 Auto-tuning in Computer Systems . . . 11
- 3.2 Hyper-parameter tuning . . . 12
- 3.3 Workload Modeling . . . 12

**4 Experimental Settings** 13
- 4.1 Hardware . . . 13
- 4.2 Workload . . . 13
- 4.3 Parameter Space . . . 15
- 4.4 Experiments and Implementations . . . 16

**5 Towards Better Understanding of Black-box Auto-Tuning: A Comparative Analysis for Storage Systems** 18
- 5.1 Overview of Datasets . . . 19
- 5.2 Comparative Analysis . . . 20
- 5.3 Impact of Hyper-Parameters . . . 24
- 5.4 Peering into the Black Box . . . 25
- 5.5 Limitations . . . 28

**6 On the Performance Variation in Modern Storage Systems** 29
- 6.1 Motivations . . . 29
- 6.2 Background . . . 31

6.2.1 Measures of Variation . . . 31
6.3 Methodology . . . 32
6.4 Related Work . . . 35
6.5 Experimental Setup and Workloads . . . 35
6.6 Evaluation . . . 36
6.6.1 Variation at a Glance . . . 37
6.6.2 Case Study: Ext4 . . . 39
6.6.3 Temporal Variation . . . 43
6.6.3.1 Throughput over Time . . . 43
6.6.3.2 Latency Variation . . . 46

**7 A Practical Auto-Tuning Framework for Storage 49**
7.1 Motivations . . . 49
7.2 Problem Statement . . . 50
7.3 Proposed Auto-Tuning Framework . . . 50
7.3.1 Workload Modeling . . . 51
7.3.2 Optimizer . . . 52
7.3.3 (Re-)Initialization . . . 52
7.3.4 Stopping Criteria . . . 53
7.3.5 Penalty Functions . . . 54
7.3.6 Machine Learning . . . 55
7.3.7 Visualizer . . . 56

**8 Proposed and Future Work 57**
8.1 Proposed Work . . . 57
8.2 Future Work . . . 59

**9 Conclusions 60**

# List of Figures

2.1 Storage systems are non-linear . . . . . . . . . . . . . . . . . . . . . . . . . . . . 5
2.2 Evaluation results depend on workloads . . . . . . . . . . . . . . . . . . . . . . . . 6
2.3 Crossover and mutation in a Genetic Algorithm . . . . . . . . . . . . . . . . . . . . 8

5.1 Throughput CDF with different hardware and workloads, with symbols marking the default configurations. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 19
5.2 Highest throughput found over time, zooming in the $Y \in [15 : 19]$ range. The blue number (15.2) on the Y axis shows the default, and the red one (18.7) shows the optimal. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 21
5.3 Comparing optimization methods' efficacy in finding near-optimal configurations. The Y axis shows the percentage of total runs (1,000) that found near-optimal configurations within certain time (X axis). . . . . . . . . . . . . . . . . . . . . . 22
5.4 Comparing optimization methods' instantaneous performance (Y axis) over time (X axis). . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 23
5.5 Impact of mutation rates on GA. . . . . . . . . . . . . . . . . . . . . . . . . . . . 24
5.6 Number of alleles (parameter values) in the first 10 generations from one GA experiment run, with more frequent ones colored with darker colors. . . . . . . . . . 25
5.7 Scatter plot for all Ext3-SSD configurations under fileserver-def workload, with one dot corresponding to one configuration. . . . . . . . . . . . . . . . . . . . . . 26

6.1 Cumulative throughput over time for one Ext4 configuration under multiple workloads. Each workload ran for 7,200s; only the first 3,000s are plotted. . . . . . . . 36
6.2 Overview of performance and its variation with different storage configurations under three workloads: (a) maileserver-heavy, (b) fileserver-heavy, and (c) webserver-heavy. The X axis represents the mean of throughput over 10 runs; the Y axis shows the relative range of cumulative throughput. Ext4 configurations are represented with squares, XFS with circles, and Btrfs with triangles. HDD configurations are shown with filled symbols, and SSDs with hollow ones. . . . . . . . . . . 37
6.3 Storage system performance variation with 20 sampled Ext4-HDD configurations under three workloads. The range is computed among 10 experiment runs, and is represented as bars corresponding to the Y1 (left) axis. The mean of throughput among the 10 runs is shown with symbols (squares, circles, and triangles), and corresponds to the Y2 (right) axis. The X axis represents configurations formatted by $\langle$block size - inode size - journal - atime - I/O scheduler - device$\rangle$. . . . . . . . . 38

6.4 Performance variation for 2 Ext4-HDD configurations with several diagnoses. Each experiment is shown as one box, representing a throughput distribution for 10 identical runs. The top border line of each box marks the $1^{st}$ quartile; the bottom border marks the $3^{rd}$ quartile; the line in the middle is the median throughput; and the whiskers mark maximum and minimum values. The dots to the right of each box show the exact throughputs of all 10 runs. The percentage numbers below each box are the relative range values. The bottom label shows configuration details for each figure. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 39
6.5 Performance variation for Ext4-HDD configuration under the Fileserver workload with different partition sizes from inner tracks of disks . . . . . . . . . . . . . . . . 41
6.6 Physical blocks of allocated files in Ext4 under the Fileserver workload. The X axis represents the physical block number of each file in the dataset. Since the *Fileserver* workload consists of small files, and one extent per file, we use the starting block number for each file here. The Y axis is the final cumulative throughput for each experiment run. Note that the Y axis does not start from 0. Lines marked with solid circles are experiment runs with the default setting; lines with triangles represent experiment runs where we set the field s_hash_seed in Ext4s's superblock to null. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 42
6.7 *Throughput-120* over time for Btrfs, XFS, and Ext4 HDD configurations under the Fileserver workload. Each configuration was evaluated for 10 runs. Two lines were plotted connecting maximum and minimum throughput values among 10 runs. We fill in colors between two lines, green for Btrfs, red for Ext4, and blue for XFS. We also plotted the average *Throughput-120* among 10 runs as a line running through the band. The maximum relative range values of *Throughput-120* for Ext4, Btrfs, and XFS are 43%, 23%, and 65%, while the minimum values are 14%, 2%, and 7%, respectively. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 43
6.8 CDFs for relative range of throughput under Fileserver workload with different window sizes. For window size N, we calculated the relative range values of throughput for all configurations within each file system type, and then plotted the corresponding CDF. . . . . . . . . . . . . . . . . . . . . . . . . . . . . 45
6.9 Normalized instantaneous throughput (*Throughput-10*) over time for experiments with various workloads, file systems, and devices. The Y axis shows the normalized values divided by the maximum instantaneous throughput through the experiment. Only the first 500s are presented for brevity. . . . . . . . . . . . . . . . . . 46
6.10 Latency CDF of one Ext4-HDD configuration under *Fileserver* workload. . . . . . 47
6.11 Pearson Correlation Coefficient (PCC) between throughput range and operation types, for three workloads and three file systems. The horizontal dashed red line at Y=0.7 marks the point above which a strong correlation is often considered to exist. 48

7.1 Auto-tuning Framework . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 51
7.2 Work flow for an enhanced Optimizer (GA). . . . . . . . . . . . . . . . . . . . . 52
7.3 Comparison of different initialization methods. . . . . . . . . . . . . . . . . . . 53
7.4 Time window based stopping criteria. . . . . . . . . . . . . . . . . . . . . . . . 54

8.1 Auto-tuning Framework. Components are re-colored based on our project timeline. We plan to completely finish work colored by green; partially finish work colored by yellow. Components colored with red are left for future work beyond this thesis. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 58
8.2 Work flow for an enhanced Optimizer (GA). Components are re-colored based on our project timeline. We plan to completely finish work colored by green; partially finish work colored by yellow. Components colored with red are left for future work beyond this thesis. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 59

# List of Tables


2.1 Comparison and summaries of optimization techniques . . . . . . . . . . . . . . . . . . . . 9

4.1 Details of experiment machines. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 14
4.2 Filebench workload characteristics. . . . . . . . . . . . . . . . . . . . . . . . . . . . . 14
4.3 Details of Parameter Spaces . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 16

5.1 Global optimal configurations with different settings and workloads. Workloads are abbreviated. Db: dbserver-def; File: fileserver-def; Mail: mailserver-def; Web: webserver-def. . 20
5.2 Importance of parameters (measured by $R^2$) among SSD configurations, with the most important one colored in yellow and second in green. . . . . . . . . . . . . . . 27

6.1 Comparison for parameter spaces. Time is computed by assuming 15 minutes per experimental run, 10 runs per configuration and 3 workloads in total. . . . . . . . . 33
6.2 List of parameters and value ranges. . . . . . . . . . . . . . . . . . . . . . . . . . . . 34

7.1 Categories of parameter penalties . . . . . . . . . . . . . . . . . . . . . . . . . . . . 55

# Acknowledgments

# Chapter 1

# Introduction

Storage is a critical element of computer systems and key to data-intensive applications. Storage systems come with a vast number of configurable parameters that control a system's behavior. Ext4 alone has around 60 parameters with whopping $10^{37}$ unique combinations of values. Default parameter settings provided by vendors are often suboptimal for a specific user deployment; previous research showed that tuning even a small subset of parameters can improve power and performance efficiency of storage systems by as much as $9\times$ [132].

Traditionally, system administrators pick parameter settings based on their expertise and experience. Due to the increased complexity of storage systems, however, manual tuning becomes intractable, error-prone, and has a low chance of finding an optimal configuration. A myriad of file systems with diverse goals and designs have been developed [47, 83, 87, 124, 144]. Newer types of devices (SSDs [64, 108], SMR drives [2, 3], PCM [81, 163]) and more layers (LVM, RAID) are added. Storage systems expand from one or few identical nodes to hundreds of highly heterogeneous environments [55, 129]. Tuning results from one workload are often inapplicable in another [24, 155]. Furthermore, the composition of hardware and workload in a modern environment changes at a fast pace that prohibits timely manual tuning.

In recent years, several attempts were made to automate the tuning of computer systems in general and storage systems in particular [141, 155]. Black-box auto-tuning is an especially popular approach thanks to its obliviousness to system's internals [170]. The basic mechanism behind black-box auto-tuning is to iteratively try different configurations, measure an objective function's value—and based on the previously learned information—select the next configurations to try. For storage systems, objective functions can be throughput, I/O latency, energy consumption, purchase cost, or even a formula combining multiple metrics [97, 141]. Many black-box auto-tuning algorithms exist and some were applied to systems. Genetic Algorithms (GA) were applied to optimize the I/O performance of HDF5-based applications [12]. Bayesian Optimization (BO) was used to find a near-optimal configuration for Cloud VMs [5]. Other methods include Evolutionary Strategies [126], Smart Hill-Climbing [164], and Simulated Annealing [42]. Although these methods were originally proposed in different scientific disciplines, they all maintain a trade-off among three behavioral dimensions: (1) *Exploration*: how much the technique searches the space randomly. (2) *Exploitation*: how much the technique leverages the "neighborhood" of the current candidate or previous search history to find even better configurations. (3) *History*: how much data from previous evaluations is kept and utilized in the overall search process. For this dissertation, we propose to investigate and design a general framework based on black-box optimization, which

can efficiently auto-tune storage systems in real-time.

To demonstrate black-box optimization's ability to find optimal (or at least near-optimal) storage configurations, we started by exhaustively evaluating several storage systems under four workloads on two servers with different hardware and storage devices; the largest system consisted of 6,222 unique configurations. Over a period of 2+ years, we executed 450,000+ experimental runs, with 18 different combination of workload and hardware settings. We stored all data points in a relational database for query convenience, including hardware and workload details, throughput, energy consumption, running time, etc. In this thesis proposal, we mainly focused on optimizing for 0throughput, but our methodology and observations are applicable to other metrics as well. We plan to release our dataset publicly to facilitate more research into auto-tuning and better understanding of storage systems.

Despite some appealing results in auto-tuning, there is no deep understanding how exactly these black-box optimization methods work, their efficacy and efficiency, and which methods are more suitable for which problems. Previous works picked algorithms somewhat arbitrarily and evaluated only one algorithm at a time. Therefore, in this proposal, for the first time and to the best of our knowledge, we apply and analytically compare *multiple* black-box optimization techniques on storage systems. We applied several popular techniques to the collected dataset to find optimal configurations under various hardware and workload settings: Simulated Annealing (SA), Genetic Algorithms (GA), Bayesian Optimization (BO), and Deep Q-Networks (DQN). We also tried Random Search (RS) in our experiments, which showed surprisingly good results in previous research [15]. We compared these techniques from various aspects, such as their ability to find near-optimal configurations, convergence time, and instantaneous system throughput during auto-tuning. For example, we found that several techniques were able to converge to good configurations given enough time, but their efficacy differed a lot. GA and BO outperformed SA and DQN on our parameter spaces, both in terms of convergence time and instantaneous throughputs. Surprisingly, RS was also able to identify good configurations, sometimes even more efficiently than sophisticated optimization methods. We further compared the techniques across the aforementioned three behavioral dimensions: *exploration*, *exploitation*, and *history*. Based on our experimental results and domain expertise, we also provide explanations of efficacy of such black-box optimization methods from a storage perspective. We observed that certain parameters would have a greater effect on system performance than others, and the set of dominant parameters depends on file systems and workloads.

During our auto-tuning experiments, we noticed that sometimes multiple runs of the same workload—in a carefully controlled environment—produced widely different performance results. In one experiment setting, over 18% of 6,222 different storage configurations that we tried exhibited a standard deviation of performance larger than 5% of the mean, and a range value (maximum minus minimum performance, divided by the average) that exceeding 9%. In a few extreme cases, the standard deviation exceeded 40% even with numerous repeated experiments. This motivated us to conduct a more detailed study of storage system performance variation and seek its root causes, as performance stability is important for the success of auto-tuning and more broadly is critical in modern storage systems. Therefore, in this proposal we conducted experiments on three local file systems (Ext4, XFS, and Btrfs) which are used in many modern local and distributed environments. We benchmarked over 100 configurations using different workloads and repeated each experiment 10 times to balance the accuracy of variation measurement with the total time taken to complete these experiments. We then characterized performance variation from several angles: throughput,

latency, temporally, spatially, and more. We found that performance variation depends heavily on the specific configuration of the storage system. We then further dove into the details, analyzed and explained certain performance variations. For example, we found that unpredictable layouts in Ext4 could cause over 16–19% of performance variation in some cases. Finally, we analyzed latency variations from various aspects, and proposed a novel approach for quantifying the impacts of each operation type on overall performance variation.

Despite some promising preliminary results, we believe traditional black-box optimization techniques still lack several critical features to achieve practical, real-time auto-tuning in storage systems. Our own experiments demonstrate that auto-tuning sometimes can be slow in finding near-optimal configurations, especially when the evaluation of even a single configuration takes long time (e.g., due to slow I/O). Worse, when each experiment itself takes a long time (e.g., due to slower I/Os), using such techniques alone can take even longer. Moreover, there is no implicit mechanism to stop the search when it reaches a sufficiently good configuration (and restart it later on as needed); little is known on how to initialize the search and give it a good starting point; and there is no accounting for the cost of moving from one configuration to another, which is critically important in some production settings.

Therefore, for this dissertation we propose to investigate and develop a more intelligent and practical auto-tuning framework, intended to dynamically optimize storage systems. We are exploring techniques that add vital missing features from existing optimization methods: **(1)** A criteria when the optimization algorithm should *stop searching*, having reached a "good enough" system configuration. **(2)** A similar criteria when the search algorithm should be *restarted*, useful when the environment conditions (e.g., workload) have changed enough to take the system off of its optimal point. **(3)** A workload modeler, which can extract features from collected system metrics and characterize the running workload based on them. This is useful in determining when to restart the auto-tuning process and how to "transfer" evaluation results from one workload to another. **(4)** A mechanism to *pick an initial set of search space locations*, as well as *re-initialize* the search space after restarting a search—which we have found to have a big impact on the efficacy of any search [23, 43]. **(5)** A *penalty function* to assign a (weighted) cost to any new configuration based on the current system state, to account for costly configuration changes (e.g., a simple run-time changeable parameter vs. one that requires a system reboot and some downtime). We describe some preliminary results exploring these proposed ideas, and discuss how these components could help us achieve the goal of auto-tuning storage systems in real-time.

The rest of this dissertation proposal is organized as follows. Chapter 2 describes challenges of auto-tuning storage systems and background knowledge on black-box optimization. Chapter 3 discusses related work. We list our experimental settings in Chapter 4. In Chapter 5 we perform a comparative analysis on multiple optimization methods. Chapter 6 provides our characterization work on performance variation in modern storage stocks. Chapter 7 discusses several missing yet important components from traditional black-box optimization. Based on it, we propose to investigate and design a more intelligent and practical auto-tuning framework for storage systems. Chapter 9 concludes this proposal.

# Chapter 2

# Background

In this thesis proposal we use "storage systems" to refer file systems, underlying storage hardware and any layers between them. Storage systems have always been a critical component of most computer systems, and are the foundation for many data-intensive applications. Usually they come with a large number of configurable options that could affect or even determine the systems' performance [24, 146], energy consumption [132], and other aspects [94, 141]. Here we define a *parameter* as one configurable option, and a *configuration* as a certain combination of parameter values. For example, the *journal_mode* is one *parameter* for Ext4, with 3 possible values: *data=writeback*, *data=ordered*, and *data=journal*. Two other common parameters are *block_size* and *inode_size* with several possible numeric values (e.g., 4K, 8K). *[journal_mode="data=writeback", block_size=4K, inode_size=4K]* is one *configuration* with 3 specific parameters: *journal mode*, *block size*, and *inode size*. All possible configurations form a *parameter space*.

When configuring storage systems, users often stick with the default configurations provided by vendors because

- it is nearly impossible to know the impact of every parameter across multiple layers; and
- vendors' default configurations are trusted to be safe and "good enough".

However, previous studies [132] showed that tuning even a tiny subset of parameters could improve the performance and energy efficiency for storage systems by as much as 9×. As Moore's law slows down, it becomes even more important to squeeze every bit of performance out of deployed storage systems.

The rest of this chapter is organized as follows. We first discuss the challenges of storage system tuning in Chapter 2.1. Then, Chapter 2.2 briefly introduces several black-box optimization techniques that we explore in this proposal. Chapter 2.3 discusses how Machine Learning (ML) techniques can help in auto-tuning storage systems. Chapter 2.4 provides a unified view of these optimization methods.

## 2.1 Problem Statement

The tuning task for storage systems is difficult, due to the following four challenges.

**(1) Large parameter space** Modern storage systems are fairly complex and easily come with hundreds or even thousands of tunable parameters. This makes it impossible to explore even a small fraction of the parameter space exhaustively. Even human experts or file-system developers cannot know the exact impact of every parameter and thus have little insight into how to optimize them. For example, Ext4 + NFS alone would result in a parameter space consisting of more than $10^{22}$ unique configurations. IBM's General Parallel File System (GPFS) [129] contains more than 100 tunable parameters, and hence $10^{40}$ configurations. From the hardware perspective, which also constitutes part of parameter space, SSDs [64,108,115,130], SMRs [2,3,68,92], and PCM [81,163] are gaining popularity and more layers (LVM, RAID) are added to storage systems.

**(2) Discrete and non-numeric parameters** Among storage system parameters, some can take a continuous spectrum of values, while many others are discrete and take only a limited set of values. Some parameters do not even have numeric values (e.g., I/O scheduler name or file system type). These types of parameters make gradient-based information for objective functions (e.g., linear regression) unavailable.

**(3) Non-linearity** A system is *non-linear* when the output is not directly proportional to the input. Many computer systems are non-linear [32], including storage systems [146]. For example, Figure 2.1 shows the average operation latency of GPFS under a typical database server workload while changing only the value of the parameter *pagepool* and setting all the others to their default. We changed the *pagepool* size from 32MB to 128MB in steps of 8MB. Clearly the average latency is not directly proportional to the *pagepool* size. In fact, through our experiments, we have seen many more parameters with similar behavior. Worse, parameter spaces for storage systems are often sparse, irregular, and contains multiple peaks. This makes optimization even more challenging, as it has to avoid getting stuck in a local optima [74].

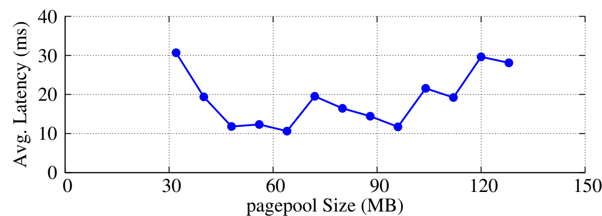


Figure 2.1: Storage systems are non-linear

**(4) Non-reusable results** Previous studies have shown that evaluation results of storage systems [24, 132] and databases [155] are dependent on the specific hardware and workloads. One good configuration might perform poorly when the environment changes. Figure 2.2 shows the I/O throughput under 4 different workloads with default configurations for Ext4, XFS, Btrfs, and Reiserfs—all on the same hardware. Under the *Mail Server* workload, the default XFS configuration performs best among these four configurations; but with the *Database Server*, Btrfs produces the highest throughput. In addition, these four configurations show similar results under the *Web Server* workload. We observed similar behavior when the hardware changed.

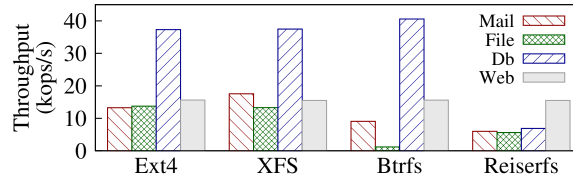


Figure 2.2: Evaluation results depend on workloads

Given these challenges, manual tuning of storage systems becomes nearly impossible while efficient automatic tuning is challenging. In this thesis we propose to design a practical auto-tuning framework for storage systems. We treat auto-tuning storage configurations as an optimization problem, and use the terms *"auto-tune"* and *"optimize"* interchangeably. Our framework is general enough to optimize for any user-specified objective, as long as possible outputs of the objective function form a totally ordered set. Examples of optimization objectives include maximizing throughput, minimizing average latency, minimizing energy consumption, etc. It can even be a complex formula combining several metrics together [97]. In this proposal we will mainly focus on auto-tuning storage systems for maximizing throughput, but our methodology and observations are applicable to other objectives as well.

Many previous efforts have been made and various techniques have been applied to parameter tuning problems. Control Theory (CT) was historically used to manage linear system parameters. CT builds a controller for a system, called the plant, so its output follows a desired control signal, called the reference [69,88]. CT has been applied to database systems [39] and storage systems [77, 89] to provide QoS guarantees. However, CT has been shown to have the following three problems: 1) CT tends to be unstable in controlling non-linear systems [95,96]. Although some variants were proposed for non-linear ones, they do not scale well. 2) CT cannot handle non-numeric parameters; and 3) CT requires an expensive learning phase, called *identification* to build a good controller, which requires having lots of data to learn from.

Supervised Machine Learning (ML) have been applied in black-box storage device modeling and prediction [158]. However, a well-known problem for supervised ML techniques is that they usually require a long training period and a large amount of data to build models; the models' quality depends heavily on the quality and amount of training data [158]. This data is not available or impossible to collect for large parameter spaces such as ours. Moreover, once the environment changes, the training data collected before it becomes invalid.

Based on the above reasons, we feel that neither CT nor supervised ML can be *directly and efficiently* applied for auto-tuning storage systems in its current state. Still, it was shown that many optimization techniques share some similarities with supervised Machine Learning [170]. Moreover, sub-disciplines of ML, including Online Learning [7, 137] and Active Learning [134], are evolving and gaining interests. They are practical and useful in solving certain problems where data becomes available incrementally. We believe ML techniques can still play an important role in our auto-tuning framework. Therefore, We provide a general introduction to ML in Chapter 2.3 and discuss how we plan to apply them in Chapter 7.

## 2.2 Black-box Optimization

Several classes of algorithms have been proposed for optimization tasks, including automated tuning of hyper-parameters of machine learning systems [14, 15, 119] and optimization of physical systems [5, 155]. Examples include Genetic Algorithms (GA) [35, 70], Simulated Annealing (SA) [27, 82], Bayesian Optimization (BO) [19, 136], etc. Although these methods were proposed originally in different scholarly fields, they can all be characterized as black-box optimizations. In this section we introduce several of these techniques that we successfully applied in auto-tuning storage systems.

**Simulated Annealing (SA)** is inspired by the annealing process in metallurgy. Annealing involves the heating and controlled cooling of a material to get to a state with minimum thermodynamic free energy to enhance, e.g., metal conductivity. When applied to storage systems, a *state* corresponds to one *configuration*. *Neighbors* of a state refer to new configurations achieved by altering only one parameter value of the current state. The thermodynamic free energy is analogous to user-defined optimization objectives. SA works by maintaining the *temperature* of the system, which determines the probability of accepting a certain move. Instead of always moving towards better states as hill-climbing methods do, SA defines an *acceptance probability distribution*, which allows it to accept some bad moves in the short run, that can lead to even-better moves later on. The system is initialized with a high temperature, and thus has high probability of accepting worse states in the beginning. The temperature is gradually reduced based on a pre-defined *cooling schedule*, thus reducing the probability of accepting bad states over time. SA has been applied in various areas and proved efficient in solving different types of problems, including the Traveling Salesman Problem (TSP) [1, 104, 156], Very Large Scale Integration (VLSI) design [131, 162], and network design [51, 52, 73].

**Genetic Algorithms (GA)** were proposed in 1975 [70] and inspired by the process of natural selection. GA maintains a population of *chromosomes* (configurations) and applies several genetic operators to them. *Crossover* takes two parent chromosomes and generates new ones. As Figure 2.3(a) illustrates, two parent Nilfs2 configurations are cut at the same *crossover* point, and then the subparts after the crossover point are exchanged between them to generate two new child configurations. Better chromosomes will have a higher probability to "survive" in future *selection* phases. *Mutation* randomly picks a chromosome and mutates one or more parameter values, which produces a completely different chromosome. Figure 2.3(b) illustrates such mutation, where the journal option is randomly mutated from *writeback* to *journal*. GA and its variants have been widely applied to various areas including the Traveling Salesman Problem (TSP) [59,62,86,113,122,140], VLSI Design [16,33,99,105], High-Performance Computing [11], and system design [30, 36, 101].

**Bayesian Optimization (BO)** [19,136] is a popular framework to solve optimization problems. It models the objective function as a stochastic process, with the argument corresponding to one storage configuration. In the beginning, a set of prior points (configurations) are given to the algorithm to get a fair estimate of the entire parameter space. BO works by computing the *confidence interval* of the objective function according to previous evaluation results. Here the *confidence interval* is the range of values that the evaluation result is most likely to fall into (e.g., with 95% probability). The next configuration is selected based on a pre-defined *acquisition function*. Both confidence intervals and the acquisition function are updated with each new evaluation. BO has been successfully applied in various areas, including hyper-parameter optimization [34] and sys-

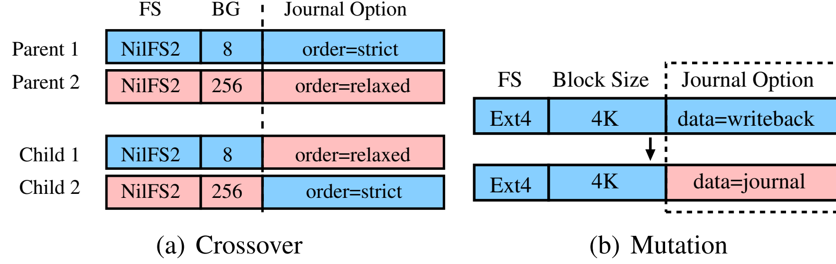


(a) Crossover (b) Mutation

Figure 2.3: Crossover and mutation in a Genetic Algorithm

tem configuration optimization [5]. BO and its variants differ mainly in their form of probabilistic models and acquisition functions. In this thesis proposal our evaluation results focus mainly on Gaussian priors and an Expected Improvement acquisition function [136].

Other promising black-box optimization techniques include Tabu Search [56–58], Particle Swarm Optimization [31, 79, 80], Ant Colony Optimization [40, 41], and Memetic Algorithms [84, 107], etc. Most of them are nature-inspired as they have been developed based on the successful evolutionary behavior of natural systems. In the current stage of our project, we focused on several representative algorithms, SA, GA, and BO. We plan to experiment with more techniques in the future (part of our future work). In fact, as detailed in §2.4, most of these techniques actually share similar traits.

## 2.3 Machine Learning

As we enter the era of big data, Machine Learning (ML) has becoming more popular in the last few decades. We can define ML as a set of methods that can automatically detect patterns in data, and then use the discovered patterns to predict future behavior, or to perform other kinds of decision making under uncertainty [114]. Generally, there are three types of ML techniques: *Supervised Learning*, *Unsupervised Learning*, and *Reinforcement Learning*.

**Supervised Learning** Supervised Learning is sometimes also called *predictive learning*, and its goal is to learn the mapping from the inputs $\overrightarrow{x}$ to outputs $y$, based on a labeled set of input-output pairs $D = \{(\overrightarrow{x_i}, y_i)\}_{i=1}^{N}$. $D$ is often referred as a *training set* consisting of $N$ training examples. In the training set, each input $\overrightarrow{x_i}$ is usually a multi-dimensional vector, and the elements in the vector are called *features* or *attributes*. Depending on whether the output $y$ is *categorical* or *real-valued*, supervised learning can be further classified into two categories, *classification* and *regression*.

**Unsupervised Learning** Unsupervised Learning (or *descriptive learning*) is another main type of Machine Learning, where the dataset is unlabeled: $D = \{(\overrightarrow{x_i})\}_{i=1}^{N}$. The goal of Unsupervised Learning is often to find certain patterns existing on the dataset; that is why it is also called *knowledge discovery*. Unsupervised Learning is arguably more typical of human and animal learning behaviors. It is also more widely applicable than supervised learning, since it does not require a human expert to manually label the data [114]. Approaches of Unsupervised Learning include *clustering*, *Latent Variable Modeling*, etc.

**Reinforcement Learning** Reinforcement Learning (RL) [143] is an area of machine learning inspired by behaviorist psychology. RL explores how software agents take actions in an environment to maximize the defined cumulative rewards. Most RL algorithms can be formulated as a model consisting of: (1) A set of environment states; (2) A set of agent actions; and (3) A set of scalar rewards. In case of storage systems, *states* correspond to *configurations*, *actions* mean changing to a different configuration, and *rewards* are differences in evaluation results. The agent records its previous experience (history), and makes it available through a *value function*, which can be used to predict the expected reward of state-action pairs. The *policy* determines how the agent takes action. A simple example is $\epsilon$-policy. For each action the agent may can take a random action with probability $\epsilon$; otherwise it will exploit the current value function and take the best action to maximize the rewards. The value function's history can be stored in a tabular form, but this does not scale well to many dimensions. Function approximation is one way for generalization when the state and/or action spaces are large or continuous. However, most approximation methods are still known to be unstable or even divergent. With recent advances in Deep Learning [61], deep convolutional neural networks, termed Deep Q-Networks (DQN), were proposed to parameterize the value function, and have been successfully applied in solving various problems [110, 111]. Many variants of DQN have been proposed [93]; in this proposal we applied its original version [111]. Another interesting fact here is that many RL algorithms, including DQN, also maintains a trade-off between exploitation, exploration, and history. In the early stages of execution, when the agent knows little about the environment, it will explore the space and try unknown actions. When it interacts enough with the environment, it will tend to choose the actions that it knows will receive the higher rewards.

## 2.4 Unified Framework

| Algorithm | Origin | Exploration | Exploitation | History |
|---|---|---|---|---|
| **Simulated Annealing (SA)** | Annealing technology in metallurgy | Allowing moving to worse neighbor states | Neighbor function | N/A |
| **Genetic Algorithms (GA)** | Natural evolution | Mutation | Crossover and selection | Current population |
| **Deep Q-Networks (DQN)** | Behaviorist psychology and neuroscience | Taking random actions | Taking actions based on action-reward function | Deep convolutional neural network |
| **Bayesian Optimization (BO)** | Statistics and experimental design | Selecting samples with high variances | Selecting samples with high mean values | Acquisition function & probabilistic model |

Table 2.1: Comparison and summaries of optimization techniques

Most optimization techniques are known to follow the *exploration-exploitation* dilemma [45, 93, 136, 157]. Here we summarize the aforementioned methods by extending the unified framework

with a third factor, the *history*. Our unified view thus defines three factors or dimensions:

- **Exploration** defines how the technique searches unvisited areas. This often includes a combination of pure random and also guided search.
- **Exploitation** defines how the technique leverages current neighborhood or *history* to find next sample.
- **History** defines how much data from previous evaluations is kept. History information can be used to help guide both future exploration and exploitation (e.g., avoiding less promising regions, or selecting regions that have never been explored before).

Table 2.1 summarizes how the aforementioned techniques work by maintaining the balance among these three key factors. For example, GA keeps the evaluation results from the last generation, which corresponds to the concept of *history* in our unified framework. GA then *exploits* the stored information, applying selection and crossover to search nearby areas and pick the next generation. Occasionally, it also randomly mutates some chosen parameters, which is the idea of *exploration*. The trade-off among exploration, exploitation, and history largely determines the effectiveness and efficiency of these optimization techniques.

# Chapter 3

# Related Work

This chapter describes related previous work and compare them with our project.

## 3.1 Auto-tuning in Computer Systems

In recent years, several attempts were made to automate the tuning of storage systems. Gaonkar et al. [53] apply GAs to design dependable data storage systems for multi-application environments, with the goal of minimizing the overall cost of the system while meeting business requirements. Strunk et al. [141] proposed to use utility functions combining different system metrics and applied GA to automate storage system provisioning. Babak et al. [12] utilized GA to optimize I/O performance of HDF5 applications. Kimberly et al. [78] formulate the data recovery scheduling problem as an optimization problem. They aim at finding the schedule that minimizes the financial penalties due to downtime, data loss, and vulnerability to subsequent failures. GAs are applied and compared with several other heuristics. Xue et al. [166, 167] propose an autonomic technique that learns the intensity patterns of user workload in tiered storage systems over long time-scales using a probabilistic model. They use the model to predict the coming workload patterns and proactively stop/start bulky internal system work. MINERVA [6], is a suite of tools for automating storage system design, which uses declarative specifications of application requirements and device capabilities; constraint-based formulations of the various sub-problems; and simple bin-packing heuristics to explore the search space of possible solutions. More recently, Deep Q-Networks has been successfully applied in optimizing performance for Lustre [168].

Auto-tuning is also a hot topic in other computer systems: Bayesian Optimization was applied to find near-optimal configurations for databases [155] and Cloud VMs [5]. Other applied techniques include Evolutionary Strategies [126], Simulated Annealing [51, 73], Tabu Search [127], and more.

However, previous work all focused on a single algorithm or technique. One contribution of our work is to provide the first comparative study of multiple, applicable optimization methods and compare them for their efficacy in auto-tuning storage systems from various aspects. We also provide some insights into the working mechanism of auto-tuning. More importantly, we propose to design a more intelligent and practical framework for auto-tuning storage systems in real-time.

## 3.2 Hyper-parameter tuning

Esteban et al. [119] applied Evolutionary Algorithms to hyper-parameter optimization for neural networks, and achieved state-of-art results on certain data-sets. Bergstra and Bengio [15] found that randomly chosen trials are more efficient for hyper-parameter optimization than trials on a grid, and explained the cause as the objective function having a low effective dimensionality. In addition, Reinforcement Learning [13] and Bayesian Optimization [44] were also applied to hyper-parameter optimization. Another direction of research focuses on eliminating all hyper-parameters and tries to propose non-parametric versions of optimization methods. Examples of this include GA [66, 100] and BO [136]

In this work, we will investigate the impacts of hyper-parameters on various optimization techniques, when applied to auto-tune storage systems.

## 3.3 Workload Modeling

A few efforts have been made on modeling or characterizing storage workloads. Bumjoon et al. [133] tried to model storage workloads on HDDs from a data-mining point of view. They use a unique clustering method for feature selection that reduces computational time on a list of 20 features available through blktrace and use a hierarchy of clustering and classification to label a workload based on access patterns. Busch et al. [21] proposed to design an automated approach for extracting workload models in virtualized environments. Features used include average file size, file set size, average request size, etc. Li et al. [90] attempted to better define sequential I/O. They focused on LBA and I/O size, and concluded that "consecutive bytes accessed" should be taken into consideration. Shen et al. [138] characterized workloads with the goal of improving performance debugging by separating their model into OS caching, prefetching, OS I/O Scheduling, and storage devices. Wang et al. [158] used CART models to predict per-request response time based on workload characteristics, and provides detailed explanations about how CART models work and why they are suitable for this problem. Riska et al. [123] tried to characterize workloads based on their environment: enterprise, desktop, or consumer electronics.

We feel that most previous work were either vague on what features to pick for characterizing workload, or they limited the model built to one or few use cases. In this work, we target at finding the minimum set, or a small-enough set of features (out of many), which is general and can characterize most storage workload. The feature engineering work will utilize cutting-edge ML and data mining techniques, but we will explain out observations from storage perspective as well.

# Chapter 4

# Experimental Settings

In this chapter we detail the experimental environments, parameter spaces, and our implementations of several optimization algorithms.

## 4.1 Hardware

We performed experiments on two sets of machines with different hardware categorized as low-end (*S1*) and mid-range (*S2*). We list the details of these two sets of machines in Table 4.1. We also use Watts Up Pro ES power meters to measure the energy consumption. During our experiments on characterizing storage performance variation (Chapter 6), to maintain realistically high ratio of the dataset size to the RAM size and ensure that our experiments produce enough I/O, we limited the RAM size on all machines to 4GB. We denote this hardware setting as *S3*. We have one type of storage device on S1 and four others on S2 and S3, which will be denoted as *HDD1*, *HDD2*, *HDD3*, *HDD4*, and *SSD* for short in this proposal.

## 4.2 Workload

We used *Filebench* [50,149] to generate various workloads in our experiments. In each experiment, if not stated otherwise, we formatted and mounted the storage devices with a file system and then ran Filebench. We mainly experimented with the four pre-configured Filebench macro-workloads that exhibit the following significantly different I/O properties:

- **Mailserver** emulates the I/O workload of a multi-threaded email server. It generates sequences of I/O operations that mimic the behavior of reading emails (open, read the whole file, and close), composing emails (open/create, append, close, and fsync) and deleting emails. It uses a flat directory structure with all the files in a single directory, and thus exercises the ability of file systems to support large directories and fast lookups.

- **Fileserver** emulates the I/O workload of a server that hosts users' home directories. Here, each thread represents a user, which performs create, delete, append, read, write, and stat operations on a unique set of files. It exercises both the metadata and data paths of the targeted file system.

| Setting | S1 | S2 | S3 |
|---|---|---|---|
| **Model** | Dell PowerEdge SC1425 | Dell PowerEdge R710 | Dell PowerEdge R710 |
| **CPU** | Intel Xeon single-core 2.8GHz CPU × 2 | Intel Xeon quad-core 2.4GHz CPU × 2 | Intel Xeon quad-core 2.4GHz CPU × 2 |
| **Memory** | 2GB | 24GB | 4GB (set by *mem=* in */etc/default/grub*) |
| **Storage** | HDD1 (73GB Seagate ST373207LW SCSI drive) × 2 | HDD2 (146GB Seagate ST9146853SS SAS HDD), HDD3 (500GB Seagate ST9500430SS SAS HDD), HDD4 (200GB Intel SSDSC2BA200G3 SATA HDD), SSD (250GB Fujitsu MHZ2250BKG2 SATA HDD) | HDD2 (146GB Seagate ST9146853SS SAS HDD), HDD3 (500GB Seagate ST9500430SS SAS HDD), HDD4 (200GB Intel SSDSC2BA200G3 SATA HDD), SSD (250GB Fujitsu MHZ2250BKG2 SATA HDD) |
| **Partition** | 100GB | 100GB | Full size |
| **OS** | Ubuntu 14.04 with kernel 3.13 | Ubuntu 14.04 with kernel 3.13 | Ubuntu 14.04 with kernel 4.4 |

Table 4.1: Details of experiment machines.

- **Webserver** emulates the I/O workload of a typical static Web server with a high percentage of reads. Files (Web pages) are read sequentially by multiple threads (users); each thread appends to a common log file (Web log). This workload exercises fast lookups, sequential reads of small files and concurrent data and metadata management.
- **Dbserver** mimics the behaviors of Online Transaction Processing (OLTP) databases. It mainly consists of random asynchronous writes, random asynchronous reads and moderate synchronous writes to the log file. It exercises the ability of large file management, extensive concurrency, and random read/write operations.

| Workload | Avg. File Size | Avg. Dir Width | # Files | Running Time (s) | Num. of Threads | R/W Ratio | Filebench Version |
|---|---|---|---|---|---|---|---|
| fileserver-def | 128KB | 20 | 10,000 | 100 | 50 | 1:2 | 1.4.9 |
| mailserver-def | 16KB | 1,000 | 1,000 | 100 | 16 | 1:1 | 1.4.9 |
| webserver-def | 16KB | 20 | 1,000 | 100 | 100 | 10:1 | 1.4.9 |
| dbserver-def | 10MB | 1,024 | 10 | 100 | 10 + 1 | 10:1 | 1.4.9 |
| fileserver-heavy | 128KB | 20 | 80,000 | 800 | 50 | 1:2 | 1.5.0 |
| mailserver-heavy | 16KB | 1,000,000 | 640,000 | 2,000 | 16 | 1:1 | 1.5.0 |
| webserver-heavy | 16KB | 20 | 640,000 | 800 | 100 | 10:1 | 1.5.0 |

Table 4.2: Filebench workload characteristics.

Table 4.2 shows the detailed settings of our workloads. The first four workloads (named as **-def*) are used in our auto-tuning experiments, while the last three were mainly applied in the performance variation study.

**Auto-tuning Experiments** It is well known that the working set size has a significant impact on the duration of an experiment [146]. In our auto-tuning experiments, the goal was to explore a large set of parameters and values quickly (though it still took us over two years to search some spaces exhaustively). We therefore decided to trade the working set size in favor of increasing the number of configurations we could explore in a practical time period. We mainly experimented with the default settings provided by Filebench. We did not perform a separate cache warm-up phase, since performance usually become relatively stable within a short time given the default dataset size.

**Performance Variation Study** For studying performance variations, nearly all workload characteristics were set to Filebench's default values, except for the number of files and the running time. As the average file size is an inherent property of a workload and should not be changed [149], the dataset size is determined by the number of files. We increased the number of files such that the dataset size is 10GB—or 2.5× the machine RAM size (*S3* in Table 4.1). By fixing the dataset size, we normalized the experiments' set-size and run-time, and ensured that the experiments run long enough to produce enough I/O. With these settings, our experiments exercise both in-memory cache and persistent storage devices [147]. We did not perform a separate cache warm-up phase in our experiments because in this study we were interested in performance variation that occurred *both* with cold and warm caches [147]. The default running time for Filebench is too short to warm the cache up. We therefore conducted a calibration phase to pick a running time that was long enough for the cumulative throughput to stabilize. We ran each workload for up to two hours for testing purposes, and finally picked the running time as shown in Table 4.2. We also let Filebench output the throughput (and other performance metrics) every 10 seconds, to capture and analyze performance variation at a finer time granularity.

## 4.3 Parameter Space

To test the efficacy of auto-tuning algorithms, ideally we wanted our storage parameter spaces to be large and complex enough. Alas, evaluations for storage systems take a long time. Considering experimentation on multiple hardware settings and workloads, we decided to experiment with a reasonable subset of the most relevant storage system parameters. We selected parameters in close collaboration with several storage experts that have either contributed to storage system designs or have spent years tuning storage systems in the field. We experimented with 7 Linux file systems that span a wide range of designs and features: *Ext2* [25], *Ext3* [153], *Ext4* [47], *XFS* [144], *Btrfs* [124], *Nilfs2* [83], and *Reiserfs* [121].

Our experiments were mainly conducted on two sets of parameters, termed as *Storage V1* and *Storage V2*. We started with a relatively smaller set of 7 parameters, and refer it as *Storage V1*. It contains the following common file system parameters: *file system type*, *block size*, *inode size*, *blocks per group*, *mount options*, *journal options*, and *special options*. We tested *Storage V1* with *Setting S1*. After some preliminary experiments, we extended our search space with one

| Param. | Abbr. | Values |
|---|---|---|
| File System | FS | Ext2, Ext3, Ext4, XFS, Btrfs, Nilfs2, Reiserfs |
| Block Size, Leaf Size | BS | 1K, 2K, 4K |
| Inode Size, Sector Size | IS | n/a, 128, 256, 512, 1024, 2048, 4096, 8192 |
| Block Group, Alloc. Group | BG | n/a, 2, 4, 8, 16, 32, 64, 128, 256 |
| Journal Option | JO | n/a, order=strict, order=relaxed, data=journal, data=ordered, data=writeback |
| Atime Option | AO | relatime, noatime |
| Special Option | SO | n/a, compress, nodatacow, nodatasum, notail |
| I/O Scheduler | I/O | noop, cfq, deadline |

Table 4.3: Details of Parameter Spaces

more parameter, the *I/O Scheduler*, and refer it as *Storage V2*. Experiments with *Storage V2* were conducted with *Setting S2*. We list all the aforementioned parameters and their values in Table 4.3. Note that certain combinations of parameter values could produce invalid configurations. For example, for Ext2, the journaling options make no sense because Ext2 does not have a journal. To handle this, we added a value *n/a* to the existing range of parameters. Any parameter with *n/a* value is considered invalid. Invalid configurations will always come with evaluation results of zero (i.e., no throughput); this ensures they are purged in an upcoming optimization process. There are 2,074 valid configurations in *Storage V1* and 6,222 in *Storage V2*.

## 4.4 Experiments and Implementations

Our experiments and implementation consist of two parts. First, we exhaustively ran all configurations for each workload on the S1 and S2 machines, and stored the results in a relational database. We collected the throughput in terms of I/O operations per second, as reported by Filebench, the running time (including setup time), as well as power and energy consumption. To acquire more accurate and stable results, we evaluated each configuration under the same environment for at least 3 runs, resulting in more than 450,000 total experimental runs. This data collection benefited our evaluation on auto-tuning as we can simply simulate a variety of algorithms by just querying the database for the evaluation results for different configurations, without having to rerun slow I/O experiments. The exhaustive search also let us know exactly what the global optimal configurations are, so that we can better understand how each optimization method performs.

Second, we simulated the process of auto-tuning storage systems by running the desired optimization method and querying the database for the evaluation results of the targeted storage configurations. We focused on optimizing for throughput in this proposal. Our implementations of optimization methods are mostly based on open-source publicly-available libraries. We use Pyevolve [118] for Genetic Algorithms, Scikit-Optimize [139] for Bayesian Optimization, and TensorFlow [150] for the DQN implementation. We implemented a simple version of Simulated Annealing, with both linear and geometric cooling schedules. (We also fixed bugs in Pyevolve and plan to release our patches.) Most of our implementation was done by applying storage-related concepts into algorithm-specific ones. For example, for GA, we defined each storage parameter as a *gene*, and each configuration as a *chromosome*. For DQN we provided storage-specific defi-

nitions for states, actions, and rewards. The complete implementation uses around 8,000 lines of code, consisting of Python and Shell scripts.

# Chapter 5

# Towards Better Understanding of Black-box Auto-Tuning: A Comparative Analysis for Storage Systems

In this chapter we apply several popular techniques to the collected dataset to find optimal configurations under various hardware and workload settings: Simulated Annealing (SA), Genetic Algorithms (GA), Bayesian Optimization (BO), and Deep Q-Networks (DQN). We also tried Random Search (RS) in our experiments, which showed surprisingly good results in previous research [15]. We compared these techniques from various aspects, such as the ability to find near-optimal configurations, convergence time, and instantaneous system throughput during auto-tuning. We also showed that hyper-parameter settings of these optimization algorithms, such as mutation rate in GA, could affect the tuning results. We compared the techniques across three behavioral dimensions: (1) *Exploration*: how much the technique searches the space randomly. (2) *Exploitation*: how much the technique leverages the "neighborhood" of the current candidate or previous search history to find even better configurations. (3) *History*: how much data from previous evaluations is kept and utilized in the overall search process. Based on our evaluation results, we show that all techniques employ these three key concepts to varying degrees and the trade-off among them plays an important role in the effectiveness and efficiency of the algorithms.

Most black-box optimization methods lack solid theoretical understanding, partially due to the large variety of problems that they were proposed to solve [170]. Based on our experimental results and domain expertise, we provide explanations of efficacy of such black-box optimization methods from a storage perspective. We observed that certain parameters would have a greater effect on system performance than others, and the set of dominant parameters depends on file systems and workloads. This allows us to provide more insights into the auto-tuning process.

Part of the results from this chapter will be published in ATC 2018.

The chapter is organized as follows. Chapter 5.1 overviews the datasets that we collected for over two years. Chapter 5.2 compares five popular optimization techniques from several aspects. Chapter 5.3 uses GA as a case study to show that hyper-parameters of these methods could also impact the auto-tuning results.

## 5.1 Overview of Datasets

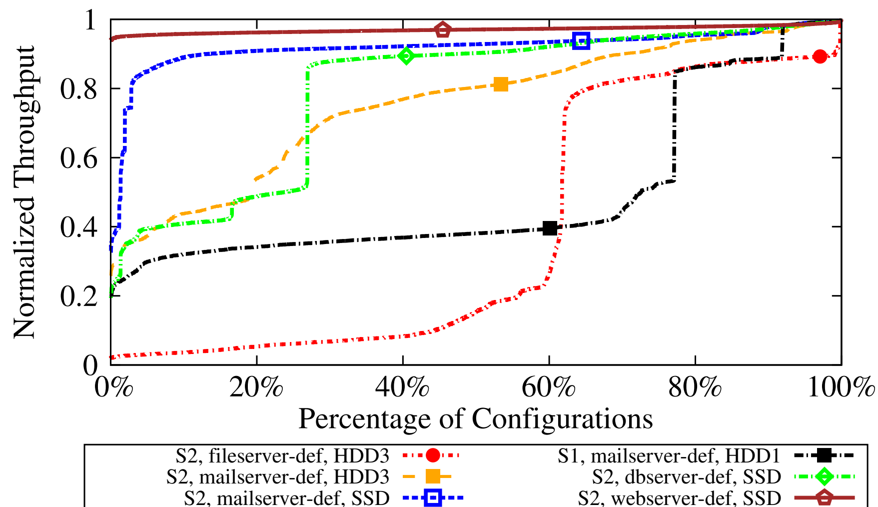


Figure 5.1: Throughput CDF with different hardware and workloads, with symbols marking the default configurations.

As per Chapter 4, our experimental methodology is to first exhaustively run all configurations under different workloads and test machines. We stored the results in a database for future use. This data collection benefits future experiments as we can simulate a variety of algorithms by querying the database for the evaluation results of different configurations.

Figure 5.1 shows the throughput CDF among all configurations for each hardware setting and workload. Due to space limits, we show only 6 representative datasets out of 18 here. The Y axis is normalized by the maximum throughput under each experiment setting. The symbols on each line mark the default configurations. As seen, for most settings, throughput values vary across a wide range. The ratios of the worst throughput to the best one are mostly between 0.2–0.4. In one extreme case, for *fileserver-def* on *S1* machines and with HDD1 device, the worst configuration only produces $1\%$ I/O operations per unit time, compared with the global optimal one. This underlines the importance of tuning storage systems: an improperly configured system could be remarkably under-utilized, and thus wasting a lot of resources. However, *S2, webserver-def, SSD* shows a much narrower range of throughput, with the worst-to-best ratio close to 0.9. This is attributed mainly to the fact that *webserver-def* consists of mostly sequential read operations that are processed similarly by different I/O stack configurations. Another useful observation from Figure 5.1 is that default configurations are always sub-optimal and, under most settings, ranked lower than the top 40% configurations. For *S1, fileserver-def, HDD1*, the default configuration shows a normalized throughput of 0.39, which means that the optimal configuration performs 2.5 times better.

We list the optimal configurations for each hardware setting and workload from our datasets in Table 5.1. As we can see, optimal configurations depend on the specific hardware as well as the running workload. For *mailserver-def* with *S1* machines and the *HDD1*, the global best is a *Nilfs2* configuration. However, if we fix the workload and change the hardware to *S2-HDD3*, the optimum becomes an *Ext4* configuration. Similarly, fixing the hardware to *S2-SSD* and experimenting under

| Hardware Workload-Device | File System | Block Size | Inode Size | BG Count | Journal Options | Atime Options | Special Options | I/O Scheduler | Through-put (IOPS) |
|---|---|---|---|---|---|---|---|---|---|
| **S1-Mail-HDD1** | Nilfs2 | 2K | n/a | 256 | order=relaxed | relatime | n/a | - | 3,677 |
| **S2-Mail-HDD3** | Ext2 | 4K | 256 | 32 | n/a | relatime | n/a | noop | 18,744 |
| **S2-Mail-SSD** | Ext2 | 4K | 256 | 8 | n/a | relatime | n/a | noop | 18,845 |
| **S2-File-SSD** | Btrfs | 4K | 4,096 | n/a | n/a | relatime | nodatacow | deadline | 16,587 |
| **S2-DB-SSD** | Ext4 | 1K | 128 | 2 | data=ordered | noatime | n/a | noop | 41,948 |
| **S2-Web-SSD** | Ext4 | 4K | 128 | 4 | data=ordered | noatime | n/a | noop | 16,185 |

Table 5.1: Global optimal configurations with different settings and workloads. Workloads are abbreviated. Db: dbserver-def; File: fileserver-def; Mail: mailserver-def; Web: webserver-def.

different workloads leads to different optimal configurations. This proves our early claim that performance (and other metrics) are sensitive to the environment (i.e., hardware, configuration, and workloads); this actually complicates the problem as results from one environment cannot be directly applied in another.

It is known that the working set size has a significant impact on the duration of an experiment [146]. Our goal in this study was to explore a large set of parameters and values quickly (though it still took us over two years). We therefore decided to trade the working set size in favor of increasing the number of configurations we could explore in a practical time period. In our experimental results, this trade-off sometimes manifests itself since SSD configurations produce comparable throughputs as HDD ones (see Table 5.1). The experiments, however, do demonstrate a wide range of performance numbers and, therefore, are valid for evaluating different optimization methods. We plan to include the working set size in the set of optimization parameters in the future.

## 5.2 Comparative Analysis

Many optimization techniques have been applied to various auto-tuning tasks [141, 155]. However, previous efforts picked algorithms somewhat arbitrarily and evaluated only one algorithm at a time. Here we provide the first comparative study of multiple black-box optimization techniques on auto-tuning storage systems. As discussed in §2.2, we focus our evaluations on a representative set of optimization methods, and their common hyper-parameter settings, including 1) Simulated Annealing (SA), with a linear cooling schedule; 2) Genetic Algorithms (GAs) with population size of 8, mutation rate of 2%; 3) Deep Q-Networks (DQN) with experience replay [111] and $\epsilon = 0.2$; and 4) Bayesian Optimization (BO) with Expected Improvement (EI) and Gaussian prior. 5) Random Search (RS), which merely performs random selection without replacement. We provide more discussion on the impact of hyper-parameters in Chapter 5.3. Note that SA, DQN, and RS experiments start with the default Ext4 configuration. GA and BO require several initial configurations (*prior points*), which we set to default configurations of all seven file systems. This allows us to simulate real-world use cases, where users often deploy their system with the default settings (and may manually optimize starting from the defaults).

Figure 5.2 presents one simulated run of each optimization method on *S2, mailserver-def, HDD3*; the Y axis shows the throughput value of the best configuration found so far, and the

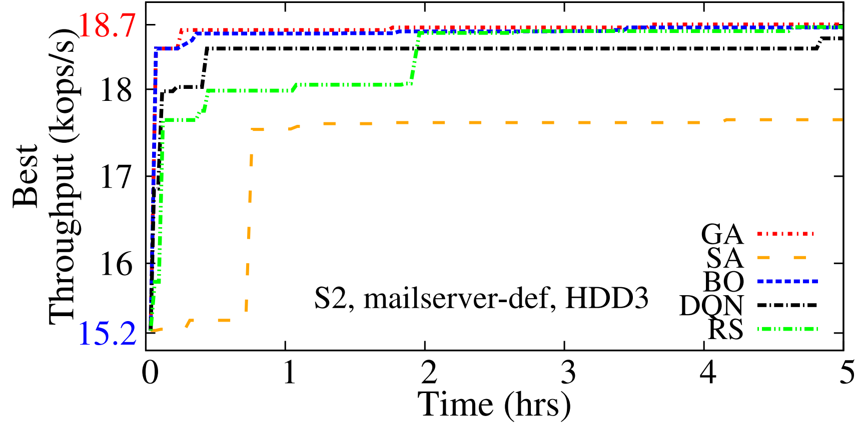


Figure 5.2: Highest throughput found over time, zooming in the $Y \in [15 : 19]$ range. The blue number (15.2) on the Y axis shows the default, and the red one (18.7) shows the optimal.

X axis is the running time. All time-related metrics in this chapter are based on the actual running time of evaluating each storage configuration, which is stored in our database. This includes both setup time and benchmarking time. We are not comparing the running costs (including any necessary training phases) for optimization methods here, which is our future work. Figure 5.2 is plotted by zooming in the range of $Y \in [15 : 19]$, with the blue number (15.2) on Y axis represents the default, while the red one (18.7) shows the global optimal. It shows that all five methods were able to gradually find better configurations, but their effectiveness and efficiency differed a lot. SA performed the worst, and got stuck in a configuration with throughput value of less than 18K IOps. DQN was able to converge to a good configuration, but spent more time to achieve that than RS. GA and BO performed best out of these five tested optimization methods. They both successfully identified a near-optimal configuration within one hour. Interestingly, we observed that pure Random Search (RS) produced better results than some other optimization methods. This is because not all storage parameters have significant impact on system performance, resulting in an *effective* search space that is much smaller than the original one. Similar results were observed in hyper-parameter optimization for neural networks [14]. We discuss this further in §5.4.

Since exploration is one critical component of all optimization methods (see §2.4), their evaluation results could also exhibit some degree of randomness. To compare them more thoroughly, we ran each optimization technique on the same environment (*S2, HDD3*) for 1,000 runs. Figure 5.3 shows the results, which evaluate the techniques' probability to find good and near-optimal configurations. Here we define a near-optimal configuration as one with throughput higher than **99%** of the global optimal value. The Y axis shows the percentage of total runs that found a near-optimal configuration within a certain time (X axis). Under *mailserver-def* workload, seen in the upper part of Figure 5.3, SA had the lowest probability among 5 algorithms.

Even after 5 hours, only around 80% of its runs found one near-optimal configuration, which suggests that SA can sometimes get stuck in a local optima. For other optimization methods, given enough time, over 90% of their runs converged to a near-optimal configuration, with BO outperforming GA, and GA outperforming DQN. RS shows the highest probability of finding near-optimal configurations when approaching 5 hours. This is reasonable because given enough time, a random selection will eventually hit near-optimal points. However, when conducting the

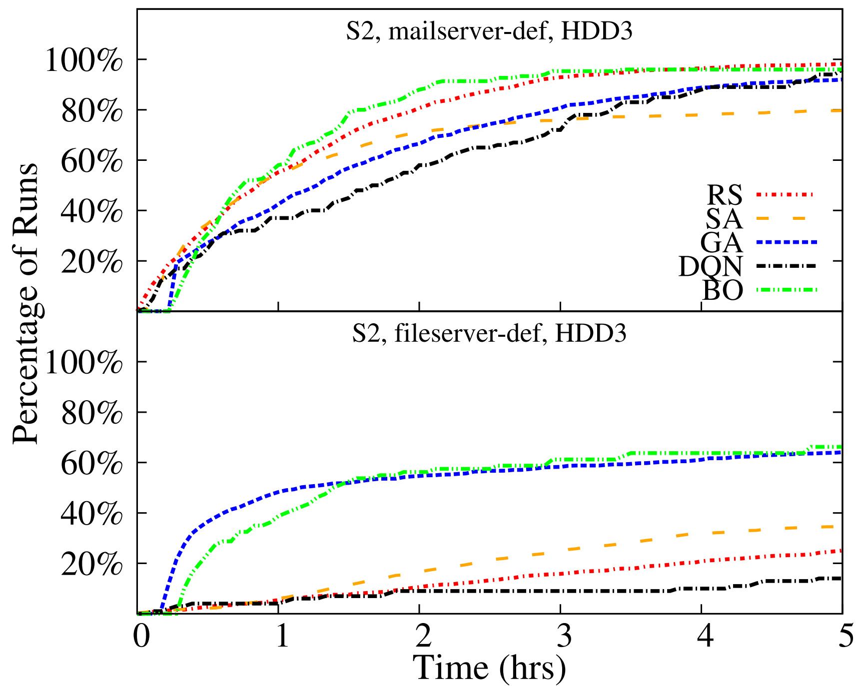


Figure 5.3: Comparing optimization methods' efficacy in finding near-optimal configurations. The Y axis shows the percentage of total runs (1,000) that found near-optimal configurations within certain time (X axis).

same experiments under the *fileserver-def* workload, it becomes more difficult to find near-optimal configurations. GA and BO are still the best, though only 65% of their runs were able to find near-optimal configurations within 5 hours. SA, RS, and DQN have a probability of lower than 40% to do so, with DQN perform the worst. This is because the global optimum under *fileserver-def* is a Btrfs configuration (see Table 5.1). It is more difficult for optimization algorithms to pick such configurations for the following reasons: 1) Few Btrfs configurations reside in the neighborhood of the default Ext4 configurations; 2) Fewer than 2 % of all valid configurations are Btrfs ones, which make them less likely to be selected through mutation.

The above results all focused on finding near-optimal configurations. However, another important aspect to compare is the system's performance *during* the auto-tuning process. This is especially important if the targeted system is deployed and online. Some randomness (exploration) is necessary when searching a complex parameter space, but ideally optimization algorithms should spend less time on bad configurations. To compare this, in Figure 5.4 we plotted the instantaneous throughput (Y axis) over time (X axis) for one run with each method under *S2, mailserver-def, HDD3*.

BO and GA are still the best two methods in terms of instantaneous throughput.

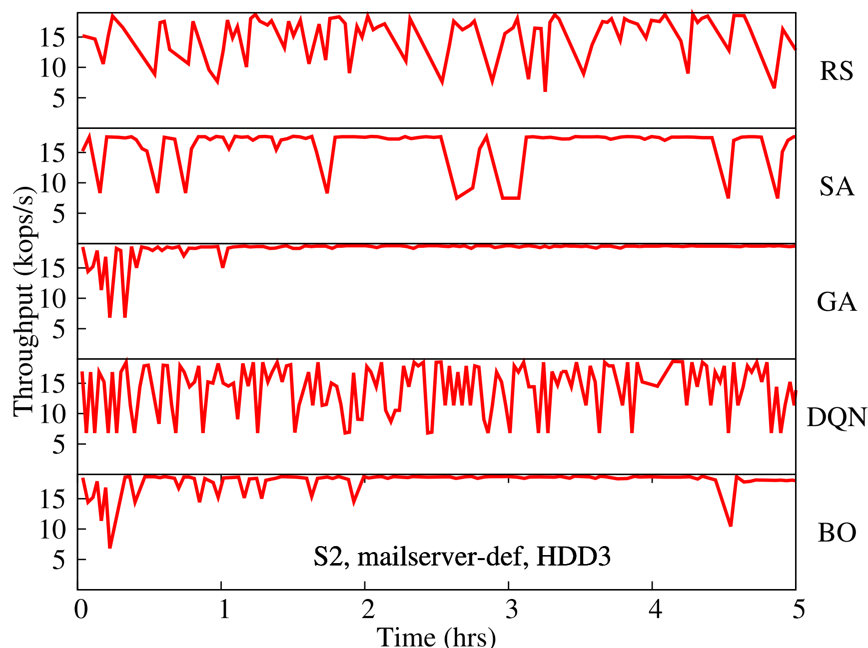


Figure 5.4: Comparing optimization methods' instantaneous performance (Y axis) over time (X axis).

During the tuning process, occasionally they will pick a worse configurations than the current one. However, they both possess the ability to quickly discard these unpromising configurations. GA achieves this by assigning the probability of surviving to next generation based on the fitness values (i.e., throughput). Configurations with low throughput values have a lower chance to be picked as parents, and thus their genes (parameter values) have a lower chance of appearing in configurations of the next generation (i.e., "survival of the fittest"). The reason for stable instantaneous throughputs with BO is that it uses an intelligent acquisition function to guide the selection

of the next generation, with the goal of maximizing the potential gain; this makes BO less likely to choose a bad configuration. In contrast, SA performs poorly possibly because it lacks a history to guide the exploitation and exploration phases, and only uses its neighborhood information (and current temperature) to pick the next configuration. DQN shows similar results with RS, which is likely caused by the fact that DQN was originally designed as an agent interacting with an unknown environment, and thus a lot of exploration (randomness) occurs in the training phase [111, 168].

In conclusion, BO and GA perform best among the 5 tested methods, on either the ability to converge to near-optimal configurations or in maintaining stable instantaneous performance during the tuning process. DQN and SA can find good configurations, although they were less efficient and less stable. Surprisingly, Random Search sometimes can produce better results than some traditional optimization methods, given enough time. We provide more explanations on these methods in Chapter 5.4.

## 5.3 Impact of Hyper-Parameters

Many optimization methods' efficacy depend on the specific hyper-parameter settings, and choosing the right hyper-parameters has caused headache to researchers for a long time [14, 15]. In this section we use GA as a case study, and show the impact of one hyper-parameter, the *mutation rate*, on auto-tuning results.

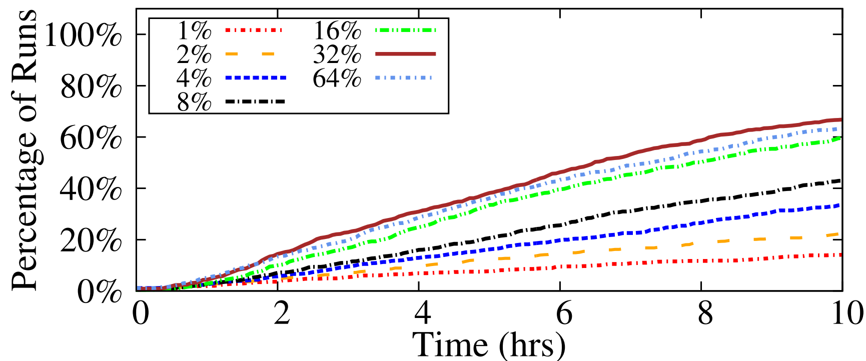


Figure 5.5: Impact of mutation rates on GA.

The mutation rate controls the probability of randomly mutating one parameter to a different value, and aligns with the idea of *exploration*, as per §2.4.

Figure 5.5 shows the results from 7 sets of GA experiments with different mutation rates (from 1% to 64%) under *S2, mailserver-def, HDD3*. Each experiment was repeated for 1,000 runs.

It is similar to Figure 5.3, but with the goal of finding near-optimal configurations whose throughput values are higher than **99.5%** of the global optimal. This makes the optimization more challenging, as GA already performs quite well on easier tasks (Chapter 5.2). As shown in the figure, when increasing the mutation rate, GA has a higher probability to converge to near-optimal configurations within a shorter time period. This is because GA works by identifying promising combination of alleles (parameter values) for the subset of *effective* genes (parameters). We define effective parameters as those having a higher impact on performance than all others. A higher mutation rate means a higher chances of exploration, and thus finding combinations of effective

alleles within a shorter time. We explain this effect more in Chapter 5.4. However, a mutation rate of 64% actually performs worse than 32%. This is because in order to reach near-optimal configurations, GA needs both exploration and exploitation. Exploration lets GA identify processing subspaces (i.e., combinations of certain parameter values) while exploitation helps GA search within promising subspaces. In this case, with a mutation rate of 64%, GA spends too much time on exploration (too much randomness), resulting in fewer chances for exploitation.

Note that in this section we are only using GA mutation rates as an example showing the impact of hyper-parameters on the efficacy of optimization methods. There are other hyper-parameters for nearly all techniques, such as the cooling schedule and initial temperature in SA, the acquisition function in BO, the population size and selection method in GA, etc. In the future, we plan to conduct more experiments on all these hyper-parameters.

## 5.4 Peering into the Black Box

Despite some successful applications of black-box optimization on auto-tuning system parameters, few have explained how and why some techniques work better than others for certain problems. Here we take the first step towards unpacking the "black box" and provide some insights into their internals based on our evaluation results and storage domain knowledge.

Our attempts for explanations stem from a somewhat unexpected but beneficial behavior of GA in the experiments. We found that as GA runs, there is often a small set of alleles (parameter values) that dominate the current population and are unlikely to change. We present and explain

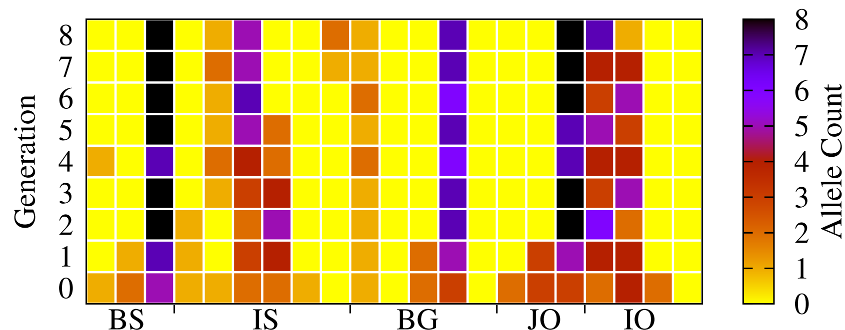


Figure 5.6: Number of alleles (parameter values) in the first 10 generations from one GA experiment run, with more frequent ones colored with darker colors.

this observation in Figure 5.6. The experiment was conducted on a parameter space consisting of 2,208 Ext3 configurations under *S2, fileserver-def, SSD*. The X axis shows 5 genes (parameters) separated by major ticks, while one cell represents one allele (parameter value). The parameters are denoted with their abbreviations from Table 4.3. The Y axis shows the generation number, and we only plotted the first 10 generations. Cells were colored based on the number of alleles in each generation. More frequent alleles are colored with darker colors. In the first generation, the gene's alleles (parameter values) were quite diverse. For example, there were 3 alleles (1K, 2K, 4K) for the *Block Size* gene, and 3 alleles (journal, ordered, writeback) for the *Journal Option* gene. However, the diversity of alleles decreased in later generations, and several genes began to dominate and even converged to a single allele. For the *Block Size* gene, only the 4K allele

survived and other two became extinct. Since GA was proposed by simulating the process of natural selection, where alleles with better fitness are more likely to survive, this suggests that GA works by identifying the combination of good alleles (storage parameter values), and producing offspring with these alleles. As shown in Figure 5.6, in the $10^{th}$ generation, all configurations have a *Block Size* of 4K and *Journal Option* of writeback.

To confirm the above observations, in Figure 5.7 we plotted all Ext3-SSD configurations under *fileserver-def* workload, with one dot corresponding to one configuration. Configurations are separated based on the *Journal Option*, shown as the X axis, and colored based on their *Block Size*. To clearly see all points within each X-axis section, we ordered configurations by their unique identification number in our database. The Y axis represents throughput values. This resulted in the formation of nine "clusters" on the graph, each corresponding to a fixed ⟨*Journal Option*, *Block Size*⟩ pair. We can see that configurations with *data=ordered* tend to produce higher throughput than those with *data=journal*, and *data=writeback* produces the best throughput. This is somewhat expected from a storage point of view, as Ext3's more fault tolerant journal option (*data=journal*) may hurt throughput by writing data as well as meta-data to the journal first.

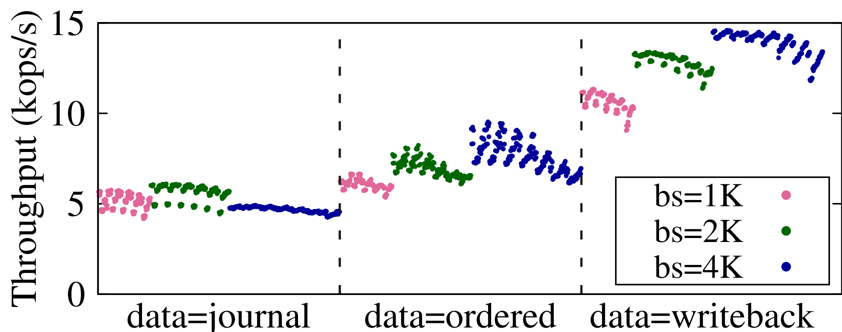


Figure 5.7: Scatter plot for all Ext3-SSD configurations under fileserver-def workload, with one dot corresponding to one configuration.

Moreover, among journal configurations with *data=writeback*, those with a 4K *Block Size* turn out to produce the highest throughput. This aligns with our observation from Figure 5.6 that GA works by identifying a subset of genes that have a greater impact on performance—*Block Size* and *Journal Option*—and finding the best alleles for them ([*4K, data=writeback*]).

Based on these observations, one interesting question to ask is whether the conclusion that a subset of parameter have greater impact on performance than other parameters, also holds for other file systems and workloads. To answer this question, we quantified the correlation between parameter values and the throughput. As most of our parameters are categorical or discrete numeric, whereas the throughput is continuous, we took a common approach to quantify the correlation between categorical and continuous variables [26]. We illustrate with the *Block Size* parameter as an example. Since it can take 3 values, we convert this parameter to three binary variables $x_1$, $x_2$, and $x_3$. If the *Block Size* is 1K, we assign $x_1 = 1$ and $x_2$ and $x_3$ are set to 0. Let Y represent the throughput values. We then do a linear regression with ordinary least squares (OLS) on Y and $x_1, x_2, x_3$. $R^2$ is a common metric in statistics to measure how the data fits a regression line. In our approach, $R^2$ actually quantifies the correlation between the selected parameter and throughput. We consider $R^2 > 0.6$ as an indication that the parameter has significant impact on performance, as is common in statistics [26]. The same calculation is applied to all parameters among SSD

configurations under the *fileserver-def* and *dbserver-def*. Parameters with the highest $R^2$ values are colored in yellow background in Table 5.2. If all $R^2$ values are below 0.6, we simply leave the entries blank, meaning no highly correlated parameters were found. To find the second important parameter, the same process is applied to the remaining parameters, but with the value of the most important one fixed (to isolate its effect on the remaining parameters' importance). Taking Ext4 as an example, we calculate $R^2$ values for all other parameters among configurations with the same *Journal Option*. For one parameter, 3 *Journal Options* lead to three $R^2$ values; we then take the maximum one as the $R^2$ value for this parameter. We color the parameter with the highest $R^2$ in Table 5.2 with a green background.

| Workload | FS | BS | IS | BG | JO | AO | SO | I/O |
|---|---|---|---|---|---|---|---|---|
| **fileserver-def** | **Ext2** | - | - | - | - | - | - | 0.68 |
| | **Ext3** | 0.84 | - | - | 0.90 | - | - | - |
| | **Ext4** | 0.92 | - | - | 0.99 | - | - | - |
| | **XFS** | 0.94 | - | 0.82 | - | - | - | - |
| | **Btrfs** | - | - | - | - | - | - | - |
| | **Nilfs2** | 0.99 | - | - | - | - | - | 0.94 |
| | **Reiserfs** | - | - | - | 0.74 | - | - | 0.99 |
| **dbserver-def** | **Ext2** | - | - | - | - | - | - | - |
| | **Ext3** | 0.72 | - | - | 0.96 | - | - | - |
| | **Ext4** | - | - | - | 0.96 | 0.68 | - | - |
| | **XFS** | - | - | - | - | - | - | - |
| | **Btrfs** | - | - | - | - | - | - | - |
| | **Nilfs2** | 0.62 | - | - | - | - | - | 0.80 |
| | **Reiserfs** | - | - | - | 0.99 | - | - | - |

Table 5.2: Importance of parameters (measured by $R^2$) among SSD configurations, with the most important one colored in yellow and second in green.

We can see that the correlated parameters are quite varied, and depend a lot on file systems. For example, under *fileserver-def*, the two most important parameters for Ext3 (in descending order) are *Journal Option* and *Block Size*; this aligns with our observation in Figure 5.6 and 5.7. However, for Reiserfs, the top 2 changes to *I/O Scheduler* and *Journal Option*. Interestingly, all parameters for Btrfs come with low $R^2$ values, which indicates that no parameter has significant impact on system performance under *fileserver-def* with Btrfs. Correlation of parameters can also depend on the workloads. For instance, the two dominant parameters for XFS under *fileserver-def* are *Block Size* and *Allocation Group*. When the workload changes to *mailserver-def*, all parameters for XFS seem to have minor impact on performance. Note that here we are isolating the impact of each parameter, thus assuming that their effect on throughput is independent; in future work we plan to investigate whether parameters have inter-dependencies.

The fact that parameters have varied impact on performance can also help explain the auto-tuning results in Chapter 5.2. Although our parameter space comes with 8 parameters, only a subset of them are correlated with performance. The number of dominant parameters is termed as *effective dimension*, and has also been observed in hyper-parameter optimization problems [14]. In our experiments (Chapter 5.2), Random Search (RS) is actually searching in a smaller effective space than the original one, and thus can find good configurations within a short time. GA's

efficacy comes from assigning a higher chance of survival to configurations with a certain combination of values for the effective parameters. BO stores its previous search experience (history) in a probabilistic surrogate model that it is building, which eventually encodes the combination of dominant parameter values that can result in good throughput values. SA does not work as well because it lacks history information to identify the dominant parameters: it wastes time on changing less useful parameters and converges slowly. Similarly, DQN also spends lots of its effort on exploring unpromising spaces, which slows its ability to find near-optimal configurations.

## 5.5 Limitations

In this chapter we provided the first comparative analysis of applying multiple optimization methods on auto-tuning storage systems. However, auto-tuning is a complex topic and more effort is required. We list some limitations of this comparative work below. ■ **(1)** We assume that changing parameter values come at no cost. In reality, parameters like *Block Size* may need re-formatting file systems. We propose to address this in Chapter 7. One possible solution is to add a *penalty function* to optimization algorithms. ■ **(2)** Previous studies [15], as well as our results from Chapter 5.3, suggest that the choice of hyper-parameter settings could have a significant impact on the efficacy of optimization algorithms. We plan to further explore the impact of hyper-parameters on optimization algorithms.

# Chapter 6

# On the Performance Variation in Modern Storage Systems

## 6.1 Motivations

Predictable performance is critical in many modern computer environments. For instance, to achieve good user experience, which notably impacts the revenues, interactive Web services require stable response time [37, 72, 91]. In cloud environments users pay for computational resources. Therefore, achieving predictable system performance, or at least establishing the limits of performance variation, is of utmost importance for the clients' satisfaction [142, 165]. In a broader sense, humans generally expect repetitive actions to yield the same results and take the same amount of time to complete; conversely, the lack of performance stability, is fairly unsatisfactory to humans.

Performance variation is a complex issue and can arise from nearly every layer in a computer system. At the hardware level, CPU, main memory, buses, and secondary storage can all contribute to overall performance variation [37,91]. At the OS and middleware level, when background daemons and maintenance activities are scheduled, they impact the performance of deployed applications. More performance disruptions come into play when considering distributed systems, as applications on different machines have to compete for heavily shared resources, such as network switches [37].

In this chapter we focus on characterizing and analyzing performance variations arising from benchmarking a typical modern storage system that consists of a file system, a block layer, and storage hardware. Storage have been proven to be a critical contributor to performance variation [67, 128, 146]. Furthermore, among all system components, the storage system is the cornerstone of data-intensive applications, which become increasingly more important in the big data era [29,75]. Although our main focus here is reporting and analyzing the variations in benchmarking processes, we believe that our observations pave the way for understanding stability issues in production systems.

Historically, many experienced researchers noticed how workloads, software, hardware, and the environment—even if reportedly "identical"—exhibit different degrees of performance variations in repeated, controlled experiments [28, 37, 46, 91, 98]. We first encountered such variations in exhaustive search experiments (see Chapter 4) with Ext4: multiple runs of the same workload in

a carefully controlled environment produced widely different performance results. Over a period of two years of collecting performance data, we later found that such high performance variations were not confined to Ext4. Over 18% of 6,222 different storage configurations on 4 different storage devices that we tried exhibited a standard deviation of performance larger than 5% of the mean, and a range value (maximum minus minimum performance, divided by the average) exceeding 9%. In a few extreme cases, standard deviation exceeded 40% even with numerous repeated experiments. The observation that some configurations are more stable than others motivated us to conduct a more detailed study of storage system performance variation and seek its root causes, as performance stability is critical for storage systems and important in achieving the success of auto-tuning.

To the best of our knowledge there are no systematic studies of performance variation in storage systems. Thus, our first goal was to characterize performance variation in different storage configurations. However, measuring this for even a single storage configuration is time consuming; and measuring all possible configurations is time-prohibitive. Even with our Storage V2 (see Chapter 4.3), it could take more than 2 years of evaluation time. Therefore, in this study we combined two approaches to reduce the configuration space and therefore the amount of time to run the experiments: (1) we used domain expertise to select the most relevant parameters, and (2) we applied a Latin Hypercube Sampling (LHS) to the configuration space. Even for the reduced space, it took us over 33 clock days to complete these experiments alone.

We focused on three local file systems (Ext4, XFS, and Btrfs) which are used in many modern local and distributed environments. Using our expertise, we picked several widely used parameters for these file systems (e.g., block size, inode size, journal options). We also varied the Linux I/O scheduler and storage devices, as they can have significant impact on performance. We benchmarked over 100 configurations using different workloads and repeated each experiment 10 times to balance the accuracy of variation measurement with the total time taken to complete these experiments. We then characterized performance variation from several angles: throughput, latency, temporally, spatially, and more. We found that performance variation depends heavily on the specific configuration of the system. We then further dove into the details, analyzed and explained certain performance variations. For example: we found that unpredictable layouts in Ext4 could cause over 16–19% of performance variation in some cases. We discovered that the magnitude of variation also depends on the observation window size: in one workload, 40% of XFS configurations exhibited higher than 20% variation with a window size of 60s, but almost all of them stabilized when the window size grew to 400s. Finally, we analyzed latency variations from various aspects, and proposed a novel approach for quantifying the impacts of each operation type on overall performance variation.

We summarize key contributions of our performance variation study as follows: ■ **(1)** To the best of our knowledge, we are the first to provide a detailed characterization of performance variation occurring in benchmarking a typical modern storage system. We believe our study paves the way towards the better understanding of complex storage system performance variations, in both experimental and production settings. ■ **(2)** We conducted a comprehensive study of storage system performance variation. Our analysis includes throughput and latency, and both spatial and temporal variations. ■ **(3)** We offer insights into the root causes of some performance variations, which could help anyone who seeks stable results from benchmarking storage systems, and encourage more follow-up work in understanding variations in production systems.

This study has been published in FAST 2017 [24]. The rest of the chapter is organized as

follows. Chapter 6.2 explains background knowledge. Chapter 6.3 describes our experimental methodology. Chapter 6.4 covers related work on storage performance variation. We list our experimental settings in Chapter 6.5. Chapter 6.6 evaluates performance variations from multiple dimensions. .

## 6.2 Background

The storage system is an essential part of modern computer systems, and critical to the performance of data-intensive applications. Often, the storage system is the slowest component and thus is one of the main contributors to the overall variability in a system's performance. Characterizing this variation in storage system performance is therefore essential for understanding overall system-performance variation.

We first define common performance metrics and notations used in this chapter. *Throughput* is defined as the average number of I/O operations completed per second. Here we use a "*Throughput-N*" notation to represent the throughput within the last N seconds of an observation. There are two types of throughput that are used most frequently in our analysis. One is *cumulative* throughput, defined as the throughout from the beginning to the end of the experiment. In this chapter, cumulative throughput is the same as *Throughput-800* or *Throughput-2000*, because the complete runtime of a single experiment was either 800 or 2,000 seconds, depending on the workload. The other type is called *instantaneous* throughput, which we denote as *Throughput-10*. Ten seconds is the smallest time unit we collected performance for, in order to avoid too much overhead.

### 6.2.1 Measures of Variation

Since our goal is to characterize and analyze collected experimental data, we mainly use concepts from *descriptive statistics*. *Statistical variation* is closely related to *central tendency*, which is an estimate of the *center* of a set of values. *Variation* (also called *dispersion* or *variability*), refers to the spread of the values around the central tendency. We considered the most commonly used measure for central tendency—the *mean*.

$$\bar{x} = \sum_{i=1}^{N} x_i. \tag{6.1}$$

Here, $x_i$ is the value number $i$ and we have $N$ such values in total (e.g., collected from experiments).

In descriptive statistics, a measure of variation is usually a non-negative real number that is zero if all readings are the same and increases as the measurements become more dispersed. To reasonably compare variations across datasets with different mean values, it is common to normalize the variation by dividing any absolute metric of variation by the mean value. There are several different metrics for variation. We initially considered two that are most commonly used in descriptive statistical analysis:

- *Relative Standard Deviation (RSD)*: the RSD, (or *Coefficient of Variation (CV)*) is

$$RSD = \frac{\sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \bar{x})^2}}{\bar{x}} \tag{6.2}$$

- *Relative Range*: this is defined as the difference between the smallest and largest values:

$$RelativeRange = \frac{max(X) - min(X)}{\bar{x}} \tag{6.3}$$

Because a range uses maximum and minimum values in its calculation, it is more sensitive to outliers. We did not want to exclude or otherwise diminish the significance of performance outliers. We found that even a few long-running I/O operations can substantially worsen actual user experience due to outliers (which are re-producible). Such outliers have real-world impact, especially as more services are offloaded to the cloud, and customers demand QoS guarantees through SLAs. That is one reason why researchers recently have begun to focus on tail latencies [37, 65, 67]. In considering the two metrics above, we felt that the RSD hides some of the magnitudes of these variations—because using square root tends to "compress" the outliers' values. We therefore decided to use the *Relative Range* as our main metric of variation in the rest of this chapter.

## 6.3 Methodology

Although we encountered storage system performance variations in past projects, we were especially struck by this issue in our recent experiments on automated recognition of optimal storage configurations. We found that multiple runs of the same workload in a carefully controlled environment could sometimes produce quite unstable results. We later observed that performance variations and their magnitude depend heavily on the specific configuration of the storage system. Over 18% of 24,888 different storage configurations that we evaluated (repeatedly over several workloads) exhibited results with a relative range higher than 9% and relative standard deviation higher than 5%.

Workloads also impact the degree of performance variation significantly. For the same configuration, experiments with different workloads could produce different magnitudes of variation. For example, we found one Btrfs configuration produces variation with over 40% relative range value on one workload but only 6% for another. All these findings led us to study the characteristics and analyze performance variations in benchmarking various storage configurations under multiple workloads. Due to the high complexity of storage systems, we have to apply certain methodologies in designing and conducting our experiments.

**Reducing the parameter space** In this chapter we focus on evaluating *local* storage systems (e.g., Ext4, Linux block layer, SSD). This is a useful basis for studying more complex distributed storage systems (e.g., Ceph [159], Lustre [109], GPFS [129], OpenStack Swift [117]). Even a small variation in local storage system performance can result in significant performance fluctuations in large-scale distributed system that builds on it [37, 103, 112].

Despite its simple architecture, a local storage system can still have a large number of parameters at every layer, resulting in a vast number of possible configurations. For instance, common

| Parameter Space | # Unique Parameters | # Unique Configurations | Time (years) |
|---|---|---|---|
| **Ext4** | 59 | $2.7 \times 10^{37}$ | $7.8 \times 10^{33}$ |
| **XFS** | 37 | $1.4 \times 10^{19}$ | $4.1 \times 10^{15}$ |
| **Btrfs** | 54 | $8.8 \times 10^{26}$ | $2.5 \times 10^{23}$ |
| **Expert Space** | 10 | 1,782 | 1.52 |
| **Sample Space** | 10 | 107 | 33.4 days |

Table 6.1: Comparison for parameter spaces. Time is computed by assuming 15 minutes per experimental run, 10 runs per configuration and 3 workloads in total.

parameters for a typical local file system include block size, inode size, journal options, and many more. It is prohibitively time consuming and impractical to evaluate every possible configuration exhaustively. As shown in Table 6.1, Ext4 has 59 unique parameters that can have anywhere from 2 to numerous allowed values each. If one experiment runs for 15 minutes and we conduct 10 runs for each configuration, it will take us $7.8 \times 10^{33}$ years of clock time to finish evaluating all Ext4 configurations.

Therefore, our first task was to reduce the parameter space (as compared with *Storage V2* in Table 4.3) for our experiments by carefully selecting the most relevant storage system parameter.. This selection was done in close collaboration with several storage experts that have either contributed to storage system designs or have spent years tuning storage systems in the field. We experimented with three popular file systems that span a range of designs and features. ■ **(1) Ext4** [47] is a popular file system that inherits a lot of internal structures from Ext3 [22] and FFS [106]) but enhances performance and scalability using extents and delayed allocation. ■ **(2) XFS** [135, 144] was initially designed for SGI's IRIX OS [144] and was later ported to Linux. It has attracted users' attention since the 90s thanks to its high performance on new storage devices and its high scalability regarding large files, large numbers of files, and large directories. XFS uses B+ trees for tracking free extents, indexing directory entries, and keeping track of dynamically allocated inodes. ■ **(3) Btrfs** [20, 124] is a complex file system that has seen extensive development since 2007 [124]. It uses copy-on-write (CoW), allowing efficient snapshots and clones. It has its own LVM and uses B-trees as its main on-disk data structure. These unique features are garnering attention and we expect Btrfs to gain even greater popularity in the future.

For the three file systems above we experimented with the following nine parameters. ■ **(1) Block size**. This is a group of contiguous sectors and is the basic unit of space allocation in a file system. Improper block size selection can reduce file system performance by orders of magnitude [67]. ■ **(2) Inode size**. This is one of the most basic on-disk structures of a file system [9]. It stores the metadata of a given file, such as its size, permissions, and the location of its data blocks. The inode is involved in nearly every I/O operation and thus plays a crucial role for performance, especially for metadata-intensive workloads. ■ **(3) Journal mode**. Journaling is the write-ahead logging implemented by file systems for recovery purposes in case of power losses and crashes. In Ext4, three types of journaling modes are supported: *writeback*, *ordered*, and *journal* [48]. The *writeback* mode journals only metadata whereas the *journal* mode provides full data and metadata journaling. In *ordered* mode, Ext4 journals metadata only, but all data is forced directly out to the disk prior to its metadata being committed to the journal. There is a trade-off between file system consistency and performance, as journaling generally adds I/O overhead. In comparison, XFS implements metadata journaling, which is similar to Ext4's *writeback* mode, and there is no need for

| File System | Parameter | Value Range |
|---|---|---|
| **Ext4** | Block Size | 1024, 2048, 4096 |
| | Inode Size | 128, 512, 2048, 8192 |
| | Journal Mode | data=journal, ordered, writeback |
| **XFS** | Block Size | 1024, 2048, 4096 |
| | Inode Size | 256, 512, 1024, 2048 |
| | AG Count | 8, 32, 128, 512 |
| **Btrfs** | Node Size | 4096, 16384, 65536 |
| | Special Options | nodatacow, nodatasum, default |
| **All** | atime Options | relatime, noatime |
| | I/O Scheduler | noop, deadline, cfq |
| | Storage Devices | HDD (SAS, SATA), SSD (SATA) |

Table 6.2: List of parameters and value ranges.

journaling in Btrfs because of its CoW nature. ■ **(4) Allocation Group (AG) count**. This parameter is specific to XFS which partitions its space into regions called Allocation Groups [144]. Each AG has its own data structures for managing free space and inodes within its boundaries. ■ **(5) Nodatacow** is a Btrfs mount-time option that turns the CoW feature on or off for data blocks. When data CoW is enabled, Btrfs creates a new version of an extent or a page at a newly allocated space [124]. This allows Btrfs to avoid any partial updates in case of a power failure. When data CoW is disabled, partially written blocks are possible on system failures. In Btrfs, *nodatacow* implies *nodatasum* and compression disabled. ■ **(6) Nodatasum** is a Btrfs mount-time option and when specified, it disables checksums for newly created files. Checksums are the primary mechanism used by modern storage systems to preserve data integrity [9], computed using hash functions such as SHA-1 or MD5. ■ **(7) atime Options**. These refer to mount options that control the inode access time. We experimented with *noatime* and *relatime* values. The *noatime* option tells the file system not to update the inode access time when a file data read is made. When *relatime* is set, atime will only be updated when the file's modification time is newer than the access time or atime is older than a defined interval (one day by default). ■ **(8) I/O scheduler**. The I/O Scheduler manages the submission of block I/O operations to storage devices. The choice of I/O scheduler can have a significant impact on storage system performance [17]. We used the *noop*, *deadline*, and *Completely Fair Queuing (CFQ)* I/O schedulers. Briefly explained, the *noop* scheduler inserts all incoming I/O requests into a simple FIFO queue in order of arrival; the *deadline* scheduler associates a deadline with all I/O operations to prevent starvation of requests; and the *CFQ* scheduler try to provide a fair allocation of disk I/O bandwidth for all processes that requests I/O operations. ■ **(9) Storage device**. The underlying storage device plays an important role in nearly every I/O operation. We ran our experiments on three types of devices: two HDDs (SATA vs. SAS) and one (SATA) SSD.

Table 6.2 summarizes all parameters and the values used in our experiments.

**Latin Hypercube Sampling** Reducing the parameter space to the most relevant parameters based on expert knowledge resulted in 1,782 unique configurations ("Expert Space" in Table 6.1). However, it would still take more than 1.5 years to complete the evaluation of every configuration in that space. To reduce the space further, we intelligently sampled it using *Latin Hypercube*

*Sampling (LHS)*, a method often used to construct computer experiments in multi-dimensional parameter spaces [71, 102]. LHS can help explore a search space and discover unexpected behavior among combinations of parameter values; this suited our needs here. In statistics, a *Latin Square* is defined as a two-dimensional square grid where each row and column have only one sample; *Latin Hypercube* generalizes this to multiple dimensions and ensures that each sample is the only one in the axis-aligned hyper-plane containing it [102]. Using LHS, we were able to sample 107 representative configurations from the Expert Space and complete the evaluation within 34 days of clock time (excluding lengthy analysis time). We believe this approach is a good starting point for a detailed characterization and understanding of performance variation in storage systems.

## 6.4 Related Work

To the best of our knowledge, there are no systematic studies of performance variation of storage systems. Most previous work focuses on long-tail I/O latencies. Tarasov et al. [146] observed that file system performance could be sensitive to even small changes in running workloads. Arpaci-Dusseau [8] proposed an I/O programming environment to cope with performance variations in clustered platforms. Worn-out SSDs exhibit high latency variations [38]. Hao et al. [64] studied device-level performance stability, for HDDs and SSDs.

For long-tail latencies of file systems, He et al. [67] developed Chopper, a tool to explore a large input space of file system parameters and find behaviors that lead to performance problems; they analyzed long-tail latencies relating to block allocation in Ext4. In comparison, our goal is broader: a detailed characterization and analysis of several aspects of storage system performance variation, including devices, block layer, and the file systems. We studied the variation in terms of both throughput and latency, and both spatially and temporally. Tail latencies are common in network or cloud services [37, 91]: several tried to characterize and mitigate their effects [65, 72, 142, 165], as well as exploit them to save data center energy [154]. Li et al. [91] characterized tail latencies for networked services from the hardware, OS, and application-level sources. Dean and Barroso [37] pointed out that small performance variations could affect a significant fraction of requests in large-scale distributed systems, and can arise from various sources; they suggested that eliminating all of them in large-scale systems is impractical. We believe there are possibly many sources of performance variation in storage systems, and we hope this work paves the way for discovering and addressing their impacts.

## 6.5 Experimental Setup and Workloads

All experiments from this chapter were conducted on *S3* machines (see Table 4.1). We characterized variations on three storage devices, HDD2, HDD4, and SSD in Table 4.1. We use *SAS-HDD* to refer HDD2, and *SATA-HDD* for *HDD4*. When discussing results on both HDD devices, we just refer them together as *HDD* for short. Workload settings were described in Table 4.2, denoted as "*-heavy". As the average file size is an inherent property of a workload and should not be changed [149], the dataset size is determined by the number of files. We increased the number of files such that the dataset size is 10GB—2.5× the machine RAM size. By fixing the dataset size, we normalized the experiments' set-size and run-time, and ensured that the experiments run

long enough to produce enough I/O. With these settings, our experiments exercise both in-memory cache and persistent storage devices [147].

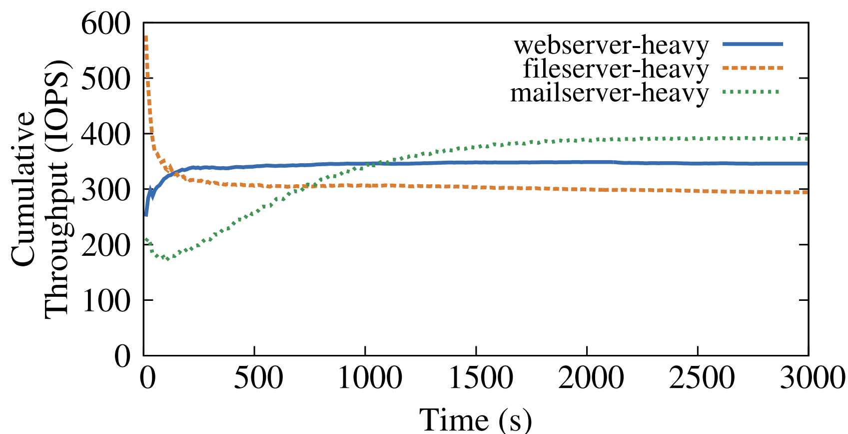


Figure 6.1: Cumulative throughput over time for one Ext4 configuration under multiple workloads. Each workload ran for 7,200s; only the first 3,000s are plotted.

We did not perform a separate cache warm-up phase in our experiments because in this study we were interested in performance variation that occurred *both* with cold and warm caches [147]. The default running time for Filebench is set to 60 seconds, which is too short to warm the cache up. We therefore conducted a "calibration" phase to pick a running time that was long enough for the cumulative throughput to stabilize. We ran each workload for up to 2 hours for testing purposes. To find a suitable run time, we ran each workload for 7,200 seconds, and measured its cumulative throughput. Figure 6.1 shows the first 3,000 seconds for Ext4 configurations. In this chapter we define the cumulative throughput as the average number of I/O operations completed per second since the start of the experiment. We can see that Fileserver and Webserver took around 600 seconds to achieve stable cumulative throughputs, and Mailserver took about 1,800 seconds. We ran the same experiments multiple times, for all file systems (Ext4, XFS, and Btrfs), and we found similar behavior. Therefore, if not stated otherwise, we set the default running time to 800 seconds for Fileserver and Webserver, and to 2,000 seconds for Mailserver. We have other choices of running time in several supplement experiments as well. We also let Filebench output the throughput (and other performance metrics) every 10 seconds, to capture and analyze performance variation from a short-term view.

## 6.6 Evaluation

In this chapter we are characterizing and analyzing storage performance variation from a variety of angles. These experiments represent a large amount of data, and therefore, we first present the information with brief explanations, and in subsequent subsections we dive into detailed explanations. Chapter 6.6.1 gives an overview of performance variations found in various storage configurations and workloads. Chapter 6.6.2 describes a case study by using Ext4-HDD configurations with the Fileserver workload. Chapter 6.6.3 presents temporal variation results. Here, temporal variations consist of two parts: changes of throughput over time and latency variation.

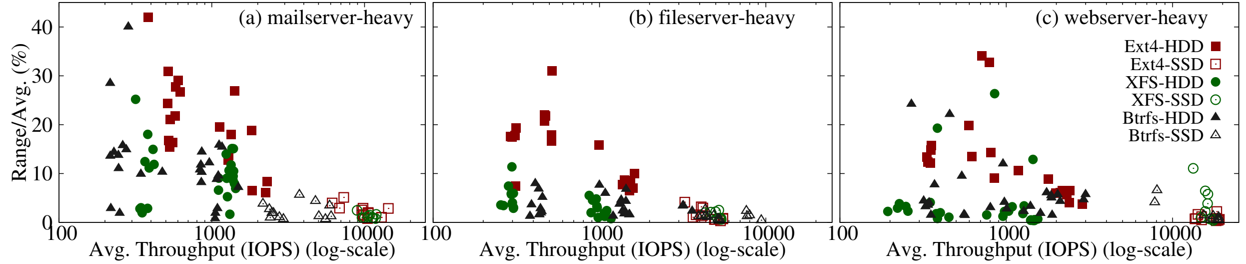


Figure 6.2: Overview of performance and its variation with different storage configurations under three workloads: (a) maileserver-heavy, (b) fileserver-heavy, and (c) webserver-heavy. The X axis represents the mean of throughput over 10 runs; the Y axis shows the relative range of cumulative throughput. Ext4 configurations are represented with squares, XFS with circles, and Btrfs with triangles. HDD configurations are shown with filled symbols, and SSDs with hollow ones.

### 6.6.1 Variation at a Glance

We first overview storage system performance variation and how configurations and workloads impact its magnitude. We designed our experiments by applying the methodology described in Chapter 6.3. We benchmarked configurations from the Sample Space (see Table 6.1) under three representative workloads from Filebench. The workload characteristics are shown in Table 4.2. We repeated each experiment 10 times in a carefully-controlled environment in order to get unperturbed measurements.

Figure 6.2 shows the results as scatter plots broken into the three workloads: *mailserver-heavy* (Figure 6.2(a)), *fileserver-heavy* (Figure 6.2(b)), and *webserver-heavy* (6.2(c)). Each symbol represents one storage configuration. We use squares for Ext4, circles for XFS, and triangles for Btrfs. Hollow symbols are SSD configurations, while filled symbols are for HDD. We collected the cumulative throughput for each run. As described in Chapter 6.2, we define the cumulative throughput as the average number of I/O operations completed per second throughout each experiment run. This can also be represented as *Throughput-800* for *fileserver-heavy* and *webserver-heavy*, and *Throughput-2000* for *mailserver-heavy*, as per our notation. In each subfigure, the Y axis represents the relative range of cumulative throughputs across the 10 runs. As explained in Chapter 6.2, here we use the relative range as the measure of variation. A higher relative range value indicates higher degree of variation. The X axis shows the mean cumulative throughput across the runs; higher values indicate better performance. Since performance for SSD configurations is usually much better than HDD configurations, we present the X axis in $log_{10}$ scale.

Figure 6.2 shows that HDD configurations are generally slower in terms of throughput but show a higher variation, compared with SSDs. For HDDs, throughput varies from 200 to around 2,000 IOPS, and the relative range varies from less than 2% to as high as 42%. Conversely, SSD configurations usually have much higher throughput than HDDs, ranging from 2,000 to 20,000 IOPS depending on the workload. However, most of them exhibit variation less than 5%. The highest range for any SSD configurations we evaluated was 11%.

Ext4 generally exhibited the highest performance variation among the three evaluated file systems. For the *mailserver-heavy* workload, most Ext4-HDD configurations had a relative range higher than 12%, with the highest one being 42%. The *fileserver-heavy* workload was slightly better, with the highest relative range being 31%. Half of the Ext4-HDD configurations show variation

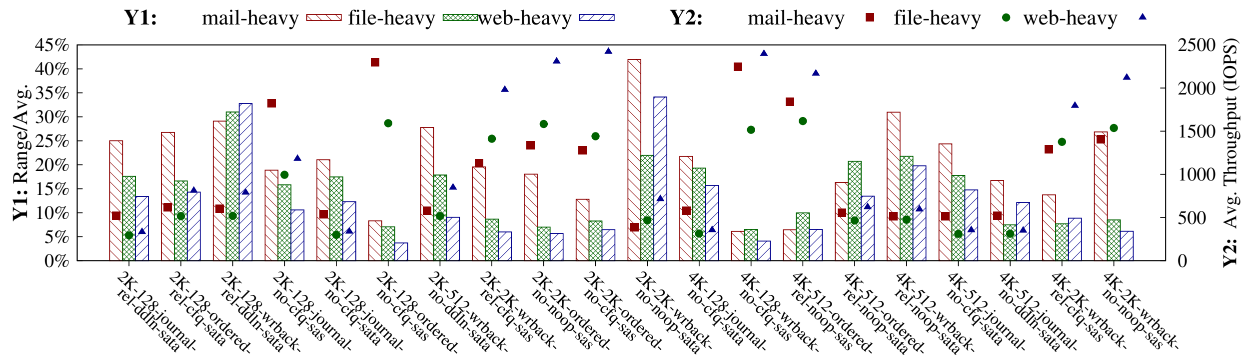


Figure 6.3: Storage system performance variation with 20 sampled Ext4-HDD configurations under three workloads. The range is computed among 10 experiment runs, and is represented as bars corresponding to the Y1 (left) axis. The mean of throughput among the 10 runs is shown with symbols (squares, circles, and triangles), and corresponds to the Y2 (right) axis. The X axis represents configurations formatted by ⟨block size - inode size - journal - atime - I/O scheduler - device⟩.

higher than 15% and the rest between 5–10%. For *webserver-heavy*, the Ext4-HDD configuration varies between 6–34%. All Ext4-SSD configurations are quite stable in terms of performance variation, with less than 5% relative range.

Btrfs configurations show a moderate level of variation in our evaluation results. For *mailserver-heavy*, two Btrfs-HDD configurations exhibited 40% and 28% ranges of throughput, and all others remained under 15%. Btrfs was quite stable under the *fileserver-heavy* workload, with the highest variation being 8%. The highest relative range value we found for Btrfs-HDD configurations under *webserver-heavy* is 24%, but most of them were below 10%. Similar to Ext4, Btrfs-SSD configurations were also quite stable, with a maximum variation of 7%.

XFS had the least amount of variation among the three file systems, and is fairly stable in most cases, as others have reported before, albeit with respect to tail latencies [67]. For *mailserver-heavy*, the highest variation we found for XFS-HDD configurations was 25%. In comparison, Ext4 was 42% and Btrfs was 40%. Most XFS-HDD configurations show variation smaller than 5% under *fileserver-heavy* and *webserver-heavy* workloads, except for one with 11% for *fileserver-heavy* and three between 12–23% for *webserver-heavy*. Interestingly, however, across all experiments for all three workloads conducted on SSD configurations, the highest variation was observed on one XFS configuration using the *webserver-heavy* workload, which had a relative range value of 11%.

Next, we decided to investigate the effect of workloads on performance variation in storage systems. Figure 6.3 compares the results of the same storage configurations under three workloads. These results were extracted from the same experiments shown in Figure 6.2. Although we show here only all Ext4-HDD configurations, we have similar conclusions for other file systems and for SSDs. The bars represent the relative range of 10 repeated runs, and correspond to the left Y1 axis. The average throughput of 10 runs for each configuration is shown as symbols, and corresponds to the right Y2 axis. The X axis consists of configuration details, and is formatted as the six-part tuple ⟨*block size* - *inode size* - *journal option* - *atime option* - *I/O scheduler* - *device*⟩. We can see that some configurations remain unstable in all workloads. For example, the configuration *2K-128-writeback-relatime-deadline-SATA* exhibited high performance variation (around 30%) under all three workloads. However, for some configurations, the actual work-

load played an important role in the magnitude of variation. For example, in the configuration *2K-2K-writeback-noatime-noop-SATA*, the *mailserver-heavy* workload varies the most; but in the configuration *4K-512-ordered-relatime-noop-SATA*, the highest range of performance was seen on *fileserver-heavy*. Finally, configurations with SAS HDD drives tended to have a much lower range variation but higher average throughput than SATA drives.

### 6.6.2 Case Study: Ext4

Identifying root causes for performance variation in the storage system is a challenging task, even in experimental settings. Many components in a modern computer system are not isolated, with complex interactions among components. CPU, main memory, and secondary storage could all contribute to storage variation. Our goal was not to solve the variation problem completely, but to report and explain this problem as thoroughly as we could. We leave to future work to address these root causes from the source code level [152]. At this stage, we concentrated our efforts solely on benchmarking local storage systems, and tried to reduce the variation to an acceptable level. In this section we describe a case study using four Ext4 configurations as examples. We focused on Ext4-HDD (SATA) here, as this combination of file systems and device types produced the highest variations in our experiments (see Figures 6.2 and 6.3).

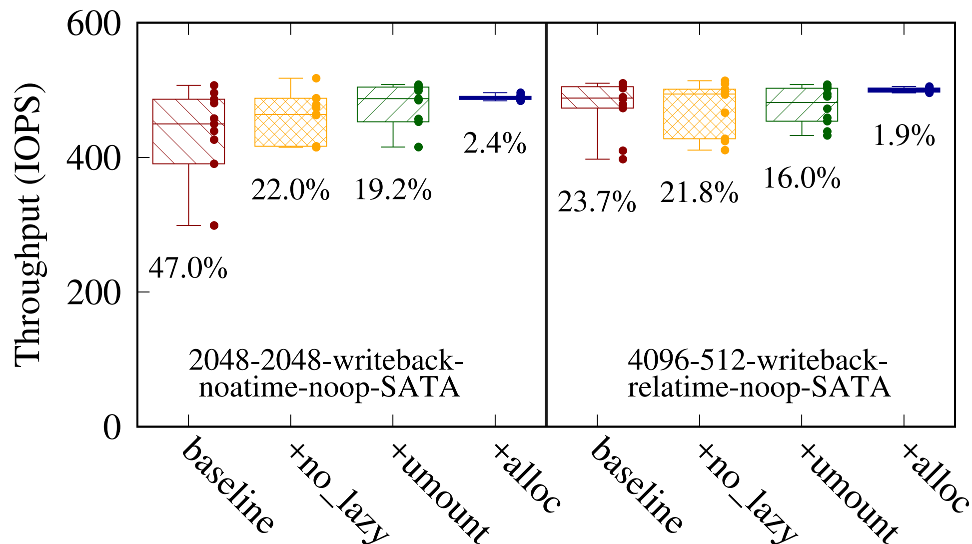


Figure 6.4: Performance variation for 2 Ext4-HDD configurations with several diagnoses. Each experiment is shown as one box, representing a throughput distribution for 10 identical runs. The top border line of each box marks the $1^{st}$ quartile; the bottom border marks the $3^{rd}$ quartile; the line in the middle is the median throughput; and the whiskers mark maximum and minimum values. The dots to the right of each box show the exact throughputs of all 10 runs. The percentage numbers below each box are the relative range values. The bottom label shows configuration details for each figure.

Figure 6.4 shows results as two boxplots for the *fileserver-heavy* workload, where each box plots the distribution of throughputs across the 10 runs, with the relative range shown below. The top border represents the $1^{st}$ quartile, the bottom border the $3^{rd}$ quartile, and the line in the middle

is the median value. Whiskers show the maximum and minimum throughputs. We also plotted one dot for the throughput of each run, overlapping with the boxes but shifted to the right for easier viewing. The X axis represents the relative improvements that we applied based on our successive investigations and uncovering of root causes of performance variation, while the Y axis shows the cumulative throughput for each experiment run. Note that the improvement label is prefixed with a "+" sign, meaning that an additional feature was added to the previous configuration, cumulatively. For example, +*umount* actually indicates *baseline* + *no_lazy* + *umount*. We also added labels on the bottom of each subfigure showing the configuration details, formatted as ⟨*block size* - *inode size* - *journal option* - *atime option* - *I/O scheduler* - *device*⟩.

After addressing all causes we found, we were able to reduce the relative range of throughput in these configurations from as high as 47% to around 2%. In the rest of this section, we detail each root cause and how we addressed it.

**Baseline** The first box for each subfigure in Figure 6.4 represents our original experiment setting, labeled *baseline*. In this setting, before each experimental run, we format and mount the file system with the targeted configuration. Filebench then creates the dataset on the mounted file system. After the dataset is created, Filebench issues the *sync* command to flush all dirty data and metadata to the underlying device (here, SATA HDD); Filebench then issues an *echo 3 > /proc/sys/vm/drop_caches* command, to evict non-dirty data and metadata from the page cache. Then, Filebench runs the Fileserver workload for a pre-defined amount of time (see Table 4.2). For this baseline setting, both Ext4-HDD configurations show high variation in terms of throughput, with range values of 47% (left) and 24% (right).

**Lazy initialization** The first contributor to performance variation that we identified in Ext4-HDD configurations is related to the lazy initialization mechanism in Ext4. By default, Ext4 does not immediately initialize the complete inode table. Instead, it gradually initializes it in the background when the created file system is first mounted, using a kernel thread called *ext4lazyinit*. After the initialization is done, the thread is destroyed. This feature speeds up the formatting process significantly, but also causes interference with the running workload. By disabling it during format time, we reduced the range of throughput from 47% to 22% for Configuration *2048-2048-writeback-noatime-noop-SATA*. This improvement is labelled +*no_lazy* in Figure 6.4.

**Sync then umount** In Linux, when *sync* is called, it only guarantees to *schedule* the dirty blocks for writing: there is often a delay until all blocks are actually written to stable media [116, 145]. Therefore, instead of calling *sync*, we *umount* the file system each time after finishing creating the dataset and then *mount* it back, which is labelled as +*umount* in Figure 6.4. After applying this, both Ext4-HDD configurations exhibited even lower variation than the previous setting (disabling lazy initialization only).

**Block allocation and layout** After applying the above improvements, both configurations still exhibited higher than 16% variations, which could be unacceptable in settings that require more predictable performance. This inspired us to try an even more strictly-controlled set of experiments. In the *baseline* experiments, by default we re-created the file system before each run and then Filebench created the dataset. We assumed that this approach would result in identical

datasets among different experiment runs. However, block allocation is not a deterministic procedure in Ext4 [67]. Even given the same distribution of file sizes and directory width, and also the same number of files as defined by Filebench, multiple trials of dataset creation on a freshly formatted, clean file system did not guarantee to allocate blocks from the same or even near physical locations on the hard disk. To verify this, instead of re-creating the file system before each run, we first created the file system and the desired dataset on it. We then dumped out the entire partition image using *dd*. Then, before each run of Filebench, we used *dd* to restore the partition using the image, and mounted the file system back. This approach guaranteed an identical block layout for each run. Figure 6.4 shows these results using *+alloc*. We can see that for both Ext4-HDD configurations, we were able to achieve around 2% of variation, which verified our hypothesis that block allocation and layout play an important role in the performance variation for Ext4-HDD configurations.

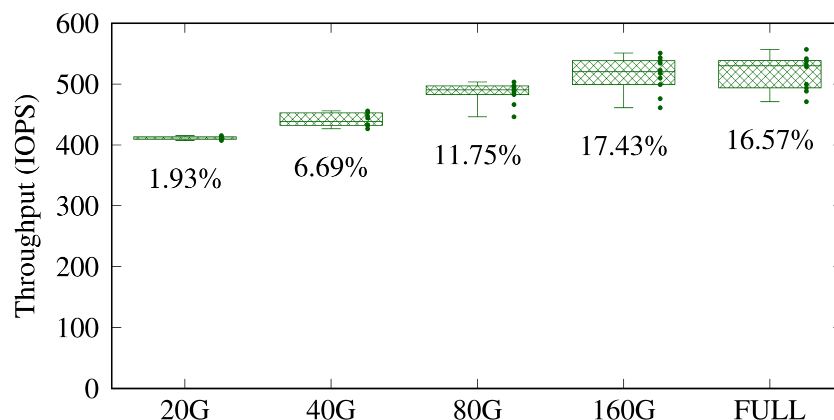


Figure 6.5: Performance variation for Ext4-HDD configuration under the Fileserver workload with different partition sizes from inner tracks of disks

After further investigation, we found this nondeterminism for Ext4 block allocation was caused by the fact that *Ext4 always tries to spread first-level directories* [49]. In the meanwhile, Filebench puts its dataset in one directory (with pre-defined directory width and depth distribution), directly under the mount point of the targeted file system. To prove this, we conducted a set of experiments by varying the partition size of the underlying hard disk. As shown in Figure 6.5, we experimented with *20G*, *40G*, *80G*, *160G* and *Full-disk* partitions. All partitions start from inner tracks of disks. We repeated each experiment for 10 runs. The meanings of boxes, whiskers, and dots are the same with those of Figure 6.4. Remember the dataset size in our experiments is 10G (see Chapter 6.5). When the partition size is 20G, the difference in physical positions of allocated files among 10 experiment runs could be quite small, which results in a relative range of $1.9\%$ in final throughput values. They all clustered in inner tracks of disks. As we increase the partition size, the relative range of throughput also increases. This is because with larger partition sizes, in different experiment runs Ext4 could allocate all blocks in physically different "clusters" across the disks. Datasets allocated in the outer tracks will result in higher final throughputs, while inner tracks produce lower results. This also explains the increasing trend of the average throughput among these experiments.

Storing the images of file systems using the *dd* command, however, could be too costly in prac-

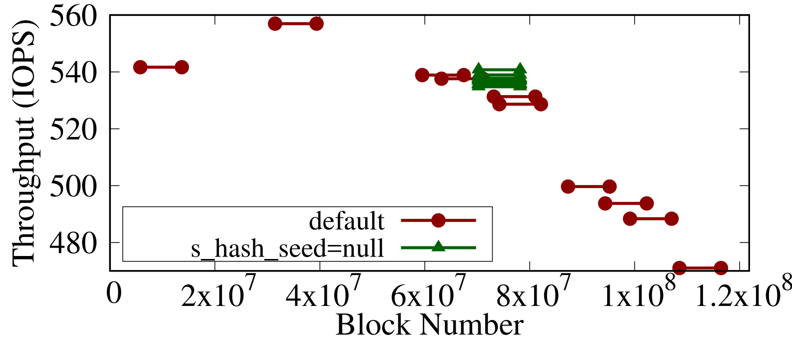


Figure 6.6: Physical blocks of allocated files in Ext4 under the Fileserver workload. The X axis represents the physical block number of each file in the dataset. Since the *Fileserver* workload consists of small files, and one extent per file, we use the starting block number for each file here. The Y axis is the final cumulative throughput for each experiment run. Note that the Y axis does not start from 0. Lines marked with solid circles are experiment runs with the default setting; lines with triangles represent experiment runs where we set the field s_hash_seed in Ext4s's superblock to null.

tice, taking hours of clock time. We found a faster method to generate reproducible Ext4 layouts by setting the *s_hash_seed* field in Ext4's superblock to *null* before mounting. Figure 6.6 shows the distribution of physical blocks for allocated files in two sets of *fileserver-heavy* experiments on Ext4. This workload consists of only small files, resulting in exactly one extent for each file in Ext4, so we used the starting block number (X axis) to represent the corresponding file. The Y axis shows the final cumulative throughput for each experiment run. Here the lines starting and ending with solid circles are 10 runs from the experiment with the full-disk partition. The lines with triangles represent the same experiments, but here we set the *s_hash_seed* field in Ext4's superblock to *null*. We can see that files in each experiment run are allocated into one cluster within a small range of physical block numbers. In most cases, experimental runs with their dataset allocated near the outer tracks of disks, which correspond to smaller block numbers, tend to produce higher throughput. As shown in Figure 6.6, with the default setting, datasets of 10 runs clustered in 10 different regions of the disk, causing high throughput variation across the runs. By setting the Ext4 superblock parameter *s_hash_seed* to *null*, we can eliminate the non-determinism in block allocation. This parameter determines the group number of top-level directories. By default, *s_hash_seed* is randomly generated during format time, resulting in distributing top-level directories all across the LBA space. Setting it to *null* forces Ext4 to use the hard-coded default values, and thus the top-level directory in our dataset is allocated on the same position among different experiment runs. As we can see from Figure 6.6, for the second set of experiments, the ranges of allocated block numbers in all 10 experiment runs were exactly the same. When we set the *s_hash_seed* parameter to *null*, the relative range of throughput dropped from and 16.6% to 1.1%. Therefore, setting this parameter could be useful when users want stable benchmarking results from Ext4. In addition to the case study we conducted on Ext4-HDD configurations, we also observed similar results for Ext4 on other workloads, as well as for Btrfs. For two of the Btrfs-HDD configurations, we were able to reduce the variation to around 1.2%, by using *dd* to store the partition image. We did not try to apply any improvements on XFS, since most of its configurations were already quite stable

(in terms of cumulative throughput) even with the *baseline* setting, as shown in Figure 6.2.

### 6.6.3 Temporal Variation

In Chapter 6.6.1 and 6.6.2, we mainly presented and analyzed performance variation among repeated runs of the same experiment, and only in terms of throughput. Variation can actually manifest itself in many other ways. We now focus our attention on *temporal* variations in storage system performance—the variation related to time. Chapter 6.6.3.1 discusses temporal throughput variations and Chapter 6.6.3.2 focuses on latency variations.

#### 6.6.3.1 Throughput over Time

After finding variations in cumulative throughputs, we set out to investigate whether the performance variation changes over time within single experiment run.

To characterize this, we calculated the throughput within a small time window. As defined in Chapter 6.2, we denote throughput with window size of N seconds as *Throughput-N*. Figure 6.7 shows the *Throughput-120* value (Y axis) over time (X axis) for Btrfs-HDD, XFS-HDD, and Ext4-HDD configurations using the *Fileserver* workload.

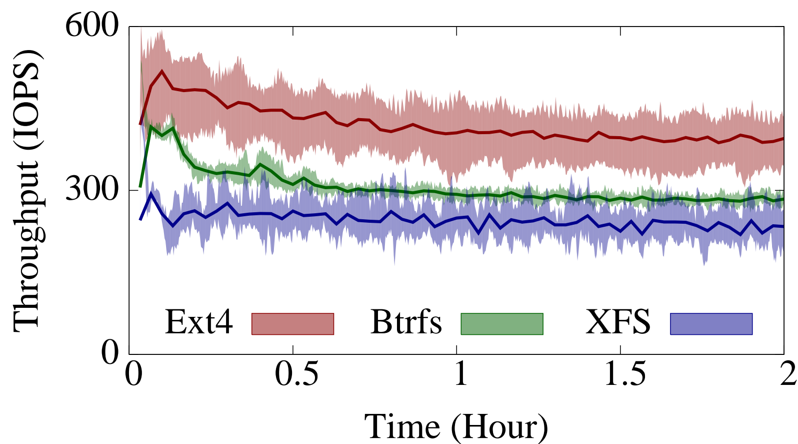


Figure 6.7: *Throughput-120* over time for Btrfs, XFS, and Ext4 HDD configurations under the Fileserver workload. Each configuration was evaluated for 10 runs. Two lines were plotted connecting maximum and minimum throughput values among 10 runs. We fill in colors between two lines, green for Btrfs, red for Ext4, and blue for XFS. We also plotted the average *Throughput-120* among 10 runs as a line running through the band. The maximum relative range values of *Throughput-120* for Ext4, Btrfs, and XFS are 43%, 23%, and 65%, while the minimum values are 14%, 2%, and 7%, respectively.

Here we use a window size of 120 seconds, meaning that each throughput value is defined as the average number of I/O operations completed per second with the latest 120 seconds. We also investigated other window sizes, which we discuss later. The three configurations shown here exhibited high variations in the experiments discussed in Chapter 6.6.1. Also, to show the temporal aspect of throughput better, we extended the running time of this experiment set to 2 hours, and we repeated each experiment 10 times. Two lines are plotted connecting the maximum and minimum

throughput values among 10 runs. We fill in colors between two lines, this producing a color band: green for Btrfs, red for Ext4, and blue for XFS. The line in the middle of each band is plotted by connecting the average *Throughput-120* value among 10 runs. We observed in Figure 6.2(b) that for the *fileserver-heavy* workload, Ext4-HDD configurations generally exhibited higher variations than XFS-HDD or Btrfs-HDD configurations in terms of final cumulative throughput. However, when it comes to *Throughput-120* values, Figure 6.7 leads to some different conclusions. The Ext4-HDD configuration still exhibited high variation in terms of short-term throughout across the 2 hours of experiment time, while the Btrfs-HDD configuration is much more stable. Surprisingly, the XFS-HDD configuration has higher than 30% relative range of *Throughput-120* values for most of the experiment time, while its range for cumulative throughput is around 2%. This suggests that XFS-HDD configurations might exhibit high variations with shorter time windows, but produces more stable results in longer windows. It also indicates that the choice of window sizes matters when discussing performance variations.

We can see from the three average lines in Figure 6.7 that performance variation exists even within one single run—the short-term throughput varies as the experiment proceeds. For most experiments, no matter what the file system type is, performance starts slow and climbs up quickly in the beginning phase of experiments. This is because initially the application is reading cold data and metadata from physical devices into the caches; once cached, performance improves. Also, for some period of time, dirty data is kept in the cache and not yet flushed to stable media, delaying any impending slow writes. After an initial peak, performance begins to drop rapidly and then declines steadily. This is because the read performance already reached its peak and cached dirty data begins to be flushed out to slower media. Around several minutes in, performance begins to stabilize, as we see the throughput lines flatten.

The unexpected difference in variations for short-term and cumulative throughput of XFS-HDD configurations lead us to investigate the effects of the time window size on performance variations. We calculated the relative range of throughput with different window sizes for all configurations within each file system type. We present the CDFs of these range values in Figure 6.8. For example, we conducted experiments on 39 Btrfs configurations. With a window size of 60 seconds and total running time of 800 seconds, the corresponding CDF for Btrfs is based on $39 \times \frac{800}{60} = 507$ relative range values. We can see that Ext4's unstable configurations are largely unaffected by the window size. Even with *Throughput-400*, around 20% of Ext4 configurations produce higher than 20% variation in terms of throughput. Conversely, the range values for Btrfs and XFS are more sensitive to the choice of window size. For XFS, around 40% of the relative range values for *Throughput-60* are higher than 20%, whereas for *Throughput-400*, nearly all XFS values fall below 20%. This aligns with our early conclusions in Chapter 6.6.1 that XFS configurations are relatively stable in terms of cumulative throughput, which is indeed calculated based on a window size of 800 seconds; whereas XFS showed the worst relative range for *Throughput-60*, it stabilized quickly with widening window sizes, eventually beating Ext4 and Btrfs.

All the above observations are based on the throughput within a certain window size. Another approach is to characterize the *instant throughput* within an even shorter period of time. Figure 6.9 shows the instantaneous throughput over time for various configurations under the *fileserver-heavy* workload. We collected and calculated the throughput every 10 seconds. Therefore we define instantaneous throughput as the average number of I/O operations completed in the past 10 seconds. This is actually *Throughput-10* in our notation. We normalize this by dividing each value by the maximum instantaneous throughput value for each run, to compare the variation across multiple

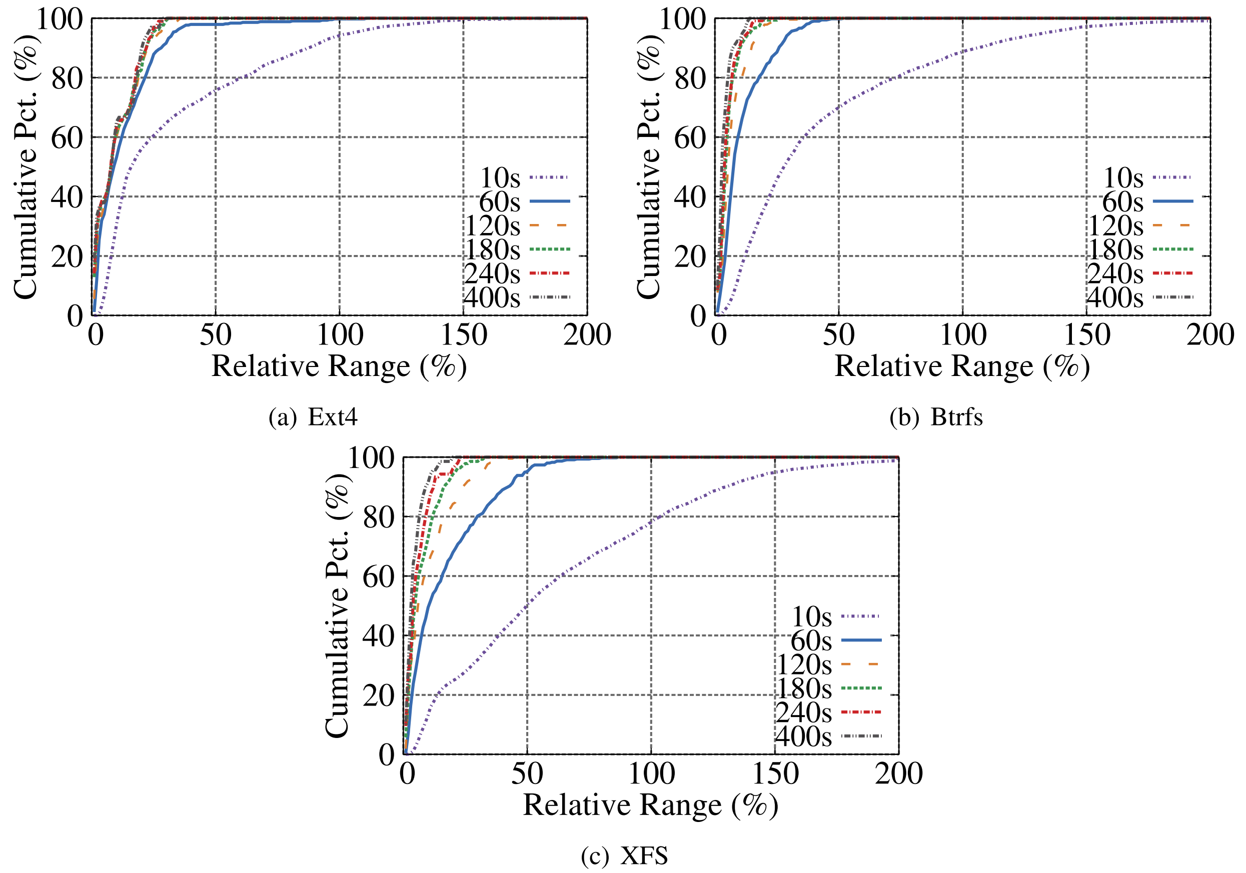


(a) Ext4

(b) Btrfs

(c) XFS

Figure 6.8: CDFs for relative range of throughput under Fileserver workload with different window sizes. For window size N, we calculated the relative range values of throughput for all configurations within each file system type, and then plotted the corresponding CDF.

experimental runs. The X axis still shows the running time.

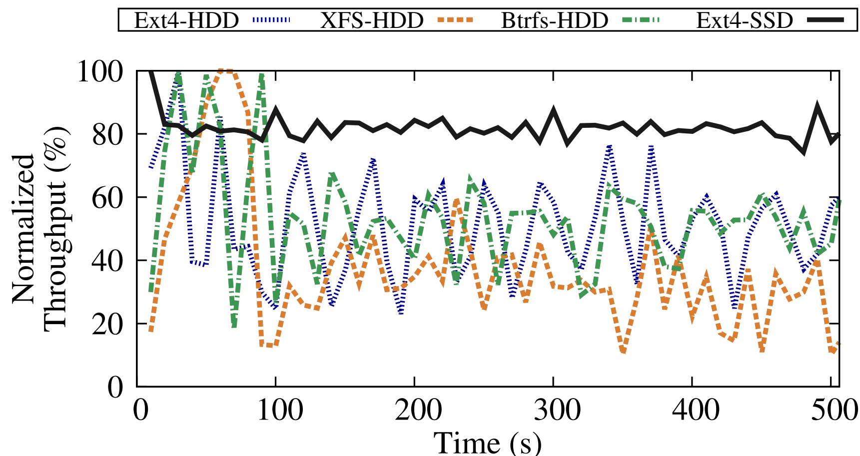


Figure 6.9: Normalized instantaneous throughput (*Throughput-10*) over time for experiments with various workloads, file systems, and devices. The Y axis shows the normalized values divided by the maximum instantaneous throughput through the experiment. Only the first 500s are presented for brevity.

We picked one illustrative experiment run for each configuration (Ext4-HDD, XFS-HDD, Btrfs-HDD, and Ext4-SSD). We can see from Figure 6.9 that for all configurations, instantaneous performance fluctuated a lot throughout the experiment. For all three HDD configurations, the variation is even higher than 80% in the first 100 seconds. The magnitude for variation reduces later in the experiments, but stays around 50%.

The throughput spikes occur nearly every 30 seconds, which could be an indicator that the performance variation in storage systems is affected by some cyclic activity (e.g., kernel flusher thread frequency). For SSD configurations, the same up-and-down pattern exists, although its magnitude is much smaller than for HDD configurations, at only around 10%. This also confirms our findings from Chapter 6.6.1 that SSDs generally exhibit more stable behavior than HDDs.

#### 6.6.3.2 Latency Variation

Another aspect of performance variation is latency, defined as the time taken for each I/O request to complete. Much work has been done in analyzing and taming long-tail latency in networked systems [72, 91] (where $99.9^{th}$ percentile latency is orders of magnitude worse than the median), and also in local storage systems [67]. Throughout our experiments, we found out that long-tail latency is not the only form of latency variation; there are other factors that can impact the latency distribution for I/O operations.

A *Cumulative Distribution Function (CDF)* is a common approach to present latency distribution. Figure 6.10(a) shows the latency CDFs for 6 I/O operations of one Ext4-HDD configuration under the *fileserver-heavy* workload. The X axis represents the latency in $log_{10}$ scale, while the Y axis is the cumulative percentage. We can see that for any one experimental run, operations can have quite different latency distribution. The latencies for *read*, *write*, and *create* form two clusters. For example, about 20% of the *read* operation has less than 0.1ms latency while the other

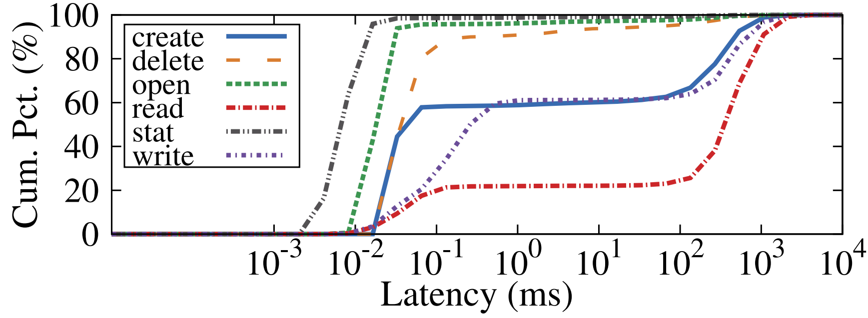


(a) CDFs of operations within one single experiment run

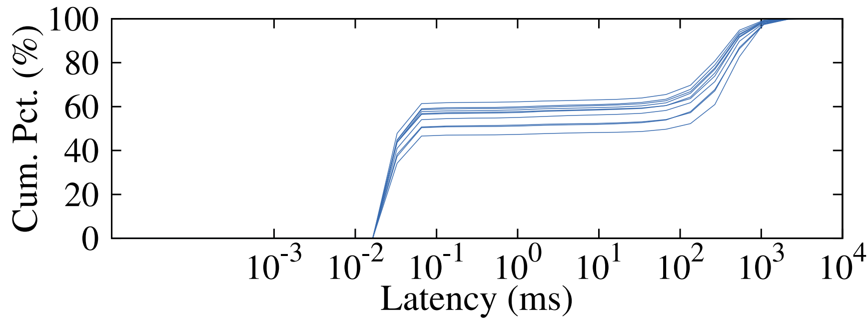


(b) CDFs of *create* operation among repeated experiment runs

Figure 6.10: Latency CDF of one Ext4-HDD configuration under *Fileserver* workload.

80% falls between 100ms and 4s. Conversely, the majority of *stat*, *open*, and *delete* operations have latencies less than 0.1ms. The I/O operation type is not the only factor that impacts the latency distribution. Figure 6.10(b) presents 10 CDFs for *create* from 10 repeated runs of the same experiment. We can see for the $60^{th}$ percentile, the latency can vary from less than 0.1ms to over 100ms.

Different I/O operations and their latencies impact the overall workload throughput to a different extent. With the empirical data that we collected—per-operation latency distributions and throughput—we were able to discover correlations between the speed of individual operations and the throughput. We first defined a metric to quantify the difference between two latency distributions. We chose to use the Kolmogorov-Smirnov test (K-S test), which is commonly used in statistics to determine if two datasets differ significantly [151]. For two distributions (or discrete dataset), the K-S test uses the maximum vertical deviation between them as the distance. We further define the range for a set of latency distributions as the maximum distance between any two latency CDFs. This approach allows us to use only one number to represent the latency variation, as with throughput. For each operation type, we calculated its range of latency variation for each configuration under all three workloads. We then computed the Pearson Correlation Coefficient (PCC) between the relative range of throughput and the range of latency variation.

Figure 6.11 shows our correlation results. The PCC value for any two datasets is always between [-1,+1], where +1 means total positive correlation, 0 indicates no correlation, and –1 means total negative correlation. Generally, any two datasets with PCC values higher than 0.7 are considered to have a strong positive correlation [125], which we show in Figure 6.11 with a horizontal dashed red line. The Y axis represents the PCC value while the X axis is the label for each opera-

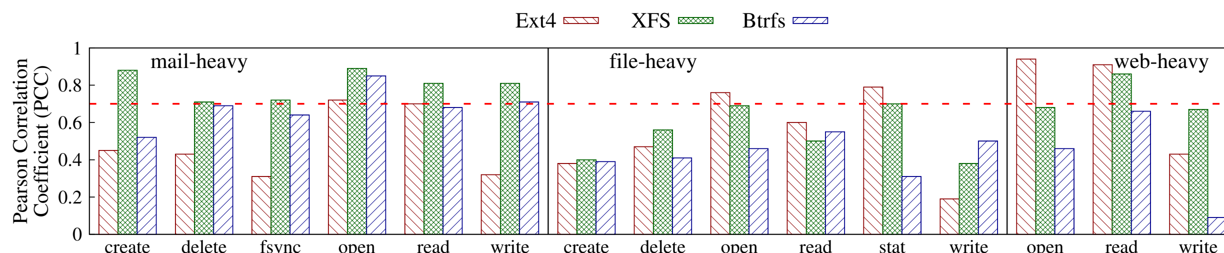


Figure 6.11: Pearson Correlation Coefficient (PCC) between throughput range and operation types, for three workloads and three file systems. The horizontal dashed red line at Y=0.7 marks the point above which a strong correlation is often considered to exist.

tion. We separate workloads with vertical solid lines. As most SSD configurations are quite stable in terms of performance, we only considered HDD configurations here. For Ext4 configurations, *open* and *read* have the highest PCC values on both *mailserver-heavy* and *webserver-heavy* workloads; however, on *fileserver-heavy*, *open* and *stat* have the strongest correlation. These operations could possibly be the main contributors to performance variation on Ext4-HDD configurations under each workload; such operations would represent the first ones one might tackle in the future to help stabilize Ext4's performance on HDD. In comparison, *write* has a PCC value of only around 0.2, which indicates that it may not contribute much to the performance variation. Most operations show PCC values larger than 0.4, which suggest weak correlation. This is possibly because I/O operations are not completely independent with each other in storage systems.

For the same workload, different file systems exhibit different correlations. For example, under the *webserver-heavy* workload, Ext4 show strong correlation on both *read* and *open*; but for XFS, *read* shows a stronger correlation than *open* and *write*. For Btrfs, no operation had a strong correlation with the range of throughput, with only *read* showing a moderate level of correlation.

Although such correlations do not always imply direct causality, we still feel that this correlation analysis sheds light on how each operation type might contribute to the overall performance variation in storage systems.

# Chapter 7

# A Practical Auto-Tuning Framework for Storage

## 7.1 Motivations

Despite some promising results in applying black-box optimization techniques for auto-tuning storage systems, we believe these techniques still lack several critical features to achieve practical, real-time auto-tuning. Our own experiments demonstrate that auto-tuning sometimes can be slow in finding near-optimal configurations, especially when the evaluation of even a single configuration takes long time (e.g., due to slow I/O). Moreover, there is no implicit mechanism to stop the search when it reaches a sufficiently good configuration (and restart it later on as needed). Little is known on how to initialize the search and give it a good starting point. There is no accounting for the cost of moving from one configuration to another, which is critically important in some production settings.

In this proposal we propose to investigate and develop a more intelligent and practical auto-tuning framework, intended to dynamically optimize storage systems. We are exploring techniques that add vital missing features from existing optimization methods:

- A criteria when the optimization algorithm should *stop searching*, having reached a "good enough" system configuration.
- A similar criteria when the search algorithm should be *restarted*, useful when the environment conditions (e.g., workload) have changed enough to take the system off of its optimal point.
- A workload modeller, which can extract features from system collected metrics and characterize the running workload based on them. This is useful in determining when to restart the auto-tuning process and how to "transfer" evaluation results from one workload to another.
- A mechanism to *pick an initial set of search space locations*, as well as *re-initialize* the search space after restarting a search—which we have found to have a big impact on the efficacy of any search [23, 43].

- A *penalty function* to assign a (weighted) cost to any new configuration based on the current system state, to account for costly configuration changes (e.g., a simple run-time changeable parameter vs. one that requires a system reboot and some downtime).

We will describe some preliminary results exploring these proposed ideas in this chapter, and discuss how these components could help us achieve the goal of auto-tuning storage systems in real-time.

## 7.2 Problem Statement

Results from Chapter 5 show that several popular black-box optimization methods were all able to gradually find better configurations; GA and BO successfully found near-optimal configurations. However, our preliminary results also indicate some limitations when simply applying these traditional black-box optimization. For example, in Figure 5.2 after around 3.5 hours, GA already found a near-optimal configuration, but it spent a lot of additional rounds and resources, yet not improving overall performance much if at all; Moreover, Figure 2.2 and Table 5.1 showed that storage evaluation results depend heavily on the hardware and running workloads. Our previous work reported similar observations [24, 132]. Therefore, our auto-tuning framework also needs to react to environment changes (e.g., hardware, workload). In Figure 5.2, SA got stuck in a configuration with throughput value of less than 18K IOps. More experiments we conducted suggest that the quality of **initialization** has a large effect on the convergence time and final optimization results. Lastly, our preliminary experiments assume that all configurations have identical cost: that moving from one configuration to another has the same (low) cost. For many storage systems, however, it is not true: for example, changing the *block size* of a file system may require costly and time-consuming reformat and data migration.

## 7.3 Proposed Auto-Tuning Framework

To address the limitations discussed in Chapter 7.2, we propose our enhanced auto-tuning framework, as shown in Figure 7.1. It consists of 6 components.

- **Monitor**, which collects and processes system metrics for other components' use.

- **Workload Modeler**, which utilizes metrics collected by the *Monitor*, to identify a running workload on the current system.

- **Optimizer**, which includes the core auto-tuning algorithm with newly added features, further detailed in §7.3.2. It calculates the optimization objective for the current system configuration, based on metrics collected by the *Monitor*. Our framework is general enough to optimize for *any objective* that can be quantitatively measured. Examples are I/O throughput or latency, energy consumption, or even an economic cost function comprising multiple metrics [97, 141].

- **Controller**, which is responsible for changing the system settings based on the configuration picked by the *Optimizer*.

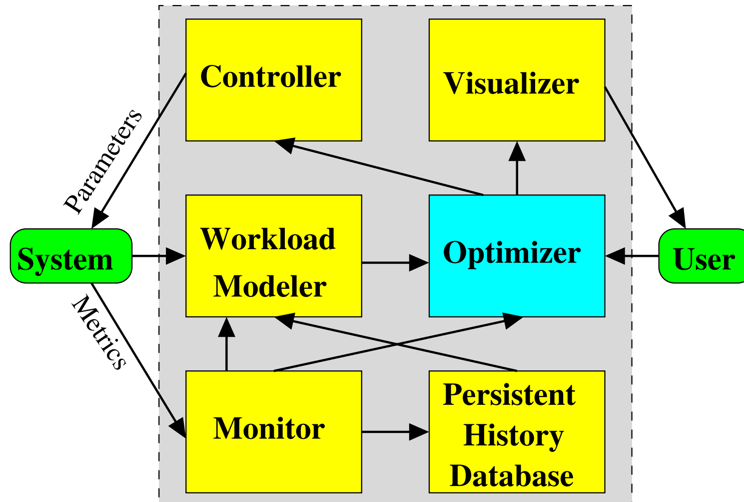


Figure 7.1: Auto-tuning Framework

- **Persistent History Database**, which stores previous evaluation results persistently. The auto-tuning algorithm can use part (or all) of this history to direct the search or build predictive ML models. A practical implementation may also periodically purge older or less valuable database entries to reduce storage costs.

- **Visualizer**, which provides the user with real time visualizations and insights into complex n-D spaces.

### 7.3.1 Workload Modeling

In this chapter we use *parameters* to refer to system factors whose values can be manually set, and *features* or *metrics* for values that can only be measured. The workload modeler's role is to find a set of features that is sufficient to differentiate workloads and quantify their changes. How to characterize a system workload remains an open problem. A few efforts were made in specific types of applications [85, 90, 155], but there is no general and well-accepted solution yet. We plan to first collect various system metrics, conducted by the *Monitor*. Example metrics include the ratio of I/O operation types (read, write, open, etc.) and how many operations are sequential vs. random. Randomness can be inferred from the difference in file offsets between consecutive I/O requests. Advanced features could be based on time series data. Our proposed modeler will then perform feature selection or clustering analysis on all extracted features and remove redundant ones.

For the selected list of metrics, we consider a distance function to quantify the similarities between workloads. Example distance functions include the earth-mover-distance (EMD) function [74, 152]. Workload similarity is useful when the *Optimizer* wants to re-utilize past evaluation results for a new workload—by finding the closest known workload for which we have a high-performing system configuration (see Chapter ).

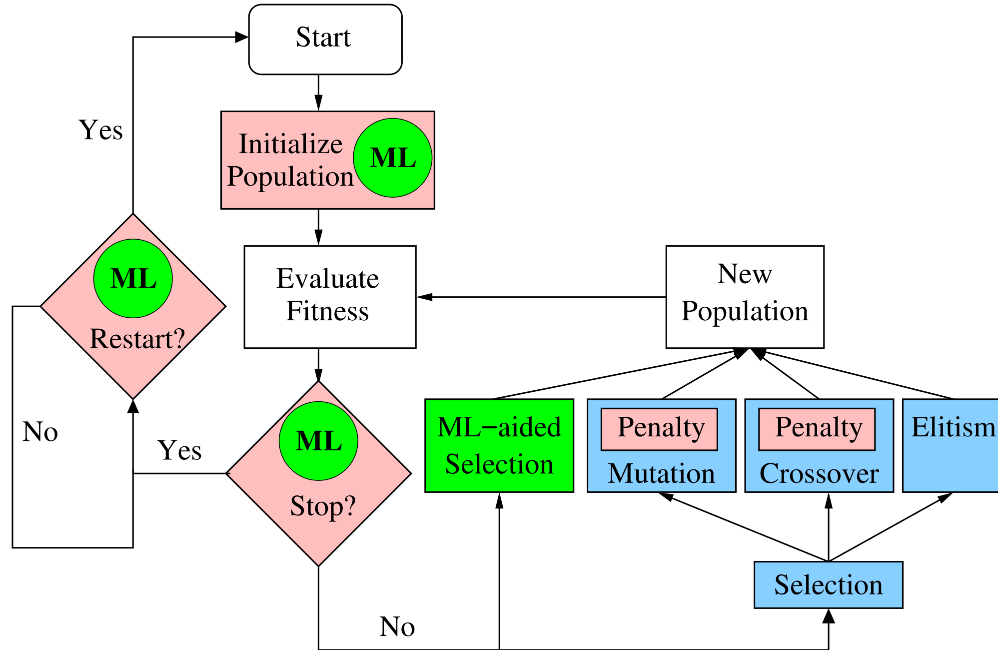


Figure 7.2: Work flow for an enhanced Optimizer (GA).

### 7.3.2 Optimizer

Our proposed *Optimizer* is designed to address the issues observed in our preliminary experiments. Here we use GA as a case study explaining how it works, but all the new components are applicable to other black-box optimization algorithms as well. As shown in Figure 7.2, white boxes represent GA's original optimization loop components and blue ones relate to GA's selection process; pink ones are new components in our hybrid optimization algorithm; and green ones show new, ML-aided components.

### 7.3.3 (Re-)Initialization

As discussed in Chapter 7.2, the quality of initialization has a large impact on the convergence time and final results of optimization. Much work have been done on proposing and analyzing different initialization methods for various optimization algorithms [54, 63, 136]. We are investigating the following initialization methods to design the best one for our needs: **(1) Simple Random Sampling**, where each configuration is chosen entirely by chance and has an equal chance of being included in the sample [26]. It is the default for many optimization techniques [54]. Although we expect it to be inefficient, it serves as a useful baseline for more intelligent methods. **(2) Stratified Random Sampling**, which divides the whole space into sub-spaces, and takes samples from each sub-space. It is quite useful when we expect the measurement of interest to vary among the different sub-spaces [26]. In case of optimizing for storage configurations, since parameters directly impact performance, an ideal initialization method should cover each parameter value more uniformly. In fact, Latin Hypercube Sampling (LHS) [102, 161], which belongs to this type of sampling, are proved to be effective in black-box optimization [54, 120]. **(3) Including domain knowledge**. Domain experts may already know some good configurations for certain

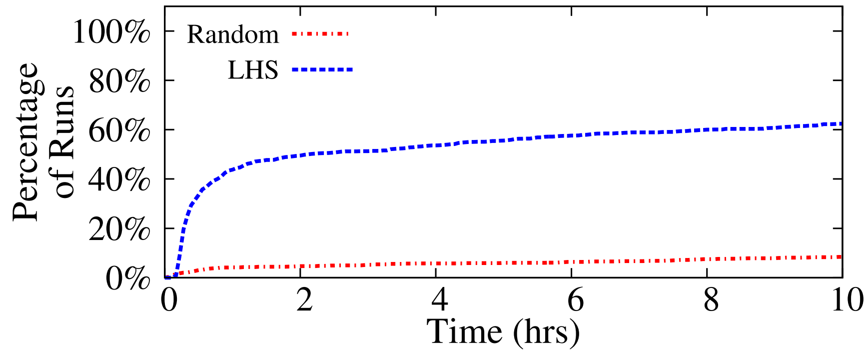


Figure 7.3: Comparison of different initialization methods.

workloads. Including them has been proved to increase the search efficacy [4,18,76,160]. Another good example here is if experts know the impact of several common parameters on overall system performance, the initialization method could try to sample the preferred parameter values more frequently. Interestingly, our automated techniques can also be used to evaluate the accuracy of domain experts' actual recommendations.

Figure 7.3 shows our preliminary results of the efficacy of two GA initialization methods: *Simple Random Sampling* and *LHS*. Since exploration is one critical component of all optimization methods (see Chapter 2.4), We repeated experiments on each initialization method 1,000 times, conducted for optimizing SSD configurations for *Fileserver* workload. We compare different initialization methods for their probability of finding near-optimal configurations. Here we define a near-optimal configuration as one with throughput higher than **99%** of the global optimal value. The Y axis shows the percentage of total runs that found a near-optimal configuration within a certain time (X axis). Clearly LHS outperforms simple random sampling, with a higher chance to find near-optimal configurations and in less time. We believe this is because different parameters have a different level of impact on performance (see Chapter 7.3.6), and GA's efficacy comes from assigning higher chances of survival to configurations with a certain combination of more effective parameter values. Initialization through stratified random sampling can let GA find these effective parameter values earlier. We plan to conduct experiments with more initialization methods and using larger search spaces.

When the environment changes, the optimizer should restart and re-optimize a targeted system. For this re-initialization phase, we are investigating how to utilize previous evaluation results, based on our quantitative analysis and characterization of environment changes (Chapter 7.3.1).

### 7.3.4 Stopping Criteria

As shown in Chapter 7.2, some stopping criteria should be included in an auto-tuning algorithm, otherwise it can spend a lot of additional rounds and resources without improving overall performance much if at all. Our proposed stopping criteria include: **(1)** *Time-based stopping criteria*, which let the optimization algorithm stop after a certain time or number of evaluations. **(2)** *Sliding-window (weighted) average*, which stop the optimization algorithm if it fails to find a better configuration within a certain time window. **(3)** *Performance curve stopping criteria*, which conducts

regression analysis on the performance curve of previous evaluations, and tries to predict the performance for upcoming configurations. If the probability of finding better configurations in the near future is low enough, then the algorithm stops. **(4)** *User-specified stopping criteria*. Users may want to specify that they want to achieve X IOPS in terms of throughput. After finding a configuration that meets such requirements, the optimization algorithm can safely stop running.

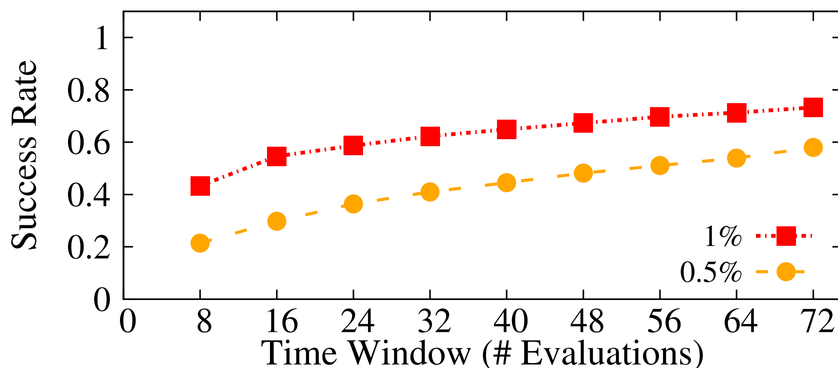


Figure 7.4: Time window based stopping criteria.

Figure 7.4 shows results using sliding-window based stopping criteria. We define a successful stop as one that stopped early and the best configuration found at that point has throughput that is at most K% lower than the best possible one. We again repeated the experiment for 1,000 runs, and the Y axis shows the percentage of runs with success stops. The X axis represents the sliding window size, which means that if the algorithm fails to find a better configuration in X consecutive evaluations, we just stop. We conducted two sets of experiments, with K percent values of 1 and 0.5. For $K = 1$ and window size of 8, around 40% of runs stopped early and successfully. Larger window sizes generally result in better success rates, with a window size of 72 evaluations reaching nearly 70% for $K = 1$ and 60% for $K = 0.5$. Despite some promising results, we believe more sophisticated criteria are needed to stop the algorithm more accurately. The key challenge would be how to determine how close the current solution is to the global best, or whether the algorithm just got stuck in a local optima and simply needs more time for (random) exploration.

In addition to a stopping criteria for the whole optimization process, we also plan to investigate early stopping criteria within each evaluation of a configuration. Evaluating a single configuration for storage systems may take several minutes or even hours. If our optimizer can recognize early that the configuration under evaluation is operating worse compared to known ones, then the optimizer can stop evaluating the current run early.

### 7.3.5 Penalty Functions

Many traditional optimization problems assume that moving from one configuration to another has the same constant cost. In practice, however, this is not always true. Imagine our optimizer finds a configuration with 10% better performance than the current one, but needs to format the underlying file system—requiring a lengthy downtime to backup the data, reformat, then restore the data. Some users may not accept such a cost to gain 10% better performance—but other users might. Therefore, we propose to include the concept of *penalty functions* into our auto-tuning

framework. This penalty roughly correlates to how much downtime the system has to endure while moving to a new configuration. We can broadly categorize system parameters into several penalty classes from least to most onerous. We expect to need only a small number of categories (shown in Table 7.1), but the exact number and granularity is still being investigated.

| Category | Penalty Description | Examples |
|---|---|---|
| 0 | Dynamically changeable | *vm.dirty_ratio* (kernel) |
| 1 | Restarting app or remounting f/s. | *journal option* (Ext4), I/O sched. (kernel) |
| 2 | Rebooting OS | new kernel image installed |
| 3 | Recreating f/s, requiring data migration | *block size* (Ext4), f/s type (kernel) |

Table 7.1: Categories of parameter penalties

To identify the penalties associated with each system parameter, administrators have to label them and also come up with weights for each category in the beginning. Such labels need only be done once for each system parameter, and can then be disseminated publicly. Users, however, would need to assign weights to the various penalty classes in each environment (e.g., more conservative in production, more aggressive in experimental systems); we are investigating several default weights.

Our optimizer can then take into account the penalties when making decisions. One approach is to include the penalty into the optimization objective function, and thus the objective becomes a complex cost formula (similar to our economics-based cost functions [97]), rather than one single system metric like I/O throughput.

### 7.3.6 Machine Learning

Despite the issues discussed in Chapter 2.1, we believe that ML can be helpful in our auto-tuning framework. Figure 7.2 highlights with green circles where we feel that ML can be useful as described next.

**Initialization** As described in §7.3.3, the *Optimizer* has to re-optimize the system when the environment changes. One applicable ML area, named *Transfer Learning*, uses data from prior studies to guide and accelerate current ones [10, 60, 169]. To guide the optimization process more effectively, we are investigating how to transfer evaluation results and workload models (§7.3.1) from one environment to another.

**Stopping criteria** In addition to the stopping criteria described in §7.3.4, we will also develop and test ML-based stopping criteria. As our optimizer picks configurations to evaluate, over time it collects more and more useful data, which can be used for incremental model training. We can use the data to train ML models to predict the performance of system configurations. If the probability of getting a better configuration than the current best one is low, the optimizer can stop the optimization process.

**Selection** ML can help the selection phase in several ways. For example, ML can predict the performance of the next configuration picked, and thus save time on evaluating unpromising ones. In the beginning, the prediction accuracy would be low due to lack of training data; hence the optimizer will need to determine a confidence level in the ML model's accuracy. ML can also help guide the optimizer to pick configurations from unexplored sub-spaces; it can label each configuration according to how many in its neighborhood have been explored. Moreover, ML can help our auto-tuning framework to identify important parameters and automatically select them. In our past work, we observed that parameters have different importance to the optimization objective [43, 94, 132, 148]. By finding out unrelated or less important parameters and removing them during the optimization process, we can *exponentially reduce the search space* [35, 70].

**Restarting criteria** The ML-aided restarting criteria is related to the *Workload Modeler*. Our optimizer needs to restart when the environment changes enough. We propose to develop an intelligent approach with the help of ML techniques, to identify changes in the environment automatically.

### 7.3.7 Visualizer

Visual Analytics (VA) techniques include a mix of automated and user-driven actions to improve the understanding of large-dimensional data sets. The human eye and brain can recognize patterns that computers are yet unable to. By visualizing N-dimensional spaces into 2D/3D, we will discover patterns that will help us (domain experts) develop more effective optimization techniques. VA will help us discover complex correlations among search parameters (i.e., dimensions, attributes, features), enabling us to reduce the overall search space size. Our search spaces are so large because each additional parameter effectively increases the search space exponentially; therefore, reducing the search space must also be done exponentially. We will integrate feature reduction methods to reduce the search space: e.g., parameters that have little impact on overall performance can be removed, and highly correlated parameters can be combined. VA will also help us visualize and hence evaluate the efficacy of any of our proposed methods and options therein: seeing an n-D space in 2D/3D, as its landscape is gradually revealed via more search points; identifying the path any search technique takes through the space and the path's efficiency; enabling domain experts optionally to provide hints to the search process (e.g., to refocus a search in an unexplored area the domain expert believes is promising); visualizing the cumulative and current penalty to move to a new configuration; the location and distribution of initially selected configurations; and more.

# Chapter 8

# Proposed and Future Work

In this chapter, we provide a summary of work that we propose to accomplish in this thesis, as well as future work beyond this thesis.

## 8.1 Proposed Work

Our proposed work consist of three parts: 1) Experimenting with larger parameter space; 2) Design and implement new components for a practical auto-tuning framework; and 3) Design and implement a workload modeler. We re-colored Figure 7.1 and Figure 7.2 here as Figure 8.1 and Figure 8.2, respectively, to provide a clearer overview of our proposed and future work. Components are re-colored based on our project timeline. We plan to completely finish work colored by green and partially finish work colored by yellow. Components colored with red are left for future work beyond this thesis.

**Experimenting with larger and more complex parameter spaces** Our current experiments were conducted on a 9-parameter space with 6,222 unique configurations. We plan to extend our experiments to larger and more complex parameter spaces, which we denote as *Storage V3* and *Storage V4*.

- *Storage V3*: It will consist of all categories of parameters (as shown in Table 7.1) carefully chosen from several representative local file systems. Unlike the default workloads used in our previous experiments (see Table 4.2), which are assigned with relatively small dataset size to speed up the whole exhaustive search process, we will apply various workloads with larger dataset size (at least $2\times$ the RAM size) on *Storage V3*. We will exhaustively evaluate all configurations in *Storage V3*. The collected datasets will allow us to design and test our practical framework.

- *Storage V4*: It will be an even larger parameter space, where exhaustive search is impossible. We will use it to test the efficacy of our auto-tuning framework in real-time.

**Practical auto-tuning framework** We propose to design and implement a practical real-time auto-tuning framework for storage systems, with new features that are missing from traditional black-box optimization techniques. The details are listed as follows:

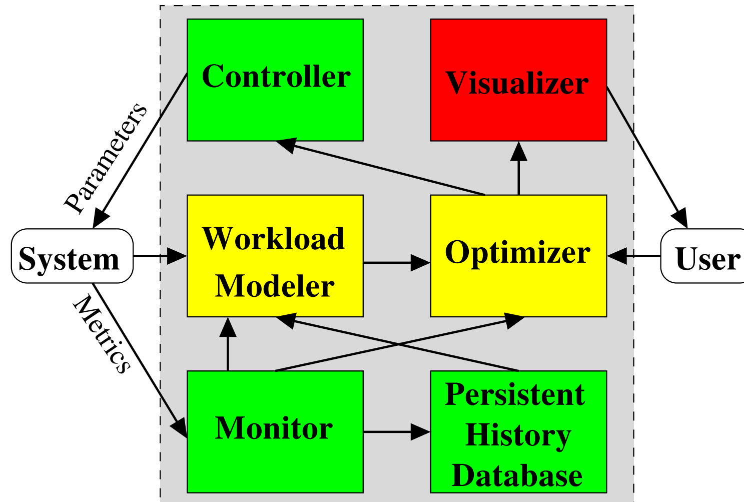


Figure 8.1: Auto-tuning Framework. Components are re-colored based on our project timeline. We plan to completely finish work colored by green; partially finish work colored by yellow. Components colored with red are left for future work beyond this thesis.

- Design efficient initialization methods and test their efficiency on Storage V2 and V3; come up the final solution for use in Storage V4 (real-time). This corresponds to the green box *"Initialize Population"* in Figure 8.2.
- Design practical stopping criteria and test their efficiency on Storage V2 and V3; come up the the final solution for use in Storage V4 (real-time). This corresponds to the green rhombus *"Stop?"* in Figure 8.2.
- Design practical penalty functions using storage domain expertise; design experiments to test their practicality. This corresponds to the green boxes *"penalty"* in Figure 8.2.
- Investigate approaches to help improve the overall efficiency of optimization algorithms on auto-tuning storage systems. One example is to do feature selection on the fly to help reduce parameter spaces and thus make optimization more efficient. This part is denoted as yellow circles *"ML"* and yellow box *"ML-aided Selection"* in Figure 8.2.
- The *Controller*, *Monitor*, *Persistent History Database* will be implemented.

**Workload modeling** We propose to design and implement the workload modeler (see Chapter 7.3.1).

- Run different workloads (ideally the more the better) and collect system traces Feature engineering: extract features from system traces (e.g., blktrace) that can characterize a storage workload.
- Design the restarting criteria, which requires the workload modeler to sense the changes in environment. This corresponds to the green rhombus *"Restart?"* in Figure 8.2.

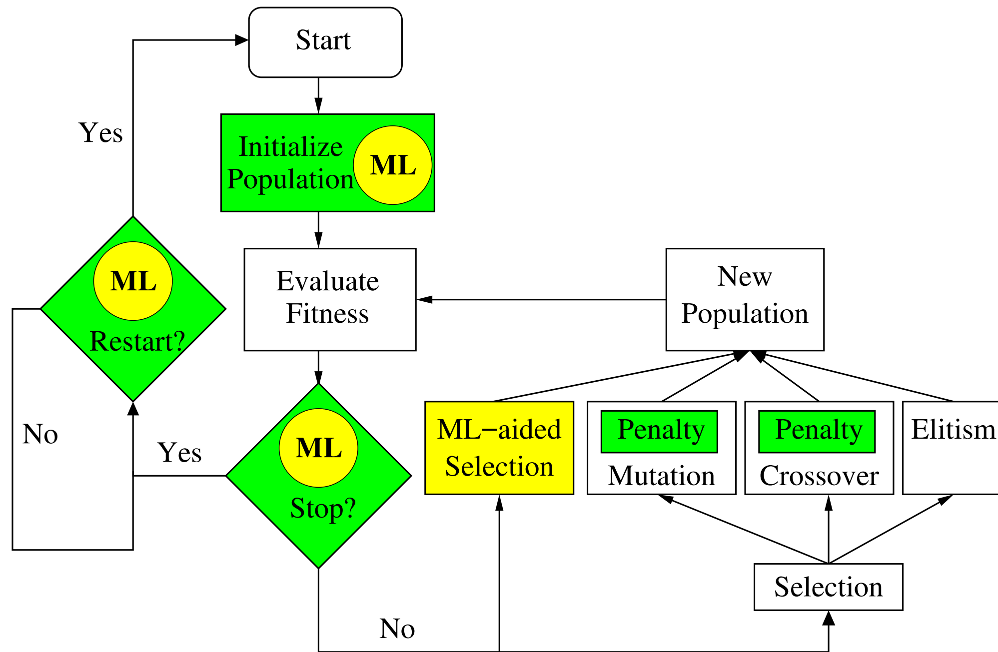


Figure 8.2: Work flow for an enhanced Optimizer (GA). Components are re-colored based on our project timeline. We plan to completely finish work colored by green; partially finish work colored by yellow. Components colored with red are left for future work beyond this thesis.

- If time permits, quantify the similarities between workloads and come up with solutions to transfer and reuse previous evaluation results. This is why the *"Workload Modeler"* is marked as yellow (partially finished by thesis defense) in Figure 8.2.

## 8.2 Future Work

This work can be extended further beyond the scope of the thesis. We see at least the following interesting and promising directions.

- Use visual analytics to help understand our parameter spaces and datasets. It might also be useful in showing intermediate status of optimization algorithms to storage domain experts, who then can make difficult decisions for the auto-tuning framework. This corresponds to the red box *"Visualizer"* in Figure 8.1.

- Find the minimum set of system metrics that can characterize a storage workload. This is not only useful for our auto-tuning framework, but could have broader impacts on understanding and optimizing performance for all storage systems. This will be part of the *"Workload Modeler"* in Figure 8.1.

- Design hybrid algorithms that combine traditional optimization algorithms and ML techniques. This is part of the *"Optimizer"* in Figure 8.1.

# Chapter 9

# Conclusions

Optimizing storage systems can provide significant benefits especially in improving I/O performance. Alas, storage systems are getting more complex, contain many parameters and an immense number of possible configurations; manual tuning is therefore impractical. Worse, many of those parameters are non-linear or non-numeric; traditional linear-regression-based optimization techniques do not work well for such problems. Therefore, in this work, we propose to auto-tune storage system configurations in real-time.

We first performed a comparative study on various black-box optimization algorithms. **(1)** We evaluated *five* popular but different auto-tuning techniques, varied some of their hyper-parameters, and applied them to storage and file systems. **(2)** We show that the speed at which the techniques can find optimal or near-optimal configurations (in terms of throughput) depends on the hardware, software, and workload; this means that no single technique can "rule them all." **(3)** We explain why some techniques appear to work better than others.

In our auto-tuning experiments, we observed that repeated experiments in well-controlled, identical environments could produce results with high variations. Therefore, we then provided the first systematic study on performance variation in benchmarking a modern storage system. We showed that variation is common in storage systems, although its magnitude depends heavily on specific configurations and workloads. Our analysis revealed that block allocation is a major cause of performance variation in Ext4-HDD configurations. From the temporal view, the magnitude of throughput variation also depends on the window size and changes over time. The latency distribution for the same operation type could also vary even over repeated runs of the same experiment. We quantified the correlation between performance and latency variations using a novel approach.

Our preliminary results show that applying black-box optimization techniques to auto-tune storage system configurations are feasible and promising, but also indicate several limitations. Therefore, we propose to design a more intelligent and practical framework for real-time auto-tuning storage systems, with several new components: a workload modeler to characterize a running workload dynamically; stopping criteria to save resources; restarting criteria in response to environment changes; and effective initialization and ML-aided selection methods, including dimensionality reduction, to speed up the auto-tuning process. We conducted simple experiments to demonstrate that our framework is the right direction for auto-tuning storage system in real-time.

Another big contribution of our project is that we are collecting a lot of data on evaluating different storage configurations on various workloads. For more than two years, we have collected a large data-set of over 450,000 data points. In our proposed work, we will conduct experiments

on even more complex parameter space in the near future. All our results will be stored in a carefully designed database. We plan to release our datasets in the future, to facilitate research on understanding and optimizing storage performance.

Finally, it is our thesis that real-time auto-tuning storage systems is important, promising, and feasible with a carefully designed framework to include missing yet critical features. We hope our auto-tuning framework can improve systems' performance efficiency, and save energy and human resources in the long term.

# Bibliography

[1] Emile Aarts and Jan Korst. *Simulated annealing and Boltzmann machines*. New York, NY; John Wiley and Sons Inc., 1988.

[2] Abutalib Aghayev, Mansour Shafaei, and Peter Desnoyers. Skylight—a window on shingled disk operation. *Trans. Storage*, 11(4):16:1–16:28, October 2015.

[3] Abutalib Aghayev, Theodore Ts'o, Garth Gibson, and Peter Desnoyers. Evolving ext4 for shingled disks. In *Proceedings of the 15th USENIX Conference on File and Storage Technologies (FAST)*, pages 105–120, Santa Clara, CA, February/March 2017. USENIX Association.

[4] Ravindra K Ahuja and James B Orlin. Developing fitter genetic algorithms. *INFORMS Journal on Computing*, 9(3):251–253, 1997.

[5] Omid Alipourfard, Hongqiang Harry Liu, Jianshu Chen, Shivaram Venkataraman, Minlan Yu, and Ming Zhang. Cherrypick: Adaptively unearthing the best cloud configurations for big data analytics. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pages 469–482. USENIX Association, 2017.

[6] Guillermo A. Alvarez, Elizabeth Borowsky, Susie Go, Theodore H. Romer, Ralph Becker-Szendy, Richard Golding, Arif Merchant, Mirjana Spasojevic, Alistair Veitch, and John Wilkes. Minerva: An automated resource provisioning tool for large-scale storage systems. *ACM Trans. Comput. Syst.*, 19(4):483–518, November 2001.

[7] Terry Anderson. *The theory and practice of online learning*. Athabasca University Press, 2008.

[8] R. H. Arpaci-Dusseau, E. Anderson, N. T., D. E. Culler, J. M. Hellerstein, D. Patterson, and K. Yelick. Cluster I/O with river: making the fast case common. In *Workshop on Input/Output in Parallel and Distributed Systems*, pages 10–22, Atlanta, GA, May 1999.

[9] Remzi H. Arpaci-Dusseau and Andrea C. Arpaci-Dusseau. *Operating Systems: Three Easy Pieces*. Arpaci-Dusseau Books, 0.91 edition, May 2015.

[10] Rémi Bardenet, Mátyás Brendel, Balázs Kégl, and Michele Sebag. Collaborative hyperparameter tuning. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 199–207. PMLR, 17–19 Jun 2013.

[11] Babak Behzad, Joey Huchette, Huong Luu, Ruth Aydt, Quincey Koziol, Mr Prabhat, Suren Byna, Mohamad Chaarawi, and Yushu Yao. Auto-tuning of parallel io parameters for hdf5 applications. In *Proceedings of the 2012 SC Companion: High Performance Computing, Networking Storage and Analysis*, SCC '12, pages 1430–, Washington, DC, USA, 2012. IEEE Computer Society.

[12] Babak Behzad, Huong Vu Thanh Luu, Joseph Huchette, Surendra Byna, Prabhat, Ruth Aydt, Quincey Koziol, and Marc Snir. Taming parallel i/o complexity with auto-tuning. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, SC '13, pages 68:1–68:12, New York, NY, USA, 2013. ACM.

[13] Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*, 2016.

[14] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb):281–305, 2012.

[15] James S. Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems 24*, pages 2546–2554, 2011.

[16] MC Bhuvaneswari. *Application of Evolutionary Algorithms for Multi-objective Optimization in VLSI and Embedded Systems*. Springer, 2015.

[17] D. Boutcher and A. Chandra. Does virtualization make disk scheduling passé? In *Proceedings of the 1st USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage '09)*, October 2009.

[18] Mark F Bramlette. Initialization, mutation and selection methods in genetic algorithms for function optimization. In *ICGA*, pages 100–107, 1991.

[19] Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

[20] BTRFS. `http://btrfs.wiki.kernel.org/`.

[21] Axel Busch, Qais Noorshams, Samuel Kounev, Anne Koziolek, Ralf Reussner, and Erich Amrehn. Automated workload characterization for i/o performance analysis in virtualized environments. In *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering*, pages 265–276. ACM, 2015.

[22] M. Cao, T. Y. Ts'o, B. Pulavarty, S. Bhattacharya, A. Dilger, and A. Tomas. State of the art: Where we are with the Ext3 filesystem. In *Proceedings of the Linux Symposium*, Ottawa, ON, Canada, July 2005.

[23] Zhen Cao. Parametric optimization of storage systems. Technical Report FSL-16-01, Computer Science Department, Stony Brook University, January 2016.

[24] Zhen Cao, Vasily Tarasov, Hari Raman, Dean Hildebrand, and Erez Zadok. On the performance variation in modern storage stacks. In *Proceedings of the 15th USENIX Conference on File and Storage Technologies (FAST)*, pages 329–343, Santa Clara, CA, February/March 2017. USENIX Association.

[25] R. Card, T. Ts'o, and S. Tweedie. Design and implementation of the second extended filesystem. In *Proceedings to the First Dutch International Symposium on Linux*, Seattle, WA, December 1994.

[26] George Casella and Roger L Berger. *Statistical Inference*, volume 2. Duxbury Pacific Grove, CA, 2002.

[27] Vladimír Černỳ. Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. *Journal of optimization theory and applications*, 45(1):41–51, 1985.

[28] Kevin K. Chang, Abhijith Kashyap, Hasan Hassan, Saugata Ghose, Kevin Hsieh, Donghyuk Lee, Tianshi Li, Gennady Pekhimenko, Samira Khan, and Onur Mutlu. Understanding latency variation in modern DRAM chips: Experimental characterization, analysis, and optimization. In *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*, SIGMETRICS'16, pages 323–336, New York, NY, USA, 2016. ACM.

[29] Min Chen, Shiwen Mao, and Yunhao Liu. Big data: A survey. *Mobile Networks and Applications*, 19(2):171–209, 2014.

[30] Y. Chen, M. Winslett, Y. Cho, and S. Kuo. Automatic parallel i/o performance optimization using genetic algorithms. In *Proceedings of the 7th IEEE International Symposium on High Performance Distributed Computing*, HPDC '98, pages 155–, Washington, DC, USA, 1998. IEEE Computer Society.

[31] Maurice Clerc. *Particle swarm optimization*, volume 93. John Wiley & Sons, 2010.

[32] Yvonne Coady, Russ Cox, John DeTreville, Peter Druschel, Joseph Hellerstein, Andrew Hume, Kimberly Keeton, Thu Nguyen, Christopher Small, Lex Stein, and Andrew Warfield. Falling off the cliff: When systems go nonlinear. In *Proceedings of the 10th Conference on Hot Topics in Operating Systems - Volume 10*, HOTOS'05, 2005.

[33] James Cohoon, John Kairo, and Jens Lienig. Evolutionary algorithms for the physical design of vlsi circuits. In *Advances in evolutionary computing*, pages 683–711. Springer, 2003.

[34] Valentin Dalibard, Michael Schaarschmidt, and Eiko Yoneki. Boat: Building auto-tuners with structured bayesian optimization. In *Proceedings of the 26th International Conference on World Wide Web*, WWW '17, pages 479–488. International World Wide Web Conferences Steering Committee, 2017.

[35] Kenneth Alan De Jong. *Analysis of the behavior of a class of genetic adaptive systems*. PhD thesis, University of Michigan, Ann Arbor, Ann Arbor, MI, USA, 1975.

[36] Pablo de Oliveira Castro, Yuriy Kashnikov, Chadi Akel, Mihail Popov, and William Jalby. Fine-grained benchmark subsetting for system selection. In *Proceedings of Annual IEEE/ACM International Symposium on Code Generation and Optimization*, page 132. ACM, 2014.

[37] Jeffrey Dean and Luiz André Barroso. The tail at scale. *Commun. ACM*, 56(2):74–80, February 2013.

[38] Peter Desnoyers. Empirical evaluation of nand flash memory performance. In *HotStorage '09: Proceedings of the 1st Workshop on Hot Topics in Storage*. ACM, 2009.

[39] Y. Diao, J. L. Hellerstein, A. J. Storm, M. Surendra, S. Lightstone, S. Parekh, and C. Garcia-Arellano. Using MIMO linear control for load balancing in computing systems. In *2004 American Control Conferences*, 2004.

[40] Marco Dorigo and Mauro Birattari. Ant colony optimization. In *Encyclopedia of machine learning*, pages 36–39. Springer, 2010.

[41] Marco Dorigo, Mauro Birattari, and Thomas Stützle. Ant colony optimization. *Computational Intelligence Magazine, IEEE*, 1(4):28–39, 2006.

[42] Fred Douglis, Deepti Bhardwaj, Hangwei Qian, and Philip Shilane. Content-aware load balancing for distributed backup. In *LISA*, 2011.

[43] E. Zadok and A. Arora and Z. Cao and A. Chaganti and A. Chaudhary and S. Mandal. Parametric optimization of storage systems. In *HotStorage '15: Proceedings of the 7th USENIX Workshop on Hot Topics in Storage*, Santa Clara, CA, July 2015. USENIX, USENIX.

[44] Katharina Eggensperger, Matthias Feurer, Frank Hutter, James Bergstra, Jasper Snoek, Holger Hoos, and Kevin Leyton-Brown. Towards an empirical foundation for assessing bayesian optimization of hyperparameters. In *NIPS workshop on Bayesian Optimization in Theory and Practice*, volume 10, 2013.

[45] A.E. Eiben and C.A. Schippers. On evolutionary exploration and exploitation. *Fundam. Inf.*, 35(1-4):35–50, January 1998.

[46] Nosayba El-Sayed, Ioan A. Stefanovici, George Amvrosiadis, Andy A. Hwang, and Bianca Schroeder. Temperature management in data centers: Why some (might) like it hot. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS'12, pages 163–174, New York, NY, USA, 2012. ACM.

[47] Ext4. `http://ext4.wiki.kernel.org/`.

[48] Ext4 documentation. `https://www.kernel.org/doc/Documentation/filesystems/ext4.txt`.

[49] Linux/fs/ext4/ialloc.c. `http://lxr.free-electrons.com/source/fs/ext4/ialloc.c`.

[50] Filebench, 2016. https://github.com/filebench/filebench/wiki.

[51] Terry L Friesz, Hsun-Jung Cho, Nihal J Mehta, Roger L Tobin, and G Anandalingam. A simulated annealing approach to the network design problem with variational inequality constraints. *Transportation Science*, 26(1):18–26, 1992.

[52] RA Gallego, AB Alves, A Monticelli, and R Romero. Parallel simulated annealing applied to long term transmission network expansion planning. *Power Systems, IEEE Transactions on*, 12(1):181–188, 1997.

[53] Shravan Gaonkar, Kimberly Keeton, Arif Merchant, and William H. Sanders. Designing dependable storage solutions for shared application environments. *IEEE Trans. Dependable Secur. Comput.*, 7(4):366–380, October 2010.

[54] Michel Gendreau and Jean-Yves Potvin. *Handbook of metaheuristics*, volume 2. Springer, 2010.

[55] S. Ghemawat, H. Gobioff, and S. T. Leung. The Google file system. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP '03)*, pages 29–43, Bolton Landing, NY, October 2003. ACM SIGOPS.

[56] F. Glover. Tabu Search – Part II. *ORSA Journal on Computing*, 2:4–32, 1990.

[57] Fred Glover. Tabu search: A tutorial. *Interfaces*, 20(4):74–94, 1990.

[58] Fred Glover and Manuel Laguna. *Tabu Search*. Springer, 2013.

[59] David E Goldberg and Robert Lingle. Alleles, loci, and the traveling salesman problem. In *Proceedings of the first international conference on genetic algorithms and their applications*, pages 154–159. Lawrence Erlbaum Associates, Publishers, 1985.

[60] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and D. Sculley. Google vizier: A service for black-box optimization. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, pages 1487–1495. ACM, 2017.

[61] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. http://www.deeplearningbook.org.

[62] John Grefenstette, Rajeev Gopal, Brian Rosmaita, and Dirk Van Gucht. Genetic algorithms for the traveling salesman problem. In *Proceedings of the first International Conference on Genetic Algorithms and their Applications*, pages 160–168. Lawrence Erlbaum, New Jersey (160-168), 1985.

[63] Lov K Grover. A new simulated annealing algorithm for standard cell placement. In *Proceedings of the International Conference on Computer-Aided Design*, pages 378–380, 1986.

[64] Mingzhe Hao, Gokul Soundararajan, Deepak Kenchammana-Hosekote, Andrew A Chien, and Haryadi S Gunawi. The tail at store: a revelation from millions of hours of disk and ssd deployments. In *14th USENIX Conference on File and Storage Technologies (FAST 16)*, pages 263–276, 2016.

[65] Md E. Haque, Yong hun Eom, Yuxiong He, Sameh Elnikety, Ricardo Bianchini, and Kathryn S. McKinley. Few-to-many: Incremental parallelism for reducing tail latency in interactive services. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS'15, pages 161–175, New York, NY, USA, 2015. ACM.

[66] Georges R Harik and Fernando G Lobo. A parameter-less genetic algorithm. In *GECCO*, volume 99, pages 258–267, 1999.

[67] Jun He, Duy Nguyen, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. Reducing file system tail latencies with Chopper. In *Proceedings of the 13th USENIX Conference on File and Storage Technologies*, FAST'15, pages 119–133, Berkeley, CA, USA, 2015. USENIX Association.

[68] Weiping He and David H.C. Du. Smart: An approach to shingled magnetic recording translation. In *Proceedings of the 15th USENIX Conference on File and Storage Technologies (FAST)*, pages 121–134, Santa Clara, CA, February/March 2017. USENIX Association.

[69] J. L. Hellerstein, Y. Diao, S. Parekh, and D. M. Tibury. *Feedback Control of Computing Systems*. Wiley-IEEE Press, 2004.

[70] J. H. Holland. *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence.* U. Michigan Press, 1975.

[71] Ronald L Iman, Jon C Helton, James E Campbell, et al. An approach to sensitivity analysis of computer models, part 1. introduction, input variable selection and preliminary variable assessment. *Journal of quality technology*, 13(3):174–183, 1981.

[72] Myeongjae Jeon, Saehoon Kim, Seung-won Hwang, Yuxiong He, Sameh Elnikety, Alan L. Cox, and Scott Rixner. Predictive parallelization: Taming tail latencies in web search. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR'14, pages 253–262, New York, NY, USA, 2014. ACM.

[73] Young-Jae Jeon, Jae-Chul Kim, Jin-O Kim, Joong-Rin Shin, and Kwang Y Lee. An efficient simulated annealing algorithm for network reconfiguration in large-scale distribution systems. *Power Delivery, IEEE Transactions on*, 17(4):1070–1078, 2002.

[74] N. Joukov, A. Traeger, R. Iyer, C. P. Wright, and E. Zadok. Operating system profiling via latency analysis. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI 2006)*, pages 89–102, Seattle, WA, November 2006. ACM SIGOPS.

[75] Karthik Kambatla, Giorgos Kollias, Vipin Kumar, and Ananth Grama. Trends in big data analytics. *Journal of Parallel and Distributed Computing*, 74(7):2561–2573, 2014. Special Issue on Perspectives on Parallel and Distributed Processing.

[76] A Kapsalis, Vic J Rayward-Smith, and George D Smith. Solving the graphical steiner tree problem using genetic algorithms. *Journal of the Operational Research Society*, pages 397–406, 1993.

[77] M. Karlsson, C. Karamanolis, and X. Zhu. Triage: Performance differentiation for storage systems using adaptive control. *ACM Trans. Storage*, 1(4), 2005.

[78] Kimberly Keeton, Dirk Beyer, Ernesto Brau, Arif Merchant, Cipriano Santos, and Alex Zhang. On the road to recovery: Restoring data after disasters. In *Proceedings of the 1st ACM SIGOPS/EuroSys European Conference on Computer Systems 2006*, EuroSys'06, pages 235–248, New York, NY, USA, 2006. ACM.

[79] James Kennedy. Particle swarm optimization. In *Encyclopedia of Machine Learning*, pages 760–766. Springer, 2010.

[80] James Kennedy and Russell C. Eberhart. Particle swarm optimization. In *Proceedings of the IEEE International Conference on Neural Networks*, pages 1942–1948, 1995.

[81] H. Kim, S. Seshadri, C. L. Dickey, and L. Chiu. Evaluating phase change memory for enterprise storage systems: A study of caching and tiering approaches. In *Proceedings of the 12th USENIX Conference on File and Storage Technologies*, pages 33–45, Berkeley, CA, 2014. USENIX.

[82] S. Kirkpatrick, C D. Gelatt, M. P Vecchi, et al. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.

[83] Ryusuke Konishi, Yoshiji Amagai, Koji Sato, Hisashi Hifumi, Seiji Kihara, and Satoshi Moriai. The linux implementation of a log-structured file system. *ACM SIGOPS Operating Systems Review*, 40(3):102–107, 2006.

[84] Natalio Krasnogor and Jim Smith. A tutorial for competent memetic algorithms: model, taxonomy, and design issues. *Evolutionary Computation, IEEE Transactions on*, 9(5):474–488, 2005.

[85] Sajib Kundu, Raju Rangaswami, Ajay Gulati, Ming Zhao, and Kaushik Dutta. Modeling virtualized applications using machine learning techniques. In *Proceedings of the 8th ACM SIGPLAN/SIGOPS Conference on Virtual Execution Environments*, VEE '12, pages 3–14. ACM, 2012.

[86] Pedro Larrañaga, Cindy M. H. Kuijpers, Roberto H. Murga, Inaki Inza, and Sejla Dizdarevic. Genetic algorithms for the travelling salesman problem: A review of representations and operators. *Artificial Intelligence Review*, 13(2):129–170, 1999.

[87] Changman Lee, Dongho Sim, Jooyoung Hwang, and Sangyeun Cho. F2fs: A new file system for flash storage. In *Proceedings of the 13th USENIX Conference on File and Storage Technologies (FAST)*, pages 273–286, Santa Clara, CA, February 2015. USENIX Association.

[88] Ernest Bruce Lee and Lawrence Markus. Foundations of optimal control theory. Technical report, DTIC Document, 1967.

[89] H. D. Lee, Y. J. Nam, K. J. Jung, S. G. Jung, and C. Park. Regulating I/O performance of shared storage with a control theoretical approach. In *NASA/IEEE Conference on Mass Storage Systems and Technologies (MSST)*. IEEE Society Press, 2004.

[90] Cheng Li, Philip Shilane, Fred Douglis, Darren Sawyer, and Hyong Shim. Assert(!defined(sequential i/o)). In *6th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 14)*, Philadelphia, PA, 2014. USENIX Association.

[91] Jialin Li, Naveen Kr. Sharma, Dan R. K. Ports, and Steven D. Gribble. Tales of the tail: Hardware, os, and application-level sources of tail latency. In *Proceedings of the ACM Symposium on Cloud Computing*, SoCC'14, pages 9:1–9:14, New York, NY, USA, 2014. ACM.

[92] Yin Li, Hao Wang, Xuebin Zhang, Ning Zheng, Shafa Dahandeh, and Tong Zhang. Facilitating magnetic recording technology scaling for data center hard disk drives through filesystem-level transparent local erasure coding. In *Proceedings of the 15th USENIX Conference on File and Storage Technologies (FAST)*, pages 135–148, Santa Clara, CA, February/March 2017. USENIX Association.

[93] Yuxi Li. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*, 2017.

[94] Z. Li, M. Chen, A. Mukker, and E. Zadok. On the trade-offs among performance, energy, and endurance in a versatile hybrid drive. *ACM Transactions on Storage (TOS)*, 11(3), July 2015.

[95] Z. Li, K. M. Greenan, A. W. Leung, and E. Zadok. Power consumption in enterprise-scale backup storage systems. In *Proceedings of the Tenth USENIX Conference on File and Storage Technologies (FAST '12)*, San Jose, CA, February 2012. USENIX Association.

[96] Z. Li, R. Grosu, K. Muppalla, S. A. Smolka, S. D. Stoller, and E. Zadok. Model discovery for energy-aware computing systems: An experimental evaluation. In *Proceedings of the 1st Workshop on Energy Consumption and Reliability of Storage Systems (ERSS'11)*, Orlando, FL, July 2011.

[97] Z. Li, A. Mukker, and E. Zadok. On the importance of evaluating storage systems' $costs. In *Proceedings of the 6th USENIX Conference on Hot Topics in Storage and File Systems*, HotStorage'14, 2014.

[98] Chieh-Jan Mike Liang, Jie Liu, Liqian Luo, Andreas Terzis, and Feng Zhao. RACNet: A high-fidelity data center sensing network. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, SenSys'09, pages 15–28, New York, NY, USA, 2009. ACM.

[99] Jens Lienig and James P Cohoon. Genetic algorithms applied to the physical design of vlsi circuits: A survey. In *Parallel Problem Solving from NaturePPSN IV*, pages 839–848. Springer, 1996.

[100] Fernando G Lobo and David E Goldberg. The parameter-less genetic algorithm in practice. *Information Sciences*, 167(1):217–232, 2004.

[101] Christoffer Loffler, Christopher Mutschler, and Michael Philippsen. Evolutionary algorithms that use runtime migration of detector processes to reduce latency in event-based systems. In *Adaptive Hardware and Systems (AHS), 2013 NASA/ESA Conference on*, pages 31–38. IEEE, 2013.

[102] W. J. Conover M. D. McKay, R. J. Beckman. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979.

[103] Pradipta De Vijay Mann and Umang Mittaly. Handling OS jitter on multicore multithreaded systems. In *Parallel & Distributed Processing Symposium (IPDPS), 2009 IEEE International*, IPDPS'09, pages 1–12. IEEE, 2009.

[104] Olivier Martin, Steve W Otto, and Edward W Felten. Large-step markov chains for the tsp incorporating local search heuristics. *Operations Research Letters*, 11(4):219–224, 1992.

[105] Pinaki Mazumder and Elizabeth M. Rudnick, editors. *Genetic Algorithms for VLSI Design, Layout &Amp; Test Automation*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1999.

[106] M. K. McKusick, W. N. Joy, S. J. Leffler, and R. S. Fabry. A fast file system for UNIX. *ACM Transactions on Computer Systems*, 2(3):181–197, August 1984.

[107] Peter Merz. Memetic algorithms for combinatorial optimization problems: Fitness landscapes and effective search strategies, 2001.

[108] Justin Meza, Qiang Wu, Sanjev Kumar, and Onur Mutlu. A large-scale study of flash memory failures in the field. In *Proceedings of the 2015 ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 2015)*, pages 177–190, Portland, OR, June 2015. ACM.

[109] Sun Microsystems. Lustre file system: High-performance storage architecture and scalable cluster file system white paper. `www.sun.com/servers/hpc/docs/lustrefilesystem_wp.pdun.com/servers/hpc/docs/lustrefilesystem_wp.pdf`, December 2007.

[110] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[111] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[112] Alessandro Morari, Roberto Gioiosa, Robert W Wisniewski, Francisco J Cazorla, and Mateo Valero. A quantitative analysis of OS noise. In *Parallel & Distributed Processing Symposium (IPDPS), 2011 IEEE International*, IPDPS'11, pages 852–863. IEEE, 2011.

[113] Heinz Muhlenbein. Evolution in time and space-the parallel genetic algorithm. In *Foundations of genetic algorithms*. Citeseer, 1991.

[114] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.

[115] Iyswarya Narayanan, Di Wang, Myeongjae Jeon, Bikash Sharma, Laura Caulfield, Anand Sivasubramaniam, Ben Cutler, Jie Liu, Badriddine Khessib, and Kushagra Vaid. Ssd failures in datacenters: What? when? and why? In *Proceedings of the Second ACM Israeli Experimental Systems Conference (SYSTOR '16)*, pages 7:1–7:11, Haifa, Israel, May 2016. ACM.

[116] E. B. Nightingale, K. Veeraraghavan, P. M. Chen, and J. Flinn. Rethink the sync. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI 2006)*, pages 1–14, Seattle, WA, November 2006. ACM SIGOPS.

[117] OpenStack Swift. `http://docs.openstack.org/developer/swift/`.

[118] Christian S. Perone. Pyevolve: A python open-source framework for genetic algorithms. *SIGEVOlution*, 4(1):12–20, November 2009.

[119] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Quoc Le, and Alex Kurakin. Large-scale evolution of image classifiers. *arXiv preprint arXiv:1703.01041*, 2017.

[120] Jackie Rees and Gary J Koehler. An investigation of ga performance results for different cardinality alphabets. In *Evolutionary Algorithms*, pages 191–206. Springer, 1999.

[121] H. Reiser. ReiserFS v.3 whitepaper. `http://web.archive.org/web/20031015041320/http://namesys.com/`.

[122] Bernd Reisleben and Peter Merz. A genetic local search algorithm for solving symmetric and asymmetric traveling salesman problems. In *Evolutionary Computation, 1996., Proceedings of IEEE International Conference on*, pages 616–621. IEEE, 1996.

[123] Alma Riska and Erik Riedel. Disk drive level workload characterization. In *USENIX Annual Technical Conference*, volume 2006, pages 97–102, 2006.

[124] Ohad Rodeh, Josef Bacik, and Chris Mason. BTRFS: The Linux B-tree filesystem. *Trans. Storage*, 9(3):9:1–9:32, August 2013.

[125] Richard P Runyon, Kay A Coleman, and David J Pittenger. *Fundamentals of behavioral statistics* . McGraw-Hill, 2000.

[126] Anooshiravan Saboori, Guofei Jiang, and Haifeng Chen. Autotuning configurations in distributed systems for performance improvements using evolutionary strategies. In *Proceedings of the 2008 The 28th International Conference on Distributed Computing Systems*, ICDCS '08, pages 769–776, Washington, DC, USA, 2008. IEEE Computer Society.

[127] Sadiq M Sait, Mahmood R Minhas, Junhaid Khan, et al. Performance and low power driven vlsi standard cell placement using tabu search. In *Evolutionary Computation, 2002. CEC'02. Proceedings of the 2002 Congress on*, volume 1, pages 372–377. IEEE, 2002.

[128] Ricardo Santana, Raju Rangaswami, Vasily Tarasov, and Dean Hildebrand. A fast and slippery slope for file systems. In *Proceedings of the 3rd Workshop on Interactions of NVM/FLASH with Operating Systems and Workloads*, INFLOW '15, pages 5:1–5:8, New York, NY, USA, 2015. ACM.

[129] F. Schmuck and R. Haskin. GPFS: A shared-disk file system for large computing clusters. In *Proceedings of the First USENIX Conference on File and Storage Technologies (FAST '02)*, pages 231–244, Monterey, CA, January 2002. USENIX Association.

[130] Bianca Schroeder, Raghav Lagisetty, and Arif Merchant. Flash reliability in production: The expected and the unexpected. In *Proceedings of the 14th USENIX Conference on File and Storage Technologies (FAST)*, pages 67–80, Santa Clara, CA, February 2016. USENIX Association.

[131] Carl Sechen. *VLSI placement and global routing using simulated annealing*, volume 54. Springer Science & Business Media, 2012.

[132] P. Sehgal, V. Tarasov, and E. Zadok. Evaluating performance and energy in file system server workloads. In *Proceedings of the USENIX Conference on File and Storage Technologies (FAST)*, pages 253–266, San Jose, CA, February 2010. USENIX Association.

[133] Bumjoon Seo, Sooyong Kang, Jongmoo Choi, Jaehyuk Cha, Youjip Won, and Sungroh Yoon. Io workload characterization revisited: A data-mining approach. *IEEE Transactions on Computers*, 63(12):3026–3038, 2014.

[134] Burr Settles. *Active Learning*. Morgan & Claypool Publishers, 2012.

[135] SGI. XFS filesystem structure. `http://oss.sgi.com/projects/xfs/papers/xfs_filesystem_structure.pdf`.

[136] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando de Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.

[137] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.

[138] Kai Shen, Ming Zhong, and Chuanpeng Li. I/o system performance debugging using model-driven anomaly characterization. In *Proceedings of the USENIX Conference on File and Storage Technologies (FAST)*, San Francisco, CA, December 2005. USENIX Association.

[139] Scikit-Optimize. https://scikit-optimize.github.io/.

[140] T Starkweather, S Mcdaniel, D Whitley, K Mathias, D Whitley, et al. A comparison of genetic sequencing operators. In *Proceedings of the fourth International Conference on Genetic Algorithms*, 1991.

[141] John D. Strunk, Eno Thereska, Christos Faloutsos, and Gregory R. Ganger. Using utility to provision storage systems. In *Proceedings of the 6th USENIX Conference on File and Storage Technologies*, FAST'08, pages 21:1–21:16, Berkeley, CA, USA, 2008. USENIX Association.

[142] Lalith Suresh, Marco Canini, Stefan Schmid, and Anja Feldmann. C3: Cutting tail latency in cloud data stores via adaptive replica selection. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation*, NSDI'15, pages 513–527, Berkeley, CA, USA, 2015. USENIX Association.

[143] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press Cambridge, 1998.

[144] A. Sweeney, D. Doucette, W. Hu, C. Anderson, M. Nishimoto, and G. Peck. Scalability in the XFS file system. In *Proceedings of the Annual USENIX Technical Conference*, pages 1–14, San Diego, CA, January 1996.

[145] *sync(8) - Linux manual page*. https://linux.die.net/man/8/sync.

[146] V. Tarasov, S. Bhanage, E. Zadok, and M. Seltzer. Benchmarking file system benchmarking: It *is* rocket science. In *Proceedings of HotOS XIII:The 13th USENIX Workshop on Hot Topics in Operating Systems*, Napa, CA, May 2011.

[147] Vasily Tarasov, Zhen Cao, Ming Chen, and Erez Zadok. The dos and don'ts of file system benchmarking. *FreeBSD Journal*, January/February, 2016.

[148] Vasily Tarasov, Abhishek Gupta, Kumar Sourav, Sagar Trehan, and Erez Zadok. Terra incognita: On the practicality of user-space file systems. In *HotStorage '15: Proceedings of the 7th USENIX Workshop on Hot Topics in Storage*, Santa Clara, CA, July 2015. USENIX, USENIX.

[149] Vasily Tarasov, Erez Zadok, and Spencer Shepler. Filebench: A flexible framework for file system benchmarking. *;login: The USENIX Magazine*, 41(1):6–12, March 2016.

[150] TensorFlow. https://www.tensorflow.org/.

[151] Olivier Thas. *Comparing distributions*. Springer, 2010.

[152] A. Traeger, I. Deras, and E. Zadok. DARC: Dynamic analysis of root causes of latency distributions. In *Proceedings of the 2008 International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 2008)*, pages 277–288, Annapolis, MD, June 2008. ACM.

[153] Stephen Tweedie. Ext3, journaling filesystem. In *Ottawa Linux Symposium*, July 2000. http://olstrans.sourceforge.net/release/OLS2000-ext3/OLS2000-ext3.html.

[154] Balajee Vamanan, Hamza Bin Sohail, Jahangir Hasan, and T. N. Vijaykumar. TimeTrader: Exploiting latency tail to save datacenter energy for online search. In *Proceedings of the 48th International Symposium on Microarchitecture*, MICRO'48, pages 585–597, New York, NY, USA, 2015. ACM.

[155] Dana Van Aken, Andrew Pavlo, Geoffrey J. Gordon, and Bohan Zhang. Automatic database management system tuning through large-scale machine learning. In *Proceedings of the 2017 ACM International Conference on Management of Data*, SIGMOD '17, pages 1009–1024, 2017.

[156] Peter J Van Laarhoven and Emile H Aarts. *Simulated annealing: theory and applications*, volume 37. Springer Science & Business Media, 1987.

[157] Matej Črepinšek, Shih-Hsi Liu, and Marjan Mernik. Exploration and exploitation in evolutionary algorithms: A survey. *ACM Comput. Surv.*, 45(3):35:1–35:33, July 2013.

[158] Mengzhi Wang, Kinman Au, Anastassia Ailamaki, Anthony Brockwell, Christos Faloutsos, and Gregory R. Ganger. Storage device performance prediction with cart models. In *The IEEE Computer Society's 12th Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems. (MASCOTS)*, pages 588–595, 2004.

[159] S. Weil, S. Brandt, E. Miller, D. Long, and C. Maltzahn. Ceph: A scalable, high-performance distributed file system. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI 2006)*, pages 307–320, Seattle, WA, November 2006. ACM SIGOPS.

[160] Darrell Whitley, Keith Mathias, and Patrick Fitzhorn. Delta coding: An iterative search strategy for genetic algorithms. In *ICGA*, volume 91, pages 77–84. Citeseer, 1991.

[161] Latin Hypercube Sampling. https://en.wikipedia.org/wiki/Latin_hypercube_sampling.

[162] DF Wong, Hon Wai Leong, and HW Liu. *Simulated annealing for VLSI design*, volume 42. Springer Science & Business Media, 2012.

[163] H-S Philip Wong, Simone Raoux, SangBum Kim, Jiale Liang, John P Reifenberg, Bipin Rajendran, Mehdi Asheghi, and Kenneth E Goodson. Phase change memory. *Proceedings of the IEEE*, 98(12):2201–2227, Dec 2010.

[164] Bowei Xi, Zhen Liu, Mukund Raghavachari, Cathy H. Xia, and Li Zhang. A smart hill-climbing algorithm for application server configuration. In *Proceedings of the 13th International Conference on World Wide Web*, WWW '04, pages 287–296, New York, NY, USA, 2004. ACM.

[165] Yunjing Xu, Zachary Musgrave, Brian Noble, and Michael Bailey. Bobtail: Avoiding long tails in the cloud. In *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation*, NSDI'13, pages 329–342, Berkeley, CA, USA, 2013. USENIX Association.

[166] Ji Xue, Feng Yan, A. Riska, and E. Smirni. Proactive management of systems via hybrid analytic techniques. In *Cloud and Autonomic Computing (ICCAC), 2015 International Conference on*, pages 137–148, Sept 2015.

[167] Ji Xue, Feng Yan, Alma Riska, and Evgenia Smirni. Storage workload isolation via tier warming: How models can help. In *11th International Conference on Autonomic Computing (ICAC 14)*, pages 1–11, Philadelphia, PA, June 2014. USENIX Association.

[168] Oceane Bel Ethan L. Miller Darrell D. E. Long Yan Li, Kenneth Chang. Capes: Unsupervised system performance tuning using neural network-based deep reinforcement learning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '17, 2017.

[169] Dani Yogatama and Gideon Mann. Efficient transfer learning method for automatic hyperparameter tuning. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33 of *Proceedings of Machine Learning Research*, pages 1077–1085, Reykjavik, Iceland, 22–25 Apr 2014. PMLR.

[170] Yang Yu, Hong Qian, and Yi-Qi Hu. Derivative-free optimization via classification. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, pages 2286–2292. AAAI Press, 2016.