

# Failure Analysis of Hadoop Schedulers using an Integration of Model Checking and Simulation

Mbarka Soualhia

Concordia University,  
Montréal, Canada  
soualhia@ece.concordia.ca

Foutse Khomh

Polytechnique Montréal,  
Montréal, Canada  
foutse.khomh@polymtl.ca

Sofiène Tahar

Concordia University,  
Montréal, Canada  
tahar@ece.concordia.ca

**Abstract:** The Hadoop scheduler is a centerpiece of Hadoop, the leading processing framework for data-intensive applications in the cloud. Given the impact of failures on the performance of applications running on Hadoop, testing and verifying the performance of the Hadoop scheduler is critical. Existing approaches such as performance simulation and analytical modeling are inadequate because they are not able to ascertain a complete verification of a Hadoop scheduler. This is due to the wide range of constraints and aspects involved in Hadoop. In this paper, we propose a novel methodology that integrates and combines simulation and model checking techniques to perform a formal verification of Hadoop schedulers, focusing on the following properties: schedulability, fairness and resources-deadlock freeness. We use the CSP language to formally describe a Hadoop scheduler, and the PAT model checker to verify its properties. Next, we use the proposed formal model to analyze the scheduler of OpenCloud, a Hadoop-based cluster that simulates the Hadoop load, in order to illustrate the usability and benefits of our work. Results show that our proposed methodology can help identify several tasks failures (up to 78%) early on, i.e., before the tasks are executed on the cluster.

## 1 Introduction

Motivated by the reasonable prices and the good quality of cloud services, several enterprises and governments are deploying their applications in the cloud. Hadoop [1] has become enormously popular for processing data-intensive applications in the cloud. A core constituent of Hadoop is the scheduler. The Hadoop scheduler is responsible for the assignment of tasks belonging to applications across available computing resources. Scheduling data-intensive applications' tasks is a crucial problem, especially in the case of real-time systems, where several constraints should be satisfied. Indeed, an effective scheduler must assign tasks to meet applications' specified deadlines and to ensure their successful completions. In the cloud, failures are the norm rather than the exception and they often impact the performance of tasks running in Hadoop. Consequently, the Hadoop scheduler may fail to achieve its typical goal for different reasons such as resources-deadlock, task starvation, deadline non-satisfaction and data loss [8] [11]. This results in extra delays that can be added and propagated to the overall completion time of a task, which could lead to resources wastage, and hence significantly decreasing the performance of the applications and increasing failure rates. For instance, Dinu and Ng [7] analyzed the performance of the Hadoop framework under different types of failures. They reported that the Hadoop scheduler experiences several failures due to the lack of information sharing about its environment (e.g., scheduler constraints, resources availability), which can lead to poor scheduling decisions. For instance, they claim that the Hadoop scheduler experiences several failures because of prioritized executions of some long tasks before small ones, long delays of the struggling tasks or unexpected resources contentions [7] [14]. The widespread use of the Hadoop framework in safety and critical applications, such as healthcare [4] and

aeronautics [3], poses a real challenge to software engineers for the designing and testing of such framework to meet its typical goal and avoid poor scheduling decisions. Although several fault-tolerance mechanisms are integrated in Hadoop to overcome and recover from these failures, several failures still occur when scheduling and executing tasks [14]. Indeed, these failures can occur because of poor scheduling decisions (e.g., resources deadlock, task starvation) or constraints related to the environment where they are executed (e.g., data loss, network congestion). As such, verifying the design of Hadoop scheduler is an important and open challenge to identify the circumstances leading to task failures and performance degradation so that software engineers studying Hadoop can anticipate these potential issues and propose solutions to overcome them. This is to improve the performance of the Hadoop framework.

Different approaches have been adopted by the research community to verify and validate the behavior of Hadoop with respect to scheduling requirements. Traditionally, simulation and analytical modeling have been widely used for this purpose. However, with many Hadoop-nodes being deployed to accommodate the increasing number of demands, they are inadequate because they are not efficient in exploring large clusters. In addition, given the highly dynamic nature of Hadoop systems and the complexity of its scheduler, they cannot provide a clear understanding and exhaustive coverage of the Hadoop system especially when failures occur. Particularly, they are not able to ascertain a complete verification of the Hadoop scheduler because of the wide range of constraints and unpredictable aspects involved in the Hadoop model and its environment (e.g., diversity of running loads, availability of resources and dependencies between tasks). Formal methods have recently been used to model and verify Hadoop. However, to the best of our knowledge very few efforts have been invested in applying formal methods to verify the performance of Hadoop (e.g., data locality, read and write operations, correctness of running application on Hadoop) [15] [12] [9]. In addition, there is no work that addresses the formal verification of Hadoop schedulers. Considering the above facts, we believe that it is important to apply formal methods to verify the behavior of Hadoop schedulers. This is because formal methods are more able to model complex systems and thoroughly analyze their behavior than empirical testing. Our aim is to formally analyze the impact of the scheduling decisions on the failures rate and avoid simulation cost in terms of execution time and hardware cost (especially for large clusters). This allows to early identify circumstances leading to potential failures, in a shorter time compared to simulation, and prevent their occurrences. In contrast to simulation and analytical modeling, knowing these circumstances upfront would help practitioners better select the cluster settings (e.g., number of available resources, type of scheduler) and adjust the scheduler design (e.g., number of queues, priority handling strategies, failures recovery mechanisms) to prevent poor scheduling decisions in Hadoop.

In this paper, we present a novel methodology that combines simulation and model checking techniques to perform a formal analysis of Hadoop schedulers. Particularly, we study the feasibility of integrating model checking techniques and simulation to help in formally verifying some of the functional scheduling properties in Hadoop including schedulability, resources-deadlock freeness, and fairness [5]. To this aim, we first present a formal model to construct and analyze the three mentioned properties within the Hadoop scheduler. We use the Communicating Sequential Processes (CSP) language [16] to model the existing Hadoop schedulers (FIFO, Fair and Capacity schedulers [14]), and use the Process Analysis Toolkit (PAT) model checker [18] to verify the schedulers' properties. Indeed, the CSP language has been successfully used to model the behavior of synchronous and parallel components in different distributed systems [16]. Furthermore, PAT is a CSP-based model checker that has been widely used to simulate and verify concurrent, real-time systems, etc. [18]. Based on the generated verification results in PAT, we explore the relation between the scheduling decisions and the failures rate while simulating different load scenarios running on the Hadoop cluster. This is in order to propose possible scheduling strategies to reduce the number of failed tasks and improve the overall cluster performance

(resources utilization, total completion time, etc). In order to illustrate the usability and benefits of our work to identify failures that could have been prevented using our formal analysis methodology, we apply our approach on the scheduler of OpenCloud, a Hadoop-based cluster that simulates the real Hadoop load [10]. Using our proposed methodology, developers would identify up to 78% of tasks failures early on before the deployment of the application. Based on the performed failures analysis, our methodology could provide potential strategies to overcome these failures. For instance, our solution could propose new scheduling strategies to adjust the Hadoop cluster settings (e.g., size of the queue, allocated resources to the queues in the scheduler, etc.) and reduce the failures rate when compared to the real-execution simulation results. To the best of our knowledge, the present work is different from existing research in applying and integrating both formal methods and simulation for the analysis of Hadoop schedulers and formally analyzing the impact of the scheduling decisions on the failures rate in Hadoop. Furthermore, our proposed methodology to integrate model checking and simulation to verify Hadoop schedulers could be also applied to formally analyze other schedulers, than Hadoop, such as Spark [2], which has become one of the key cluster-computing framework that can be running on Hadoop.

The rest of the paper is organized as follows: Section 2 describes basics of Hadoop architecture, CSP language, and the tool PAT. Section 3 presents the proposed methodology for the formal analysis of the Hadoop scheduler. We describe the analysis of the scheduler of the OpenCloud a Hadoop cluster in Section 4. Section 5 summarizes the most relevant related work. Finally, Section 6 concludes the paper and highlights some future directions.

## 2 Preliminaries

In this section we briefly present the Hadoop architecture and some of the basic of the CSP language and the tool PAT, which will be used in the rest of the paper. This is in order to better understand the different steps of our proposed methodology.

### 2.1 Hadoop Architecture

Hadoop [14] is an open source implementation of the MapReduce programming model [8]. MapReduce is designed to perform parallel processing of large datasets using a large number of computers. A MapReduce job is comprised of two functions: a map and reduce, and the input data [14]. Hadoop has become the *de facto* standard for processing large data in today's cloud environment. It is a master-slave-based framework, the master node consists of a JobTracker and NameNode. The worker/slave node consists of a TaskTracker and DataNode. The Hadoop Distributed File System (HDFS) is the storage unit responsible for controlling access to data in Hadoop. Hadoop is equipped with a centerpiece; the scheduler which distributes tasks across worker nodes according to their availability. The default scheduler in Hadoop is based on the First In First Out (FIFO) principle. Facebook and Yahoo! have developed two new schedulers for Hadoop: Fair scheduler and Capacity scheduler, respectively [14].

### 2.2 CSP and PAT

CSP is a formal language used to model and analyze the behavior of processes in concurrent systems. It has been practically applied in modeling several real time systems and protocols [16]. In the sequel, we present a subset of the CSP language, which will be used in this work, where  $P$  and  $Q$  are processes,  $a$  is an event,  $c$  is a channel, and  $e$  and  $x$  are values:

$$P, Q ::= \text{Stop} \mid \text{Skip} \mid a \rightarrow P \mid P ; Q \mid P \parallel Q \mid c!e \rightarrow P \mid c?x \rightarrow P$$

- **Stop**: indicates that a process is in the state of deadlock.
- **Skip**: indicates a successfully terminated process.
- $a \rightarrow P$ : means that an object first engages in the event  $a$  and then behaves exactly as described by  $P$ .
- $P ; Q$ : denotes that  $P$  and  $Q$  are sequentially executed.
- $P \parallel Q$ : denotes that  $P$  and  $Q$  are processed in parallel. The two processes are synchronized with the same communication events.
- $c!e \rightarrow P$ : indicates that a value  $e$  was sent through a channel  $c$  and then a process  $P$ .
- $c?x \rightarrow P$ : indicates a value was received through a channel  $c$  and stored in a variable  $x$  and then a process  $P$ .

PAT [18] is a CSP-based tool used to simulate and verify concurrent, real-time systems, etc. [17]. It implements different model checking techniques for system analysis and properties verification in distributed systems (e.g., deadlock-freeness, reachability, etc.). Different advanced optimizations techniques, such as partial order reduction, symmetry reduction, etc., are available in PAT to reduce the number of explored states and CPU time.

### 3 Formal Analysis Methodology

In this section, we first present a general overview of our methodology followed by a description of each step.

#### 3.1 Methodology Overview

Figure 1 provides an overview of the main idea behind our formal analysis of the Hadoop scheduler when integrating model checking and simulation, to early identify failure occurrences. The inputs of our proposed methodology are (1) the description of the Hadoop scheduler, (2) the specification of the properties to be verified, and (3) the scheduling metrics (e.g., type of scheduler, number of nodes, workload and failure distributions, schedulability rate). Our proposed methodology is comprised of three main steps, including (a) the formal modeling of the Hadoop scheduler and its properties in CSP and LTL, (b) the quantitative analysis of failures using model checking in PAT, and (c) the qualitative analysis of the failures using simulation of the proposed scheduling strategies and real-simulation traces. The outputs of our methodology are the rate of failures of the verified scheduler, and a set of possible scheduling strategies to overcome these failures.

We have chosen the CSP language to formally describe the scheduler as it allows to model the behavior of processes in concurrent systems. It has been successfully applied in modeling synchronous and parallel components for several real-time and distributed systems [16] (which is the case of Hadoop). The properties we aim to verify are written in Linear Temporal Logic (LTL). Thereafter, we use the PAT model checker to perform the formal quantitative analysis of failures in Hadoop scheduler. PAT is based on CSP and implements various model checking techniques to analyze and simulate the behavior of several distributed systems (e.g., transportation management system, Web authentication protocols, wireless sensor networks, etc.) [18]. Furthermore, it allows to model timed and probabilistic processes in the studied systems [17]. Based on the generated results from PAT, we perform a qualitative failures analysis to determine the circumstances and specifications leading to tasks' failures in the scheduler. The remainder of this section elaborates more on each of the steps of our methodology.

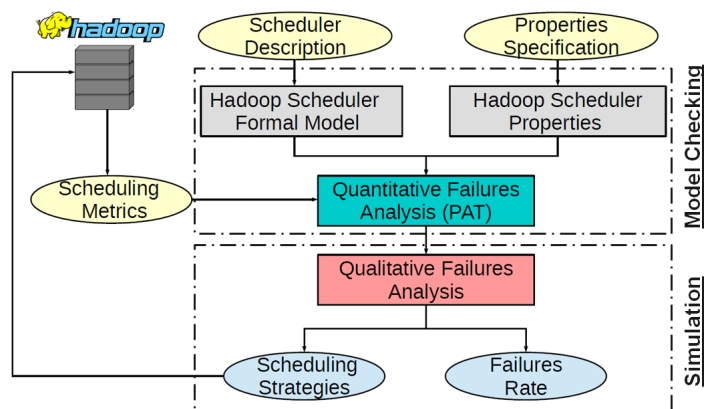


Figure 1: Overview of the Formal Analysis of Hadoop Schedulers Methodology

### 3.2 Hadoop Scheduler Formal Model: Model Checking

The first step in conducting the proposed formal analysis of the Hadoop scheduler using a model checker is to construct a formal model of the main components in Hadoop responsible for the scheduling of tasks, using the CSP language. To do so, we start by writing a formal description of the Hadoop master node: the JobTracker and NameNode. At the master level, we model the scheduler and the main entities responsible for task assignment and resources allocation in Hadoop. Next, we model the TaskTracker and the DataNode including the entity responsible for the task execution at the worker nodes. In addition, we integrate some of the important scheduling constraints in Hadoop in the model of the TaskTracker nodes including the data locality, data placement, and speculative execution of tasks [14]. Here, we selected these three constraints because of their direct impact on the scheduling strategies and the performance of executed tasks [14]. At this level, we should mention that the provided model of the Hadoop scheduler represents a close representation of the actual ones (e.g., FIFO, Fair, and Capacity) because it includes the main functionalities responsible for assigning the received tasks on available nodes. Furthermore, we checked the correspondence between the provided model and the real Hadoop scheduler by comparing the scheduling outcomes of scheduled tasks when using the formal model and the existing ones. More details about this comparison is given in Section 4.2.2. The obtained results showed a good matching between the two models. Nevertheless, further functionalities can be added to the presented model in order to improve its results (*i.e.*, recovery mechanisms, efficient resources assignment). In the following, we present a formal description of the steps to model a Hadoop cluster. Then, we present examples of implemented CSP processes<sup>1</sup> to describe the *Hadoop scheduler*, *TaskTracker activation* and *task assignment*, where “*Cluster()*” is the main process and  $N$  is the number of available TaskTracker nodes in the cluster:

```
Cluster() = initialize(); NameNode_activate() || JobTracker_activate() ||
  (|| i:{0..(N-1)}@DataNode_activate(i)) ||
  (|| i:{0..(N-1)}@TaskTracker_activate(i)) ||
  Hadoop_Scheduler();
```

The following process presents a formal description of the steps in “*Hadoop\_Scheduler()*” to check scheduling constraints of map/reduce tasks. First, it checks the availability of resource slots ( $slotTT[i] > 0$ ). Then, it checks the type of task to be scheduled, either a map ( $Queue[index] == 1$ ) or reduce ( $Queue[index] == 2$ ) task. It also checks whether a task is speculatively executed or not (e.g., map:

<sup>1</sup>The entire CSP script is available at:  
<http://hvg.ece.concordia.ca/projects/cloud/fvhs.html>

$Queue[index] == 3$  or reduce:  $Queue[index] == 4$ ). Next, it assigns the received task to the TaskTracker node where it will be executed ( $signedtask?i \rightarrow signedtask.i \rightarrow Task\_Assignment(location,type);$ ).

```
Hadoop.Scheduler() =
{
  if( (slotTT[i]>0) )
    {while( (found==0) && (index < maxqueue ) )
      {
        if (Queue[index] == 1) //it is a map task
          {
            schedulable = 1; found =1; location = index;
            type = MapTask; IDjob_ task = IDJob[index];}
          if((Queue[index] == 2)) //it is a reduce task
            {
              if(FinishedMap[IDJob[index]] == Map[IDJob[index]] )
                {
                  schedulable = 1; found =1; location = index;
                  type = ReduceTask; IDjob_ task = IDJob[index];}
            }
          if(Queue[index] == 3) //it is a speculated map task
            {
              schedulable = 1; found =1; location = index; type = MapTask;
              IDjob_ task = IDJob[index]; SpeculateTask[location] = 1;}
          if(Queue[index] == 4) //it is a speculated reduce task
            {
              schedulable = 1; found =1; location = index; type = ReduceTask;
              IDjob_ task = IDJob[index]; SpeculateTask[location] = 1;}
          ...
        }
      }
    }
} → signedtask?i → signedtask.i → Task.Assignment(location, type);
```

“*TaskTracker\_activate(i)*” presents our proposed process to activate the TaskTracker, after checking that the JobTracker was already activated ( $JobTracker == ON$ ) and this TaskTracker was not already activated ( $TaskTracker[i] == OFF$ ). Next, it activates this TaskTracker ( $TaskTracker[i] == ON$ ) and the number of slots specified to this TaskTracker ( $slotsnb$ ). These activated slots are ready to execute tasks.

```
TaskTracker_activate(i) = activate_jt_success → ifa(TaskTracker[i] ==
OFF && JobTracker == ON) {activate_tt.i{
TaskTracker[i] = ON; trackercount++;}
→ atomic{activate_tt_success.i →
(|| j:{1..(slotsnb)}@TaskTracker_sendready(i))}};
```

The following “*TaskTracker\_execute(i)*” process presents an example of executing a task after checking its locality in Hadoop. For instance, it checks the availability of the slot assigned to a given task by the scheduler (if slot is free then  $task\_running[nbTT][k]$  is equal to 0, where  $nbTT$  is the ID of the TaskTracker and  $k$  is the ID of the assigned slot). Next, it checks the locality of the task by checking whether the node where to execute the task ( $selectedTT$ ) is the same node where its data is located ( $Data-LocalTT[idtask]$ ).

```
TaskTracker_execute(i) =
{
  var nbTT =i; var found = 0; var k = 0;
  while((k<slotsnb) && (found == 0) ){
    if(task_at_tasktracker[nbTT][k]==1 && task_running[nbTT][k]== 0)
      {selectedslot =k; found = 1; }
  }
```

```

k++; } ...
if(Data-LocalTT[idtask] == selectedTT) //check locality of the task
{locality = locality + 1; Locality[idtask] = 1;}
else {nonlocality = nonlocality + 1; Locality[idtask] = 0; }
...}
} → if(pos== -1) {TaskTracker_ execute(i)}
      else {execute(i,selectedslot)};

```

### 3.3 Hadoop Scheduler Properties: Model Checking

The three selected properties we aim to verify in our work are the schedulability, fairness and resources-deadlock freeness. We select these properties because they represent some of the main critical properties affecting the execution of tasks in real-time systems (e.g., task outcome, delays, resources utilization) [5]. The three properties can be described as follows: the *schedulability* checks whether a task is scheduled and satisfies its specified deadline when scheduled according to a scheduling algorithm. The *fairness* checks whether the submitted tasks are served (e.g., receiving resources slots that ensure their processing) and there are no tasks that will never be served or will wait in the queue more than expected. The *resources-deadlock* checks whether two tasks are competing for the same resource or each is waiting for the other to be finished.

To better illustrate the properties to be verified, we explain as an example the schedulability property and its corresponding states in our approach. For the schedulability, a task can go from state: *submitted* to *scheduled* to *processed* then to *finished-within-deadline* or *finished-after-deadline* or *failed*. Let  $X$  be the total number of scheduled tasks and  $Y$  be the number of tasks finished within their expected deadlines. The schedulability rate can be defined as the ratio of  $X$  over  $Y$ . The property “*schedulabilityrate* > 80” means checking whether the scheduler can have a total of 80% of tasks finished within their deadlines.

To verify above properties in PAT, we need to provide their descriptions in LTL. For example, the following LTL formulas check the schedulability and resource-deadlock freeness of given tasks. The first example checks whether a given task eventually goes from the state submitted to the state finished within the deadline. The second example checks whether a given task should not go to a state of waiting-resources. Here,  $\diamond$ ,  $\models$ ,  $\rightarrow$ , and  $\neg$  represent *eventually*, *satisfy*, *imply*, and *not*, respectively, in the LTL logic.

```

# assert  $\diamond$  (task  $\models$  (submitted  $\rightarrow$  finished-within-deadline) );
# assert  $\neg$  (task  $\models$  (submitted  $\rightarrow$  waiting-resources) );

```

### 3.4 Quantitative Failures Analysis using Model Checking

Provided the CSP model of the Hadoop scheduler and the properties to be verified, we perform the formal analysis of the scheduler performance using the PAT model checker while simulating different Hadoop workload scenarios. Here, we can vary the property requirements to assess the performance of the scheduler under different rates and evaluate their impact on the failures rates of the cluster and the simulated scenarios. For example, we can define “*goal0*” to check whether all the submitted tasks (“*workload*”) are successfully scheduled. Using PAT, we can verify if the modeled cluster, “*cluster1*”, can reach this goal or not. Another example could be to check if “*cluster1*” meets “*goal0*” with a “*schedulabilityrate*” of 80%. The following examples present some of the properties that can be verified using our approach.

```

#define goal0 completedscheduled == workload && workload >0;
#define goal1 fairnessrate ==50;
#define goal2 resourcedeadlockrate ==50;
#assert cluster1 reaches goal0 && schedulabilityrate >80;
#assert cluster1 reaches goal0 && goal1;
#assert cluster1 reaches goal0 && goal1 && goal2;

```

### 3.5 Qualitative Failures Analysis using Simulation

The last step in our proposed methodology is to use the traces simulated and generated by the PAT model checker to extract information about the applied scheduling strategies and explore their impact on tasks' failures using simulation. For instance, we can investigate the states where the scheduler does (not) meet a given property and map these states to the obtained scheduler performance and to the input cluster settings. This step allows us to find a possible correlation between the cluster settings, the applied scheduling strategies, and the failures rate. Next, we use these correlations to suggest scheduling strategies to overcome these failures by either (1) recommending to the Hadoop developers to change the scheduling decisions (e.g., delay long tasks, wait for a local task execution) or (2) to customers to change and adjust their cluster settings (e.g., number of nodes, number of allowed speculative executions). In next section, we propose a case study to evaluate and simulate the suggested scheduling strategies on a Hadoop environment and measure their impact on the failure rates and the Hadoop cluster performance. Generally, our solution could propose new scheduling strategies to adjust the Hadoop cluster settings and hence reduce failures rate when compared to the real-execution simulation results.

## 4 Case Study: Formal Analysis of OpenCloud Scheduler

In this section, we illustrate the practical usability and benefits of our work by formally analyzing the scheduler of OpenCloud, a Hadoop-based cluster [10] that simulates the Hadoop load.

### 4.1 Case Study Description

To apply our formal approach on a case study, we investigated existing Hadoop case studies in the open literature. Overall, we found three main public case studies including: Google [19], Facebook [20], and OpenCloud [10] traces. We found out that the Google traces do not provide data about the cluster settings, which is an important factor affecting the analysis results in our approach. For the Facebook traces, we noticed that they do not provide data about the cluster settings, capacity of nodes, failures rate, etc., which are essential for our approach. Whereas, we found that the OpenCloud traces provide the required inputs of our verification approach (e.g., # nodes, capacity of nodes, workload). Therefore, we choose to formally analyze the scheduler of this cluster because it provides public traces of real-executed Hadoop workload for more than 20 months. OpenCloud [10] is an open research cluster managed by Carnegie Mellon University. It is used by several groups in different areas such as machine learning, astrophysics, biology, cloud computing, etc.

Based on the modeled system of the OpenCloud's scheduler and the specified properties, we start the verification by parsing the trace files to extract the input information needed in our methodology. Specifically, we use the description of the workload included in the first month traces, and the first six months traces together. This allows us to evaluate the scalability of our methodology, in terms of number of visited states and execution time, using different traces. Although these traces provide the required inputs for our methodology, we did not find any information describing the type of scheduler



used in the cluster. Since the type of the scheduler is an important factor that impacts the performance of the cluster, we evaluate the performance of the modeled cluster for the three existing schedulers of Hadoop (FIFO, Fair and Capacity). We vary the property requirements to assess the performance of the used scheduler under different rates and evaluate their impact on the failures rate in the cluster while executing the same Hadoop workload. For the search engines in PAT, we use the “First Witness Trace using Depth First Search” in order to perform an analysis without reduction, and the “First Witness Trace with Zone Abstraction” for the analysis with symmetry reduction. The testbed is a workstation with Intel i7-6700HQ (2.60GHz\*8) CPU and 16 GB of RAM.

## 4.2 Verification and Evaluation

In this section, we present the obtained scalability results of our methodology along with the results of the formal quantitative and qualitative analyses of failures in Hadoop schedulers.

### 4.2.1 Properties Verification and Scalability Analysis

Tables 1 and 2 summarize the verification results of the first month trace, and the first six months traces together, respectively. The first trace provides information about 1,772,144 scheduled tasks, whereas the six files describing the executed workload for six months contain information about 4,006,512 scheduled tasks. In the sequel, we discuss the results of the performed analysis.

The results presented in Table 1 show that only the Fair scheduler satisfies the schedulability property (for the two given rates), meaning that up to 80% are scheduled and executed within their expected deadlines. The Capacity scheduler does not meet the schedulability rate of 80%. However, the FIFO does not satisfy the schedulability properties for the two input rates, meaning that more than 50% of scheduled tasks are exceeding their expected deadlines. Hence, these tasks are using their assigned resources more than expected, which can affect the overall performance of the cluster.

For the fairness property, only the Fair scheduler satisfies the property of fairness with a rate of 50% and 80%, meaning that at most 80% of tasks get served and executed on time. However, both the FIFO and the Capacity schedulers violate the fairness property for the two given values. Therefore, we can claim that more than 80% of the submitted tasks are waiting longer than expected in the queue before getting executed. This may lead to a problem of task starvation.

For the resources-deadlock, we can report that the Capacity scheduler is characterized by more tasks that experience a resources-deadlock compared to the FIFO and the Fair schedulers. This is because it satisfies the property of resources-deadlock for the two given rates (e.g., 10% and 30%). This can be explained by the fact that the Capacity scheduler suffers from the problem of miscalculation of resources (headroom calculation). The FIFO scheduler shows that only 10% of the scheduled tasks experience the problem of resources-deadlock. The Fair scheduler does not satisfy the resources-deadlock property for two given rates (10% and 30%) for the first trace, which means that less than 10% of tasks may experience an issue of resources-deadlock. Hence, we can conclude that the Fair scheduler generates less scheduling decisions leading to resources-deadlock situations for the OpenCloud cluster.

Overall, our proposed methodology is able to verify the given three properties for the three existing Hadoop schedulers, while exploring on average 11086 K states in 3724 seconds without symmetry reduction, and 742 K states in 1619 seconds with symmetry reduction, as shown in Table 1.

In order to evaluate the scalability of our methodology, we perform the formal analysis of a cumulative workload. This is done by incrementally adding the Hadoop workload of each month to the previous trace(s) to be analyzed. We start the evaluation by analyzing the trace of the first month, and add the traces up to the trace where the tool cannot perform the analysis. This is in order to check the bounds of the analysis performed in terms of explored states. The obtained results, presented in Table 2, show

Table 1: Verification Results: Trace for the First Month (1,772,144 Tasks)

Property	Scheduler	Results without SR			Results with SR		
		Valid?	#States	Time(s)	Valid?	#States	Time(s)
Schedulability = 50%	FIFO	No	11086 K	3869	No	742 K	1648
	Fair	Yes	11086 K	3524	Yes	742 K	1597
	Capacity	Yes	11086 K	3601	Yes	742 K	1604
Schedulability = 80%	FIFO	No	11086 K	3789	No	742 K	1650
	Fair	Yes	11086 K	3676	Yes	742 K	1614
	Capacity	No	11086 K	3721	No	742 K	1602
Fairness = 50%	FIFO	No	11086 K	3714	No	742 K	1594
	Fair	Yes	11086 K	3759	Yes	742 K	1615
	Capacity	No	11086 K	3736	No	742 K	1612
Fairness = 80%	FIFO	No	11086 K	3855	No	742 K	1675
	Fair	Yes	11086 K	3795	Yes	742 K	1642
	Capacity	No	11086 K	3761	No	742 K	1619
Resources- Deadlock = 10%	FIFO	Yes	11086 K	3615	Yes	742 K	1602
	Fair	No	11086 K	3698	No	742 K	1610
	Capacity	Yes	11086 K	3704	Yes	742 K	1632
Resources- Deadlock = 30%	FIFO	No	11086 K	3742	No	742 K	1596
	Fair	No	11086 K	3733	No	742 K	1618
	Capacity	Yes	11086 K	3752	Yes	742 K	1623

SR = Symmetry Reduction

that our approach could formally analyze the first six traces (combined) describing the workload (about 4,006,512 tasks) for the first six months. It could analyze the first five traces together with and without symmetry reduction, however, it could analyze the six traces together only with symmetry reduction by exploring on average 17,692 K states in 4346 seconds.

Table 2: Verification Results: Trace for the 1-6 Months (4,006,512 Tasks)

Property	Scheduler	Results without SR			Results with SR		
		Valid?	#States	Time(s)	Valid?	#States	Time(s)
Schedulability = 50%	FIFO	**	**	**	Yes	17692K	4350
	Fair	**	**	**	Yes	17692K	4362
	Capacity	**	**	**	Yes	17692K	4359
Schedulability = 90%	FIFO	**	**	**	No	17692K	4346
	Fair	**	**	**	No	17692K	4341
	Capacity	**	**	**	No	17692K	4367
Fairness = 30%	FIFO	**	**	**	Yes	17692K	4377
	Fair	**	**	**	Yes	17692K	4312
	Capacity	**	**	**	Yes	17692K	4335
Fairness = 90%	FIFO	**	**	**	No	17692K	4352
	Fair	**	**	**	No	17692K	4328
	Capacity	**	**	**	No	17692K	4369
Resources- Deadlock = 10%	FIFO	**	**	**	Yes	17692K	4322
	Fair	**	**	**	No	17692K	4360
	Capacity	**	**	**	Yes	17692K	4354
Resources- Deadlock = 50%	FIFO	**	**	**	No	17692K	4338
	Fair	**	**	**	No	17692K	4342
	Capacity	**	**	**	No	17692K	4328

SR = Symmetry Reduction, \*\* = Not Available

#### 4.2.2 Quantitative Failures Analysis

We use the traces generated by the PAT model checker during the verification of the three properties described in Section 4.2.1, to explore states where the scheduler did not satisfy a given property value. This step is important to map these states to the failed tasks, if found, and examine the relationship between

the verified property and the task failure. We apply this step on the first trace file because it contains an important number of scheduled tasks (i.e., 1,772,144 tasks). To do so, we first identify the tasks that did not satisfy a given property, for example schedulability = 80% or resources-deadlock = 10%. Then, we check whether these tasks were failed when they were executed in the real cluster. Next, we classify the observations into four main categories: *True Positive (TP)*, *True Negative (TN)*, *False Positive (FP)*, and *False Negative (FN)*. *TP* represents the successfully executed tasks, based on the simulation traces, that are identified as finished tasks using our approach. Likewise, *TN* represents the failed tasks, based on the simulation traces, that are identified as failed tasks. While *FP* denotes the amount of identified finished tasks that failed during the real simulation. Finally, *FN* denotes the number of identified failed tasks that are finished in reality. These four metrics are calculated by comparing the status of the tasks in the generated traces using our methodology, specifically from PAT, and the simulation traces across the total number of scheduled tasks in the input trace (1,772,144 tasks).

Overall, we observe that our formal analysis approach could identify up to 61.49% of the finished tasks (*TP*, Fair, schedulability = 50%), and up to 4.62% of the failed tasks (*TN*, Fair, schedulability = 80%) without symmetry reduction. Here, we notice that the *TN* values are small compared to the *TP* ones. This can be explained by the fact that the first trace contains more than 94% of successful tasks. For this reason, we computed another metric that we call the *Detected Failures (DF)*. The *DF* is calculated by mapping the *TN* rate over the total number of failed tasks in the input trace. Our aim here is to quantify the amount of detected failures using our formal verification approach when compared to the given OpenCloud traces (obtained from real-execution simulation). This is to answer the question of how many failures could be identified by our formal verification approach before the application is deployed in the field?

At this level, we can claim that our approach is able to catch between 42% and 78.57% of the total failed tasks without symmetry reduction. While it can catch between 42% and 78.91% of the failures when using the symmetry reduction. On the other hand, we notice that the *FN* is in the order of 40%, which means that more than 40% of tasks are finished, in the simulation traces, but our methodology indicates their failures. This can be explained by the internal mechanisms implemented in Hadoop internally to recover these tasks in case of a failure. For example, the mechanism of pausing and resuming running tasks allows higher priority tasks to be executed without killing the lower priority ones [11]. Furthermore, we notice that the false positive rate varies from 1.26% and 3.41% and indicates failed tasks that are identified as finished using our methodology. A possible explanation for these false positive results can be the lack of some real-world constraints that affect the execution of tasks (e.g., network congestion, lack of resources). In addition, failed tasks identified as finished can generate more false positive results due to the tight dependency between the map and the reduce tasks (the reduce tasks will be launched when all the map tasks are successfully executed).

Next, we check the states of the failed tasks and analyze the factors that may cause the failure of these tasks. We find that the scheduler of OpenCloud experiences several failures, up to 32%, because of long delays that exceed the maximum timeout for a task to be finished (property “mapred.task.timeout”: defining the maximum timeout for a task to be finished). We carefully checked the states of these tasks and found that these delays are mainly caused by data locality constraints [14] and that they can reach 10 minutes (for small and medium tasks). Moreover, we found out that about 40% of these straggling tasks cause the failure of the job to which they belong, resulting in a waste of resources. Moreover, we noticed that many failures are cascaded from one job to another, especially in the long Hadoop chains. A Hadoop chain is a connection between more than one Hadoop job to execute a specific workload. This may result in a degradation in performance and many failures. We can claim that the Hadoop scheduler lacks mechanisms to share information about these failures between its components to avoid their occurrence.

Table 3: Coverage Results(%): Trace for the First Month (1,772,144 Tasks)

Property	Scheduler	Results without SR					Results with SR				
		TP	TN	FP	FN	DF	TP	TN	FP	FN	DF
Schedulability = 50%	FIFO	56.03	3.2	2.68	38.09	54.76	47.29	2.47	3.41	46.83	42.00
	Fair	61.49	4.53	1.35	32.63	77.04	55.38	3.82	2.06	38.74	64.96
	Capacity	49.16	2.47	3.41	44.96	42.00	46.09	2.85	3.03	48.03	48.46
Schedulability = 80%	FIFO	50.14	3.46	2.42	43.98	58.84	47.98	2.79	3.09	46.14	47.44
	Fair	59.83	4.62	1.26	34.29	78.57	56.82	4.21	1.67	37.3	71.59
	Capacity	43.61	3.05	2.83	37.73	51.87	42.18	3.03	2.85	51.94	51.53
Fairness = 50%	FIFO	49.28	2.92	2.96	44.84	49.65	47.84	2.92	2.96	46.28	49.65
	Fair	55.36	3.89	1.99	38.76	66.15	49.63	3.86	2.02	44.49	65.64
	Capacity	47.03	2.47	3.41	47.09	42.00	44.11	3.04	2.84	50.01	51.70
Fairness = 80%	FIFO	44.88	3.76	2.12	49.29	63.94	49.77	3.32	2.56	44.35	56.46
	Fair	57.21	4.07	1.81	36.91	69.21	50.14	4.64	1.24	43.98	78.91
	Capacity	43.14	3.01	2.87	50.98	51.19	48.29	3.48	2.4	45.83	76.18
Resources-Deadlock = 10%	FIFO	43.73	3.06	2.82	50.39	52.04	46.21	3.19	2.69	47.92	54.25
	Fair	49.99	3.82	2.06	44.13	64.96	51.44	3.63	2.25	42.68	61.73
	Capacity	46.07	3.19	2.69	48.05	54.25	43.77	2.84	3.04	50.35	48.29
Resources-Deadlock = 30%	FIFO	50.22	3.53	2.35	43.9	60.03	48.54	2.91	2.97	45.58	49.48
	Fair	52.39	4.17	1.71	41.73	70.91	55.26	3.75	2.13	38.86	63.77
	Capacity	49.02	3.16	2.72	45.1	53.74	49.71	3.19	2.69	44.41	54.25

SR = Symmetry Reduction, DF = Detected Failures  
 TP = True Positive, TN = True Negative, FP = False Positive, FN = False Negative

On the other hand, we found out that 26% of tasks failed because they exceeded the maximum number of allowed speculative execution (e.g., property “mapred.map.tasks.speculative.execution”: defining the maximum number of allowed speculative execution, they have likely more chances to fail). When checking the tasks waiting for a long time in the queue before being served, we observed that this is due to the fact that the long tasks are scheduled first and they occupy their assigned resources for a long time (e.g., a large input file to be processed, a job composed of more than 1000 tasks). Consequently, we can conclude that knowing these factors and issues, one can adapt the scheduler system of the Hadoop framework to overcome these failures and improve its performance.

### 4.2.3 Qualitative Failures Analysis

To show the benefits of our formal analysis approach, we propose to integrate and simulate some guidelines or strategies to adapt the scheduling decisions of the created Hadoop cluster and evaluate their impact on the failures rate. This is based on the generated traces and the performed failures analysis to check whether our work can help early identify the occurrence of failures and then propose appropriate mechanisms to overcome them. For instance, one possible strategy could be to change the cluster settings by adding more resources on its nodes or adding the number of nodes on it. This could solve for example the fairness and resources-deadlock issues. Another scheduling strategy could be to change the value of timeout of scheduled task, which represents the number of milliseconds before terminating a task. Another strategy could be to adjust the type of the scheduler used in the cluster (e.g., FIFO, Fair, Capacity, etc.). In this paper, we evaluate the impact of the strategy to change the cluster settings by adding more resources on the OpenCloud cluster where we change the number of nodes from 64 to 100 and simulate the same Hadoop workload. Figure 2 gives an overview of the impact of adding more resources in the cluster for the three schedulers considering a failures’ rate of 5.88%; the identified failure rate from the first trace file. Overall, we noticed that adding more nodes in the cluster could reduce the failure rates by up to 2.34% (Fair scheduler), which means a reduction rate of 39.79%. This was expected since when we analyzed the traces we found out that several tasks are straggling for more than 10 minutes, waiting for other resources to be released.

Given the obtained findings, we can claim that our solution could identify new scheduling strategies to adjust the Hadoop cluster to reduce failures rates by combining simulation and model checking techniques and formally verifying the functional scheduling properties.

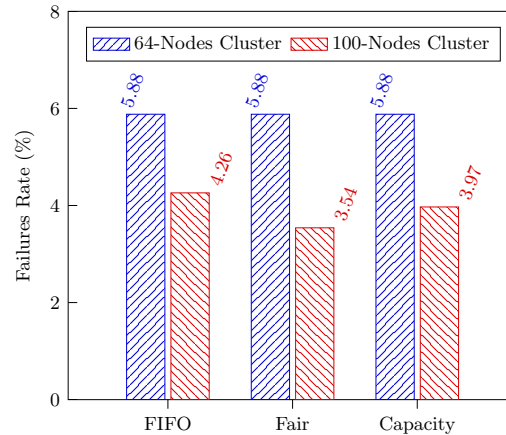


Figure 2: Impact of Resources Adding on Failures Rate

## 5 Related Work

In this section, we discuss the most relevant work that apply formal methods in the context of Hadoop framework. In [6] and [13], the authors analyzed the behavior of MapReduce using Stochastic Petri Nets and Timed Colored Petri Nets, respectively. They modeled the mean delay in each time transition of the scheduled tasks as formulas, and the Hadoop jobs were simulated based on the used Petri Nets. The proposed approaches could evaluate the correctness of the system and analyze the performance of MapReduce. But, they lack several constraints about the scheduling of the jobs and cannot cover larger Hadoop clusters. Su *et al.* [15] used the CSP language to formalize four key components in MapReduce. Specifically, they formally modeled the master, mapper, reducer and file system while considering the basic operations in Hadoop (e.g., task state storing, error handling, progress controlling). However, none of the properties of the Hadoop framework is verified using the formalized components. Xie *et al.* [22] address the formal verification of the HDFS reading and writing operations using CSP and the PAT model checker. For instance, they formally modeled the reading and writing operations for the HDFS based on the formalized components proposed in [15]. Moreover, they verified some of the HDFS properties including the deadlock-freeness, minimal distance scheme, mutual exclusion, write-once scheme and robustness. While this approach allows to detect unexpected traces generating errors and verify data consistency in the HDFS, a limitation of this work is that it only models the reading and writing operations for just one file system and requires to investigate the validity of these operations for distributed files as in HDFS. In [12], Reddy *et al.* propose an approach to model Hadoop's system using the PAT model checker. They used CSP to model Hadoop software components including the "NameNode", "DataNode", task scheduler and cluster setup. They identified the benefits of some properties like data locality, deadlock-freeness and non-termination among others and proved the correctness of these properties. However, their proposed model is evaluated on a small workload, and none of the properties is verified to check the performance of the Hadoop scheduler. Kosuke *et al.* [9] used the proof assistant

Coq to write an abstract computation model of MapReduce. They modeled the mapper and reducer as Coq functions and proved that the implementation of the mapper and reducer satisfies the specification of some applications such as WordCount [21]. The authors present an abstracted model of MapReduce, where many details are not included such as task assignment or resources allocation. These issues can affect the performance of applications running on Hadoop.

## 6 Conclusion and Future Work

Given the dynamic behavior of the cloud environment, the verification of the Hadoop scheduler has become an open challenge especially with many Hadoop-nodes being deployed to accommodate the increasing number of demands. Existing approaches such as performance simulation and analytical modeling are not able to ascertain a complete verification of the Hadoop scheduler because of the wide range of constraints involved in Hadoop. In this paper, we presented a novel methodology that combines and integrates simulation and model checking techniques to perform a formal analysis of Hadoop schedulers. Particularly, we studied the feasibility of integrating model checking techniques and simulation to help in formally verifying some of the functional scheduling properties. So, we formally verified some of the most important scheduling properties on Hadoop including the schedulability, resources-deadlock freeness, and fairness properties, and analyzed their impact on the failure rates. The ultimate goal of this work is to be able to propose possible scheduling strategies to reduce the number of failed tasks and improve the overall cluster performance and practitioner better assign and configure resources to mitigate the failures that a Hadoop scheduler may experience while simulating the Hadoop workload. We used CSP to model Hadoop schedulers, and the PAT model checker to verify the mentioned properties. To show the practical usefulness of our work, we applied our approach on the scheduler of OpenCloud, a real Hadoop-based cluster. The obtained results show that our approach is able to formally verify the selected properties and that such analysis can help identify up to 78% of failures before they occur in the field.

An important direction of future work is to verify other properties that can impact the failures rate in Hadoop, like the resource assignment, load balancing, and modeling the internal recovery mechanisms of Hadoop. Another direction can be the use of our work to automatically generate scheduling guidelines to improve the performance of the Hadoop framework and reduce the failures rate. Finally, the failure analysis approach and findings of this paper can be extended and evaluated on Spark [2], which has become one of the key cluster-computing framework that can be running on Hadoop. In fact, one can easily adapt the proposed methodology according to the architecture of Spark.

## References

- [1] Apache Hadoop: <http://hadoop.apache.org/>.
- [2] Apache Spark: <http://spark.apache.org/>.
- [3] Applying Apache Hadoop to NASA's Big Climate Data: [http://events.linuxfoundation.org/sites/events/files/slides/ApacheCon\\_NASA\\_Hadoop.pdf](http://events.linuxfoundation.org/sites/events/files/slides/ApacheCon_NASA_Hadoop.pdf).
- [4] D. Chen, Y. Chen, B. N. Brownlow, P. P. Kanjamala, C. A. G. Arredondo, B. L. Radspinner & M. A. Raveling (2017): *Real-Time or Near Real-Time Persisting Daily Healthcare Data Into HDFS and Elastic-Search Index Inside a Big Data Platform*. *IEEE Transactions on Industrial Informatics* 13(2), pp. 595–606, doi:10.1109/TII.2016.2645606.
- [5] A.M. K. Cheng (2002): *Real-Time Systems: Scheduling, Analysis, and Verification*. John Wiley & Sons, Inc., doi:10.1002/0471224626.

- [6] S.-T. Cheng, H.-C. Wang, Y.-J. Chen & C.-F. Chen (2015): *Performance Analysis Using Petri Net Based MapReduce Model in Heterogeneous Clusters*. In: *Advances in Web-Based Learning*, LNCS 8390, Springer, pp. 170–179, doi:10.1007/978-3-662-46315-4\_18.
- [7] F. Dinu & T.S. E. Ng (2012): *Understanding the Effects and Implications of Compute Node Related Failures in Hadoop*. In: *International Symposium on High-Performance Parallel and Distributed Computing*, pp. 187–198, doi:10.1145/2287076.2287108.
- [8] S. Li, S. Hu, S. Wang, L. Su, T. Abdelzaher, I. Gupta & R. Pace (2014): *WOHA: Deadline-Aware Map-Reduce Workflow Scheduling Framework over Hadoop Clusters*. In: *IEEE International Conference on Distributed Computing Systems*, pp. 93–103, doi:10.1109/ICDCS.2014.18.
- [9] K. Ono, Y. Hirai, Y. Tanabe, N. Noda & M. Hagiya (2011): *Using Coq in Specification and Program Extraction of Hadoop Mapreduce Applications*. In: *International Conference on Software Engineering and Formal Methods*, pp. 350–365, doi:10.1007/978-3-642-24690-6\_24.
- [10] OpenCloud: <http://ftp.pdl.cmu.edu/pub/datasets/hla/dataset.html>.
- [11] J. A. Quiané-Ruiz & et al. (2011): *RAFTing MapReduce: Fast Recovery on the RAFT*. In: *IEEE International Conference on Data Engineering*, pp. 589–600, doi:10.1109/ICDE.2011.5767877.
- [12] G.S. Reddy, F. Yuzhang, L. Yang, S.D. Jin, J. Sun & R. Kanagasabai (2013): *Towards Formal Modeling and Verification of Cloud Architectures: A Case Study on Hadoop*. In: *International World Congress on Services*, pp. 306–311, doi:10.1109/SERVICES.2013.47.
- [13] M. C. Ruiz, J. Calleja & D. Cazorla (2015): *Petri Nets Formalization of Map/Reduce Paradigm to Optimise the Performance-Cost Tradeoff*. In: *IEEE Trustcom/BigDataSE/ISPA*, 3, pp. 92–99, doi:10.1109/Trustcom.2015.617.
- [14] M. Soualhia, F. Khomh & S. Tahar (2017): *Task Scheduling in Big Data Platforms: A Systematic Literature Review*. *Journal of Systems and Software* 134, pp. 170 – 189, doi:10.1016/j.jss.2017.09.001.
- [15] W. Su, F. Yang, H. Zhu & Q. Li (2009): *Modeling MapReduce with CSP*. In: *IEEE International Symposium on Theoretical Aspects of Software Engineering*, pp. 301–302, doi:10.1109/TASE.2009.28.
- [16] J. Sun, Y. Liu, J. S. Dong & C. Chen (2009): *Integrating Specification and Programs for System Modeling and Verification*. In: *IEEE International Symposium on Theoretical Aspects of Software Engineering*, pp. 127–135, doi:10.1109/TASE.2009.32.
- [17] J. Sun, Y. Liu, J. S. Dong & J. Pang (2009): *PAT: Towards Flexible Verification under Fairness*. In: *Computer Aided Verification*, LNCS 5643, pp. 709–714, doi:10.1007/3-540-10843-2\_22.
- [18] Process Analysis Toolkit: <http://sav.sutd.edu.sg/PAT/>.
- [19] Google Traces: <https://github.com/google/cluster-data>.
- [20] Facebook Traces: <https://github.com/SWIMProjectUCB/SWIM/wiki/Workloads-repository>.
- [21] WordCount Example: <http://wiki.apache.org/hadoop/WordCount>.
- [22] W. Xie, H. Zhu, X. Wu, S. Xiang & J. Guo (2016): *Modeling and Verifying HDFS Using CSP*. In: *IEEE Annual Computer Software and Applications Conference*, 1, pp. 221–226, doi:10.1109/COMPSAC.2016.158.