

Reliable and Locality-driven scheduling in Hadoop

Phuong Tran Anh

Institute Superior Tecnico

Abstract. The increasing use of computing resources in our daily lives leads to data being generated at an astonishing rate. The computing industry is being repeatedly questioned for its ability to accommodate the unpredictable growth rate of data, and its ability to process them. This has encouraged the development of cluster based data-intensive applications. Hadoop is a popular open source framework known for its massive cluster based data processing power. Hadoop is widely used in the computer industry because of its scalability, reliability, ease of use, and low cost of implementation.

Cloud computing in the recent years has gained increasingly popularity by its cost-efficient and flexible way to leverage the power of commodity hardware. Hadoop based services on the Cloud have also emerged as one of the prominent choices for smaller businesses. However, evidence in the literature shows that faults on the Cloud do occur and normally result with performance problems. Hadoop hides the complexity of discovery and handling failures from the schedulers, but the expenses of failure recovery rest entirely on users, regardless of root causes. We systematically assess these expenses through a set of experiments, and argue that more effort to reduce this cost to users is desirable.

We also analyze the drawback of current Hadoops mechanism in prioritizing failed tasks. By trying to launch failed tasks as soon as possible regardless of locality, it significantly increases the execution time of jobs with failed tasks, due to two reasons: 1) available slots might not be free up as quickly as expected and 2) the slots might belong to machines with no data on it, introducing extra cost for data transferring through network, which is normally the most scarce resource in nowadays data centers.

This thesis then introduces a new algorithmic approach called the waste-free preemption. The waste-free preemption saves Hadoop scheduler from choosing solely between kill, which instantly re-leases the slots but is wasteful, and wait, which does not waste any previous effort but falls for the two above mentioned reasons. With this new option, a preemptive version of Hadoops default scheduler (FIFO) is implemented. The evaluation demonstrates the effectiveness of the new feature by comparing its performance with the traditional Hadoops mechanism.

Keywords: Distributed systems, Hadoop, scheduling, preemption

1 Introduction

Data insight forms an essential part in today's decision making process. The massive growth in the amount of data, and the need to analyze data motivated the development of data intensive applications. In this context, Hadoop MapReduce [6] is a big data processing framework that has rapidly gained popularity in both industry and academia. Many enterprises also claim to have Hadoop clusters of various sizes: while most of the enterprises have clusters of size less than 100 machines, some consist up to hundreds of machines. The main reasons of such popularity are the ease-of-use, scalability, and failover properties of Hadoop MapReduce.

The unprecedented growth in data center technologies and services in the recent years allow smaller companies to take advantages of the "cloud-based" infrastructure. Well-known Cloud providers such as Amazon Web Services, Google App Engine and Microsoft Azure respond to the need of data processing by equipping their software stack with MapReduce-like systems. Amazon Elastic MapReduce is a platform that facilitate large-scale data applications and perhaps a very successful example of scalable MapReduce on Cloud.

Regarding reliability, user of Cloud services often have to face the threat of failures. In fact, researchers interested in fault tolerance have accepted that failure is becoming a norm, rather than an exception. With respect to MapReduce-Cloud, Cloud providers rely completely on the fault-tolerance mechanisms which is provided by Hadoop. Currently, Hadoop handles failures by simply re-executing all the failed tasks. However, all these efforts to handle failures are entirely entrusted to the core of Hadoop and hidden from the task schedulers. To our knowledge, there has been no scheduler that explicitly copes with failure. This potentially leads to degradation in Hadoop's performance. In this study, we address the problem of failure in Hadoop, and present our approach to improve Hadoop performance under failure.

Contribution and structure

This thesis contributes to the field of data-intensive applications in several ways. First, it systematically assesses the Hadoop architecture, focusing on the fault tolerance mechanisms that Hadoop employs. A set of experiments were conducted to illustrate the limitations of its default mechanism. Second, it explores the possibility of improving the performance by introducing a new Preemption scheduling heuristic in Hadoop. This new option is accompanied with a new scheduler that explicitly deal with failures in Hadoop. Finally, it performs evaluation of the new scheduler and discusses other possibilities to fully utilize the new feature in different scenarios.

This thesis is organized as follows. Section 2 provides background into the Hadoop framework, its fault tolerance mechanism and the effect of failure on Hadoop's performance. Section 3 discusses recent works that share the same similar objectives in improving the fault tolerance mechanism in Hadoop. Chapter 4 and 5 open new possibilities to improve Hadoop performance by introducing the work-conserving preemption mechanism with its implementation details. Chapter 6 analyzes the performance of the new feature in comparison with the default scheduler in terms of execution time and data locality. Finally, chapter 7 concludes the thesis with some ideas on how to further improve Hadoop's performance with the introduced feature.

2 Background

2.1 Hadoop framework

Hadoop is an open-source implementations of the MapReduce programming paradigm which was designed to process and generate large data sets. The MapReduce abstractions allows users to express their computation as a series of Map and Reduce functions. Map function processes a series of $\langle key, value \rangle$ pairs to generate a set of intermediate $\langle key, value \rangle$ pairs, and Reduce function aggregates all intermediate values that associate to the same intermediate key to produce the final output, also in the form of $\langle key, value \rangle$ pairs.

Each Hadoop cluster contains one Job Tracker. Job Tracker is responsible for (a) querying the underlying HDFS for the block locations, (b) scheduling the tasks on the slave which is hosting the tasks blocks, and (c) monitoring the successes and failures of the tasks. Equivalent to the slave workers are Task Trackers, who executes the tasks as directed by the Job Tracker. Each Task Tracker is set up with certain numbers of Map and Reduce slots respectively i.e. a Task Tracker cannot have more than these numbers of Map or Reduce tasks running simultaneously. Each Map (or Reduce) task is a separate program that contains the user-defined map (or reduce) function and can be missing (Map-only jobs).

2.2 Fault tolerance in Hadoop

Detection of failures Hadoop employs a static timeout mechanism for the detection of fail-stop failures. It keeps track of each Task Tracker's last heartbeat, so that should a Task Tracker has not sent any heart beat in a certain amount of time, that Task Tracker will be declared Failed. Each Task Tracker sends a heart beat every 0.3s (many literatures have claimed that the heart beat interval is 3 seconds, however here we use the value we found in the source code). The Job Tracker checks every 200s for any Task Tracker that has been silent for 600s. Once found, the Task Tracker is labeled as a failed machine, and the Job Tracker will trigger the failure handling and recovery process.

Failure handling and recovery Tasks that were running on the failed Task Tracker will be restarted on other machines. Map tasks that completed on this Task Tracker will also be restarted if they belong to jobs that are still in progress and contain some reduce tasks. Completed Reduce Tasks are not restarted, as the output is supposedly stored persistently on HDFS.

Hadoop prioritizes failed tasks to run first over pending (non-running) tasks when it comes to assigning new tasks. The aim of this decision is to quickly discover any jobs' internal failures'. Hadoop jobs' internal failures' are failures that are specific to a certain job and cannot be tolerated by re-execution mechanism. One example of internal failures is corrupted input files. While each task consumes a certain amount of resources (disk space, computational cycles, memory...), faster discovery of internal failures allows Hadoop to quickly purge those failed jobs and release the resources for other waiting jobs, hence achieve better utilization of resources.

Although Hadoop pays much effort to achieve locality for Map tasks in normal situation (it tries to assign as many local tasks as possible, while only assign at most 1 non-local tasks regardless of the number of available slot at each time), it completely ignores locality when it comes to failed tasks. As long as there are failed (Map) tasks, any Task Tracker that has free slots will be assigned the maximum number of tasks that it can handle. This leads to degradation in the performance of Hadoop when there are many failed tasks, as the number of non-local tasks might become very high.

2.3 Failures effect on Hadoop execution




Fig. 1: Performance of Hadoop in 3 scenarios: Normal, Mix Stress and Failure

Figure 1a presents the total execution time in 3 different scenarios: Normal, Mix Stress and Failure in a cluster of 21 nodes with 6 Hadoop jobs of different sizes. In Normal scenario, we allow Hadoop run on dedicated nodes without introducing any stress or failure. The Mix Stress scenario has some CPU-intensive and IO-intensive processes launched on one node to simulate sporadic overloading situations. Failure scenario has 1 node fail after arbitrarily 80 seconds (by simply killing the Task Tracker process running on that node). As we can see from Figure 1a, stress slightly degrades Hadoop’s performance a little, while failures prolong the execution of Hadoop jobs by a significant amount of time as much as 56 seconds (roughly 22.5%) in case of Capacity scheduler. Fair scheduler appears to suffer the least: its execution time is prolonged for only 29 seconds (10.4%), and it also finishes the fastest among the three schedulers under failure (278 seconds compared to 294 seconds for FIFO, and 305 seconds for Capacity scheduler).

Figure 1b shows the percentage of locally executed tasks over the total number of tasks in the same scenarios. Fair scheduler still enjoys the highest number of locality, even though this number is decreasing. FIFO and Capacity scheduler shows some degradation, though this degradation is rather small compared to Fair scheduler (3% and 1% compared to 9%). To explain this phenomenon, remember that Fair scheduler was designed based on the assumption that most tasks are short and therefore, nodes will release slot quickly enough for other tasks get locally executed. However, in case of failure, the long failure detection time (expiry time) creates the illusion of long-lasting tasks on failed node. These ”fake” long tasks break the assumption of Fair scheduler, leading to high degradation.

3 Related works

There exists a problem regarding the failure handling mechanism of Hadoop that often gets overseen. When a task is declared failed, it gets ”special treatment” in the manner that, failed tasks will be launched as soon as any slot becomes available, regardless of data locality. In a cluster where Data Node and Task Tracker processes co-reside, a machines failure will reduce the replication factor for those data splits originally on that node. Providing that Hadoop tries its best to provide locality for tasks in normal

situation, it is likely that the failed tasks will have one less machine to run locally, which in turn leads to lower locality in general.

Providing locality for tasks is crucial for performance of Hadoop in large clusters because network bisection bandwidth becomes a bottleneck [1]. Besides, since most of the Hadoop usage is for small jobs (jobs with small number of map tasks), it is difficult for a small job to obtain slots on nodes with local data. Data then has to be transferred through the network, which might significantly increase the execution time if network bandwidth is scarce. Providing locality for these jobs will greatly increase the performance of Hadoop in term of time and resource preserving.

Unfortunately, achieving high locality is not easy. Zaharia et al [2] introduces the Delay technique inside the Fair scheduler to improve locality of tasks. Instead of strictly following the order of jobs, Fair scheduler allows behind jobs launch their tasks first if the head-of-line job fails to launch a local task. However, Fair scheduler relies on the assumption that tasks are mostly short and slots are freed up quickly. In case of long tasks that occupy the slots, a node may not free up quickly enough for other jobs to achieve locality.

In the effort to overcome the above-mentioned problems, we propose preemption. Preemption allows a task to quickly release the slot for more urgent task. Locality can be assured with the employment of preemption. Also preemption allows the scheduler to have better control on the resources (i.e. the task slots) so that optimal efficiency can be obtained. Shorter tasks can preempt a longer one to achieve fast response time.

Preemption in Hadoop

To our knowledge, there has been not much work aiming at providing the preemption feature for Hadoop:

Li Liu [3] introduces a Preemptive Deadline Constraint Scheduler (PDCS), which aims at minimizing the total completion time of jobs under deadlines. Traditional non-preemptive schedulers have to wait for previously assigned jobs completion or halt. This delays the execution of production jobs, sometimes renders them violating their deadlines. To avoid this, PDCS employs the Hadoop built-in preemption mechanism (kill) to provide slots for near deadline jobs. Upon submission, jobs are checked whether they can finish under its deadline or not using estimation. Jobs then are scheduled if the number of available slots meets the requirement. Otherwise, the scheduler would determine whether these jobs are legal to preempt the slots that have been already allocated.

Yandong Wang [4] introduces the Fair Completion Scheduler (FCS) that supports Reduce Task preemption. A long Reduce task would occupy the reduce slot, and significantly increase the completion time of shorter jobs. By check pointing the Reduce task, the reduce slot can be passed on to a different shorter job. After the short job finishes, the long Reduce task picks up the work from where it was left off, and continue until the end.

Pastorelli [5] proposes to leverage the already available POSIX signals such as SIGTSTP and SIGCONT to suspend running tasks. In Hadoop, Map and Reduce tasks are regular Unix processes running in child JVMs spawned by the TaskTracker. This means that they can safely be handled with the POSIX signaling infrastructure. The state of tasks is implicitly saved by the operating system, and kept in memory. If not enough physical memory is available for running tasks at any moment, the OS paging mechanisms saves the memory allocated to the suspended tasks in the swap area.

Pastorelli's approach save the state of the JVM and can be applied seamlessly to arbitrary tasks regardless of types. However a suspended process can only be resumed on the same machine it was suspended on. If the same task gets scheduled on a different machine, it has to be restarted from scratch, losing work done so far: in that case, the `suspend` is effectively analogous to a delayed `kill`.

Although very interesting, the above mentioned efforts all suffer from some drawbacks. The Preemptive Deadline Constraint Scheduler approach employs the naive Kill primitive from Hadoop, which incurs a large amount of waste work. The Fair Completion Scheduler only concerns about preempting Reduce tasks but ignores the case of Map tasks. Pastorelli's OS-assisted preemptive primitives allow a seamless preemption mechanism for all types of tasks, but do not support migration, i.e. restarting the preempted tasks on a different node. These drawbacks limit the usage of these approaches. In the next chapter, we will discuss more about the choices between preemption styles, and our approach to overcome these drawbacks.

4 Pause and Resume mechanism

Preemption is a highly desirable feature in many cases. It allows schedulers to quickly reallocate resources between jobs for fair sharing. In other cases, preemption can reclaim resources from long running tasks

and give them to shorter one. Providing local execution to short tasks helps reducing response time as well as lowering the network traffic between nodes in the same cluster. In this study, we introduce our preemption mechanism that is both fast and light-weight.

4.1 Map task preemption

Each Map task belonging to a Hadoop job is assigned a split of the data to process. Once all the input $\langle key, value \rangle$ pairs have been processed, the Map task reports back to the Job Tracker, and the map output is entrusted to the local Task Tracker to serve to Reduce tasks.

Although the name is pause and resume, our Map preemption mechanism is implemented in a structurally different manner. There is no pause in Map task preemption, but we actually split the map tasks into 2 sub-tasks at the boundary between input $\langle key, value \rangle$ pairs. The first sub-task covers all the map input $\langle key, value \rangle$ pair that the task has processed so far. This sub-task is considered Completed and will report back to the Job Tracker as a normally completed task. The second sub-task includes all the map input $\langle key, value \rangle$ pair that has not yet been processed. This sub-task will be treated as a new independent task with almost no difference compared to any other non-running tasks. The independence between $\langle key, value \rangle$ pairs of the Map Reduce programming paradigm guarantees that our mechanism works correctly and does not produce any extra (or lose any) $\langle key, value \rangle$ pair, thus ensuring correctness.

A Map task now does not simply loop through every single $\langle key, value \rangle$ pair like in the original source code. At every $\langle key, value \rangle$ pair, it needs to make sure the pair falls inside the tasks covering range. A simple comparison between starting key, ending key and the current key at each round serves the purpose. Besides, Map task now needs to actively listen to the signal from the Task Tracker in order to stop at any time. Finally, before committing the map output, Map task needs to mark the last processing input $\langle key, value \rangle$ pair, so the Job Tracker can know whether this task has completely finished or needs to be split into sub-tasks. Algorithm ?? illustrates the new Map tasks workflow.

The creation of sub-tasks does not modify the underlying data splits structure. Sub-tasks still require the same data split to process compared to its parental original task. This poses a drawback in our mechanism regarding the efficiency of data usage: the range of a sub-task might only cover, for example, 10% of the data, but the same amount of data (1 chunk = 64 MB by default) will be required to transfer. It is then advisable not to split the task into too many small sub-tasks, as the network and computational waste would be high.

4.2 Reduce task preemption

A Reduce task is divided into 3 different phases. The Shuffle phase fetches all the segments that belong to it from intermediate Map Output data. The Sort phase performs a sort operation on all the fetched data, which was kept in $\langle key, value \rangle$ pair format. Finally, the Reduce phase applies the user-defined reduce function on each of the Reduce-input $\langle key, value \rangle$ pairs. Since the 3 phases are heavily dependent on each other, a similar splitting approach like in Map task preemption is unviable. We instead choose the traditional pause and resume approach where, at the moment of pausing, all the data of the Reduce task is stored in the local storage of the Task Tracker and, at the moment of resuming, data is loaded back to the memory.

Our mechanism allows a Reduce task to preempt itself at any time during the course of the Shuffle phase, and at the boundary of other phases. The Sort phase is usually very short due to the fact that Hadoop launches a separate thread to merge the data as soon as they become available. Preempting Reduce task during Reduce phase is also feasible, however we will consider this in future work.

Pause During the Shuffle phase, a Reduce task fetches all the segments that belong to it from all intermediate map outputs. According to the sizes of the segments, Reduce Task stores them either to local disks or in memory. Meanwhile, multiple merging threads merge fetched data into larger segments and store them in an ordered structure that later can be popped out to feed to Reduce function. Preserving the state of Shuffle phase means to keep track of the shuffling status of all segments. Upon receiving preemption request, Reduce task stops all the fetching and merging thread gracefully: it allows the threads to finish the last unit of work they are currently on. Fetching threads can finish fetching the last segment of Map output, while Merging threads can finish merging the current segment.

After stopping all communication and sorting threads, the Reduce task flushes all the in-memory data to the disks while leave all the on-disk segments untouched. In-memory data includes in-memory segments

and other data to keep track of the progress (number of copied segments, number of sorted segments...). These data is kept in files stored in each task attempts specific folder so that later re-launched attempt can access to these files and resume operation.

Preemption at the phases boundary follows exactly the same procedure. Data is flushed to disks, and all the information needed to re-create the running task is also stored in files. After that, the Reduce task preempts itself and releases the slot. The task reports back to the Job Tracker with a new status of SUSPENDED. Suspended task will go through almost the same procedure as Failed task, except increasing the count for failed tasks. It is also added back to the pool for later resume.

Resume The information about the previously launched Task Tracker is stored inside the task. Upon re-launch, the Reduce task checks to see if it has been launched somewhere before. The location can be either local (task is re-launched on the same Task Tracker) or remote (task is re-launched on a different Task Tracker). The task then tries to fetch the already processed data from the previously launched Task Tracker before resuming to the point where it left off.

Our mechanism is incremental in the sense that it allows the task to be preempted multiple times. However, we do not allow the re-launched task to be preempted during the period when it fetches data from previously launched Task Tracker. Task needs to finish resuming to the previous state before being able to be preempted again. This is to prevent having flushed data on too many Task Trackers, which will make the resuming process complex and inefficient.

5 Preemptive locality-driven scheduler

In order to demonstrate the effectiveness of our new mechanism, we design a new scheduler, namely the Preemptive Locality-driven scheduler (PLS), that employs this new option to provide better locality for Map Reduce jobs. In normal situations without any failure, PLS behaves similarly to the default FIFO scheduler of Hadoop. However, when failures occur, PLS actively leverages the preemption mechanism to provide local execution for failed tasks.

Upon a task's failure, PLS assembles a list of Task Trackers that has local data to the failed task (if it was of Map type), or simply list of all running Task Trackers (if it was of Reduce type). If any of these task trackers has a free slot of the right task type, PLS simply allows the task to be launched from that task tracker at the next heartbeat. If no free slot is found, PLS tries to find appropriate task from all the currently running tasks on those task trackers to preempt. Figure 1 illustrates the logic of finding a task to preempt:

Algorithm 1 Finding a task to preempt

```

/*get the list of all suitable Task Trackers*/
List taskTrackers = getListOfTaskTracker();
/*get the list of all tasks running on those task trackers*/
List tasks = getListStatusesOfTasks(taskTrackers);
/*sort the tasks in the reverse job order*/
tasks.sort();
for {Task t : tasks} do
  if ( t belongs to jobs of lesser priority
  AND t.progressScore < THRESHOLD
  AND !t.isCleanUp() AND !t.isSetUp()
  AND !preemptList.contains(t) then
    preempt_task(t);
    preemptList.add(t);
  return;
  end if
end for

```

PLS only allows failed tasks to preempt tasks from later submitted jobs. To respect the FIFO order, PLS prioritizes preempting tasks that belong to the latest submitted jobs. However, not all tasks would be considered for preemption. Preempting a task at, for example, 0.99 progress score would not be advisable, as it only introduces unnecessary waste which is more costly than waiting for that task to end (and perhaps the task will be already completed by the time of the next heartbeat). We impose a

THRESHOLD of progress score so that almost complete tasks will not be considered for preemption. This THRESHOLD value is configurable via `mapred-site.xml` file, and the default value is set to 0.8.

In case PLS cannot find any task to preempt (for example, tasks all belong to the same or higher priority jobs, or all tasks are close to completion), PLS let the task to be launched at any arbitrary Task Tracker that has free slot the earliest. This is similar to the original FIFO schedulers procedure. This is to avoid having to wait for too long for a local slot.

6 Evaluation

In order to prove the efficiency of Pause and Resume preemption mechanism, we present the result from selected experiments. The goal is to demonstrate how the Preemptive Locality-driven scheduler manages to help jobs under failures. For this purpose, we concern about the execution time of jobs that suffer the failure, as well as their localities.

We performed our experiment on Grid5000 testbed, with a cluster of 21 nodes (1 Master and 20 Slave nodes). Each node has 2 Map and 2 Reduce slots. Replication factor is set to 2. For timely reason, expiry time is set at 30 seconds. Speculation is turned off for correct evaluation of Pause and Resume mechanism's efficiency.

Jobs For simplicity, we only run 2 jobs of WordCount application with the same input size of 30GB. Failure is injected at 150 seconds to maximize the number of failed tasks from job 1.

Execution time




Fig. 2: Finish time of jobs with different schedulers

Figure 2 shows the finishing time of each job in the 4 different schedulers: the three readily available schedulers of Hadoop (FIFO, Fair and Capacity scheduler) and Preemptive Locality-driven scheduler

(PLS). There is little difference in the total execution time of the two jobs (represented by the finishing time of the second job), but the trend is that PLS has to pay the most time to finish the set of jobs. However, PLS shows good improvement in the finishing time of job 1 - the job that suffers from the failure. PLS only requires 256.4 seconds to finish job 1, roughly 14 seconds (5.2%) faster than in Capacity and about 20 seconds (7.2%) when compare to FIFO scheduler. Fair scheduler performs poorly on this metrics, as jobs share the resources of the cluster. In fact, job 1 takes more time to finish than job 2 does in Fair scheduling. This can be explained by the fact that job 1 starts first and has more failed tasks than job 2.

To further understand about the efficiency of PLS, as well as the actual behavior of Pause and Resume preemption feature, we perform the same experiment with 3 different flavors of PLS. We omit completely the preemption function of PLS and mimic the FIFO scheduler in a flavor called FIFO*. FIFO* still has to pay the overhead of allowing Pause and Resume at any time (this overhead will be discussed later). The Kill-PLS uses the Kill primitive provided by Hadoop instead of Pause and Resume. Finally, the PLS that leverages the Pause and Resume function is also included.




Fig. 3: Execution with different flavors of PLS

Figure 3 compares the execution of Hadoop with different flavors of PLS. The FIFO* enjoys the best performance due to the fact that there is no killing or preemption overhead. Tasks are launched and finish naturally without any waste work (from killing) or delay (from splitting and re-launching a task).

Though witnessing some degradation in the total performance, Kill-PLS and PLS gain some improvement in term of Job 1's execution time. Both of Kill and Preemption finish the first job faster than that of FIFO*, in the order of roughly 19 (11.2%) and 7 (2.5%) seconds, respectively. Kill-PLS outperforms the other 2 competitors thanks to the fact that kill command can instantly rescue the slots from currently running tasks, while preemption command needs to pay some delay so that running tasks can save the states and release the slots. However, PLS manages to save some work from preempted tasks. PLS

only requires an extra 102 seconds to finish the second job, compared to 120 seconds in Kill-PLS (an improvement of 15%).

Data locality of the first job

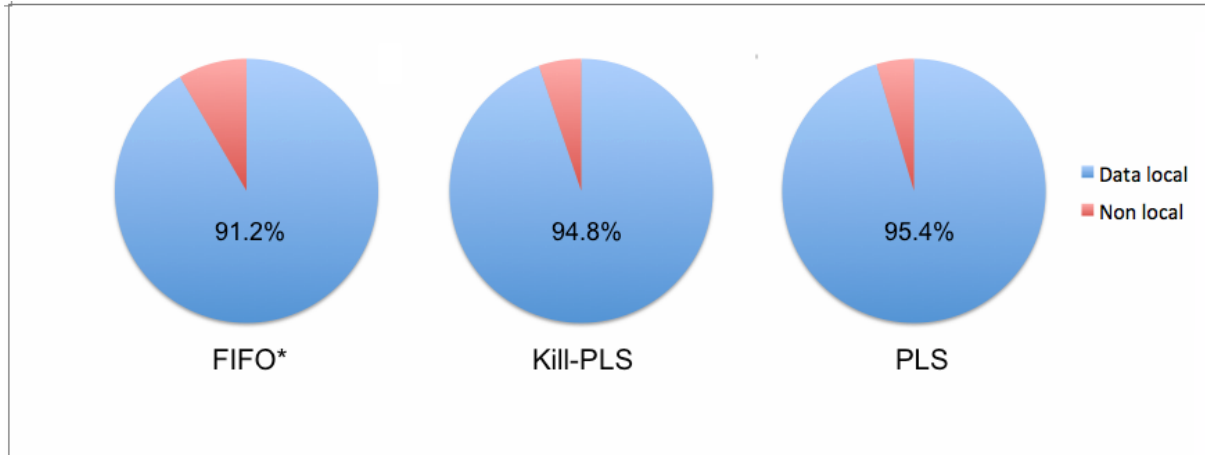


Fig. 4: Locality of the first job

Figure 4 shows the locality of the first jobs in the three different flavors. FIFO* has the lowest value of data locality (90.4%) due to the fact that failed tasks are assigned regardless of locality. Kill-PLS and PLS observe higher locality in Map tasks and equal each other with the values of 94.8% and 95.4%, respectively. This means PLS manages to improve the locality of Map tasks by almost 5%. Out of 12 failed tasks (tasks that were initially launched on the failed node), FIFO* assigns all 12 tasks remotely (rack-locally), while Kill-PLS and PLS assigns only 3 remotely (as in table 1).

FIFO*	Kill-PLS	PLS
0/12 (0%)	9/12 (75%)	9/12 (75%)

Table 1: Number of local re-executed tasks

7 Conclusion

This thesis addresses the problem by investigating the fault-tolerance mechanism of Hadoop, a popular implementation of MapReduce. It presents some results to illustrate Hadoop's behaviors under different situations. The intent was to confirm the mechanism and analyze the drawbacks of how Hadoop handles failures.

We proposed the designed of a new feature for Hadoop: the waste-free preemption function. By allowing a task to preserve its state and release the slot in a timely manner, Hadoop can have more control over resources. This in turn will help increase the performance of Hadoop under the occurrence of failure. The preemption feature was implemented not only with the intention of improving performance under failure, but also to open new possibilities for further improvement under other circumstances.

Finally, we evaluated the effectiveness of the new feature considering some basics metrics such as execution time and data locality. In this work, we compare the original Fifo scheduler with a preemptive version. Experiments show that our new feature improves the execution time of Hadoop job when failure occurs. Although the preemption mechanism imposes some overhead, if wisely used, it can greatly improve Hadoop's performance.

Future work

Our work suffers from some concerns regarding the usability, as well as efficiency. Firstly, the Reduce Pause and Resume mechanism should be further extended to allow preemption during Reduce phase. This allows a more fine-grain control over Reduce tasks. Secondly, it would be interesting to evaluate the cost of preempting a Reduce task, in comparison with Killing. Killing instantly releases the slot, but wastes the effort, while Preemption may take sometime before the slot is released. Also, upon resuming, Reduce tasks also require some extra time to load up the data from the local storage. These two delays add up and in some cases, it is more beneficial to decide to Kill rather than Preempt a Reduce task.

References

1. Dean, J., Ghemawat, S.: MapReduce: simplified data processing on large clusters. *Communications of the ACM*, 107–113 (2008)
2. Zaharia, M. , Borthakur, D., Sen, S.J., Elmeleegy, K., Shenker, S., Stoica, I.: Delay scheduling: a simple technique for achieving locality and fairness in cluster scheduling: *Proceedings of the 5th European conference on Computer systems*, 265–278 (2010)
3. Liu, L., Zhou, Y., Liu, M., Xu, G., Chen, X., Fan, D., Wang, Q.: Preemptive hadoop jobs scheduling under a deadline, In: *Semantics, Knowledge and Grids (SKG), 2012 Eighth International Conference, IEEE*, 72–79 (2012)
4. Wang, Y., Tan, J., Yu, W., Meng, X., Zhang, L.: Preemptive reducetask scheduling for fair and fast job completion In: *Proceedings of the 10th International Conference on Autonomic Computing, ICAC*, 13 (2013)
5. Pastorelli, M., Dell’Amico, M., Michiardi, OS-Assisted Task Preemption for Hadoop: IN: *arXiv preprint arXiv:1402.2107*, (2014)
6. Apache Hadoop Welcome page, <http://http://hadoop.apache.org>