

Tails in the Cloud: A Survey and Taxonomy of Straggler Management within Large-scale Cloud Datacenters

Sukhpal Singh Gill¹, Xue Ouyang², Peter Garraghan³

¹School of Electronic Engineering and Computer Science, Queen Mary University of London, UK

²School of Electronic Sciences, National University of Defense Technology, China

³School of Computing and Communications, Lancaster University, UK

s.s.gill@qmul.ac.uk, ouyangxue08@nudt.edu.cn, p.garraghan@lancaster.ac.uk

Cloud computing systems are splitting compute and data intensive jobs into smaller tasks to execute them in a parallel manner using clusters to improve execution time. However, such systems at increasing scale are exposed to stragglers, whereby abnormally slow running tasks executing within a job substantially affect job performance completion. Such stragglers are a direct threat towards attaining fast execution of data intensive jobs within cloud computing. Researchers have proposed an assortment of different mechanisms, frameworks, and management techniques to detect and mitigate stragglers both proactively and reactively. In this paper, we present a comprehensive review of straggler management techniques within large-scale cloud datacenters. We provide a detailed taxonomy of straggler causes, as well as proposed management and mitigation techniques based on straggler characteristics and properties. From this systematic review, we outline several outstanding challenges and potential directions of possible future work for straggler research.

Keywords: Computing, Stragglers, Cloud Computing, Straggler Management, Distributed Systems, Cloud Datacenters

1. INTRODUCTION AND MOTIVATION

Nowadays, applications spanning various domains including social networks, e-commerce sites, and healthcare generate vast quantities of data. The growing velocity and volume of such data generation has subsequently required the substantial computing capacity in order to store and process such data effectively [1]. Such large-scale computing systems, encompassing datacenter clusters, comprise hundreds and thousands of individual machines interconnected together that underpin application operation consumed by both businesses and consumers alike.

A combination of increasing application demand and technological innovations has resulted in greater system scale in the regions of tens of thousands of servers within an individual cluster [12]. However, such complexity has subsequently resulted in an increase in complexity within such systems, manifesting in the form of emergent phenomena whereby system operation exhibits behaviour unforeseen at design time. Such emergent phenomena manifesting within large-scale cloud datacenters has observed to negatively impact application performance. One such phenomena, known as the *Long Tail Problem*, is characterized by a minor subset of task stragglers that operate unusually slower in comparison to normal task behaviour within a job. Task stragglers occur within any highly parallelised system and become even more apparent for jobs containing many tasks executing across a large number of machines.

Frameworks such as MapReduce, Spark, and Dryad [1] [40] [41] process vast quantities of data via parallelizing jobs into a smaller sub-set of tasks, and thus makes such applications susceptible to stragglers. For example, within MapReduce, a job can only complete once all tasks have completed their execution. However, the occurrence of stragglers results in an atypically long task execution duration, thus degrading the performance of the entire job. The challenge in effectively addressing stragglers is that their root-cause is not well-understood [80] and can be resultant due to various reasons spanning daemon processes, data skew, failures, resource contention, and energy management tools [49] [42], manifesting within the application, Operating Systems (OS), or physical hardware. This can subsequently lead to subsequent applications that depend on job outputs to also fail pending on its completion [7] [11].

This has resulted in a growing body of straggler research pertaining to analysing their underlying causes [11] [28], straggler forecasting [15] [46], and straggler mitigation techniques [4] [5] [16] [21] including speculative execution [6], replication, load balancing and scheduling [18]. Each of these

works predominantly focus on a certain sub-set phenomenon within a particular context of system operation of application framework. Thus, straggler research has reached sufficient level of maturity whereby it is worthwhile to appraise the landscape of research within the field, identify cross-cutting challenges within areas, and evaluate future challenges on the horizon for future generation computing systems.

1.1 Motivation

The core motivation behind this methodical survey is to conduct a systematic review of straggler research within large-scale cloud datacenters. This systematic review encompasses clearly defining and analysing the impact of stragglers, a taxonomy of various straggler management techniques for forecasting and mitigations, as well as identify future directions within the field.

1.2 Article Organization

The rest of the article is structured as follows: Section 2 presents the background information for straggler definition as well as straggler management within large-scale systems. Section 3 presents the taxonomy of straggler causes. Section 4 explores the existing literature for straggler management techniques. Section 5 presents the comparison of straggler management techniques based on the taxonomy of straggler causes and outlines the observation, trend analysis and future research directions. Finally, Section 6 summarizes the article.

2. BACKGROUND

2.1 Straggler Definition and Impact

Applications execute within large-scale computing systems such as datacenters and clusters by submitting jobs via a resource manager (YARN, Mesos, Borg, etc). In this context, a job is composed of multiple smaller tasks (defined as the smallest unit of computation observable by the resource manager) [82]. Such jobs and subsequent tasks are scheduled onto different machines in a parallelized manner to accelerate job completion and are often divided into phases creating a Direct Acyclic Graph (DAG) [83]. Application frameworks (such as MapReduce) attempt to sub-divide jobs so that tasks will approximately complete within the same timeframe for each phase [84]. This is achieved by providing a sub-set of data (known as shards) to each task, and allocating the appropriate resources to tasks (CPU, memory, etc). This is calculated via the resource requirement module of the resource manager [85].

However, even with such measures in place, within large-scale cloud datacenters a sub-set of tasks within a job will manifest as stragglers [86] [87]. In this context, a straggler is defined as task which execute abnormally slow in comparison to the average task duration within a job [12]. The phrase '*abnormally slow*' is typically identified as any task with a task completion time 50% greater than the (average) task completion time for a job phase [14] [88]. Slowly executing tasks (stragglers) affect the performance and completion time of the entire job [5], increasing resource utilisation and performance degradation of applications at increased scale [22] [25], thus reducing system availability and incurring additional operational costs [89]. It has been identified from analysis of production systems at scale [25] that approximately 4-6% of task stragglers negatively affect over 50% of the overall jobs within the greater system.

2.2 Straggler Management

Due to the impact of long-tail problem within distributed computing systems, there has been concentrated efforts in order to effectively mitigate their effects. This has been tackled by the research community via the creation of various straggler management techniques. In this context, straggler management comprises all mechanisms that have been created in order to mitigate the effects and impact of straggler manifestation. Figure 1 shows the depiction of straggler tasks and non-straggler tasks.




Figure 1: Depiction of straggler tasks and non-straggler tasks

Such straggler management techniques can be predominantly considered into two main classes: detection and mitigation [93] [94]. *Detection* focuses on approaches to identify straggler manifestation a priori or post priori job execution within the cloud datacenter, such as offline analytics and online monitoring mechanisms [90] [91] and an example of straggler detection is NearestFit [1]. *Mitigation* approaches focus on avoiding or tolerating (detected) straggler manifestation during job execution such as scheduling, load balancing and replication [88] [71] [92]. The examples of straggler mitigation are Dolly [4], GRASS [5], LATE [21] and Wrangler [16].

2.3 Related Surveys and Our Contributions

To present data, to the best of our knowledge, only two works have conducted a survey pertaining to straggler research. Umesh and Jitendar [34] discuss an overview of straggler handling algorithms for MapReduce framework, while Ashwin [35] reviewed several straggler handling techniques. Whilst these reviews cover specific cases of stragglers related to specific frameworks and installations, they do not necessarily provide a comprehensive survey of the straggler causes and straggler management techniques which exist within the research community. Furthermore, these works do not discuss in detail the precise root-causes and analysis of straggler behaviour, which underpin the design of straggler management techniques. Therefore, this paper attempts to provide a systematic review and taxonomy of straggler causes and map them directly to straggler management techniques along with trend analysis.

3. TAXONOMY OF STRAGGLER CAUSES

As mentioned in *Section 1*, the challenge within this research area is the myriad of potential causes of straggler manifestation. According to our comprehensive appraisal of the literature, we have identified eight key causes for straggler occurrence that manifest within large-scale cloud datacenters. Figure 2 shows the taxonomy of straggler causes.

1. Data Abstraction: Stragglers can occur due to information obfuscation at different levels of the system. Literature [43] [14] [45] [95] has identified that information can be hidden at two different levels: i) OS Level and ii) Application Level. During the execution of resources, the master node (controller) hides information from workers (cluster nodes) at *OS level*. ii) at *application level*, the information regarding platform services and infrastructure services are kept hidden from the software services.

2. CPU Utilization: It has been identified that there is a strong correlation between high system CPU utilization and straggler occurrence [46] [47] [48] [49]. The reason for this occurrence is resource contention. This is further compounded due to Head-of-Line blocking (HOL blocking), task interference during execution, busy locks, queue issues, hazard rates of task execution and launching additional *speculative* replicas, which requires *additional time* for execution.

3. Scheduling: It has been identified that scheduling and resource allocation decisions also influence straggler manifestation [50] [51] [52] [53]. For *job scheduling*, stragglers can occur due to a large number of enqueued jobs within a (machine, master scheduler) that are pending for available resources to be revoked (i.e. only a portion of tasks within a job are able to successfully acquire their necessary resources to commence execution). Furthermore, straggler may occur due to the poor admission control mechanisms, which is using to submit the jobs for execution [54]. The poor admission control mechanism launches multiple tasks together resulting in resource exhaustion causing slowdown. Lastly, dynamicity of QoS requirements at runtime, result in an inability to effectively manage the resources which leads to further the straggler occurrence. In terms of *resource scheduling*, stragglers can occur in following situations [54] [55] [56] [96]: 1) when resources are allocated to the jobs in an inefficient manner without available resource optimisation, which leading to ineffective scheduling of resources for job execution and 2) sometimes resources are still in active stage even they are not utilized for execution of jobs, which wastes more energy consumption and effects the performance of other resources because some resources need more power to run continuously.

4. Inaccessible Local Disk: Stragglers may occur when a machine hard disk is not accessible to residing tasks. Such inaccessibility is predominantly caused by [8-11] [62-66]: i) Increasing Backup Tasks and ii) Failed to Store Output. Stragglers can occur, when it is difficult to find the required task due the large backlog of the tasks waiting for execution. Sometimes, an error can occur while storing the output on the disk, causing a problem when some task wants to access that data during execution.

5. Data Skew: Straggles can occur due to the data skew, caused by the different data sizes and time variation in accessing required data [62] [63] [73] [97]. With several tasks operating on a split version of a very large shared dataset, an uneven distribution of the data amongst these tasks potentially results in some tasks to progress slow in comparison to tasks within the same phase (and subsequently delays the future sub-phases and the entire job). Data non-uniformity can also impact data access and processing time data, directly affecting the timing delays between tasks, further increasing the probability of straggler occurrence. Moreover, data locality for job execution results in lower latencies whilst distant data will take longer to be accessed, incurring additional delays in task completion, again, manifesting as a straggler.

6. Resource Contention: Resource Contention occurs when the same resource is shared by multiple tasks [4-6] [8-11] [31-33] [38] [64] [65] [66] [79] [98] [99] [100]. Resource contention occurs due to conflict over task access and oversubscription to a resources within multi-tenant machines which can be exuberated within different scenarios including: 1) hardware heterogeneity, 2) poor user code, 3) extra cloning, 4) ineffective algorithm logic, 5) temporary slowdowns, 6) additional task clones requiring more resources and 7) resource usage is being higher than accepted threshold value. *Hardware heterogeneity* is the main reason of resource contention, occurs due to a mismatch between hardware specification and specified application constraints (e.g. budget, deadline etc.) leading to task performance degradation. The *source code* of scheduling algorithm also affects the performance of the scaling system due to its coding style in terms of space and time complexity. Sometimes, poorly written source code schedules resources inefficiently, which can increase resource consumption and unavailability of required resources to specific jobs [71]. The *cloning* of tasks is creating a similar of copy to task to run parallel on another resource for fast execution.

The *cloning* of tasks needs more resources (increases resource usage), which can also put tasks of other jobs on hold and when the tasks are waiting for other resources, then stragglers can occur. An *ineffective logic* in the resource scheduling algorithm can also lead to an inefficient allocation of resources and increase resource usage, which leads to resource contention for future tasks. Temporary slowdown can occur due to inefficient allocation of resources, which needs to be corrected, otherwise it will cause straggler occurrence during execution of resources.

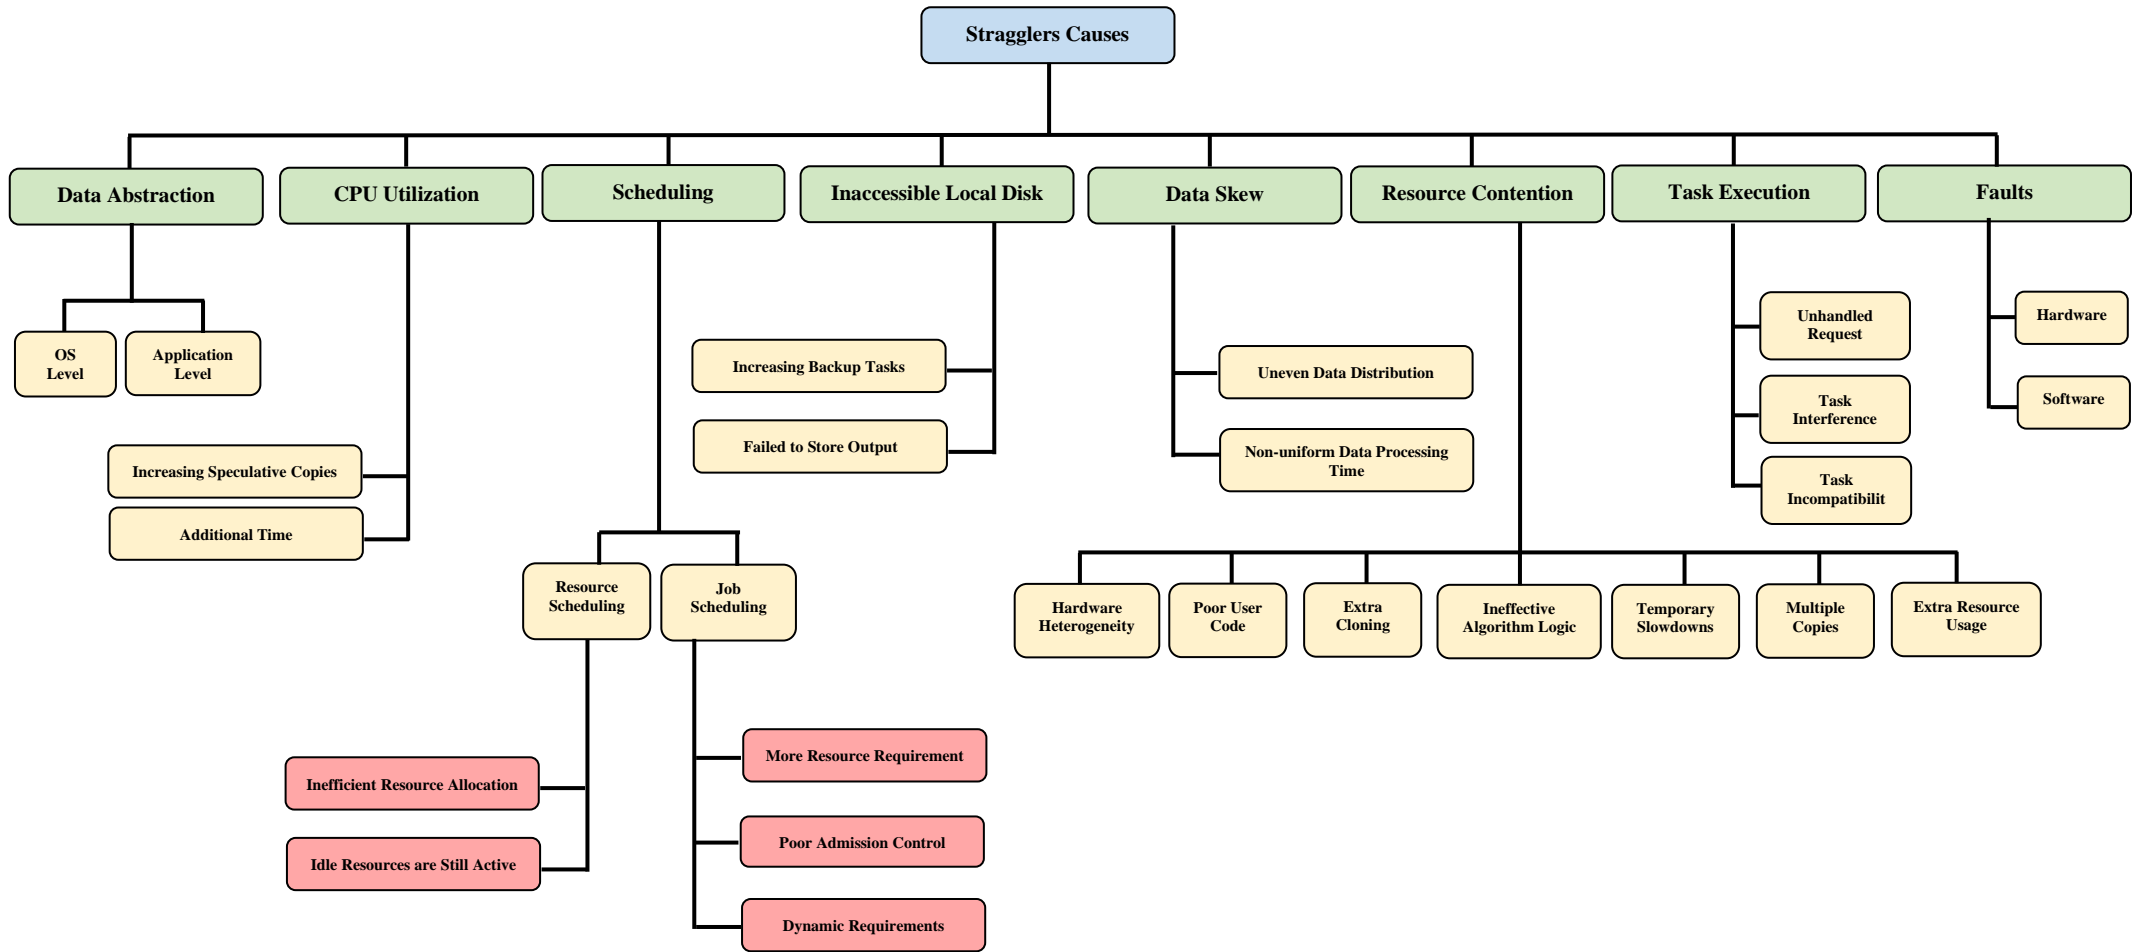


Figure 2: Taxonomy of Straggler Causes

7. Task Execution: The successful execution of a task is important to avoid straggler occurrence during execution of jobs [24-28] [67] [68]. During job execution, stragglers can occur due to unhandled requests or ineffective task interference and task incompatibility management. When a processing request is unhandled or not fully handled, tasks expecting the results of this request will have to wait until the full request output is ready, manifesting in straggling tasks. This occurs due to data dependency and task dependency. If the tasks are not oblivious to the heterogeneity of the underlying resources of the platform, their incompatibility (non-synchronization) due to different types of workloads or requirements can manifest in slower execution and ultimately straggler occurrence.

8. Faults: Faults within software and hardware resulting in to crash-stop and late-timing failure can cause straggler occurrence in large-scale systems [6] [18] [32] [33] [69] [70]. The main reasons for software-induced faults can be; development, logic or overflow errors as well as misconfigurations. In terms of hardware, the main fault occurrence reasons are: physical damage, device failures, daemon processes or power-related issues such as effective energy management. Ironically, fault-tolerance and recovery mechanisms can themselves result in straggler manifestation (for example, checkpointing introduces burst in disk access increasing resource contention resulting in a higher system hazard rate).

3.1 Relationship between Straggler Causes

Based on different types of causes of stragglers in large scale systems, we have identified the correlation among them, as described in Table 1.

Table 1: Correlation among Straggler Causes

Stragglers Causes	Dependent
Data Abstraction 1. OS Level 2. Application Level	Resource contention (1), Resource contention (6), Inaccessible Local Disk, Task Execution.
CPU Utilization 1. Increasing Speculative Copies 2. Additional Time	Resource contention, Faults (1), Scheduling (2b), Scheduling (1c)
Scheduling 1. Job Scheduling a. Number of Jobs more than available resources b. Poor admission control mechanism c. Dynamic Requirements 2. Resource Scheduling d. Inefficient resource allocation e. Idle resources are still active	Resource Contention (1), Resource Contention (2), Resource Contention (3), Resource Contention (4), Faults
Inaccessible Local Disk 1. Increasing Backup Tasks 2. Failed to Store Output	Task Execution (1), Task Execution (2), Task Execution (3), Scheduling (1c)
Data Skew 1. "Uneven data distribution among tasks" 2. "Nonuniform data processing time"	Inaccessible Local Disk (2), Data Abstraction (2), Scheduling (1b)
Resource Contention 1. Hardware Heterogeneity 2. Poor User Code 3. Extra Cloning 4. Ineffective Algorithm Logic 5. Temporary Slowdowns 6. More number of copies of same task needs more resources 7. Resource Usage is more than Threshold value	Data Abstraction (1), CPU Utilization, Inaccessible Local Disk (1), Data Skew (1), Task Execution, Scheduling (2a)
Task Execution 1. Unhandled request 2. Task interference 3. Task incompatibility	Resource Contention (1), Resource Contention (2), Scheduling (1a)
Faults 1. Hardware 2. Software	Resource Contention (6), Resource Contention (7), Task Execution (1), Task Execution (2)

As identified in [25], stragglers are not resultant of a singular cause, but can potentially be correlated. For example, data abstraction can occur due to tasks in a queue waiting for execution. Resource contention is the main reason of stragglers due to the sharing of resources among different applications, which are running on different nodes, which further affects the CPU utilization by overloading the resources. Straggler occurrence during scheduling of jobs as well as resources and the reasons of straggler occurrence during resource scheduling can be heterogenous resources, poor user code or logic error and too many copies of straggler tasks are running simultaneously. The reasons of inaccessible local disk can be large copies of backup tasks and failed to store required output, which happens due to task interference and its incompatibility with other tasks. The other reason can be requirements are changing dynamically. Data skew happens due to straggler happens at application level due to data hiding or fail to write data. The other reason can be inefficient allocation of resources for processing of data, which can increase running time of resource. The resource contention occurs at OS level, when master node hides the information from workers. Further, the overutilization of CPU causes the resource contention due to increasing speculative copies as well as when the performance of node degrades. Moreover, poor admission control can also affect the resource utilization and creates resource contention when the value of required resources is increased than the available resources. Further, resource contention affects the task execution due to unavailability of shared resources. Fault occurrences during job execution can happen due to resource failure and resource misconfiguration [39].

4. STRAGGLER MANAGEMENT TECHNIQUES: CURRENT STATUS

Straggler management techniques can be categorized into two broad categories: straggler detection and straggler mitigation. Each category can be further sub-divided into specific areas as shown in Figure 3.

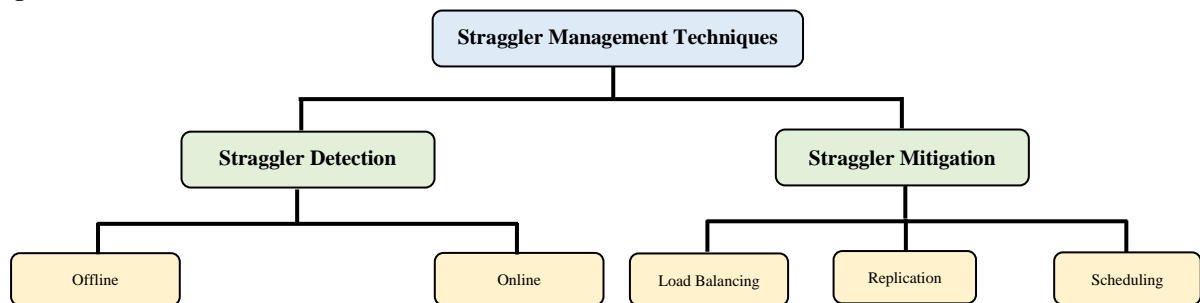


Figure 3: Taxonomy of Straggler Management Techniques

4.1 Straggler Detection Techniques

Straggler detection techniques are leveraged in order to identify straggler occurrence during job execution.

4.1.1 Offline Straggler Detection

Offline straggler detection technique attempts to identify straggler manifestation in order to enhance speculative execution via leveraging offline analytics (i.e. analysing and modelling task execution and progress patterns derived from empirical data *a priori* execution).

Coppa and Finocchi [1] identified three different challenges such as straggling tasks, load unbalancing and data skewness, which affects the performance of computing systems. To overcome these challenges, authors proposed a profile-guided progress indicator called NearestFit to gather the required combination of closest neighbour regression using statistical curve fitting approach. NearestFit is mainly suitable for long running applications and helps to identify the above discussed challenges to increase the efficiency of computing systems. Authors implemented the NodeIterator triangle counting algorithm using homogeneous clusters in Hadoop to test the capability of NearestFit dynamically in terms run time and progress.

Ouyang et al. [24] proposed a technique for Modelling and Ranking Node-Level Stragglers (MRNLS) in CDCs based on analysing the execution trace log data of parallel jobs. This was conducted by a graph-based algorithm is used to partition the server nodes into small nodes to execute more jobs in parallel. The proposed techniques improve the performance of computing systems by reducing task stragglers occurrence. Cong et al. [27] proposed a Machine Learning based Straggler Detection (MLSD) technique using unsupervised clustering method. The proposed technique effectively manages the resources while executing the jobs and diagnosing the stragglers at runtime. Wei et al. [28] proposed Straggler Detection Approach (SDA) for data-intensive computing in cloud environment to detect stragglers at early stage to preserve the efficiency of the CDC. Further, statistical method for outlier detection called Turkey is developed to detect straggler at run time because it starts the speculative execution earlier than the standard deviation method.

4.1.2 Online Straggler Detection

Online straggler detection technique detects the straggler to improve speculative execution using online monitoring tools.

Farshid [2] analysed that map phase of MapReduce (MR) framework takes longer with the increase in number of servers, which further affects negatively the execution time of MapReduce job. Moreover, authors designed an analytical model to identify the impact of stragglers on efficiency of computing system using map phase in terms of application, system and hardware parameters. Experimental results show that model reduces the execution time during execution of MapReduce applications. Zaharia et al. [3] proposed a Resilient Distributed Datasets (RDD), a distributed memory abstraction, which enables developers to provide a fault tolerant module while performing in-memory computations on a huge number of clusters. RDDs uses coarse grained transformations to offer controlled form of shared memory to perform different memory intensive computations in an iterative manner. Further, Spark is used to implement RDDs in a controlled environment to evaluate its performance.

Da et al. [6] proposed Heuristic Algorithm (HA) to search for the best replication to reduce latency in computing systems. The proposed algorithm is used to implement the proposed algorithm and experimental results demonstrate that this is capable to reduce latency and its impact on cost of execution of workloads. Jeffrey and Sanjay [9] explored Data Processing on Large Clusters (DPRC) to perform different aspects such as 1) provide fault tolerance by distributing computations, 2) optimize network bandwidth by decreasing the quantity of data transferred throughout the network and 3) decrease impact of slow machines and improve fault tolerance. In DPRC [9], speculative copy of task is executed by MapReduce on another node for increasing job completion time and reduce response time. It is challenging to select the task for which to execute speculation because it is not trivial to identify the machine or node, which is running slower than average. To implement DPRC effectively, stragglers are recognized at the earliest possible stage calculated by progress scores.

Garraghan et al. [25] explored the Root-Cause of Stragglers (RCS) and provided a method to analyse the root-cause analysis in a massive scale virtualized CDCs to solve the Long Tail challenge effectively. Authors used online analytic agents and offline execution patterns modelling for straggler detection while monitoring tasks dynamically. Heecheol et al. [29] proposed Secure Distributed Computing (SDC) approach using recovery threshold value to efficiently deal with the impact of straggling [59], which uses polynomial codes on sub-tasks allocated to nodes.

4.2 Straggler Mitigation Techniques

Straggler mitigation technique comprise all mechanisms and approaches to tolerate or avoid the impact of straggler manifestation [57]. Such techniques can be further sub-divided into three sub-categories [44] [100] [101] [81]: load balancing based, replication based, and scheduling based.

4.2.1 Load Balancing based Straggler Mitigation

Load balancing based straggler mitigation technique manages the load during mitigation of stragglers.

Ouyang et al. [12] proposed a method to reduce Late-Timing Failure (LTF) and analyse the root-cause of stragglers in Cloud Data Centres (CDC) such as server failures or task concurrency and resource contention. Further, this study identified the high temporal resource contention as a main root-cause of stragglers. Further, the output of experiments demonstrate that this technique maintains the efficiency of the computing systems while tolerating the system failures effectively. Yanfei et al. [23] proposed a user transparent task slot management approach called *FlexSlot*, which identifies the stragglers automatically and resize their slots to improve the speed of execution of task. Further, it balances the usage of resources by automatically changing the number of available slots of nodes to improve its utilization. Moreover, FlexSlot uses adaptive speculative execution approach to improve mitigation of skew data.

Neda et al. [26] proposed Log-Assisted Straggler-Aware (LASA) I/O scheduler for high-end computing to mitigate the impact of storage server stragglers. Further, a scheduling algorithm is proposed to make effective decisions to manage stragglers at runtime. The output of experiments demonstrate that LASA is performing better in load balancing while mitigating the storage server stragglers dynamically. Eman et al. [31] proposed a parallel model for straggler mitigation in distributed spatial simulation called Priority Asynchronous Parallel (PAP) to exploit data dependencies of parallel processes to be computed and synchronized based on data priority to the other workers. Moreover, load balancing and partitioning method is proposed to balance the workloads among different nodes and help to improve the performance speedup by a large extent. Haozhao et al. [37] proposed Heterogeneity-aware Gradient Coding (HGC) scheme to execute the jobs in heterogenous environment and efficiently tolerate the stragglers without degrading the effectiveness of the cloud services [58]. The output of experiments demonstrates that HGC scheme outperforms in computation time.

4.2.2 Replication based Straggler Mitigation

Replication based straggler mitigation technique replicates the adequate number of tasks during mitigation of stragglers.

Mehmet et al. [7] analysed the Trade-off between Latency and Cost (TLC) using simple replication or erasure coding for straggler mitigation in executing jobs with many tasks. Experimental results show that delaying redundancy is not effective in reducing cost. Further, Mehmet et al. [10] developed a Straggler Mitigation (SM) technique using delayed relaunch of tasks, which helps to reduce cost and latency effectively. Da et al. [11] proposed an idea of an efficient Task Replication Technique (TRT) for straggler management to improve the response time in parallel computations. Further, this technique is implemented in [13], demonstrate empirically that replicating all operations can result in significant mean and tail latency reduction in real world systems including Domain Name System (DNS) queries, database servers, and packet forwarding within networks.

Tien-Dat [15] [33] proposed Energy-Efficient Straggler Mitigation (EESM) technique for effective management of big-data applications in the cloud computing environment to optimize the energy consumption during straggler occurrence. Firstly, authors characterize the effect of straggler mitigation on energy-efficiency. Secondly, a straggler detection framework is developed, and they identified that only 12% of the detected tasks are real stragglers [33]. The usage of huge number of speculative copies is the main reason of unnecessary energy consumption. Thirdly, a reservation-based straggler handling approach is proposed to optimize the energy efficiency by allocating the required resources at runtime effectively.

Da et al. [17] analysed the trade-off between latency and cost to find out the best replication technique for straggler management based on following parameters: 1) when to perform replication

for straggling tasks, 2) number of replicas to be launched and 3) is it necessary to destroy the original copy or not. Further, a Straggler Management Approach (SMA) is proposed to calculate the value of latency-based empirical distribution of execution time of task. The output of experiments demonstrate that this work gives better for two performance parameters such as cost and latency. Lei et al. [19] proposed a straggler management technique called CREST (Combination Re-Execution Scheduling Technology) for fast speculation of straggler tasks in MapReduce framework, which further reduces the response time of MapReduce jobs. The re-execution of set of tasks on set of computing nodes in CREST improves the speed of task execution.

Radheshyam et al. [20] proposed a Job-Aware Scheduling (JAS) technique to optimize the running time of different jobs by maintaining the harmony among them, which are executing on same cluster. JAS technique is implemented using for MapReduce framework. Further, proposed algorithm selects the most compatible task with executing task to reduce more execution time. Moreover, a heuristic based load balancing technique is developed to avoid the underloading and overloading of resources. Matei et al. [21] explored the MapReduce framework for straggler management and improved its performance in heterogenous environment. Further, a resource scheduling algorithm, Longest Approximate Time to End (*LATE*) is proposed to improve the robustness in regard to heterogeneity and improves response time of tasks. *LATE* scheduling algorithm [3] estimates the longest approximate time and select the task with longest approximate time as straggler tasks and execute its speculative copy on another fast node to speed-up the job completion time. SAMR scheduling technique [18] computes the completion of tasks at runtime and discover the straggler task based on execution time. Further, historic information of node is used to detect more reliable node in SAMR and weights of reduce and map stages are updated after completion of every task.

Farhat et al. [32] proposed a Straggler Management technique for Modelling and Optimization (SMMO) of straggling mappers to show the stochastic behaviour of mapper nodes and its negative effect on completion time of MapReduce jobs. Authors identified task inter-arrival time of jobs to map the required nodes of heterogenous CDC in an optimized way. The experimental results demonstrate that the proposed technique reduced the execution time of jobs at runtime. Behrouzi-Far et al. [74] proposed an efficient straggler replication framework in large-scale parallel computing to analyze the performance of the system in terms of latency-cost trade-off. Further, it identifies the best replication technique based on different criteria such as: (i) number of replicas required, (ii) time to replicate straggling tasks and (iii) determine whether to kill the original task. Finally, performance evaluation is described that latency and cost is reduced in Google Cluster Trace as compared to MapReduce.

4.2.3 Scheduling based Straggler Mitigation

Scheduling based straggler mitigation technique schedule the resource for jobs during mitigation of stragglers.

Ganesh et al. [4] explored the straggler mitigation techniques and identified the impact of reasons of stragglers in latency sensitive jobs. Further, authors designed workloads with small number of jobs and performed cloning of small jobs. It has been identified that the cloning of small jobs uses less resources but improves the reliability of computing services. Moreover, a system named *Dolly* is developed to generate multiple clones of jobs and execute jobs within their specified budget. Experimental results demonstrate that *Dolly* sped up jobs by 46% by using only 5% extra resources.

Ganesh et al. [5] proposed Greedy Speculative scheduling and Resource Aware Speculative scheduling (GRASS) technique, which uses speculation to mitigate the impact of stragglers in approximation jobs. GRASS uses extra resources for speculation and improves accuracy for deadline-bound jobs by 47% and speeds up error-bound jobs by 38%. Aaron et al. [8] addressed the straggler problem for Iterative Convergent Parallel (ICP) machine learning technique to identify the behaviour (in terms of delay) of the system during execution of jobs by injecting the stragglers.

Amazon EC2 and Microsoft Azure [60] is used to evaluate the performance of system in terms of execution time.

Ouyang et al. [14] proposed a Straggler Management Technique (SMT) to find the task stragglers by calculating threshold value at runtime. Further, this technique considers important key parameters such as resource utilization, task execution and job QoS timing constraints to manage straggler tasks effectively. Neeraja et al. [16] proposed straggler management technique called *Wrangler* to proactively avoid the conditions, which cause stragglers. *Wrangler* [4] uses interpretable linear modelling approach to reduce the resource wastage by eradicating the requirement for replicating tasks. It uses fewer resources to complete the job in a faster way and avoids the straggler proactively by predicting in advance. A cluster resource utilization based statistical learning technique is used for confidence measure to offer reliable task scheduling by predicting errors in advance. The output of experiments show that *Wrangler* produces improvements in terms of Job Completion Time and resource utilization as compared to speculative execution.

Quan et al. [18] proposed a Self-Adaptive MapReduce (SAMR) scheduling technique for straggler management, which estimates task progress automatically and adapts to the changing conditions of environment dynamically. SAMR uses MapReduce mechanism to divide jobs into tasks and execute on different available nodes. SAMR does not create backup tasks for regular tasks. SAMR reduces the execution time of MapReduce jobs while executing tasks in heterogenous environment. Enhanced SAMR (ESAMR) [22] uses the k-means clustering algorithm to categorize the historic data of each node into k-clusters and identifies the straggler task more accurately. Furthermore, ESAMR uses weights of reduce and map stages to find the *Time to End* on different nodes, which can easily identify the more reliable node.

Ganesh et al. [22] studied and explored the straggler management in resource aware techniques and identified the main causes of stragglers such as varying bandwidth, network congestion, workload imbalance and contention of resources (network, memory and processor). Furthermore, *Mantri* [22] is used monitor task execution and take a proactive action to sustain the efficiency of the CDC in the case of resource contention or hardware/software failure [75] [77] [78]. It uses Bing traces to evaluate the performance and it improves job completion time to a large extent.

Ouyang et al. [30] proposed a Straggler Management Mechanism (SMM) to improve the execution efficiency of Internet-ware applications by dynamically calculating the straggler threshold, considering important parameters such as optimal system resource utilization, task execution progress and job QoS timing constraints. Further, YARN architecture is used to implement dynamic straggler threshold to test the performance of the proposed mechanism and experimental results gives the better outcomes in terms of response time. Rong et al. [36] developed Large-scale Multimedia Semantic Concept (LMSC) model to improve the scalability of the computing systems with heterogenous environment. Robust Subspace Bagging algorithm is used to improve learning process and further, a task scheduling algorithm is proposed to improve the scalability by executing heterogenous tasks. Proposed model is tested on MapReduce framework and experimental results demonstrate its superiority.

Figure 4 presents the evolution (2008-2019) of different types of straggler management techniques along with their focus of study and QoS.

Table 2 shows the comparison of different types of straggler management techniques based on different parameters.

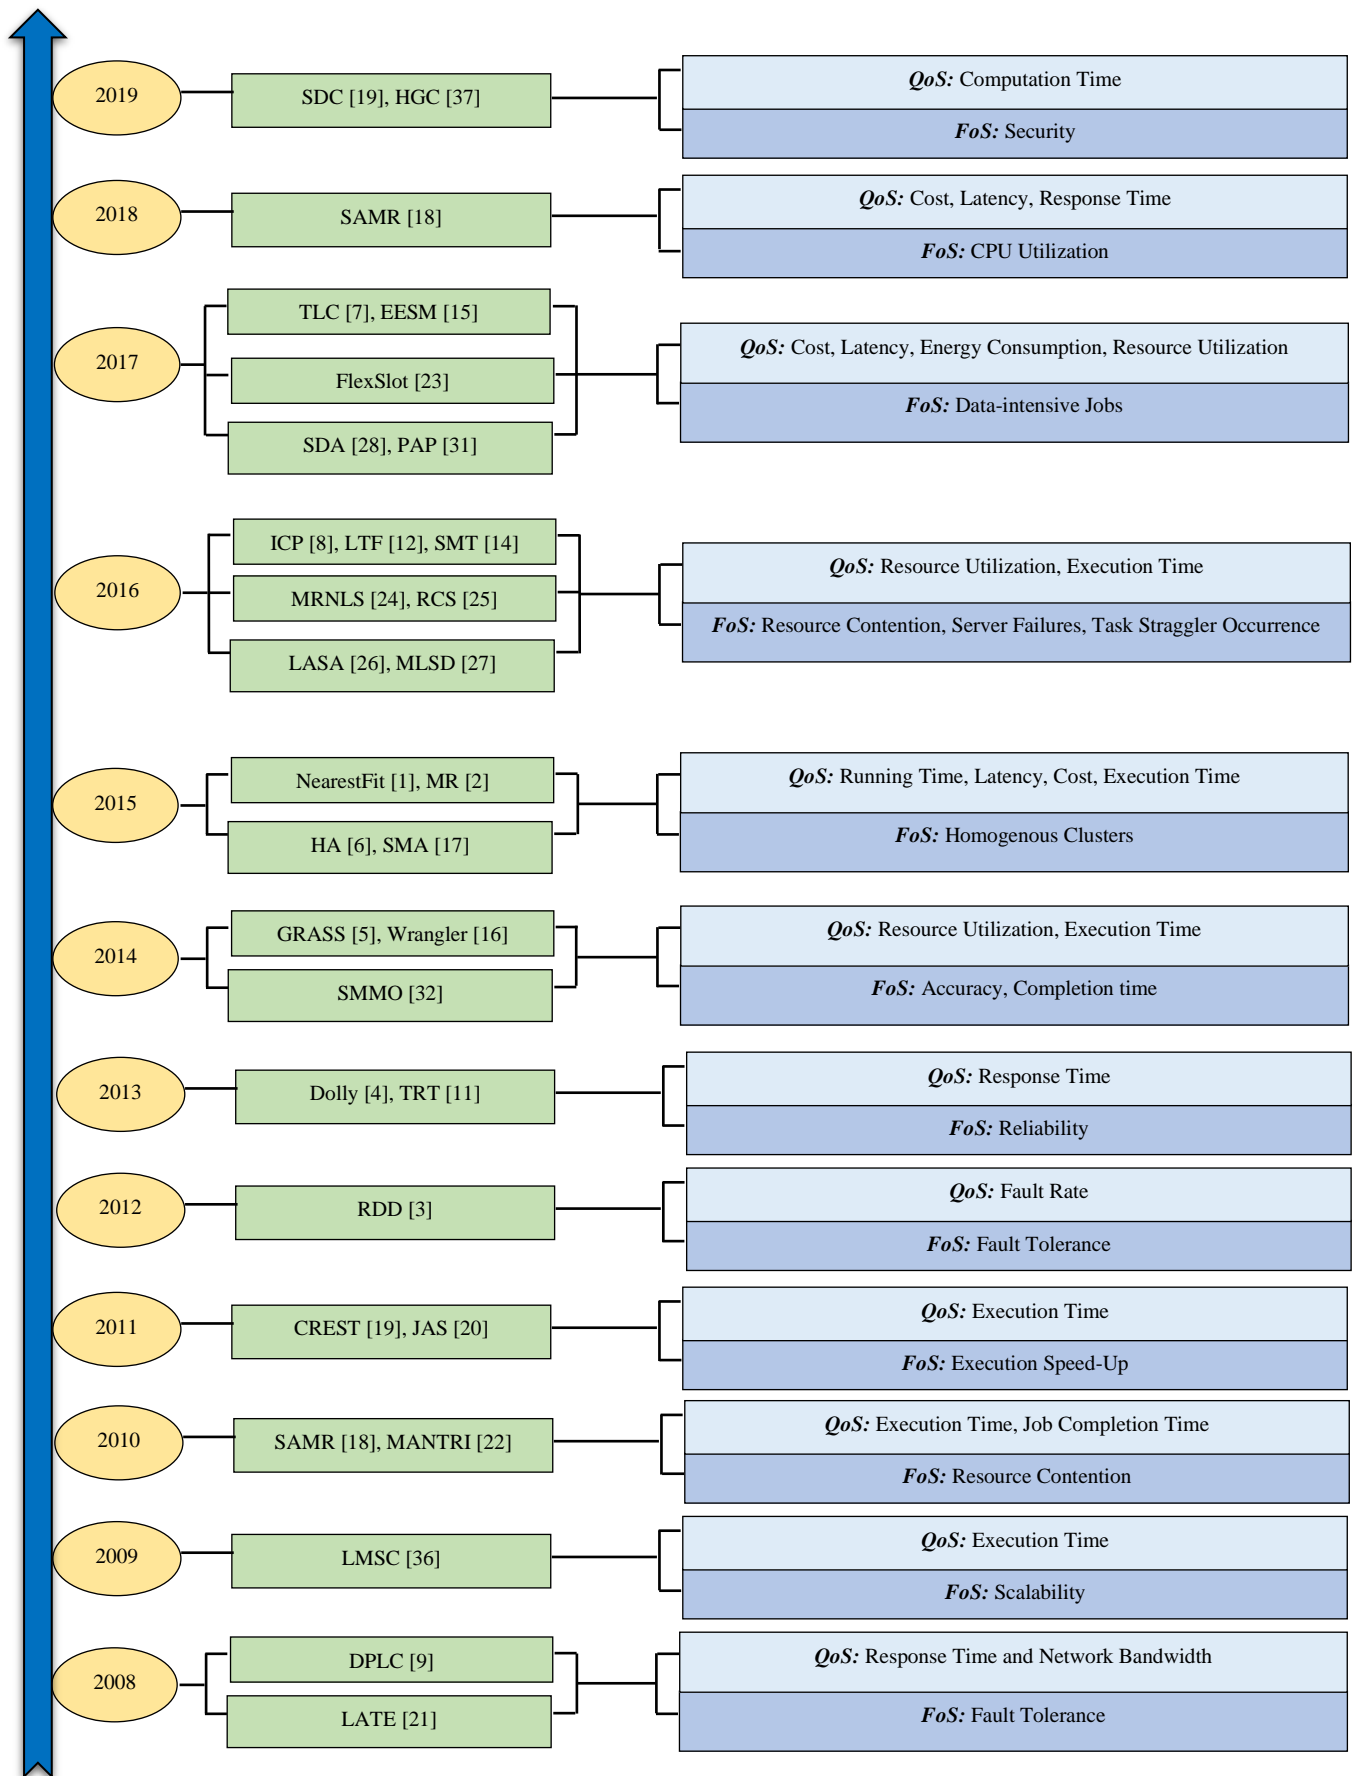


Figure 4: Evolution of Straggler Management Techniques

Table 2: Comparison of Straggler Management Techniques

Work	Straggler type	Cause	Environment	Type	Dataset	Delay	Merits	Open Challenges
[1]	Slow tasks, non-local task, data skew	Wall-clock times of slow running tasks far from prediction from linear progress assumption	Hadoop	Pro	50GB archive Wikipedia articles	Low	Higher accuracy in profile optimization	If a job is too short, unable to collect enough profiling data to obtain accurate prediction
[2]	Delay time is larger	Shuffle phase overlaps with the Map phase of MapReduce	Hadoop	Rea	Word Count Data Set (3GB)	High	Optimized execution time	Inaccurate estimation of time to execution
[3]	Memory utilization	Lack of data sharing abstractions	Spark	Pro	1TB Wikipedia logs (2 years).	Low	Efficient data recovery	Memory requirements grow as the dataset size increases
[4]	Resource contention	Extra clones cause contention for intermediate data and shared resources	Hadoop	Pro	Facebook and Bing traces	High	46% job speedup, 5% extra CPU	1. Homogenous job sizes required for optimization 2. Extensive analysis needed to determine straggler chance
[5]	CPU Utilization	Spawning speculative copy leads to task using two (or more) resources simultaneously	Spark and Hadoop	Pro & Rea	Facebook and Bing traces	Med	Improves deadline job accuracy 47%; error jobs 38%,	Weigh the impact of speculating a running task with scheduling a new task of any job.
[6]	Larger response time	Resource contention due to resource sharing (deadlock)	Hadoop	Rea	Google Trace Data	High	“Reduced latency and cost of computing resources”	Develop online strategy to learn execution time distribution to launch replica, instead of using historical traces.”
[7]	Delaying redundancy	Delaying redundancy increases latency	NA	Rea	Google Trace Data	High	Analysed trade-off between cost and latency	Degree of redundancy can be reduced without affecting latency.
[8]	Transient slowdown of worker thread	Temporary slowdowns (due to resource contention with a background activity) often occurs on non-blacklisted machines	Amazon EC2 and Microsoft Azure	Pro	Netflix dataset (480k-by-18k sparse matrix with 100m elements)	High	Reduced execution time	Resource overloading
[9]	Inaccessible local disk	Completed map tasks are re-executed on failure because output is stored on local disk of failed machine, thus inaccessible.	Hadoop	Rea	Google Zeitgeist data	Low	Fault-tolerance, locality, optimization load balancing	Redundant execution can be used to reduce the impact of slow machines, but consumes more energy
[10]	Ineffective CPU Utilization	Relaunching tasks before minimum task completion time causes work loss and latency	Hadoop	Pro	Google Trace Data	High	Reduced redundancy, cost, latency	Need of empirical evidence to accurately model execution time for resource load
[11] [13]	Larger response time	Latency in locating small disk file, time needed to load file from disk	Apache Cassandra	Pro	Facebook and Google traces	High	Reduced latency using redundancy	Ignoring data locality for launching speculative map tasks leads to performance degradation
[12]	Resource Contention	Unbalanced workload aggregation and poor user code	Hadoop	Pro	NA	Low	Reduced late-timing failures	Request handling inefficiency is due to overloaded file request
[14]	CPU and Memory Utilization	Contention of shared resources, node disk failures, and imbalanced task workloads	SEED	Rea	Microsoft Bing traces	High	Improved job completion time	To design a cost function for further optimization of proposed method
[15] [33]	Energy consumption, resource contention	Large resource cost with many speculated tasks, increased energy usage	Hadoop	Pro	MapReduce Application (Wordcount, Cloud Burst)	Med	Improved energy-efficiency	Speculative copies can be launched adaptively to improve further the output accuracy

[16]	Resource contention	Local conditions exceed fixed thresholds defined over set of resource usage statistics	Hadoop	Pro	Facebook 2009, Cloudera Hadoop trace	High	Improved resource utilization, job completion time	Reduce time spent for capturing training data per node in a cluster
[17]	Long task delay	Additional time needed for original copy to finish more likely to be shorter than new copy execution time	Hadoop	Pro	Google Cluster Trace	High	Latency and cost are improved	Technique can be tested on heterogeneous servers, identify task dependencies (some tasks need to complete prior to others)
[18]	Task Replication	Larger number of task replication resulting in high overhead	Hadoop	Pro	MapReduce Application (Wordcount and Sort)	Low	Reduced execution time	Ignoring of data locality for launching speculative map tasks may lead to severe performance degradation
[19]	High speculative map task execution	Ignoring data locality for launching speculative map task may lead to performance degradation	Hadoop	Rea	Specific gene segment from 50GB DNA	Low	Reduced running time of a speculative map task	Nonuniform data processing time; performance interference from co-running jobs.
[20]	No synchronization between tasks	Task incompatibility due to their different requirements	Hadoop	Rea	Terasort, grep, webcrawling, WordCount, video conversion	Low	Reduced execution time	Scale up/down cluster by switching on/off virtual machines or nodes based on resource usage to reduce energy
[21]	Performance degradation in heterogeneous environment	Many speculative tasks may launched leading to resource exhaustion	Hadoop and EC2	Pro	Facebook and Yahoo traces	Low	Improved task response time	Workload imbalance between workers
[22]	Varying bandwidth, network congestion, workload imbalance and resource contention (network, memory, CPU)	Disk-intensive tasks scheduled to a node with slow disk	Hadoop	Pro	Bing traces	Med	Improved job completion time	Longer communication delays
[23]	High resource contention (CPU utilization)	Uneven data distribution among tasks, nonuniform data processing time; performance interference from co-running jobs	Hadoop	Pro	Wikipedia traces	Med	Data skew mitigation, reduced job completion time	Adaptive speculation execution approach can be developed for resource management in YARN.
[24]	Task-level stragglers during job execution	Lack of correlation between straggler tasks and available slots at system run-time	Hadoop	Rea	Google trace	Low	Reduced task straggler occurrence	Performance-aware algorithm can improve straggler mitigation, node-level identification.
[25]	Resource contention	Hardware heterogeneity	SEED	Pro	Google trace	Low	Monitored and detected straggler tasks at runtime	Root-cause analysis via machine learning to cross-correlate heterogeneous system traces for intelligent failure detection.
[26]	Storage server stragglers	Workload imbalance between workers	Simulator	Rea	Synthetic workloads generated based on real-world traces	Low	Mitigate storage server stragglers dynamically	Develop prototype of integrating proposed I/O scheduling into existing parallel file systems
[27]	Workload imbalance	Datasets on stragglers are expected to demonstrate imbalanced label distribution	Spark	Rea	Google traces	Low	Run-time straggler detection	Fine-grain straggler identification by per slowest node and cluster

[28]	High job completion time	Late speculative execution on straggler tasks does not reduce job completion time.	Simulator	Rea	NA	Low	Run-time straggler diagnosis	Slow task execution (e.g., non-local task and data skew)
[29]	Data abstraction obfuscation	Master node wants to hide input data from workers	NA	Rea	NA	Med	Reduced straggling effects on sub-tasks	Communication delays and computation delays
[30]	High CPU utilization	Imbalanced workload or uneven input data Sizes	SEED	Pro	Sort, Wordcount, Hive query	Med	Improved response time	Design a cost function beyond CPU and memory utilization, including disk volume and network speed
[31]	Data dependencies of parallel process	Workload imbalance between workers in cloud datacenters	SMARTS	Pro	OpenStreetMap dataset	High	Improved performance speedup	High resource overhead associated with increased backups
[32]	Resource contention	Hardware heterogeneity	Hadoop	Rea	Bing, Facebook traces	High	Reduced job execution time	Hardware/software faults identified and removed from task scheduling
[36]	Task-level stragglers during execution of jobs	Hardware heterogeneity	Hadoop	Pro	Google traces	Med	Improved Scalability	Computation delay is larger
[37]	Resource contention	Hardware heterogeneity	Simulator	Rea	NA	High	Reduced job completion time	Overloading of resources

Pro = Proactive, Rea = Reactive, High means the value is more than 80%, Medium (Med) means the value between 50% to 80% and Low means the value less than 50%

5. COMPARISON OF STRAGGLER MANAGEMENT TECHNIQUES BASED ON TAXONOMY

Table 3 shows the comparison of straggler management techniques based on taxonomy of straggler causes from Figure 1 and Table 2.

Table 3: Comparison of Straggler Management Techniques based on Taxonomy of Straggler Causes

Work	Data Abstraction	CPU util.	Scheduling		Inaccessible Local Disk	Data Skew	Resource Contention	Task Execution	Faults
			Resource Schedule	Job Schedule					
[1]	OS Level	NA	Inefficient Resource Allocation	NA	Increased Backup Tasks	Uneven Data Distribution	Extra Cloning	Task Interference	NA
[2]	NA	Additional Time	NA	More Resource Requirement	NA	NA	Hardware Heterogeneity	Task Interference	Software
[3]	Application Level	NA	Idle Resources	High resource requirement	Failed to store output	Uneven data distribution	Extra cloning	Task incompatibility	NA
[4]	OS Level	NA		NA	NA		Extra Cloning	NA	Hardware
[5]	NA	Increased speculative copies	Idle Resources still Active	Poor admission control	NA	Uneven Data Distribution	Hardware Heterogeneity	Unhandled Request	Software
[6]	Application Level	NA	Idle Resources Still Active	Poor admission control	Failed to store output	Non-uniform data processing time	Ineffective Algorithm Logic	Unhandled Request	NA
[7]	NA	NA	NA	NA	NA	Non-uniform Data Processing Time	Temporary Slowdowns	Task Incompatibility	Hardware
[8]	OS Level	NA	NA	NA	NA	NA	Temporary Slowdowns	NA	Hardware
[9]	NA	NA	Inefficient Resource Allocation	Dynamic Requirements	Failed to Store Output	Uneven Data Distribution	Ineffective Algorithm Logic	Task Interference	NA
[10]	NA	Additional Time	NA	NA	NA	NA	Extra Cloning	NA	Hardware

[11] [13]	Application Level	Additional Time	NA	NA	Increasing Backup Tasks	NA	Hardware Heterogeneity	NA	Hardware
[12]	Application Level	NA	NA	NA	NA	NA	Poor User Code	NA	NA
[14]	NA	Increasing Speculative Copies	Idle Resources still Active	Dynamic Requirements	Increasing Backup Tasks	NA	Ineffective Algorithm Logic	Unhandled Request	Hardware
[15] [33]	OS Level	NA	NA	Dynamic Requirements	NA	Uneven Data Distribution	Multiple Copies	NA	NA
[16]	NA	NA	NA	NA	NA	NA	Extra Resource Usage	NA	Software
[17]	NA	Additional Time	NA	NA	NA	NA		NA	Software
[18]	Application Level	NA	Idle Resources are Still Active	More Resource Requirement	Increasing Backup Tasks	NA	Extra Cloning	Unhandled Request	
[19]	OS Level	Increasing Speculative Copies	NA	NA	NA	Non-uniform Data Processing Time	Extra Cloning	NA	Hardware and Software
[20]	OS Level	NA	NA	NA	NA	NA	Multiple Copies	Task Incompatibility	NA
[21]	OS Level	Increasing Speculative Copies	NA	Poor Admission Control	Failed to Store Output	NA	Multiple Copies	NA	NA
[22]	Application Level	NA	NA	NA	NA	NA	Temporary Slowdowns	Task Incompatibility	Hardware
[23]	NA	NA	Inefficient Resource Allocation	Dynamic Requirements	Increasing Backup Tasks	Uneven Data Distribution	NA		Software
[24]	NA	Additional Time	NA	NA	NA	NA	NA	Task Interference	Hardware
[25]	NA	NA	NA	NA	NA	NA	Hardware Heterogeneity	NA	Hardware
[26]	NA	NA	Idle Resources are Still Active	More Resource Requirement	Increasing Backup Tasks	NA	NA	Task Interference	Software
[27]	Application and OS Level	Increasing Speculative Copies	Idle Resources are Still Active	Poor Admission Control	Failed to Store Output	Non-uniform Data Processing Time	Extra Cloning	Unhandled Request	NA
[28]	OS Level	Increasing Speculative Copies	NA	Dynamic Requirements	NA	NA	Hardware Heterogeneity	NA	Hardware
[29]	OS level	NA	NA	NA	NA	NA	Poor User Code	Task Incompatibility	NA
[30]	Application Level	Increasing Speculative Copies	Inefficient Resource Allocation	More Resource Requirement	NA	Uneven Data Distribution	Ineffective Algorithm Logic	NA	Software
[31]	NA	NA	Inefficient Resource Allocation	Poor Admission Control	Failed to Store Output	Non-uniform Data Processing Time	Temporary Slowdowns	Task Interference	NA
[32]	Application Level	Additional Time	NA	More Resource Requirement	Increasing Backup Tasks	Non-uniform Data Processing Time	Hardware Heterogeneity	Task Incompatibility	Hardware
[36]	OS Level	Additional Time	NA	Poor Admission Control	Failed to Store Output	NA	Hardware Heterogeneity	NA	Hardware
[37]	Application Level	NA	Inefficient Resource Allocation	Dynamic Requirements	NA	NA	Hardware Heterogeneity	NA	NA

NA: Not Applicable

5.1 Analysis of Experimental Results: Practical Use-Case

The existing straggler management techniques have been categorized into two categories i.e. straggler detection and mitigation techniques. Table 4 shows the analysis of experimental results of straggler detection and mitigation techniques in the context of different performance parameters. Future researchers can use Table 4 to validate their research work based on the values of various performance parameters identified from existing literature. Literature reported that there are four types of data abstraction levels (OS, application, server and VM), where straggler can occur.

Table 4: Analysis of experimental results of straggler detection and mitigation techniques

Study	Type	Data Abstraction Level	Context																			
			Number of Nodes/Workers	Energy Consumption	CPU Utilization	Disk Utilization	Data Transfer	Memory Utilisation	Deadline	Execution Time	Number of Tasks	Number of Jobs	Data Size	Latency	Failure Prediction Accuracy	Number of Failures	Average Error	Slowdown/Delay	Interval Arrival Rate	Running Time	Response Time	
[1]	Straggler Detection Techniques	Application	8-32	NA	>61%	NA	NA	NA	NA	NA	NA	50-150	NA	NA	NA	NA	2.5-5%	5-15 %	NA	0-30 m	NA	
[2]		Server	40-200	NA	NA	NA	NA	>71%	NA	0.5s – 1000s	NA	NA	1-3 GB	6.5-9 s	NA	NA	NA	0.5 s – 2 s	10-30 J/s	NA	NA	
[3]		Application	25-100	NA	NA	NA	NA	NA	NA	0-184s	NA	NA	100 GB – 1 TB	NA	NA	1-119	NA	NA	NA	NA	1.7 – 6.6 s	
[6]		OS	NA	90 kWh	>82%	NA	NA	NA	NA	NA	400-488	1017	NA	4-8 s	NA	NA	NA	NA	NA	NA	0-3k s	NA
[9]		Server	NA	NA	NA	NA	0 – 20k MB/s	NA	NA	0-634s	3351	29423	3228 TB	NA	NA	NA	NA	NA	NA	NA	0-1000s	NA
[24]		Server	132	NA	NA	>40%	NA	NA	NA	50-300 s	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA
[28]		Application	NA	NA	>80%	NA	NA	NA	NA	10-14 m	252950	3043	NA	NA	NA	NA	NA	NA	NA	NA	29 days	NA
[25]		Application	NA	NA	>80%	NA	NA	90%	NA	10-14 m	1233879	875	NA	1200-1500 ms	NA	NA	NA	6%	NA	14 days	30 m	
[27]		Server	12000	NA	NA	NA	NA	NA	NA	NA	NA	417	NA	NA	62.92%	NA	NA	NA	NA	NA	NA	NA
[4]		VM	48	NA	NA	>60%	NA	NA	NA	0-35 s	1-5	1-500	NA	NA	NA	NA	NA	NA	NA	NA	1-6 h	NA
[5]	VM	NA	NA	NA	NA	NA	NA	2-20%	30-50%	NA	50-500	NA	NA	20-52%	NA	5-30%	NA	NA	NA	NA	NA	
[7]	Application	NA	NA	NA	NA	NA	NA	NA	50-110s	NA	NA	NA	0-120s	NA	NA	1.5-3.5%	NA	NA	NA	NA	NA	
[8]	Server	90	NA	NA	NA	NA	NA	NA	0-100s	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	5-32s	NA	
[10]	Application	NA	NA	NA	NA	NA	NA	NA	50-110s	NA	NA	NA	60-1600s	70%	NA	1.5-3.5%	NA	NA	NA	NA	NA	
[12]	Application	60	NA	>80%	>80%	NA	NA	NA	NA	NA	NA	NA	NA	NA	6-22	NA	8s	NA	NA	NA	NA	
[11]	Servers	1-10	NA	NA	NA	NA	NA	NA	0-12s	NA	700	NA	0-1sec	NA	NA	NA	NA	0-12 req/s	NA	1-1000s	NA	
[13]	Application	20-80	NA	27-135%	NA	NA	27-135%	NA	100-1100	NA	NA	24.6 GB	NA	NA	NA	2-3%	4-6s	NA	NA	680-1080ms	NA	
[14]	Application	NA	14-18 10 ⁹ J	NA	NA	NA	NA	NA	12-52 s	1759434	1735	20 GB	NA	40-60%	NA	NA	NA	NA	7-14 10 ³ s	NA	NA	
[15]	Application	50-100	NA	NA	NA	NA	NA	NA	0-70 s	NA	22974	NA	NA	20-85%	NA	NA	0.22-21.84%	NA	9 days	NA	NA	
[33]	Application	NA	NA	NA	NA	2.87 MB/s	NA	NA	NA	NA	NA	NA	1750-5200 ms	NA	NA	NA	NA	NA	0-1600 s	18s	NA	
[16]	Application	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	
[17]	Application	1-8	NA	NA	NA	NA	NA	NA	210-330s	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	
[18]	VM	100	NA	NA	NA	NA	NA	NA	NA	1-8	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	
[19]	Application	12	NA	NA	NA	NA	NA	NA	NA	NA	0-35	64 MB	NA	NA	50-91	NA	11%	NA	0-125s	NA	NA	
[20]	VM	871	NA	NA	NA	52.1 MB/s	NA	NA	10-70s	NA	NA	30 GB	NA	NA	NA	NA	NA	NA	0-2.5 s	NA	NA	
[21]	Server	NA	NA	20-80%	NA	NA	NA	NA	0-300 s	NA	NA	NA	NA	NA	NA	NA	NA	20-80 req/s	NA	NA	NA	
[22]	VM	32	60 kWh	NA	NA	NA	NA	NA	0-300 s	NA	NA	150 GB	NA	NA	NA	7-9%	NA	NA	NA	NA	NA	
[23]	Application	300	NA	NA	NA	NA	NA	NA	20-200s	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	
[26]	VM	40	NA	NA	NA	11 MB/s	NA	NA	5-85s	500-10000	NA	NA	NA	NA	NA	NA	NA	NA	NA	130-213s	NA	
[30]	Application	80	NA	NA	NA	NA	NA	NA	15-80s	NA	NA	NA	NA	NA	82-98	NA	NA	NA	NA	NA	NA	
[31]	Server	NA	NA	NA	>55%	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	3-6%	4-9 s	0-25 req/s	NA	NA	NA	
[32]	OS	16	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	60-70%	NA	NA	NA	NA	NA	NA	NA	
[36]	Application	8-48	NA	>40%	NA	NA	NA	NA	0-15 s	NA	NA	NA	NA	NA	NA	<2%	4s	NA	NA	NA	50s	
[37]	Application	8-48	NA	>40%	NA	NA	NA	NA	0-15 s	NA	NA	NA	NA	NA	NA	<2%	4s	NA	NA	NA	50s	

NA: Not Available, s: Seconds, ms: Milliseconds, GB: Giga Bytes, kWh: Kilo Watt Hour, MB/s: Mega Bytes per second, req/s: Number of Requests per Second, m: minutes

5.2 Trend Analysis

Our systematic review has identified different types of result outcomes for different categories of straggler management techniques developed from year 2008 to year 2019. The scheduling-based straggler mitigation technique appears prominent across the years except year 2012. After the scheduling-based techniques, researchers focused on replication-based straggler mitigation, during the year 2013 to 2019. The offline, online and load-balancing straggler management techniques are less focused on from year 2008 to year 2019 requiring research to improve the straggler management in large scale systems. Researchers focused on scheduling and replication-based straggler management in years 2018 and 2019. Figure 5 shows the year-wise publications of straggler management techniques and it has been clearly depicted that research from year 2008




Figure 5: Publications of Straggler Management Techniques

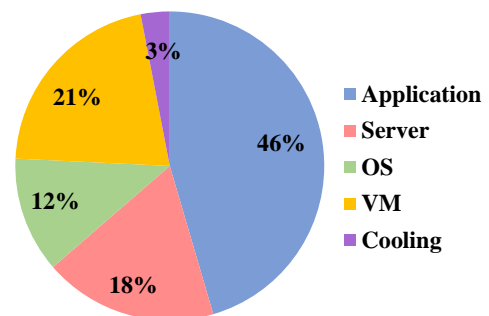


Figure 6: Straggler Type Breakdown in Literature

to 2016 was highly progressive in this area, declining after 2017 and 2018 while progressing in 2019.

Literature reports the research related to straggler management is mostly published in journals (31%), followed by conferences (28%), transactions (21%) and book chapters (10%). The rest of the research is published in symposiums, workshops, white papers and PhD thesis. Figure 6 shows the research conducted related to straggler management at different levels such as Application, Server, OS, VM and cooling. Figure 6 clearly shows that most of the research work has been done at the application level (46%) and followed by VM level (21%). Only 3% of research work has been done at cooling level.

Literature reports 44% of research work considered between 0 and 100 nodes for performance evaluation and only 7% research work considered 1000+ nodes. There are four different types of studies identified from literature: real testbed based (63%), systematic reviews (7%), conceptual models (10%) and simulation based (20%). Most of the technical research papers (63%) consider real testbeds for performance evaluation. There are only two reviews [34] [35], which have been done in this area. Table 5 shows the different research work related to different performance parameters identified from Table 4.

5.3 Observations

From the trend analysis, it is observable that current related works focus on studying and mitigating specific straggler types, ranging from resource contention to data skew as shown in Table 2. This appears to be a necessity given the complexities and management strategies appropriate for each straggler type. The challenge is that it is possible for straggler manifestation to be correlated in terms of system phenomena, but also management techniques themselves (e.g. use of speculative copies to address data skew causes increased resource contention).

The important research challenges within the large-scale cloud data centres such as latency, scalability, energy consumption and data processing are contributing to the rise in research in the field of straggler management, which can be solved by using Artificial Intelligence techniques. On the other hand, there is a need of real cloud infrastructure (at least 50 physical nodes) to test the

performance of future straggler management techniques but it would be very expensive to afford for academic institutes. To solve this problem, industries such as Facebook, Google, Amazon should collaborate with academic institutes to provide required infrastructure to do real experiments.

Table 5: Research Work Related to Performance Parameters

Performance Parameters	Study
Number of Physical Nodes/ Workers	[1] [2] [3] [11] [13] [14] [16] [18] [19] [20] [21] [23] [24] [26] [27] [28] [30] [31]
Number of Virtual Nodes/ Workers	
Energy Consumption	[15] [33]
CPU Utilization	[12] [14] [22] [25]
Disk Utilization	[12]
Data Transfer	[9] [17] [21]
Memory Utilisation	[14]
Deadline	[5]
Execution Time	[2-18] [33] [21] [22] [23] [24] [25] [26] [27] [28] [29] [30] [31] [36] [37]
Number of Tasks	[4] [6] [9] [15] [33] [25] [30]
Number of Jobs	[4] [5] [6] [8] [15] [16] [33] [21] [23]
Data Size	[33]
Latency	[2] [6] [7] [10] [11] [13] [16] [27]
Failure Prediction Accuracy	[5] [16]
Number of Failures	[3]
Average Error	[1] [5] [7] [10]
Response Time	[3] [11] [33] [14] [30]
Slowdown/ Delay	[1] [2] [16] [32]
Interval Arrival Rate	[2] [11] [13] [22] [32]
Running Time	[1] [4] [6] [8] [9] [11] [15] [16] [17] [20] [21] [25] [33]

This systematic review also identifies various research directions for perspective researcher scholars, who are working in the field of straggler management for distributed systems and searching for new research challenges to improve the performance of cloud services. The straggler management is an evolving field of research for large scale systems and it is quite challenge ring to execute user workloads without occurrence of stragglers. To solve this problem, there is a need to recognize the reasons of long-tail problem or stragglers and their correlations, which can help to find out the dependency among stragglers. This study [1] developed straggler management technique for profile guide more accurately, but accurate predication is difficult to get if job is very small to gather required profiling data. An efficient data recovery is achieved in [3], but it has been identified that the memory requirements do not grow to intolerable levels as the size of dataset is increasing, which further causes the stragglers. The jobs are increasing with time, but there is need to analyse the impact of multiple jobs on probability of stragglers [4]. Existing techniques uses historic data to estimate resource requirement [6]. However, there is a need to develop an online strategy to simultaneously learn the execution time distribution and launch replicas, instead of estimating time using historical traces. Further, the replication increases the reliability of execution of jobs, but it consumes more energy consumption, which is a global challenge to address [7]. The scale up/down infrastructure by switching on/off the virtual machines/nodes based on the resource usage of the cluster to save energy is required [20]. The dependency among tasks during task execution further effects causes the stragglers because some tasks need to complete in order to begin others [17]. Existing straggler management techniques are required to improve to attain to reduce straggler occurrence. By using this systematic review, causes of straggler can be identified easily. Therefore, an effective straggler management technique can be developed to execute the jobs without straggler occurrence while fulfilling the dynamic requirements of job, which helps to increase the efficiency of large-scale cloud data centres.

5.4 Future Research Directions

Although a substantial progress has been made in straggler management techniques for large scale systems, there are still many pressing issues and challenges in this field that need to be addressed. Based on existing research, we have identified various open issues pending in this area.

5.4.1 Data Processing

Data processing in straggler management is an important challenge [9]. It happens due to the skew in data that the computing system is able to process effectively. There are two types of problems which reduce the data processing capability of systems: 1) large variation of data size and 2) non-uniformity of data. These two reasons degrade the performance of large-scale computing systems. To improve the straggler management mechanism, there should be less variation as well as less non-uniformity of data. Tackling this challenge can further improve the processing speed of computing systems in terms of execution time and latency.

5.4.2 Heterogeneity

Hardware heterogeneity is the main reason for resource contention, which occurs due different types of resources (with different configurations, different providers etc.) being used and sometimes some resources are not compatible to execute jobs in a coordinative manner. There is a need for a single interface, which can provide a stable platform for interaction of different types of hardware in a collaborative manner.

5.4.3 Latency

The latency is another important challenge in straggler management of large-scale systems, which can affect the performance of computing systems. There are different types of reasons for latency: 1) non-uniformity of data, 2) resource contention, 3) poor user code and 4) extra cloning. To improve the processing of computing systems, there is the need to make data uniform initially. Further, efficient resource scheduling algorithms are required, which can reduce resource contention at runtime as well as reduce the latency [61]. The extra cloning of tasks to speed up the execution can increase the latency because there is a requirement for more number of resources to process more number of copies. There is a need to develop an effective straggler management technique, which schedules resources and reduces latency at runtime.

5.4.4 Scalability

To improve the performance of computing systems, the systems must be more scalable to serve the jobs within their specific deadline without further delay at runtime [72]. The scalability of the computing system can increase the capacity of the system when the load increases, which can further reduce the problem of occurrence of stragglers.

5.4.5 Resource Sharing

The sharing of resources among different jobs can improve resource utilization but it leads to resource contention, which can degrade the performance of large-scale computing systems [76]. There is a need of an effective resource contention technique, which can identify the reasons of resource contention and provide the possible solutions to avoid additional resource over-allocation, ultimately contributing to straggler occurrence.

5.4.6 Energy Management

The literature reports [45] [61] [72] [102] that the straggler management techniques create several copies of the same task to mitigate the effects of stragglers. Copying a task reserves additional resources such as the disk, memory of CPU time, increasing use of particular resource. As the resource is more continuously used, its energy consumption rises. Depending on the type of the resource, its performance can degrade as its energy consumption increases above a certain threshold level.

6. SUMMARY AND CONCLUSIONS

In this paper, we have provided a comprehensive literature review of current straggler research within Computer Science, an important problem which directly debilitates the performance of large-scale computing systems. We proposed a taxonomy of straggler causes as identified from different types of straggler management techniques. Moreover, various straggler management

techniques have been reviewed and classified into two categories: straggler detection and straggler mitigation. The comparison of straggler detection and straggler mitigation have been presented in detail, and the taxonomy mapping based comparison has been described and various result outcomes related to straggler management have been presented. Observations of interest include that the focused nature of straggler causes, and mitigation solutions may potential interfere with each other due to correlated root-causes. Hence, there is a possibility of designing a multi-purpose straggler management technique which profiles and acts based on the type of identified straggler.

ACKNOWLEDGEMENTS

This work is supported by the Engineering and Physical Sciences Research Council (EPSRC) (EP/P031617/1). We would like to thank Damian Borowiec and Shreshth Tuli for their valuable comments, useful suggestions and discussion to improve the quality of the paper. We would like to thank the Editor-in-Chief and anonymous reviewers for their valuable suggestions to help and improve paper.

REFERENCES

- [1] Coppa, Emilio, and Irene Finocchi. "On data skewness, stragglers, and MapReduce progress indicators." In Proceedings of the Sixth ACM Symposium on Cloud Computing, pp. 139-152. ACM, 2015.
- [2] Farhat, Farshid. "Stochastic modeling and optimization of stragglers in mapreduce framework." (2015). Thesis, The Pennsylvania State University.
- [3] Zaharia, Matei, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael J. Franklin, Scott Shenker, and Ion Stoica. "Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing." In Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation, pp. 2-2. USENIX Association, 2012.
- [4] Ananthanarayanan, Ganesh, Ali Ghodsi, Scott Shenker, and Ion Stoica. "Effective Straggler Mitigation: Attack of the Clones." In NSDI, vol. 13, pp. 185-198. 2013.
- [5] Ananthanarayanan, Ganesh, Michael Chien-Chun Hung, Xiaoqi Ren, Ion Stoica, Adam Wierman, and Minlan Yu. "GRASS: trimming stragglers in approximation analytics." (2014): 289-302.
- [6] Wang, Da, Gauri Joshi, and Gregory Wornell. "Using straggler replication to reduce latency in large-scale parallel computing." ACM SIGMETRICS Performance Evaluation Review 43, no. 3 (2015): 7-11.
- [7] Aktas, Mehmet Fatih, Pei Peng, and Emina Soljanin. "Effective Straggler Mitigation: Which Clones Should Attack and When?." ACM SIGMETRICS Performance Evaluation Review 45, no. 2 (2017): 12-14.
- [8] Harlap, Aaron, Henggang Cui, Wei Dai, Jinliang Wei, Gregory R. Ganger, Phillip B. Gibbons, Garth A. Gibson, and Eric P. Xing. "Addressing the straggler problem for iterative convergent parallel ML." In Proceedings of the Seventh ACM Symposium on Cloud Computing, pp. 98-111. ACM, 2016.
- [9] Dean, Jeffrey, and Sanjay Ghemawat. "MapReduce: simplified data processing on large clusters." Communications of the ACM 51, no. 1 (2008): 107-113.
- [10] Aktas, M. F., Peng, P., & Soljanin, E. (2018). Straggler mitigation by delayed relaunch of tasks. ACM SIGMETRICS Performance Evaluation Review, 45(3), 224-231.
- [11] Wang, Da, Gauri Joshi, and Gregory Wornell. "Efficient task replication for fast response times in parallel computation." In ACM SIGMETRICS Performance Evaluation Review, vol. 42, no. 1, pp. 599-600. ACM, 2014.
- [12] Ouyang, Xue, Peter Garraghan, Renyu Yang, Paul Townend, and Jie Xu. "Reducing late-timing failure at scale: Straggler root-cause analysis in cloud datacenters." In Fast Abstracts in the 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks. DSN, 2016.
- [13] Vulimiri, Ashish, Philip Brighten Godfrey, Radhika Mittal, Justine Sherry, Sylvia Ratnasamy, and Scott Shenker. "Low latency via redundancy." In Proceedings of the ninth ACM conference on Emerging networking experiments and technologies, pp. 283-294. ACM, 2013.
- [14] Ouyang, Xue, Peter Garraghan, David McKee, Paul Townend, and Jie Xu. "Straggler detection in parallel computing systems through dynamic threshold calculation." In Advanced Information Networking and Applications (AINA), 2016 IEEE 30th International Conference on, pp. 414-421. IEEE, 2016.
- [15] Phan, Tien-Dat. "Energy-efficient Straggler Mitigation for Big Data Applications on the Clouds." PhD diss., ENS Rennes, 2017.
- [16] Yadwadkar, Neeraja J., Ganesh Ananthanarayanan, and Randy Katz. "Wrangler: Predictable and faster jobs using fewer resources." In Proceedings of the ACM Symposium on Cloud Computing, pp. 1-14. ACM, 2014.
- [17] Wang, Da, Gauri Joshi, and Gregory Wornell. "Efficient Straggler Replication in Large-scale Parallel Computing." arXiv preprint arXiv:1503.03128 (2015).

- [18] Chen, Quan, Daqiang Zhang, Minyi Guo, Qianni Deng, and Song Guo. "Samr: A self-adaptive mapreduce scheduling algorithm in heterogeneous environment." In Computer and Information Technology (CIT), 2010 IEEE 10th International Conference on, pp. 2736-2743. IEEE, 2010.
- [19] Lei, Lei, Tianyu Wo, and Chunming Hu. "CREST: Towards fast speculation of straggler tasks in MapReduce." In e-Business Engineering (ICEBE), 2011 IEEE 8th International Conference on, pp. 311-316. IEEE, 2011.
- [20] Nanduri, Radheshyam, Nitesh Maheshwari, A. Reddyraja, and Vasudeva Varma. "Job aware scheduling algorithm for mapreduce framework." In 2011 Third IEEE International Conference on Cloud Computing Technology and Science, pp. 724-729. IEEE, 2011.
- [21] Zaharia, Matei, Andy Konwinski, Anthony D. Joseph, Randy H. Katz, and Ion Stoica. "Improving MapReduce performance in heterogeneous environments." In *Osd*, vol. 8, no. 4, p. 7. 2008.
- [22] Ananthanarayanan, Ganesh, Srikanth Kandula, Albert G. Greenberg, Ion Stoica, Yi Lu, Bikas Saha, and Edward Harris. "Reining in the Outliers in Map-Reduce Clusters using Mantri." In *Osd*, vol. 10, no. 1, p. 24. 2010.
- [23] Guo, Yanfei, Jia Rao, Changjun Jiang, and Xiaobo Zhou. "Moving Hadoop into the cloud with flexible slot management and speculative execution." *IEEE Transactions on Parallel & Distributed Systems* 3 (2017): 798-812.
- [24] Ouyang, Xue, Peter Garraghan, Changjian Wang, Paul Townend, and Jie Xu. "An approach for modeling and ranking node-level stragglers in cloud datacenters." In *Services Computing (SCC)*, 2016 IEEE International Conference on, pp. 673-680. IEEE, 2016.
- [25] Garraghan, Peter, Xue Ouyang, Renyu Yang, David McKee, and Jie Xu. "Straggler root-cause and impact analysis for massive-scale virtualized cloud datacenters." *IEEE Transactions on Services Computing* (2016).
- [26] Tavakoli, Neda, Dong Dai, and Yong Chen. "Log-assisted straggler-aware I/O scheduler for high-end computing." In *Parallel Processing Workshops (ICPPW)*, 2016 45th International Conference on, pp. 181-189. IEEE, 2016.
- [27] Li, Cong, Huanxing Shen, and Tai Huang. "Learning to diagnose stragglers in distributed computing." In *Many-Task Computing on Clouds, Grids, and Supercomputers (MTAGS)*, 2016 9th Workshop on, pp. 1-6. IEEE, 2016.
- [28] Dai, Wei, Ibrahim Ibrahim, and Mostafa Bassiouni. "An Improved Straggler Identification Scheme for Data-Intensive Computing on Cloud Platforms." In *Cyber Security and Cloud Computing (CSCloud)*, 2017 IEEE 4th International Conference on, pp. 211-216. IEEE, 2017.
- [29] Yang, Heecheol, and Jungwoo Lee. "Secure Distributed Computing With Stragglers Using Polynomial Codes." *IEEE Transactions on Information Forensics and Security* 14, no. 1 (2019): 141-150.
- [30] Ouyang, Xue, Peter Garraghan, Bernhard Primas, David McKee, Paul Townend, and Jie Xu. "Adaptive speculation for efficient internetware application execution in clouds." *ACM Transactions on Internet Technology (TOIT)* 18, no. 2 (2018): 15.
- [31] Bin Khunayn, Eman, Shanika Karunasekera, Hairuo Xie, and Kotagiri Ramamohanarao. "Exploiting Data Dependency to Mitigate Stragglers in Distributed Spatial Simulation." In *Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, p. 43. ACM, 2017.
- [32] Farhat, F., D. Zad Tootaghaj, A. Sivasubramaniam, M. T. Kandemir, and C. R. Das. Modeling and optimization of stragglers mappers. Technical report, Technical Report CSE-14-006, Pennsylvania State University, 2014.
- [33] Phan, Tien-Dat, Shadi Ibrahim, Amelie Chi Zhou, Guillaume Aupy, and Gabriel Antoniu. "Energy-Driven Straggler Mitigation in MapReduce." In *European Conference on Parallel Processing*, pp. 385-398. Springer, Cham, 2017.
- [34] Kumar, Umesh, and Jitendar Kumar. A comprehensive review of straggler handling algorithms for mapreduce framework. *International Journal of Grid and Distributed Computing* 7, no. 4 (2014): 139-148.
- [35] Ashwin Bhandare et al. Review and Analysis of Straggler Handling Techniques, *International Journal of Computer Science and Information Technologies*, Vol. 7, no. 5, 2016.
- [36] Yan, Rong, Marc-Olivier Fleury, Michele Merler, Apostol Natsev, and John R. Smith. "Large-scale multimedia semantic concept modeling using robust subspace bagging and MapReduce." In *Proceedings of the First ACM workshop on Large-scale multimedia retrieval and mining*, pp. 35-42. ACM, 2009.
- [37] Wang, Haozhao, Song Guo, Bin Tang, Ruixuan Li, and Chengjie Li. "Heterogeneity-aware Gradient Coding for Straggler Tolerance." *arXiv preprint arXiv:1901.09339* (2019).
- [38] Yu, Qian, Mohammad Ali Maddah-Ali, and A. Salman Avestimehr. "Straggler mitigation in distributed matrix multiplication: Fundamental limits and optimal coding." *IEEE Transactions on Information Theory* 66, no. 3 (2020): 1920-1933.
- [39] Garraghan, Peter, Renyu Yang, Zhenyu Wen, Alexander Romanovsky, Jie Xu, Rajkumar Buyya, and Rajiv Ranjan. "Emergent Failures: Rethinking Cloud Reliability at Scale." *IEEE Cloud Computing* 5, no. 5 (2018): 12-21.

- [40] Zaharia, Matei, Reynold S. Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng et al. "Apache spark: a unified engine for big data processing." *Communications of the ACM* 59, no. 11 (2016): 56-65.
- [41] Isard, Michael, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. "Dryad: distributed data-parallel programs from sequential building blocks." In *ACM SIGOPS operating systems review*, vol. 41, no. 3, pp. 59-72. ACM, 2007.
- [42] Dean, Jeffrey, and Luiz André Barroso. "The tail at scale." *Communications of the ACM* 56, no. 2 (2013): 74-80.
- [43] Eppstein, David, and Michael T. Goodrich. "Space-efficient straggler identification in round-trip data streams via newton's identities and invertible bloom filters." In *Workshop on Algorithms and Data Structures*, pp. 637-648. Springer, Berlin, Heidelberg, 2007.
- [44] Lindsay, Dominic, Sukhpal Singh Gill, and Peter Garraghan. "PRISM: An Experiment Framework for Straggler Analytics in Containerized Clusters." In *Proceedings of the 5th International Workshop on Container Technologies and Container Clouds*, pp. 13-18. 2019.
- [45] Tavakoli, Neda, Dong Dai, and Yong Chen. "Client-side straggler-aware I/O scheduler for object-based parallel file systems." *Parallel Computing* 82 (2019): 3-18.
- [46] Ozfaturay, Emre, Deniz Gunduz, and Sennur Ulukus. "Speeding up distributed gradient descent by utilizing non-persistent stragglers." *arXiv preprint arXiv:1808.02240* (2018).
- [47] Eppstein, David, and Michael T. Goodrich. "Straggler identification in round-trip data streams via Newton's identities and invertible Bloom filters." *IEEE Transactions on Knowledge and Data Engineering* 23, no. 2 (2011): 297-306.
- [48] Yu, Ze, Min Li, Xin Yang, Han Zhao, and Xiaolin Li. "Taming non-local stragglers using efficient prefetching in MapReduce." In *2015 IEEE International Conference on Cluster Computing*, pp. 52-61. IEEE, 2015.
- [49] Shen, Huanxing, and Cong Li. "Zeno: A Straggler Diagnosis System for Distributed Computing Using Machine Learning." In *International Conference on High Performance Computing*, pp. 144-162. Springer, Cham, 2018.
- [50] Sukhpal Singh and Inderveer Chana, "QoS-aware Autonomic Resource Management in Cloud Computing: A Systematic Review", *ACM Computing Surveys*, vol. 48, no. 3, 46 pages, 2016
- [51] Aaron Harlap, Henggang Cui, Wei Dai, Jinliang Wei, Gregory R. Ganger, Phillip B. Gibbons, Garth A. Gibson, and Eric P. Xing. 2016. Addressing the straggler problem for iterative convergent parallel ML. In *Proceedings of the Seventh ACM Symposium on Cloud Computing (SoCC '16)*. Association for Computing Machinery, New York, NY, USA, 98–111. DOI:<https://doi.org/10.1145/2987550.2987554>.
- [52] Ouyang, Xue, Huan Zhou, Stephen Clement, Paul Townend, and Jie Xu. "Mitigate data skew caused stragglers through ImKP partition in MapReduce." In *2017 IEEE 36th International Performance Computing and Communications Conference (IPCCC)*, pp. 1-8. IEEE, 2017.
- [53] Martha, Venkata Swamy, Weizhong Zhao, and Xiaowei Xu. "h-MapReduce: a framework for workload balancing in MapReduce." In *2013 IEEE 27th International Conference on Advanced Information Networking and Applications (AINA)*, pp. 637-644. IEEE, 2013.
- [54] Huang, Sheng-Wei, Tzu-Chi Huang, Syue-Ru Lyu, Ce-Kuen Shieh, and Yi-Sheng Chou. "Improving speculative execution performance with coworker for cloud computing." In *2011 IEEE 17th International Conference on Parallel and Distributed Systems*, pp. 1004-1009. IEEE, 2011.
- [55] Lin, Jimmy. "The curse of zipf and limits to parallelization: A look at the stragglers problem in mapreduce." In *7th Workshop on Large-Scale Distributed Systems for Information Retrieval*, vol. 1, pp. 57-62. Boston, MA, USA: ACM, 2009.
- [56] Zhou, Amelie Chi, Tien-Dat Phan, Shadi Ibrahim, and Bingsheng He. "Energy-Efficient Speculative Execution using Advanced Reservation for Heterogeneous Clusters." In *Proceedings of the 47th International Conference on Parallel Processing*, p. 8. ACM, 2018.
- [57] Li, Songze, Seyed Mohammadreza Mousavi Kalan, A. Salman Avestimehr, and Mahdi Soltanolkotabi. "Near-optimal straggler mitigation for distributed gradient methods." In *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pp. 857-866. IEEE, 2018.
- [58] Tandon, Rashish, Qi Lei, Alexandros G. Dimakis, and Nikos Karampatziakis. "Gradient coding: Avoiding stragglers in distributed learning." In *International Conference on Machine Learning*, pp. 3368-3376. 2017.
- [59] Mallick, Ankur, Malhar Chaudhari, and Gauri Joshi. "Rateless Codes for Straggler Mitigation in Distributed Computing." Online Available: https://www.andrew.cmu.edu/user/gaurij/18-847F-Lectures/rateless_codes_2018.pdf
- [60] Cipar, James, Qirong Ho, Jin Kyu Kim, Seunghak Lee, Gregory R. Ganger, Garth Gibson, Kimberly Keeton, and Eric Xing. "Solving the straggler problem with bounded staleness." In *Presented as part of the 14th Workshop on Hot Topics in Operating Systems*. 2013.
- [61] Sukhpal Singh and Inderveer Chana, "A Survey on Resource Scheduling in Cloud Computing: Issues and Challenges", *Journal of Grid Computing*, vol. 14, no. 2, pp. 217-264, 2016

- [62] Yu, Qian, Mohammad Ali Maddah-Ali, and A. Salman Avestimehr. "Straggler mitigation in distributed matrix multiplication: Fundamental limits and optimal coding." In 2018 IEEE International Symposium on Information Theory (ISIT), pp. 2022-2026. IEEE, 2018.
- [63] Baharav, Tavor, Kangwook Lee, Orhan Ocal, and Kannan Ramchandran. "Straggler-proofing massive-scale distributed matrix multiplication with d-dimensional product codes." In 2018 IEEE International Symposium on Information Theory (ISIT), pp. 1993-1997. IEEE, 2018.
- [64] Xu, Maotong, Sultan Alamro, Tian Lan, and Suresh Subramaniam. "Optimizing speculative execution of deadline-sensitive jobs in cloud." In ACM SIGMETRICS Performance Evaluation Review, vol. 45, no. 1, pp. 17-18. ACM, 2017.
- [65] Xu, Maotong, Sultan Alamro, Tian Lan, and Suresh Subramaniam. "Optimizing speculative execution of deadline-sensitive jobs in cloud." In ACM SIGMETRICS Performance Evaluation Review, vol. 45, no. 1, pp. 17-18. ACM, 2017.
- [66] Haddadpour, Farzin, Yaoqing Yang, Malhar Chaudhari, Viveck R. Cadambe, and Pulkit Grover. "Straggler-resilient and communication-efficient distributed iterative linear solver." arXiv preprint arXiv:1806.06140 (2018).
- [67] Khunayn, Eman Bin, Shanika Karunasekera, Hairuo Xie, and Kotagiri Ramamohanarao. "Straggler mitigation for distributed behavioral simulation." In 2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS), pp. 2638-2641. IEEE, 2017.
- [68] Paik, Michael. "Stragglers of the herd get eaten: security concerns for GSM mobile banking applications." In Proceedings of the Eleventh Workshop on Mobile Computing Systems & Applications, pp. 54-59. ACM, 2010.
- [69] Malewicz, Grzegorz, Marian Dvorsky, Christopher B. Colohan, Derek P. Thomson, and Joshua Louis Levenberg. "System and method for limiting the impact of stragglers in large-scale parallel data processing." U.S. Patent 8,510,538, issued August 13, 2013.
- [70] Karakus, Can, Yifan Sun, Suhas Diggavi, and Wotao Yin. "Redundancy techniques for straggler mitigation in distributed optimization and learning." arXiv preprint arXiv:1803.05397 (2018).
- [71] Ouyang, Xue, Changjian Wang, Renyu Yang, Guogui Yang, Paul Townend, and Jie Xu. "ML-NA: A machine learning based node performance analyzer utilizing straggler statistics." In 2017 IEEE 23rd International Conference on Parallel and Distributed Systems (ICPADS), pp. 73-80. IEEE, 2017.
- [72] Zheng, Pengfei, and Benjamin C. Lee. "Hound: Causal learning for datacenter-scale straggler diagnosis." Proceedings of the ACM on Measurement and Analysis of Computing Systems 2, no. 1 (2018): 17.
- [73] Zhao, Xia, Kai Kang, YuZhong Sun, Yin Song, Minhao Xu, and Tao Pan. "Insight and reduction of MapReduce stragglers in heterogeneous environment." In 2013 IEEE International Conference on Cluster Computing (CLUSTER), pp. 1-8. IEEE, 2013.
- [74] Behrouzi-Far, Amir, and Emina Soljanin. "On the Effect of Task-to-Worker Assignment in Distributed Computing Systems with Stragglers." In 2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 560-566. IEEE, 2018.
- [75] Chen, Fei, Song Wu, Hai Jin, Yin Yao, Zhiyi Liu, Lin Gu, and Yongluan Zhou. "Lever: Towards low-latency batched stream processing by pre-scheduling." In Proceedings of the 2017 Symposium on Cloud Computing, pp. 643-643. ACM, 2017.
- [76] Singh, Sukhpal, and Inderveer Chana. "Cloud resource provisioning: survey, status and future research directions." Knowledge and Information Systems 49, no. 3 (2016): 1005-1069.
- [77] Misra, Pulkit A., María F. Borge, Íñigo Goiri, Alvin R. Lebeck, Willy Zwaenepoel, and Ricardo Bianchini. "Managing Tail Latency in Datacenter-Scale File Systems Under Production Constraints." In Proceedings of the Fourteenth EuroSys Conference 2019, p. 17. ACM, 2019.
- [78] Qureshi, Nawab Muhammad Faseeh, Isma Farah Siddiqui, Asad Abbas, Ali Kashif Bashir, Keehyun Choi, Jaehyun Kim, and Dong Ryeol Shin. "Dynamic Container-based Resource Management Framework of Spark Ecosystem." In 2019 21st International Conference on Advanced Communication Technology (ICACT), pp. 522-526. IEEE, 2019.
- [79] Yang, Eunju, Dong-Ki Kang, and Chan-Hyun Youn. "BOA: batch orchestration algorithm for straggler mitigation of distributed DL training in heterogeneous GPU cluster." The Journal of Supercomputing (2019): 1-21.
- [80] Gill, Sukhpal Singh, Inderveer Chana, Maninder Singh, and Rajkumar Buyya. "RADAR: Self-configuring and self-healing in resource management for enhancing quality of cloud services." Concurrency and Computation: Practice and Experience 31, no. 1 (2019).
- [81] Kapoor, Rishi, George Porter, Malveeka Tewari, Geoffrey M. Voelker, and Amin Vahdat. "Chronos: Predictable low latency for data center applications." In Proceedings of the Third ACM Symposium on Cloud Computing, p. 9. ACM, 2012.
- [82] Sukhpal Singh Gill, Peter Garraghan, Vlado Stankovski, Giuliano Casale, Soumya K. Ghosh, Ruppa K. Thulasiram, Kotagiri Ramamohanarao and Rajkumar Buyya, "Holistic Resource Management for Sustainable and Reliable Cloud Computing: An Innovative Solution to Global Challenge", Journal of Systems and Software, Elsevier, Volume 155, Pages: 104-129, 2019.

- [83] Lama, Palden, Shaoqi Wang, Xiaobo Zhou, and Dazhao Cheng. "Performance isolation of data-intensive scale-out applications in a multi-tenant cloud." In 2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS), pp. 85-94. IEEE, 2018.
- [84] Zhou, Honggang, Yunchun Li, Hailong Yang, Jie Jia, and Wei Li. "BigRoots: An effective approach for root-cause analysis of stragglers in big data system." IEEE Access 6 (2018): 41966-41977.
- [85] Gill, Sukhpal Singh, and Rajkumar Buyya. "A taxonomy and future directions for sustainable cloud computing: 360 degree view." ACM Computing Surveys (CSUR) 51, no. 5 (2018): 104.
- [86] Mitsuzuka, Koya, Michihiro Koibuchi, Hideharu Amano, and Hiroki Matsutani. "Proxy responses by FPGA-based switch for MapReduce stragglers." IEICE TRANSACTIONS on Information and Systems 101, no. 9 (2018): 2258-2268.
- [87] Ouyang, Xue, Changjian Wang, and Jie Xu. "Mitigating stragglers to avoid QoS violation for time-critical applications through dynamic server blacklisting." Future Generation Computer Systems 101 (2019): 831-842.
- [88] Phan, Tien-Dat, Guillaume Pallez, Shadi Ibrahim, and Padma Raghavan. "A new framework for evaluating straggler detection mechanisms in mapreduce." ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS) 4, no. 3 (2019): 14.
- [89] Gill, Sukhpal Singh, Shreshth Tuli, Minxian Xu, Inderpreet Singh, Karan Vijay Singh, Dominic Lindsay, Shikhar Tuli et al. Transformative effects of IoT, Blockchain and artificial intelligence on cloud computing: evolution, vision, trends and open challenges. Internet of Things. 2019; 8:100118.
- [90] Krishna, Lolla Sai, and Lakshmi Prasad Natarajan. "Distributed Inference With Straggler Mitigation." PhD diss., Indian institute of technology Hyderabad, 2019.
- [91] Huang, Xiaohan, Chunlin Li, and Youlong Luo. "Optimized Speculative Execution Strategy for Different Workload Levels in Heterogeneous Spark Cluster." In Proceedings of the 2019 4th International Conference on Big Data and Computing, pp. 6-10. ACM, 2019.
- [92] Panda, Biswaranjan, Deepthi Srinivasan, Huan Ke, Karan Gupta, Vinayak Khot, and Haryadi S. Gunawi. "{IASO}: A Fail-Slow Detection and Mitigation Framework for Distributed Storage Services." In 2019 {USENIX} Annual Technical Conference ({USENIX}{ATC} 19), pp. 47-62. 2019.
- [93] Hamandawana, Prince, Ronnie Mativenga, Se Jin Kwon, and Tae-Sun Chung. "EPPADS: An Enhanced Phase-Based Performance-Aware Dynamic Scheduler for High Job Execution Performance in Large Scale Clusters." In International Conference on Database Systems for Advanced Applications, pp. 140-156. Springer, Cham, 2019.
- [94] Ren, Xiaoqi, Ganesh Ananthanarayanan, Adam Wierman, and Minlan Yu. "Hopper: Decentralized speculation-aware cluster scheduling at scale." In ACM SIGCOMM Computer Communication Review, vol. 45, no. 4, pp. 379-392. ACM, 2015.
- [95] Benavides, Zachary, Rajiv Gupta, and Xiangyu Zhang. "Parallel execution profiles." In Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing, pp. 215-218. ACM, 2016.
- [96] Wang, Zhigang, Lixin Gao, Yu Gu, Yubin Bao, and Ge Yu. "FSP: towards flexible synchronous parallel framework for expectation-maximization based algorithms on cloud." In Proceedings of the 2017 Symposium on Cloud Computing, pp. 1-14. ACM, 2017.
- [97] Isaacs, Katherine E., Todd Gamblin, Abhinav Bhatele, Peer-Timo Bremer, Martin Schulz, and Bernd Hamann. "Extracting logical structure and identifying stragglers in parallel execution traces." In ACM SIGPLAN Notices, vol. 49, no. 8, pp. 397-398. ACM, 2014.
- [98] Jiang, Jiawei, Bin Cui, Ce Zhang, and Lele Yu. "Heterogeneity-aware distributed parameter servers." In Proceedings of the 2017 ACM International Conference on Management of Data, pp. 463-478. ACM, 2017.
- [99] Patgiri, Ripon, and Rajdeep Das. "rTuner: A Performance Enhancement of MapReduce Job." In Proceedings of the 10th International Conference on Computer Modeling and Simulation, pp. 176-183. ACM, 2018.
- [100] Zaharia, Matei, Tathagata Das, Haoyuan Li, Timothy Hunter, Scott Shenker, and Ion Stoica. "Discretized streams: Fault-tolerant streaming computation at scale." In Proceedings of the twenty-fourth ACM symposium on operating systems principles, pp. 423-438. ACM, 2013.
- [101] Chen, Chen, Qizhen Weng, Wei Wang, Baochun Li, and Bo Li. "Fast Distributed Deep Learning via Worker-adaptive Batch Sizing." In Proceedings of the ACM Symposium on Cloud Computing, pp. 521-521. ACM, 2018.
- [102] Fuerst, Carlo, Stefan Schmid, Lalith Suresh, and Paolo Costa. "Kraken: Towards elastic performance guarantees in multi-tenant data centers." In ACM SIGMETRICS Performance Evaluation Review, vol. 43, no. 1, pp. 433-434. ACM, 2015.