What Distributed Systems Say: A Study of Seven Spark Application Logs*

Sina Gholamian
University of Waterloo
Waterloo, ON, Canada
sgholamian@uwaterloo.ca

Paul A. S. Ward
University of Waterloo
Waterloo, ON, Canada
pasward@uwaterloo.ca

arXiv:2108.08395v1 [cs.DC] 18 Aug 2021

***Abstract*—Execution logs are a crucial medium as they record runtime information of software systems. Although extensive logs are helpful to provide valuable details to identify the root cause in postmortem analysis in case of a failure, this may also incur performance overhead and storage cost. Therefore, in this research, we present the result of our experimental study on seven Spark benchmarks to illustrate the impact of different logging verbosity levels on the execution time and storage cost of distributed software systems. We also evaluate the log effectiveness and the information gain values, and study the changes in performance and the generated logs for each benchmark with various types of distributed system failures. Our research draws insightful findings for developers and practitioners on how to set up and utilize their distributed systems to benefit from the execution logs.**

***Index Terms*—logging statement, log verbosity level, log4j, logging cost analysis, information gain, entropy, distributed systems, system failure, Spark**

I. Introduction and Motivation

The rapid growth of processing requirements and data scale in computing systems has contributed to the development and adaptation of large-scale, parallel, and distributed computation and storage platforms, *e.g.*, Apache Spark and Hadoop Distributed File System (HDFS). Laterally, as the size of the data and computing systems grow, and they become more distributed in nature, evaluating their reliability and performance becomes more daunting. As such, execution log files and instrumentation of the source code are important origins of information for dependability analysis and gaining insight into the runtime state of the system. Execution logs have advantages over instrumentation, as they are readily available, do not require access to the source code, and do not introduce perturbation [30]. However, instrumentation requires access to the source code, and it incurs perturbation due to the added instrumentation code.

Logging is an important integral part of the software development process to record necessary run-time information [14, 19]. Software developers insert logging statements into the source code to record a wealth of information such as variable values, state of the system, and error messages. Developers and system operators use this information for different purposes, among them failure and performance diagnosis [12, 43]. Although logging has proven benefits, it can incur system costs. Excessive logging can cause system overhead, such as CPU and I/O consumption. Contrarily, logging too little may miss important information and degrade the usefulness of execution logs [14]. Authors of [22] described a typical online system at Microsoft that could produce execution logs in the terabyte order-of-magnitude per day. As such, this high volume of logs can impair the quality of service for such systems. To address the trade-offs associated with the overhead of logging, well-known libraries, such as Apache Log4j [5] and SLF4j [6], provide facilities for different levels and granularities of logging. The libraries provide different verbosity levels to dynamically control the number of logging statements being ultimately outputted to the log file on the storage medium. As each logging statement comes with a verbosity level, the logging library filters log messages by comparing the log statement's level with the dynamic log level specified by the user. Log4j has six verbosity levels available to the developers by default: *fatal, error, warn, info, debug, and trace*. Figure 1 shows an example of a logging statement from Spark with *info* verbosity level and its end product in the log file. In addition, each logging statement consists of a constant part, *i.e.*, *"Executor added: on with core(s)"*, and a variable part, *i.e.*, *"fullId"*.

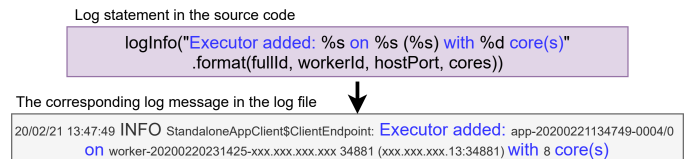

Fig. 1: Log statement and end product in the log file.

Log levels represent a measure of the importance of the messages. For example, less verbose levels (*i.e.*, *fatal*, *error*, and *warn*) are used to warn the user when a potential problem happens in the system. On the other side, more verbose levels such as *info*, *debug*, and *trace* are utilized to track more general system events and information or detailed debugging. Considering the flexibility that each log level brings, our goal in this research is to quantitatively measure the cost, in terms of storage, execution overhead, and information gain (IG) of log files while the distributed system is running under different log verbosity levels. Ultimately, we aim to reach a guideline on implications for developers and practitioners on

*We give the title's credit to the influential Oliner and Stearley's paper, "What Supercomputers Say: A Study of Five System Logs [36]."

how to utilize the logs in different verbosity level decisions while developing or operating distributed software systems in normal scenarios and in presence of failures. We guide our research with the following research questions (RQs):

- **RQ1:** *what is the quantitative cost of logging in terms of computation time (CT) and storage overhead (SO)?* (§III)
- **RQ2:** *how much information is gained from different log verbosity levels (VLs)?* (§IV)
- **RQ3:** *how the characteristics of logs change with distributed failures, i.e., distributed computation and storage failure? Does the entropy of logs increase when a failure happens?* (§V)

For each RQ, we discuss the practical findings of our analysis and their implications for developers and practitioners on how to utilize the execution logs. Our research provides insight on how to choose the level of logging, and ultimately control the amount of generated logs and the information gain, and how the failures can be detected with entropy values. In addition, we provide a discussion on our findings and the implications for future improvements in distributed systems and their scheduling in case of system failures (§V-D). With the motivation of helping developers and practitioners to gain more insight into the content of execution logs, and to make more deliberate logging level decisions, we pursue the following contributions in this paper: **(1)** We evaluate the performance and cost of logging for Spark under a set of batch and iterative workloads with different characteristics to calculate the overall execution time overhead and volume of generated logs (**RQ1**). **(2)** We calculate the information gained from the log files in different VLs based on their entropies and natural language processing (NLP) of logs with n-gram models and provide insights on how to make logging level decisions based on the observed cost and information gain from the log files (**RQ2**). **(3)** We introduce a comprehensive set of distributed system failures and evaluate the changes in execution log characteristics and entropy values, and provide insights on practical outcomes of our analysis for how to utilize execution logs to pinpoint failures (**RQ3**). Lastly, we release our labeled failure logs to encourage and enable further research in this field [1].

II. Approach and Setup

In this section, we present our approach and characterize the systems, their configurations, and the workloads that we use to conduct our study. Figure 2 outlines the steps involved in our study. We categorize the logging cost into two system aspects: 1) execution overhead and 2) storage cost. We run seven Spark benchmarks with different log verbosity levels and calculate the execution times and the size of the generated logs. We then utilize Shannon's entropy theory [40] and n-gram models [2] to measure the information gain by calculating entropies for different log levels with and without failures.

Apache Spark. Since its introduction, Spark has been widely adopted as a big-data, distributed, and parallel processing framework. Spark builds upon Hadoop's MapReduce model and brings extra flexibility and improved

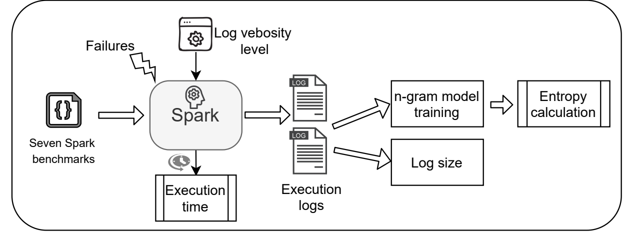

Fig. 2: Our approach for measuring the cost and effectiveness of the logs.

performance. Additionally, Spark provides interfaces to other big-data platforms such as Hadoop's distributed file system, HDFS. To achieve higher performance and as an improvement to Hadoop MapReduce, Spark utilizes Resilient Distributed Datasets (RDDs) [46], which retain the intermediate results in main memory, and therefore, reduces the overhead caused by the disk and network [47]. This optimization benefits Spark the most in iterative tasks (*e.g.*, Transitive Closure), as the following stages of the task rely on the intermediate results from the prior stages. Because of its improved performance and widespread use, we deploy a Spark cluster to perform our study.

Hardware. Table I and Fig. 3 show the main hardware characteristics and the architecture of our deployed cluster, respectively. Each node in the cluster has 12 (12*2 hyper-threaded) cores, 32 GBs of memory, and 2 storage disks of 1 TB each. We evaluate the benchmarks on a cluster of 4 commodity machines illustrated in Fig 3. Each machine is equipped with dual 2.40GHz Intel Xeon E5-2620 CPUs, supporting a total of 24 hyperthreads per machine and a 1Gbps NIC. All servers run Ubuntu Server 16.04.6 LTS 64-bit with kernel version 4.4.0-159-generic.

Name	Role	Cores	Memory	Disk	Local IP
styx01	Master/NameNode	12 (24 HT)	64 GB	2*1 TB	192.168.210.11
styx02	Worker/DataNode	12 (24 HT)	64 GB	2*1 TB	192.168.210.12
styx03	Worker/DataNode	12 (24 HT)	64 GB	2*1 TB	192.168.210.13
styx04	Worker/DataNode	12 (24 HT)	64 GB	2*1 TB	192.168.210.14

TABLE I: Styx cluster for Spark computation and HDFS.

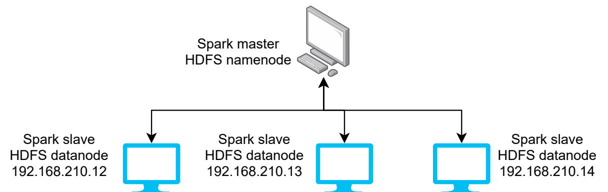

Fig. 3: Design of the distributed cluster, consisting of one master/name node and three slave/data nodes.

Framework setup. In our experiments, we use Spark 2.4.4 and Hadoop 2.9.2. We use *styx01* as a dedicated server for the Spark *master* and HDFS *name node*, while having one Spark *slave* and HDFS *data node* on each of the three other machines, *styx02, styx03, and styx04*. Hadoop file system block size is 128 MBs and replication is set to 3. We set up each Spark slave to use 12 available cores and up to half of the available memory *(i.e.*, 32 GBs out of 64 GBs) on each of the nodes. The frameworks have been carefully configured according to their corresponding user guides and the characteristics of the system (*e.g.*, number of CPU

Framework	Parameter	Value
HDFS	block size	128 MBs
HDFS	Replication factor	3
Spark	Worker/Executor per node	1
Spark	Cores per executor	24 HT
Spark	Memory per executor	32 GBs
Spark	Driver memory	10 GBs

TABLE II: Main Spark and HDFS settings.

cores and memory size). Table II summarizes the related configurations for HDFS and Spark nodes.

Benchmarks. In this research, we experiment on seven different Spark benchmarks and provide a brief explanation of each one in the following: ❶ **WordCount (WC)**, which counts the number of times each word appears in the input dataset. By applying transformations such as *'.reduceByKey()'* on RDDs, WordCount outputs a dataset of (word, value) pairs, saved to a file on HDFS. WordCount is a popular Spark's benchmark that allows us to assess CPU and I/O costs associated with different levels of logging. ❷ **TeraSort (TS),** which sorts randomly generated rows of key-value (KV) pairs with each KV being 100 bytes. The TeraSort implementation and its random data generator engine, TeraGen, are both adopted from [7]. ❸ **TransitiveClosure (TC),** checks and implements linear transitive closure (LTC) on a graph, iteratively. For example, if *x, y, and z* are three vertices, and *(x,y)* and *(y,z)* represent edges between *x and y*, and *y and z*, respectively, for satisfying the transitive closure property, a new edge is added between *x and z*. LTC grows paths by one edge, by joining the graph's edges with the already-discovered paths in each iteration. TC is an iterative CPU-intensive workload. ❹ **PageRank (PR)** is an iterative graph algorithm that ranks URLs by considering the number and rank of URLs referring to it. For example, the more URLs with higher ranks refer to a URL under consideration (URLUC), the URLUC's rank increases. For the PageRank's implementation, we use the implementation provided with the Spark's example package. ❺ **TestDFSIO (DF),** which is a benchmark designed to evaluate the I/O (read/write) performance by using Spark's tasks to read and write multiple files in parallel. The benchmark aims to read and write an even amount of data to HDFS on each node in the cluster. The implementation is adapted from [8]. ❻ **GradientBoostingClassificationTrees (GC),** which is a machine learning algorithm for classification, that generates a prediction model as an ensemble of decision trees. In this use case, the number of classes is set to two, the depth of the trees is set to five, and we perform 200 iterations for the model training. ❼ **LinearDiscriminantAnalysisClustering (LD),** which implements LDA clustering algorithm, *i.e.*, unlabeled data, that clusters the input documents into three different topics.

Table III summarizes the benchmarks used in the experiments, along with their characterization such as CPU or I/O (disk and network) intensive, and if the computation happens iteratively. The sizes of the input datasets are also shown, and we refer to the benchmarks with their abbreviation in the rest of the paper, as shown in Table III. During the benchmark selection, we were deliberate to include a variety of workloads such as Spark's conventional benchmarks, *e.g.*, WC and PR, and machine learning ones, *e.g.*, GC and LD.

III. RQ1: Cost of Logging

There exists various qualitative and quantitative metrics for assessing logging cost. Quantitative metrics consider the overhead of logging on different subsystems of the computing systems, *e.g.*, CPU and I/O overhead [12], and qualitative metrics assess the cost of logging in terms of developer and user experience, such as the cost of revealing private information through logs [27]. In our work, we focus on quantitative measurement of the logging cost and for this purpose, we conduct a set of experiments to evaluate the impact of logging verbosity level on the size of the generated logs, as well as the effect on the performance of the Spark. To measure the cost of logging, we run multiple Spark benchmarks on our commodity cluster and calculate the logging cost in terms of the size of the generated log and the execution time for each benchmark. The measured computation time and storage values in this section serve as a baseline for comparison on further RQs.

A. *Computation time (CT)*

Figures 4a-4g show the violin chart and interquartile range with mean values (noted in text) for execution time of different benchmarks. We also perform Kruskal-Wallis test [26] to reject the null hypothesis and ensure statistically significant values, *i.e.,* $p \leq 0.05$. Violin chart improves on boxplot chart by providing the width in the graph as the density of points in the experiments. The vertical axes represent the execution time in minutes, and the horizontal axes show different verbosity levels (VLs). We select to show *info* and *trace* VLs as the former is generally enabled by default and presents the normal mode of logging, and the latter represents the maximum amount of logging which is widely utilized during failure diagnosis [10, 48]. This provides a picture of the lower and upper bounds of logging. Besides, we present the data for logging *off* to present the other end of the spectrum (least amount of logging) compared to the *trace* level (most amount of logging). The overall trend shows the more verbose VLs result in higher execution times. Because *trace* VL enables more detailed logging and executes additional lines of code (and more I/O system calls), it incurs a minor but noticeable execution time overhead across different benchmarks, when compared to *info* level.

B. *Storage overhead (SO)*

Figures 4h-4n show the size of generated execution logs for different benchmarks. The vertical axes represent the size of the log file in MB, and the horizontal axes show different verbosity levels. The sizes show the aggregated logs for all the nodes (*i.e.,* master and workers) of the Spark cluster. Similar to computation times, *trace* level enables more detailed logging, which results in significantly higher volumes of log data when compared with less verbose log levels.

Results review. Based on the conducted experiments, we observe that the execution time of the benchmarks increases

Benchmark (abbrv.)	Task type	Input data size	Notes
WordCount (WC)	CPU and I/O intensive	52 GBs	The output is pairs of (word, count) written to HDFS.
TeraSort (TS)	Iterative, I/O, and CPU intensive	2 GBs	Sorts randomly generated (key, value) pairs, and the size of each pair is 100 bytes.
TransitiveClosure (TC)	Iterative and CPU intensive	small (few KBs)	Calculates the transitive closure on a randomly generated graph with 200 edges and 100 vertices.
PageRank (PR)	Iterative and CPU intensive	40 MBs	Ranks web pages based on their popularity.
DFSIO (DF)	I/O intensive	20 × 1 GB	Writes and then reads 20 files of one GB each to HDFS.
GradientBoostingClassification Trees (GC)	Iterative and CPU intensive	small (205 KBs)	Trains 200 decision trees with the depth of five for classification of a decision problem, *i.e.*, yes (1) or no (0).
LinearDiscriminantAnalysis Clustering (LD)	Iterative and CPU intensive	21 MBs	Clusters the input data into three topics using LDA.

TABLE III: Benchmark characteristics.

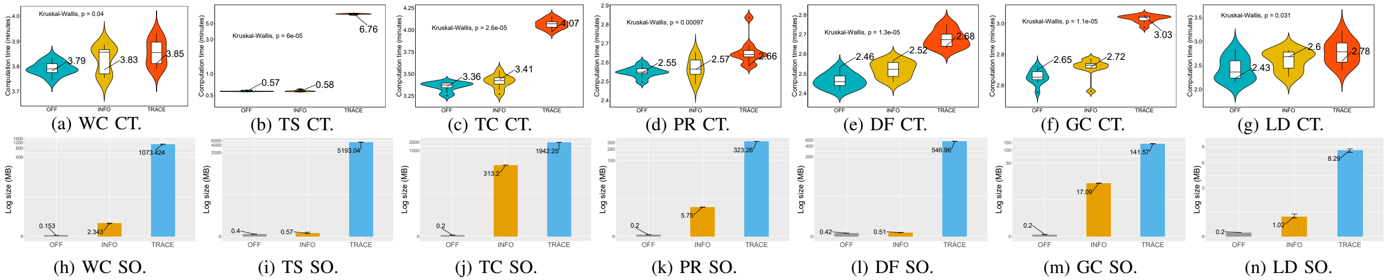

(a) WC CT. (b) TS CT. (c) TC CT. (d) PR CT. (e) DF CT. (f) GC CT. (g) LD CT.

(h) WC SO. (i) TS SO. (j) TC SO. (k) PR SO. (l) DF SO. (m) GC SO. (n) LD SO.

Fig. 4: Computation time (CT) and storage overhead (SO) for different benchmarks.

as the log level becomes more verbose. This observation is congruent with the knowledge that the execution of logging messages infers extra CPU, I/O, and storage cost. Overall, excluding TS, we see on average 8.01% overhead in execution time when the *trace* log level is enabled versus the *info* level. For the log file size, we notice on average a ∼268X increase in the volume of the generated logs for the *trace* log level versus the *info* verbosity level. We did a further investigation to better understand the ∼12X increase in computation time for TS, and we observed the significant amount of generated logs for TS in *trace* level when compared to *info* (*i.e.*, 5 GBs *vs.* 0.5 MBs). Because TS is a CPU-intensive benchmark, we rationalize that its CT suffers noticeably due to the significant amount of logs outputted in the *trace* level. Comparing CT and SO values for different benchmarks, we observe that the amount of generated logs in different VLs is benchmark dependent, and CPU-intensive applications (*e.g.*, TS, TC) observe a higher slowdown due to more verbose logging.

SO mitigation. Prior work has shown that due to the high level of repetitiveness in log files, they can benefit from large compression ratios, up to 84% [44]. Therefore, the noticeable difference in storage cost can be mitigated by the compression of log files to ∼43X.

> **Finding 1.** *Overall, we observe on average 8.01% and ∼268X overhead in the execution time and storage when the trace log level is enabled versus the info level, respectively, and CPU-intensive workloads suffer more from a higher degree of logging.*

> **Implications.** *Considering the trade-offs, if the worst-case 8.01% execution time is acceptable, by utilizing log rolling, compression, and continuous achieving, the storage overhead of more verbose logging can be further lowered.*

C. RAM Disk

We used the hard disk drive (HDD - TOSHIBA MG04ACA200E - 7,200 RPM) as the storage medium for collecting the logs. Because we observed significant degradation for the performance of some of the CPU-intensive applications such as TeraSort when the *trace* level is enabled, we further investigate the impact of utilizing faster storage systems for log collection. Because developers and practitioners mostly utilize trace log level for debugging, we rationalize that the debugging data can be saved in memory temporarily to expedite the debugging process and the final debugging outcome can be transferred to the disk when the debugging is finalized. In addition, although memory storage is volatile and there is a risk of debugging data loss due to power outage, we presume this risk is manageable as we are concerned with debugging data in contrast to the actual execution logs, and the experiments can be repeated in case the debugging data is lost. Additionally, as new storage technologies become faster, *e.g.*, solid-state drive (SSD), and its latency edges closer to the main memory speed, this data point shows the maximum potential improvement that comes in from a faster storage paradigm, *i.e.*, a latency lower bound and a *'hypothetical'* storage medium that is as fast as the main memory. Table IV compares the

CT values for HDD versus *RAM Disk* for benchmarks that we observed a noticeable increase in CT when *trace* level is enabled. RAM Disk is a utility that allows us to map a portion of RAM as disk space and redirect benchmark logs to the space on RAM. Our goal is to show how much of the extra CT introduced because of the slow storage medium can be recovered by leveraging a faster storage medium, assuming non-volatile storage mediums become as fast as RAMs.

CT (min)	TS	TC	PR	DF	GC	LD
HDD	7.40	4.07	2.66	2.68	3.03	2.78
RAM Disk	5.73	3.99	2.61	2.64	2.99	2.68
CT reduction (%)	22.62	2.05	2.02	1.38	1.19	3.60

TABLE IV: Computation time values for RAM Disk *vs.* HDD for *trace* level.

Finding 2. *TeraSort, which generates a significantly higher amount of logs in trace level compared to info level, shows the highest CT reduction while using RAM Disk.*
Implications. *Faster storage mediums can mitigate some of the overhead associated with logging for CPU-intensive workloads that generate a significant amount of logs in more verbose log levels.*

IV. RQ2: Log Effectiveness

In RQ2, we evaluate the relationship between the log verbosity levels and their effectiveness. Although more verbose logs are used generally for debugging, there has not been any effort to quantitatively assess the effectiveness of logs in more verbose levels. In other words, although the common perception is that a higher degree of logging translates to more effectiveness of the logs, this assumption might not completely hold true. For this purpose, we introduce a new metric for calculating the effectiveness of logs based on entropy values and investigate whether or not more verbose VLs are more effective.

Log effectiveness (LE). LE is a quantitative measure of logs' effectiveness in achieving their goals, which is mainly problem diagnosis and troubleshooting. For example, Yuan *et al.* [45] showed that in their experiments when log statements exist, developers could diagnose system problems *2.2X* faster compared to not having the logs. In this study, LE is directly related to Entropy (*i.e.*, $LE \propto Entropy$) that we clarify in the following. As illustrated in Figure 1, log statements consist of two parts: static and dynamic content. Static content of log statements originates from the source code, and dynamic content is the value of variables that are printed in the log files as the system is running. As such, the dynamic content of the logs can be different in each iteration, whereas the static part is unchanged and has the same value in every iteration of the program. Therefore, for more verbose VLs to be more effective than less verbose ones, they should result in higher Entropy values and information gain (IG), as this signifies more unique runtime content. In other words, higher dynamic content translates to more runtime information and value of variables, which is positively related to higher values of entropy, IG, and LE:

(↑Dynamic content) → (↑Entropy) → (↑IG) → (↑LE)

Shannon's entropy. We use entropy as a metric to measure the dynamic content and effectiveness of the log records. Shannon's entropy [40] is used to measure the amount of information that is contained in an information source (*e.g.*, a text file). Entropy is calculated as: $H = -\sum_{n=1}^{N} p(i) \log_2 p(i)$, where $p(i)$ is the probability of a possible character happening in the log data [40]. The more random (*i.e.*, less repetitive) the content of the log file, the higher the entropy. For the purpose of experimentation, we focus on the *info* and *trace* log levels as they are used during the deployment and development of the software, respectively, to gain insight from the end user and developer perspectives. Because Spark generates only a few MBs of logs in *info* level for some benchmarks and to perform the experimentation on equal log sizes, we randomly sample 1 MB size of Spark's logs for both *trace* and *info* levels for each benchmark and measure the entropies. Table V shows the entropy values per character for log files in *trace* and *info* verbosity levels. The character-level entropy values are slightly higher for the *trace* log level, which partially signifies higher information gain (IG) and less repetitiveness for this level.

Entropy	WC	TS	TC	PR	DF	GC	LD
Info	5.24	5.31	5.16	5.21	5.26	5.23	5.26
Trace	5.41	5.39	5.38	5.40	5.37	5.34	5.40

TABLE V: Shannon's entropies for *info* and *trace* for various applications.

N-gram model. Although character level entropy explains the randomness of single characters in logs, it does not provide insight on the sentence level repetitiveness of log messages. It is more reasonable to calculate the entropies for a sequence of words, as log statements are inserted as a sequence of tokens (*i.e.*, words and variables) in the source code. To accommodate for a sequence of words, which bears higher semantic meanings for log messages, entropy is also used for a sequence of grams (*i.e.*, words or tokens), such as calculating the probabilities of a sequence of tokens in the English language. For this purpose, prior research has suggested the use of n-gram models [23], to capture the repetitiveness of a sequence of words. The n-gram model captures the probability distribution of the log data, and once trained, it can predict the probability distribution of the next token in new log sequences by utilizing order-n Markov model approximation. This approximation considers the probability of i_{th} element in the sequence of n tokens to be predicted based on $n-1$ preceding tokens [25]. Therefore, we can estimate the probability of a_i succeeding tokens $a_{i-1}, a_{i-2}, ..., a_{i-n+1}$ with:

$$p(a_i|a_{i-1}a_{i-2}...a_{i-n+1}) = \frac{count(a_i a_{i-1} a_{i-2}...a_{i-n+1})}{count(a_{i-1}a_{i-2}...a_{i-n+1})} \quad (1)$$

Based on this model, the entropy for a sequence of tokens is:

$$H = -\frac{1}{N}\sum_{n=1}^{N} log\ p(a_i|a_{i-1}a_{i-2}...a_{i-n+1}) \quad (2)$$

To measure the sentence-level information gain from the logs, we evaluate the entropy of a sequence of log tokens in

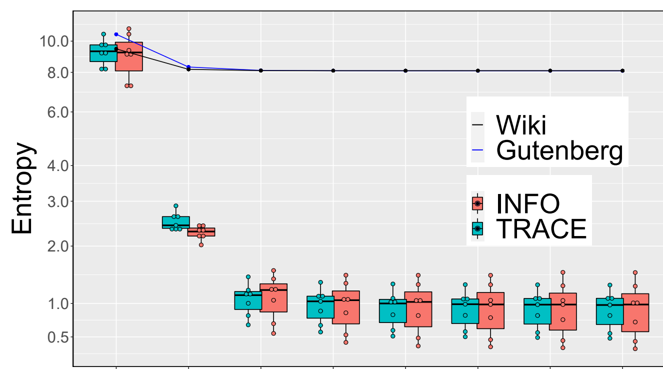

Fig. 5: Entropy for n-gram models for Spark logs and English text.

both *info* and *trace* levels with n-gram models and compare it with common English text such as *Gutenberg* [3] and *Wiki*. Gutenberg is a collection of English books, and Wiki is the English articles from Wikipedia. To train and test the n-gram models on the sequences of logs, we randomly sample 1 MB of data from each benchmark and perform a 90%-10% train-test split. We run ten-fold cross-validation to avoid overfitting [35] and plot the average entropies of 10 iterations for n-gram models in the range of $n \in (1, 8)$ in Figure 5. English text entropies stabilize around eight as the size of n-gram increases, and the median values for *trace* and *info* stabilize at 0.975 and 0.982, respectively. Our experiment reveals that English text has higher entropy than log files, and hence it has less repetitiveness, which is also observed in prior research on the naturalness of software artifacts [15, 17, 23, 42]. Lower baseline entropy of software logs compared to natural language text is beneficial as it results in *'distinguishable'* entropy changes while detecting anomalous log lines (*i.e.*, peaks in entropy values) [17], which can be utilized for log failure detection. Interestingly enough, our comparison shows, in Spark's case, the n-gram entropies are comparable for *trace* level when compared to the *info* level. This suggests that although *trace* level logging results in a higher volume of logs, *trace* log sequences are *not necessarily* less repetitive, and *trace* does not benefit from noticeable *higher information gain*, as IG from an event (*e.g.*, log event) is directly related to its entropy [38]. In addition, a higher amount of repetition that results in larger log files might decrease their effectiveness. Redundancy is an undesirable feature of logs since it adds noise to the log files and complicates the understanding of the program's behavior and hinders problem diagnosis through logs [20, 45].

Finding 3. *Although trace level generates a larger volume of logs, trace data does not provide a noticeable higher entropy, and hence, does not necessarily carry higher information gain and effectiveness when compared to less verbose log levels.*
Implications. *We presume trace logs show comparable IG to info because they contain higher repetition rather than unique dynamic values.*

V. RQ3: Failure Assessment

Logs are widely utilized in failure detection and performance diagnosis [12, 43]. Therefore, in this section, we study the effectiveness of the information gain approach in system failure detection. To evaluate the effect of system failures on the generated logs, we design a framework to inject different types of distributed failures and measure their impact on logs and how the IG approach can be applied to extract log lines related to the failures. As numerous failure scenarios exist, our goal is not to provide a comprehensive list, but to investigate common failures in a distributed environment. We categorize the distributed failures in four main categories:

1) **Compute node failure** happens when a compute resource becomes unavailable. We synthesize this scenario by terminating one of the Spark's worker nodes.
2) **Storage node failure** in a distributed environment happens when a storage medium becomes unavailable. As we utilize HDFS with the replication factor of three, the integrity of the data remains intact in case of a single node failure, however, the latency of reads and writes to the storage will increase for some compute nodes that require access to data on non-local HDFS nodes. We synthesize this failure by terminating one of the HDFS data nodes.
3) **Communication interference**, which resembles a scenario in the distributed network with variable latency and a probability of packet loss. This category can be initially observed as a performance degradation, and eventually may lead to a complete failure if the communication delay between distributed compute and storage nodes surpasses a system's predefined timeout.
4) **Combined failure** resembles a scenario in which multiple nodes become unavailable simultaneously for various reasons such as power outages. We simulate this scenario by terminating a cluster node that hosts both Spark compute and HDFS storage nodes.

With this failure categorization, our goal is to observe the changes in the content of the logs files and apply information gain approaches to detect failures. The hypothesis is that we should observe a higher information gain during a failure, as the failure related logs should resemble different dynamic content. As such, we evaluate the entropy of logs during their normal and abnormal (*i.e.*, failure) time intervals. Because we noticed comparable entropy values for *info* and *trace* log levels (Figure 5), in the following, we focus on *info* verbosity level as it is the default log level during the deployment, and, additionally, storage overhead of logs becomes more manageable. We run each Spark's benchmark in *info* level for ten iterations with and without the aforementioned failures and evaluate the changes in execution time and the storage overhead for the generated logs. In addition, we also evaluate the changes in information gain (entropy) with the normal and failure logs. For entropy calculation, the n-gram model is trained on the normal execution runs (*i.e.*, without failures) and tested on runs with failures. We choose $n = 5$ for the n-gram model, as according to Figure 5, entropy values are stabilized for $n \geq 5$.

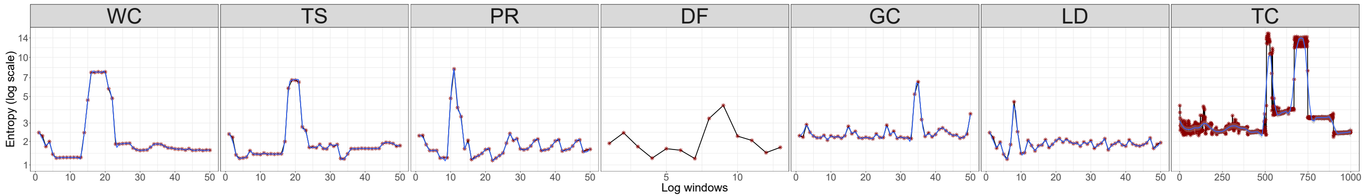

Fig. 6: Entropy values for log windows for different applications with Spark's compute node failure.

A. Compute Node Failure

Figure 6 shows the entropy values over time for different benchmarks. The x-axis shows sequential log windows of 4 KBs in size, as suggested by prior work for log analysis time window ([17]), and the y-axis shows the entropies. As the size of the generated logs varies for each benchmark, we show the timeline of entropy changes for each benchmark from the start to the end of its execution. The spikes in the entropy values are the manifestation of failures in the Spark's logs. Once a failure happens, the system first detects the failure and then plans a set of *recuperating actions* to recover from the failure. For example, for a distributed system such as Spark, the task manager resubmits the tasks previously assigned to a failed node to other nodes in the distributed cluster. This results in several log lines in the log files, which we call *failure manifestation log region*, that have higher than normal entropy values. As the failure happens, the benchmarks also show noticeable prolonged computation time and additional logs compared to the baseline scenario (*i.e.*, with no failure in Figure 4).

Detailed analysis of TransitiveClosure. We observed that failures can result in different manifestations in the execution logs, and the manifestation can be relatively benchmark dependent. To provide further insight, we review the interesting scenario of a Spark compute node (CN) failure for the TransitiveClosure benchmark, which goes through the following four failure stages: (**S1**) **Failure detection**. Upon a CN failure, the Spark's Master (SM) observes this as *"a CN has been disassociated"*, and subsequently, SM observes that the tasks associated with that CN are also lost. This results in a set of log records with high IG, and the first region of spikes in entropy values for ***TC*** in Figure 6 right after *x=500*. (**S2**) **Interleaving logs**. In addition, due to the interleaving of logs in the distributed system, other components in the system still continue to generate normal logs. This is manifested in the entropy drop after the initial spike. (**S3**) **Recovery attempts.** Happens when SM makes several unsuccessful attempts to *recover* from the failed state and reconnect to the lost CN. This is manifested as the second high entropy region in Figure 6 that ends just before *x=750*. (**S4**) **Cleaning and back to normal**. After **S3**, SM gives up attempting to reconnect to the failed CN and continues the execution by reassigning the failed tasks to the remaining CNs. In the meanwhile, it clears the data structure allocated to the failed CN. It should be noted that the outlined stages can manifest differently depending on the applications. For example, at *info* level, TS generates far fewer log lines (0.57 MBs) than TC (313 MBs). As such, TS observes less interleaving of logs compared to TC, and **S1** and **S3** manifest as one spike region in the logs.

> **Finding 4.** *A compute node failure with manifestation in log files would result in higher entropy values than normal entropies, and different runs show extended computation time and additional logs related to the failure. CT of CPU-intensive applications suffers more from compute node failure than I/O intensive benchmarks.*
> **Implications.** *Sudden changes in the entropy values of log records can signify a system failure.*

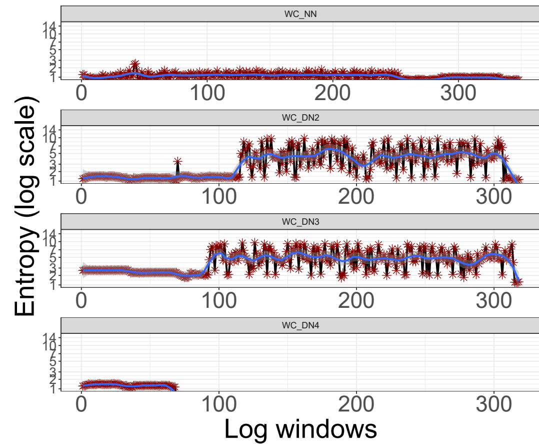

Fig. 7: Entropy values for log windows for WordCount with HDFS's data node failure.

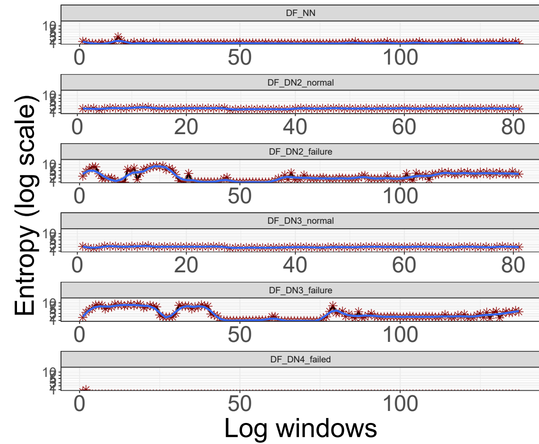

Fig. 8: Entropy values for log windows for DFSIO with HDFS's data node failure.

B. Storage Failure

To gain insight into storage failures, we investigate the entropy changes of HDFS logs. All nodes within the HDFS

file system (*i.e.*, name node and data nodes) generate logs. We perform the experimentation for all the benchmarks and review the logs from all the nodes and measure the entropy changes as the failure happens. Due to the limited space, we focus on the entropy value changes of WordCount and DFSIO as they are I/O intensive benchmarks, and they make the most use of HDFS compared to other benchmarks. Figure 7 shows the entropy values over time for the name node (WC_NN), and three data nodes (WC_DNx) as the failure happens. When DN4 fails at log window 69, we observe a delayed manifestation of entropy changes in other data nodes (DN2 and DN3) which starts at x=*100*. We observe that DN2 and DN3 directly contact DN4 (which is not available) to retrieve some blocks of data, and hence this results in failure log messages and hence higher entropies. By default, DNs are configured to send heartbeat signals to NN every 3 seconds. However, in case of a DN failure, NN marks an unresponsive DN dead after 10mm:30ss[1], which at that time manifests in high entropy values in NN logs. The log windows in Figure 7 show the entire execution span for WC, which on average finishes within four minutes, and hence, we do not observe entropy value changes in NN in the plotted timeline. Similarly, Figure 8 represents the entropies for DFSIO, another I/O intensive task in our benchmark set when a data node (DN4) fails. We have also shown the entropy values for DN2 and DN3 during a normal run for comparison. The peaks show failure log messages with some normal interleaving logs. Failure for DN4 happens very close to the start of the x-axis and thus the initial peaks for DN2 and DN3.

Finding 5. *We observe noticeable entropy changes in the HDFS logs of I/O intensive benchmarks as the storage failure occurs. CPU-intensive benchmarks that have minimal interaction with HDFS do not generate enough HDFS logs for a meaningful log analysis.*
Finding 6. *I/O intensive applications that read and write large volumes of data to the distributed storage will be negatively affected the most as a result of a storage node failure, whereas CPU-intensive applications (i.e., with minimal R/W to the storage) will be less impacted. Thus, CT of I/O tasks becomes prolonged due to the storage failure and the application's log size is also partially increased as it captures extra failure log records.*
Implications. *As failures are manifested as higher information gain, entropy-based anomaly detection approaches can be applied for online log analysis to isolate the higher entropy regions and further investigate the failures.*

C. Communication Interference Modeling

If the network connection between distributed nodes permanently disconnects, the observable failure outcome would be similar to the permanent failure of the compute/storage node, as that node becomes unreachable from the cluster manager. In contrast to permanent failures in Sections V-A-V-B, we here investigate intermittent network interference, *e.g.*, packet loss, to gain insight into non-permanent failures which are manifested as ***performance degradation***. To emulate a realistic network traffic model, we implement Gilbert-Elliot capacity modeling approach [21, 34], which is comprised of *Good* and *Bad* states (Figure 9). This model offers a more realistic emulation for network impairments, rather than simple packet loss.

[1]https://issues.apache.org/jira/browse/HDFS-3703

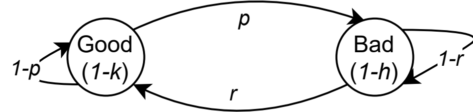

Fig. 9: Gilbert-Elliot communication interference model.

An example usage configuration for Gilbert-Elliot scheme would be as follows: ``tc qdisc add dev dev_name root netem loss gemodel 2% 15% 30% 1%.'' In this example, the error rate in Good *(1−k)* and Bad *(1−h)* states are *1%* and *30%*, respectively, and the probability of transitioning to Good *(r)* and Bad *(p)* states are *2%* and *15%*, respectively. In the following experiments, we vary the error rate in Bad state, *i.e., (1−h)*, in the range of (0%, 45%) and measure the computation time. In addition, we also evaluate the *'combined failure'* (Case **4** in Section V) by disconnecting one of the machines in the cluster that hosts both compute and data nodes.

Heterogeneous cluster. For the purpose of experimentation, we also define a new configuration that has three slave/data nodes but one node is smaller as it is using half of the cores and memory (*i.e.*, 12 cores and 16 GBs of memory instead of 24 cores and 32 GBs of memory). This is in contrast to the homogeneous cluster (Figure 3) that all the three slave/data nodes are the same size (24 cores and 32 GBs of memory). The rationale to include the heterogeneous configuration is to compare it with the performance degradation scenario that appears as a result of network interference.

Results. Figures 10a and 10b show the evaluation for **communication interference** and **combined failure** for WordCount and TransitiveClosure as examples of I/O intensive, and iterative and CPU-intensive benchmarks, respectively. We refer to the graphs by their labels in the figures, *i.e.*, (A)-(E). Graph (A) shows the computation time as a result of a combined failure for a cluster with 2 nodes, *i.e.*, two compute nodes and two storage nodes after a machine that hosts both compute and storage nodes fails. Graph (B) shows a *homogeneous* cluster with three nodes, and Graph (C) shows a *heterogeneous* cluster in which the third node is smaller ('3Nodes-1small'). No network interference is applied to (A), (B), and (C). Graphs (D) and (E) are equivalent to (B) and (C), respectively, but with added network interference. In Figure 10, as the communication interference increases, the CT time for (D) and (E) increases, and for values higher than 15% for WC and 10% for TC, the computation time of a cluster with interference surpasses a cluster with combined failure, Graph (A), *i.e.*, two nodes in the cluster. We also observe that I/O intensive benchmarks that require to transfer a large amount of data among the nodes in the cluster suffer more than CPU-intensive benchmarks that use the network to a lesser degree.

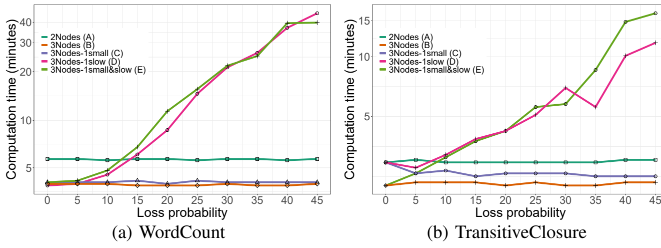

(a) WordCount (b) TransitiveClosure

Fig. 10: Execution time for WC and TC during the communication interference and combined failure.

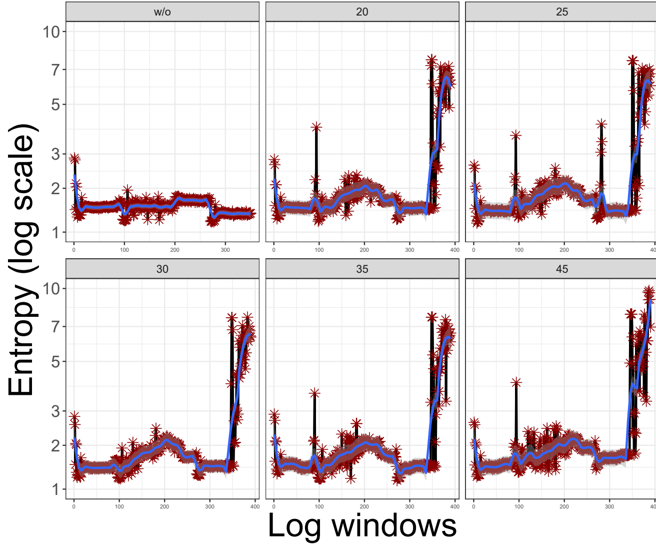

Fig. 11: Entropy values for log windows for WordCount for different values of drop rate *(1-h)*.

> **Finding 7.** *As the communication interference increases, the computation time increases, and communication interference is manifested as a performance degradation and not a complete failure.*
> **Finding 8.** *When the interference increases beyond a certain threshold, the negative impact of the performance degradation surpasses the impact of a complete failure because, for each stage of the computation, the faster nodes are awaiting the completion of the slow node.*
> **Implications.** *Distributed scheduling algorithms that can detect slow nodes in the system and remove them from the computation can benefit the entire system's performance.*

Entropy values. Figure 11 shows the effect of drop rate in Bad state *(1-h)* for WordCount logs. The top-left graph ('w/o') shows the entropy of Spark's logs with zero communication interference, and we gradually increase the drop rate from 5% to 45% (in steps of 5%) and plot the entropy values from the start to the end of the execution for selected percentages, *i.e.*, (20, 25, 30, 35, and 45)%. Our observation is that due to the non-deterministic nature of network interference, performance degradation is indirectly manifested in the logs and their corresponding entropy values, in contrast to having clear regions with high entropies (Figures 6-7). We also observe that entropy values start to climb as the interference percentage increases. In addition, higher entropy values are manifested with a delay towards the end of the execution as the system experiences timeouts and aims to reestablish the connection with the unstable node or resubmit the failing tasks to other nodes in the distributed system. Therefore, we hypothesize that a combination of execution log records and system metrics, such as average task completion time for speculative execution (Section V-D), are required to identify performance degradation cases [24].

> **Finding 9.** *As the communication interference increases, the entropy values gradually increase with a delay.*
> **Implications.** *Failures that manifest as a gradual system slowdown and performance degradation are harder to detect than complete failures solely with logs. As such, other system metrics, e.g., average task completion time, can be applied in conjugation with logs.*

D. Discussion

Speculative execution. Findings 7-9 imply that since communication interference is manifested as intermittent failures, as opposed to a complete compute or data node failure, they are harder to detect and diagnose from the log files. Therefore, we suggested the usage of distributed scheduling algorithms that can detect the slow nodes and utilization of system metrics (*e.g.*, average task's computation time) in conjunction with logs for more effective problem diagnosis. Spark provides a feature known as speculative execution (`spark.speculation`) [4] that if enabled, allows resubmitting slow tasks to other nodes in parallel, and proceeds as soon as any of the task instances completes its execution. Figure 12 shows the non-speculative ('w/o') and speculative ('w/') runs for WC, TS, TC, and DF. For WC, although the computation is moved to another node, because large amounts of data are being shuffled through the network interference to the new node, speculation would not show a noticeable improvement. For DF, we observe as the network interference increases, without speculative execution, the write pipeline[2] fails more often, and this results in a noticeable increase in the execution time of this benchmark. With speculative execution, the slow tasks are moved to other nodes, hence with a different write pipeline, which significantly helps with containing the network failures. In short, the main benefit from speculative execution in the distributed environment comes from moving tasks from the faulty node with network interference to other nodes, which helps to avoid task re-execution and data re-transfer due to the network uncertainty. Our observation is that speculative execution only marginally improves the execution time, and in general, after a certain level, even with speculative execution enabled, the distributed systems performance becomes slower than removing the node with lower performance completely from the cluster.

Cluster heterogeneity. We experimented with homogeneous and heterogeneous clusters in Section V-C. Although heterogeneity by design, *i.e.*, having a machine smaller than other machines in the cluster, is well understood [37], heterogeneity that is the result of performance

[2]https://stackoverflow.com/questions/37531946/what-is-hdfs-write-pipeline

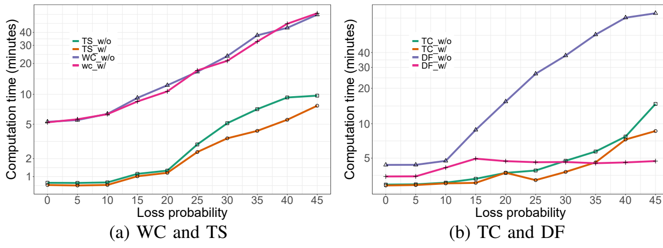

(a) WC and TS

(b) TC and DF

Fig. 12: Speculative execution for different benchmarks.

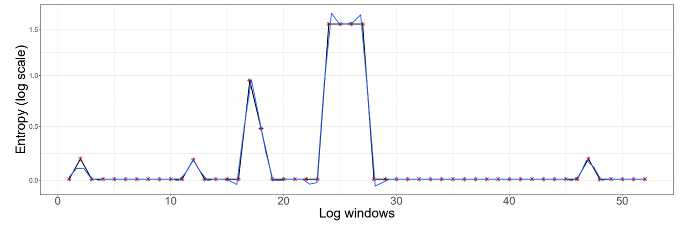

Fig. 13: OpenStack entropy values for log sequences with four anomalous VM log records.

degradation and partial failure is left untreated. In our case, the rationale to include experiments with a heterogeneous cluster is to simulate a scenario that heterogeneity is introduced in the distributed platform because of performance degradation. In other words, although the original design is homogeneous, heterogeneity can still happen due to a variety of reasons, such as communication interference, which can negatively impact the entire system's performance. Also, one of the factors that limits Spark speculation performance is its assumption about operating in a homogeneous environment, which is not the case in a performance-impaired cluster. This would encourage further research to investigate possible scenarios and solutions for failure-induced heterogeneity.

Slow distributed file system. Although with the replication factor of three in HDFS, there is no single point of failure, partial failures can negatively affect the performance of the entire file system. In our case, the slow network connection for one of the data nodes due to an induced drop rate negatively impacts the performance of the entire HDFS. A slow data node still continues to send heartbeats successfully, and the HDFS name node will keep redirecting clients to the slow DN, and therefore, degrade the performance of the entire cluster. Although HDFS provides few settings to detect and report slow data nodes, it does not provide a mechanism to automatically bypass the slow DNs, as they are still sending heartbeat signals to the name node. Thus, we foresee further research to investigate mechanisms, similar to Spark's speculative execution, to obtain data from other available data nodes in case a data node becomes intermittently unavailable or slow.

Implications on fault tolerance. Our observation is that information gain is helpful in zooming in the failure regions (*i.e.*, spikes in the entropy values), which means that these regions of logs with higher IG can be isolated and quickly reviewed by users and practitioners for failure diagnosis, which then results in faster system recovery from a failure. Additionally, we emphasize that although the spark cluster is fault-tolerant by design, its performance is impaired due to the failures. Thus, we envision that quicker detection and recovery of failed compute and storage nodes directly connects with and benefits the robustness and fault-tolerance of the system. A speedy return to the normal state will allow the system to tolerate additional failures.

VI. Case Study

In this section, we use logs from real abnormal scenarios for labeled OpenStack logs [13]. We first parse the logs to extract their templates and then group the logs based on their *instance_id*. *instance_id* serves as an identifier for a particular task execution sequence. Figure 13 shows the timeline of 52 log windows with 4 injected abnormal OpenStack VM sequences in the center for $x \in [24, 27]$. There is also one false positive high entropy value at $x = 17$. By leveraging the *Hampel* filter and outlier detection approach proposed in prior work [17], we can reach *F-Measure* ($\frac{2 \times precision \times recall}{precision + recall}$) and *Balanced Accuracy* ($\frac{TP}{2(TP+FN)} + \frac{TN}{2(TN+FP)}$) of 0.89 and 0.98, respectively. This signifies the effectiveness of the information gain approach that is achieved by measuring entropy values of log sequences in detecting OpenStack's abnormal scenarios.

Computation time. In our analysis, the training of the natural language model happens only once on normal logs. Then, testing the log records while they are generated is rather fast. We performed a quick measurement and received on average *2.4* milliseconds execution time (as a single thread executed on 2.40GHz Intel Xeon E5-2620) for a *4KB* log window. Relatively, the execution time for testing logs is faster than the rate of log generation, thus, the log records can be tested in real-time and observed for information gain and potential anomalies.

VII. Related Work

Logging cost and gain. Prior work on assessing logging cost does not quantitatively evaluate the system performance overhead contributed to logging levels. Ding et al. [12] performed a survey of logging practices among Microsoft developers and listed their findings on overheads associated with logging reported by the developers. Mizouchi et al. [33] presented PADLA, an online method to dynamically adjust the logging level, and, consequently, limit the logging cost. Miranskyy et al. [32] reviewed challenges for log analysis of big-data systems, among them *limited storage* and *unscalable log analysis*. Goshal *et al.* [18] discussed the provenance, *i.e.*, the origin, of logs and how it correlates with log levels and types of applications in large-scale systems. Different from the prior research, our study aims to quantitatively analyze the costs and benefits associated with the level of logging and the information gain.

Spark's performance evaluation. Spark [46] has been introduced for massively parallel data analytics, which improves on its predecessor, Hadoop MapReduce [11], by utilizing in-memory storage of intermediate results for iterative applications, and bringing more flexibility in performance and the programming model [39, 41]. Mavridis and Karatza [31]

evaluated various log file analyses with the cloud computational frameworks, Apache Hadoop and Apache Spark, and, experimentally, Spark achieved the best performance. Lu *et al.* [29] performed log-based anomaly detection for Spark. Our study is different as we leverage Spark's logs to evaluate the cost and IG associated with log levels.

Information gain and NLP. In this research, we propose a new perspective along with validated metrics to evaluate the impact of different verbosity levels on cost and information gain from log statements with natural language processing (NLP) approaches. Compared with related research [9, 16, 28], our approach is orthogonal to such efforts that aim to suggest the proper logging statement or its VL by extracting features from the source code. These works rely on the existing logging statements to suggest logs for newly composed instances of code. However, we aim to bring attention to the trade-offs between logging cost and the information gain, and that failures are manifested as higher information gain in logs.

VIII. Conclusions and Future Work

The goal of our work is to provide a quantitative assessment of logging cost in different verbosity levels and how that translates to information gain in distributed systems. Therefore, we evaluate the impact of log verbosity levels on performance and storage overhead, and the information gain from logs for various Spark Benchmarks. We also experiment with synthesizing various categories of distributed failures for compute and data nodes, and network interference, and measure the effect of failures in execution time, the volume of the generated logs, and the information gain. Lastly, we provide a case study of the application of our approach on OpenStack real failure logs. Our findings are helpful for developers and practitioners to better evaluate the costs and benefits of logging when choosing different verbosity levels and how failures can be tracked down with IG approaches. As future work, we will look into evaluating logs of other distributed software systems and investigate how IG can be translated to more effective troubleshooting by leveraging the execution logs and system metrics.

References

[1] Dataset for this reasearch. https://anonymous.4open.science/r/a3c64a04-501d-4296-8aa8-0f4972f9d167/.

[2] N-gram model. https://en.wikipedia.org/wiki/N-gram.

[3] Project Gutenburg. https://www.gutenberg.org/.

[4] Spark scheduling. https://spark.apache.org/docs/latest/configuration.html#scheduling.

[5] Logging services project. http://logging.apache.org/, 2019.

[6] Simple Logging Facade for Java (SLF4j). http://www.slf4j.org/, 2020.

[7] TeraSort for Spark. https://github.com/ehiggs/spark-terasort, 2021.

[8] TestDFSIO for Spark. https://github.com/BBVA/spark-benchmarks, 2021.

[9] H. Anu, J. Chen, W. Shi, J. Hou, B. Liang, and B. Qin. An approach to recommendation of verbosity log levels based on logging intention. In *ICSME*, pages 125–134. IEEE, 2019.

[10] M. Chen and et al. Failure diagnosis using decision trees. In *International Conference on Autonomic Computing, 2004. Proceedings.*, pages 36–43, 2004.

[11] J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[12] J. Ding and et al. Log2: A cost-aware logging mechanism for performance diagnosis. USENIX ATC 2015, July 2015.

[13] M. Du, F. Li, G. Zheng, and V. Srikumar. Deeplog: Anomaly detection and diagnosis from system logs through deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1285–1298, 2017.

[14] Q. Fu and et al. Where do developers log? an empirical study on logging practices in industry. In *Companion Proceedings of ICSE*, pages 24–33. ACM, 2014.

[15] E. Gabrilovich and S. Markovitch. Wikipedia-based semantic interpretation for natural language processing. *JAIR*, 34:443–498, 2009.

[16] S. Gholamian and P. A. Ward. Logging statements' prediction based on source code clones. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing*, pages 82–91, 2020.

[17] S. Gholamian and P. A. Ward. On the naturalness and localness of software logs. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*, pages 155–166. IEEE, 2021.

[18] D. Ghoshal and B. Plale. Provenance from log files: a bigdata problem. In *Proceedings of the Joint EDBT/ICDT 2013 Workshops*, pages 290–297, 2013.

[19] A. E. Hassan. The road ahead for mining software repositories. In *2008 Frontiers of Software Maintenance*, pages 48–57. IEEE, 2008.

[20] A. E. Hassan, D. J. Martin, P. Flora, P. Mansfield, and D. Dietz. An industrial case study of customizing operational profiles using log compression. In *Proceedings of the ICSE*, pages 713–723, 2008.

[21] G. Haßlinger and O. Hohlfeld. The gilbert-elliott model for packet loss in real time services on the internet. In *14th GI/ITG Conference-Measurement, Modelling and Evalutation of Computer and Communication Systems*, pages 1–15. VDE, 2008.

[22] S. He, Q. Lin, J.-G. Lou, H. Zhang, M. R. Lyu, and D. Zhang. Identifying impactful service system problems via log analysis. In *Proceedings of the 2018 26th ACM Joint ESEC/FSE*, pages 60–70. ACM, 2018.

[23] A. Hindle, E. T. Barr, Z. Su, M. Gabel, and P. Devanbu. On the naturalness of software. In *ICSE*, pages 837–847. IEEE, 2012.

[24] O. Ibidunmoye, F. Hernández-Rodriguez, and E. Elmroth. Performance anomaly detection and bottleneck identification. *CSUR*, 48(1):1–35, 2015.

[25] D. Jurafsky. *Speech & language processing*. Pearson Education India, 2000.

[26] W. H. Kruskal and W. A. Wallis. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association*, 47(260):583–621, 1952.

[27] H. Li, W. Shang, B. Adams, M. Sayagh, and A. E. Hassan. A qualitative study of the benefits and costs of logging from developers' perspectives. *IEEE TSE*, 2020.

[28] H. Li, W. Shang, and A. E. Hassan. Which log level should developers choose for a new logging statement? *EMSE*, 22(4):1684–1716, 2017.

[29] S. Lu, B. Rao, X. Wei, B. Tak, L. Wang, and L. Wang. Log-based abnormal task detection and root cause analysis for spark. In *ICWS*, pages 389–396. IEEE, 2017.

[30] A. D. Malony and et al. Performance measurement intrusion and perturbation analysis. *IEEE Computer Architecture Letters*, 3(04):433–450, 1992.

[31] I. Mavridis and H. Karatza. Performance evaluation of cloud-based log file analysis with apache hadoop and apache spark. *Journal of Systems and Software*, 125:133–151, 2017.

[32] A. Miranskyy, A. Hamou-Lhadj, E. Cialini, and A. Larsson. Operational-log analysis for big data systems: Challenges and solutions. *IEEE Software*, 33(2):52–59, 2016.

[33] T. Mizouchi, K. Shimari, T. Ishio, and K. Inoue. PADLA: a dynamic log level adapter using online phase detection. In *Proceedings of the 27th ICPC*, pages 135–138. IEEE Press, 2019.

[34] M. Mushkin and I. Bar-David. Capacity and coding for the gilbert-elliott channels. *IEEE Transactions on Information Theory*, 35(6):1277–1290, 1989.

[35] A. Y. Ng et al. Preventing" overfitting" of cross-validation data. In *ICML*, volume 97, pages 245–253. Citeseer, 1997.

[36] A. Oliner and J. Stearley. What supercomputers say: A study of five system logs. In *DSN*, pages 575–584. IEEE, 2007.

[37] A. J. Page, T. M. Keane, and T. J. Naughton. Multi-heuristic dynamic task allocation using genetic algorithms in a heterogeneous distributed system. *JPDC*, 70(7):758–766, 2010.

[38] N. R. Pal and S. K. Pal. Entropy: A new definition and its applications. *IEEE transactions on systems, man, and cybernetics*, 21(5):1260–1270, 1991.

[39] P. Petridis, A. Gounaris, and J. Torres. Spark parameter tuning via trial-and-error. In *INNS Conference on Big Data*, pages 226–237. Springer, 2016.

[40] C. E. Shannon. A mathematical theory of communication. *ACM SIGMOBILE mobile computing and communications review*, 5(1):3–55, 2001.

[41] J. Shi, Y. Qiu, U. F. Minhas, L. Jiao, C. Wang, B. Reinwald, and F. Özcan. Clash of the titans: Mapreduce vs. spark for large scale data analytics. *Proceedings of the VLDB Endowment*, 8(13):2110–2121, 2015.

[42] Z. Tu, Z. Su, and P. Devanbu. On the localness of software. In *ACM SIGSOFT FSE*, pages 269–280, 2014.

[43] W. Xu, L. Huang, A. Fox, D. Patterson, and M. I. Jordan. Detecting large-scale system problems by mining console logs. In *Proceedings of the ACM SIGOPS*, pages 117–132. ACM, 2009.

[44] K. Yao, H. Li, W. Shang, and A. E. Hassan. A study of the performance of general compressors on log files. *EMSE*, 25(5):3043–3085, 2020.

[45] D. Yuan, S. Park, and Y. Zhou. Characterizing logging practices in open-source software. In *ICSE*, pages 102–112. IEEE, 2012.

[46] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauly, M. J. Franklin, S. Shenker, and I. Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Presented as part of the 9th {USENIX} ({NSDI} 12)*, pages 15–28, 2012.

[47] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, I. Stoica, et al. Spark: Cluster computing with working sets.

[48] X. Zhou, X. Peng, T. Xie, J. Sun, C. Ji, D. Liu, Q. Xiang, and C. He. Latent error prediction and fault localization for microservice applications by learning from system trace logs. In *Proceedings of FSE/ESCE*, pages 683–694, 2019.