

A Dynamic Resource Allocation Method for Parallel Data Processing in Cloud Computing

¹Venkatesa Kumar, V. and ²S. Palaniswami

¹Department of Computer Science and Engineering,
Anna University of Technology Coimbatore, Coimbatore
Faculty of Information and Communication Engineering

²Department of Electrical Engineering,
Government College of Technology, Coimbatore
Faculty of Electrical Engineering

Abstract: Problem statement: One of the Cloud Services, Infrastructure as a Service (IaaS) provides a Compute resources for demand in various applications like Parallel Data processing. The computer resources offered in the cloud are extremely dynamic and probably heterogeneous. Nephelē is the first data processing framework to explicitly exploit the dynamic resource allocation offered by today's IaaS clouds for both, task scheduling and execution. Particular tasks of processing a job can be assigned to different types of virtual machines which are automatically instantiated and terminated during the job execution. However, the current algorithms does not consider the resource overload or underutilization during the job execution. In this study, we have focussed on increasing the efficacy of the scheduling algorithm for the real time Cloud Computing services. **Approach:** Our Algorithm utilizes the Turnaround time Utility efficiently by differentiating it into a gain function and a loss function for a single task. The algorithm also assigns high priority for task of early completion and less priority for abortions/deadlines issues of real time tasks. **Results:** The algorithm has been implemented on both preemptive and Non-preemptive methods. The experimental results shows that it outperforms the existing utility based scheduling algorithms and also compare its performance with both preemptive and Non-preemptive scheduling methods. **Conclusion:** Hence, a novel Turnaround time utility scheduling approach which focuses on both high priority and the low priority tasks that arrives for scheduling is proposed.

Key words: Cloud computing, task scheduling, resource utilization

INTRODUCTION

Cloud computing has the potential to dramatically change the landscape of the current IT industry (Armbrust *et al.*, 2009; Goldberg, 1989) For companies that only have to process large amounts of data occasionally running their own data center is obviously not an option. Instead, Cloud computing has emerged as a promising approach to rent a large IT infrastructure on a short-term pay-per-usage basis. Operators of so-called IaaS clouds, like Amazon EC2, (AWSLLC, 2011a), let their customers allocate, access and control a set of Virtual Machines (VMs) which run inside their data centers and only charge them for the period of time the machines are allocated. The VMs are typically offered in different types, each

type with its own characteristics (number of CPU cores, amount of main memory) and cost.

Since the VM abstraction of IaaS clouds fits the architectural paradigm assumed by the data processing frameworks described above, projects like Hadoop The Apache Software Foundation, 2011 (White, 2010), a popular open source implementation of Google's MapReduce framework, already have begun to promote using their frameworks in the cloud (White, 2010) Only recently, Amazon has integrated Hadoop as one of its core infrastructure services (AWSLLC, 2011b). However, instead of embracing its dynamic resource allocation, current data processing frameworks rather expect the cloud to imitate the static nature of the cluster environments (Dornemann *et al.*, 2009) they were originally designed for, e.g., at the moment the types and number of VMs allocated at the

Corresponding Author: Venkatesa Kumar, V., Department of Computer Science and Engineering,
Anna University of Technology Coimbatore, Coimbatore

beginning of a compute job cannot be changed in the course of processing, although the tasks the job consists of completely different demands on the environment. As a result, rented resources may be inadequate for big parts of the processing job, which may lower the overall processing performance and increase the cost.

One of an IaaS cloud's key feature is the provisioning of compute resources on demand. The computer resources available in the cloud are highly dynamic and possibly heterogeneous. Nephele is the first data processing framework to explicitly exploit the dynamic resource allocation offered by today's IaaS clouds for both task scheduling and execution. Particular tasks of a processing a job can be assigned to different types of virtual machines which are automatically instantiated and terminated during the job execution.

Nephele is the first data processing framework to explicitly exploit the dynamic resource allocation offered by today's IaaS clouds for both, task scheduling and execution. Particular tasks of a processing job can be assigned to different types of virtual machines which are automatically instantiated and terminated during the job execution.

While there exist different interpretations and views on cloud computing (Armbrust *et al.*, 2009) it is less disputable that being able to effectively exploit the computing resources in the clouds to provide computing service at different quality levels is essential to the success of cloud computing. For real-time applications and services, the timeliness is a major criterion in judging the quality of service. Due to the nature of real-time applications over the Internet, the timeliness here refers to more than the deadline guarantee as that for hard real-time systems. In this regard, an important performance metric for cloud computing can thus be the sum of certain value or utility that is accrued by processing all real-time service requests.

To improve the performance of cloud computing, one approach is to employ the traditional Utility Accrual (UA) approach first proposed to associate each task with a Time Utility Function (TUF), which indicates the task's importance. Specifically, the TUF describes the value or utility accrued by a system at the time when a task is completed (Li *et al.*, 2006). Based on this model, there have been extensive research results published on the topic of UA scheduling. While Jensen's definition of TUF allows the semantics of soft time constraints to be more precisely specified, all these variations of UA-aware scheduling algorithms imply that utility is accrued only when a task is successfully completed and the aborted tasks neither increase nor decrease the accrued value or utility of the system.

We believe that, to improve the performance of cloud computing, it is important to not only measure the profit when completing a job in time, but also account for the penalty when a job is aborted or discarded. Note that, before a task is aborted or discarded, it consumes system sources including network bandwidth, storage space and processing power and thus can directly or indirectly affect the system performance. This is especially true for cloud computing in considering the large possibility of migration of a task within the clouds for reasons such as the economy considerations (Casati and Shan, 2001). If a job is deemed to miss its deadline with no positive semantic gain, a better choice should be one that can detect it and discard it as soon as possible.

Recently, (Yu *et al.*, 2010) proposed a task model that considers both the profit and penalty that a system may incur when executing a task. According to this model, a task is associated with two different TUFs, a profit TUF and a penalty TUF. The system takes a profit (determined by its profit TUF) if the task completes by its deadline and suffers a penalty (determined by its penalty TUF), if it misses its deadline or is dropped before its deadline. It is tempting to use negative values for the penalties and thus combine both TUFs into one single TUF. However, a task can be completed or aborted and hence can produce either a profit value or a penalty value. Mathematically, if there exists such a single function, it would imply that a single value in its domain was mapped to two values in its range, violating that it is a function. Therefore, one utility function cannot accurately represent both the profit and penalty information when executing a task. There are also some other penalty related models proposed in the literature. For example, studied the on-line scheduling problem when penalties have to be paid for rejected jobs. This model, however, does not account for the penalty to drop the task before its deadline.

However Nephele does not consider resource overload or underutilization during the job execution automatically. In this study, a novel Turnaround time utility algorithm is proposed for scheduling the real-time cloud computing services. The most unique characteristics of this approach is that, different from traditional utility accrual approach that works under one single Time Utility Function (TUF), which have two different functions called a Gain and a loss Functions-associated with each task at the same time, to model the real-time applications for cloud computing. To compare the performance of cloud computing, the traditional Utility approach is deployed in both Non-Preemptive and Preemptive scheduling.

This study includes further details on scheduling strategies and extended experimental results. The study is structured as follows: First it starts with describing the basic concept of cloud and present the architecture of the Nephelê and outline how jobs can be described and executed in the cloud. Followed by our scheduling approach in explained in detail. Then we present the experiment setup used for the evaluation and discuss the results. Finally, we conclude the study.

MARERIALS AND METHODS

Cloud computing: The cloud is a metaphor for the Internet and is an abstraction for the complex infrastructure it conceals. Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. It define in three models Software as a Service (SaaS), Platform as a Service (PaaS), Infrastructure as a Service (IaaS). Figure 1 shows the architecture of the cloud computing. Cloud computing system scales applications by maximizing concurrency and using computing resources more efficiently One must optimize locking duration, statelessness, sharing pooled resources such as task threads network connections bus, cache reference data and partition large databases for scaling services to a large number of users. IT companies with innovative ideas for new application services are no longer required to make large capital outlays in the hardware and software infrastructures. By using clouds as the application hosting platform, IT companies are freed from the trivial task of setting up basic hardware and software infrastructures. Thus they can focus more on innovation and creation of business values for their application services. Some of the traditional and emerging Cloud- based application services include social networking, web hosting, content delivery and real time instrumented data processing.




Fig. 1: Cloud Computing Architecture

Each of these application types has different composition, configuration and deployment requirements. Quantifying the performance of provisioning (scheduling and allocation) policies in a real Cloud computing environment (AEC2, 2011), Microsoft Azure, 2011 Google App Engine (2011) for different application models.

Cloud computing also describes applications that are extended to be accessible through the Internet. These cloud applications use large data centers and powerful servers that host Web applications and Web services. Anyone with a suitable Internet connection and a standard browser can access a cloud application.

Task scheduling and load-balancing technique: A task is a (sequential) activity that uses a set of inputs to produce a set of outputs. Processes in fixed set are statically assigned to processors, either at compile-time or at start-up (i.e., partitioning) avoids overhead of load balancing using these load-balancing algorithms. The Grid computing algorithms can be broadly categorized as centralized or decentralized, dynamic or static or the hybrid policies in latest trend. A centralized load balancing approach can support larger system. Hadoop system takes the centralized scheduler architecture. In static load balancing, all information is known in advance and tasks are allocated according to the prior knowledge and will not be affected by the state of the system. Dynamic load-balancing mechanism has to allocate tasks to the processors dynamically as they arrive. Redistribution of tasks has to take place when some processors become overloaded (Zaharia *et al.*, 2009).

In cloud computing, each applications of users will run on a virtual operating systems, the cloud systems distributed resources among these virtual systems. Every application is completely different and is independent and has no link between each other whatsoever, For example, some require more CPU time to compute complex task and some others may need more memory to store data. Resources are sacrificed on activities performed on each individual unit of service.

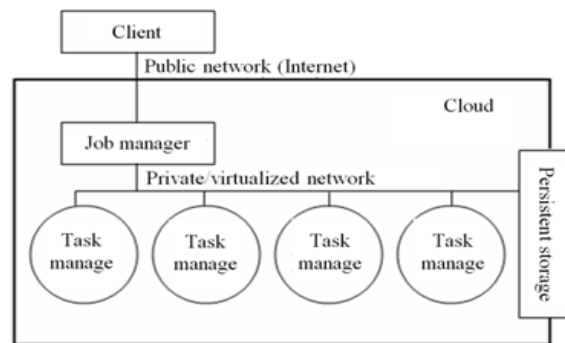


Fig. 2: Nephelê's Architecture

In order to measure direct costs of applications, every individual use of resources (like CPU cost, memory cost, I/O cost) must be measured. When the direct data of each individual resources cost has been measured, more accurate cost and profit analysis.

Overview of Nephelē architecture: Nephelē is a new data processing framework (Warneke and Kao, 2009; Ravindran *et al.*, 2005) for cloud environment that takes up many ideas of previous processing frameworks but refines them to better match the dynamic and opaque nature of a cloud.

Nephelē’s architecture follows a classic master-worker pattern as illustrated in Fig. 2.

Before submitting a Nephelē compute job, a user must start a VM in the cloud which runs the so called Job Manager (JM). The Job Manager which receives the client’s jobs, is responsible for scheduling them and coordinates their execution. It is capable of communicating with the interface the cloud operator provides to control the instantiation of VMs. We call this interface the Cloud Controller. By means of the Cloud Controller the Job Manager can allocate or deallocate VMs according to the current job execution phase. We will comply with common Cloud computing terminology and refer to these VMs as instances for the remainder of this study. The term instance type will be used to differentiate between VMs with different hardware characteristics. For example, the instance type “m1.small” could denote VMs with one CPU core, one GB of RAM and a 128 GB disk while the instance type “c1.xlarge” could refer to machines with 8 CPU cores, 18 GB RAM and a 512 GB disk.

The actual execution of tasks which a Nephelē job consists of is carried out by a set of instances. Each instance runs a so-called Task Manager (TM). A Task Manager receives one or more tasks from the Job Manager at a time, executes them and after that informs the Job Manager about their completion or possible errors. Unless a job is submitted to the Job Manager, we expect the set of instances (and hence the set of Task Managers) to be empty. Upon job reception the Job Manager then decides, depending on the job’s particular tasks, how many and what type of instances the job should be executed on and when the respective instances must be allocated/deallocated to ensure a continuous but cost-efficient processing. The newly allocated instances boot up with a previously compiled VM image. The image is configured to automatically start a Task Manager and register it with the Job Manager. Once all the necessary Task Managers have successfully contacted the Job Manager, it triggers the execution of the scheduled job.

Initially, the VM images used to boot up the Task Managers are blank and do not contain any of the

data the Nephelē job is supposed to operate on. As a result, we expect the cloud to offer persistent storage (like, e.g., Amazon S3 (Amazon Web Services)). This persistent storage is supposed to store the job’s input data and eventually receive its output data. It must be accessible for both the Job Manager as well as for the set of Task Managers, even if they are connected by a private or virtual network.

On-line non-preemptive utility:

Efficiency Scheduling: The on-line non-preemptive scheduling method which is used to maximize the efficiency gain. Since the execution of a task may gain positive profit or suffer penalty and thus degrade the overall computing performance, judicious decisions must be made with regard to executing a task, dropping or aborting a task and when to drop or abort a task. The rationale of our approach is very intuitive, i.e. a task can be accepted and executed only when it is statistically promising to bring positive gain and discarded or aborted otherwise. Before we introduce the details of our scheduling approach, we first introduce two useful concepts, the expected gain utility and the critical point.

The expected gain utility and the critical point: Since the task execution time is not known deterministically, we do not know if executing the task will lead to positive gain or loss. To solve this problem, we can employ a metric, i.e., the expected gain utility, to help us make the decision.

Given a task T with arrival time of a_i , let its predicted starting time be t_i . Then the potential Gain $P_i(t_i)$ to execute T can be represented as the integration of the summation of gain over time t_i and the difference of the starting time of the process and the arrival time of the process.

$$P_i(t_i) = \int_{\max}^{c_i - (t_i - at_i)} P_i(t + (t_i - at_i)) f_i(t) dt$$

Similarly, the potential loss $(\overline{Li}(T))$ to execute T_i can be represented as:

$$Li(T) = Li(D) \int_{c_i - (t_i - at_i)}^{\max} f_i(t) dt$$

Therefore, the expected increased efficiency $\eta(T)$ to execute T_i can be represented as:

$$\eta(T) = P_i(t_i) - Li(t_i)$$

A task can be accepted or chosen for execution when $\eta(T) > 0$, which means that the probability of to obtain positive gain is no smaller than that to incur a loss. We can further limit the task acceptance by imposing a threshold (δ) to the expected accrued utility, i.e. a task is accepted or can be chosen for execution if:
 $P_i(T) \geq \mu$

We call μ as the Efficiency threshold: Furthermore, since the task execution time is not known a priori, we need to decide whether to continue or abort the execution of a task. The longer we execute the task, the closer we are to the completion point of the task. At the same time, however, the longer the task executes the higher penalty the system has to endure if the task cannot meet its deadline. To determine the appropriate time to abort a task, we employ another metric, i.e., the critical point.

Let task T_i starts its execution at t_i , then the potential profit $T_i > t$ (i.e., $\eta(T)$) can be represented as the integration of the maximum gain and the difference of the completion of the task. The Potential loss of a function can be calculated by the integration of its completion time to the max time. Hence, the expected efficiency η is the difference between the gain of a task and the loss of a task.

If we substitute η to be equal to 0, we can see that the gains & loss are found to be equal in executing a task. As time increases, the η decrease and after a critical point at deadline more loss incurs than gain.

Algorithm 1: For non-preemptive scheduling:

Consider K accepted Task in Ready Queue and the Current Time t.
 Parameters

- 1: Accepted Task in the Queue Level. Let $\{ t_1, t_2, \dots, t_k \}$ be the Arrival Time A_T [$T=1$ to K]
- 2: Let Currently Running Task may be at $T=0$. Show the task with T and the Threshold Value T_h $A_T = A_0$.
- 3: Conditions The Current Job is in Critical, Then Abort the execution of T_0
- 4: Otherwise New Task enrolled in the end process.
- 5: Calculation of efficiency of task and reschedule the task based on the Utility value and load into the ready Queue.
- 6: Start the Execution from T_1 . The utility value is less than the Threshold value then remove the process from ready queue else the current process and start its execution

The scheduling algorithm: Our scheduling algorithm works at scheduling points that include: the arrival of a new task, the completion of the current task and the critical point of the current task. The detailed algorithm is described in algorithm 1.

In algorithm 1, when the time reaches the critical point of the current task, the current active task is immediately discarded and the task with the highest expected efficiency is selected to be executed. Upon the finish of the current task, the task with the highest expected efficiency is selected for execution. After the selection of the new task in both of the two cases, the expected efficiency for the rest of the tasks are recalculated. The tasks with the expected efficiency smaller than the threshold value are discarded.

Algorithm 2: Sort the ready queue based on the recalculated expected gain

- 1: Input: Let $Tr = \{ t_1, t_2, \dots, t_k \}$ be the accepted tasks in the ready queue, let $tr_i, i = 1, \dots, k$ represent their specific arrival times. Let current time be t and T_0 be the task currently being executed.
- 2: Output: The list of tasks in the ready queue $T'r = \{ T'1, T'2, \dots, T'k \}$ sorted based on their expected gain.
- 3: $T_{start} =$ expected finishing time of $T_0 - t$;
- 4: for $i=0$ to k do
- 5: $T'i = T_j$ where $T_j \in Tr$ is the task with the largest expected gain assuming it starts at T_{start} ;
- 6: Remove T_j from Tr ;
- 7: $T_{start} = T_{start} +$ expected execution time of $T'j$;
- 8: Calculate the following task's expected utility at time T_{start} ;
- 9: end for

When a new job comes, it is first inserted at the head of the ready queue, assuming its expected starting time would be the expected finishing time of the current active task. Based on this starting time, we then can compare its expected utility with the rest of the tasks in the queue. If its expected utility is less than that of the one following it, we re-insert this job to the queue according to its new expected utility. We calculate the new expected utility according to Algorithm 2, by estimating its new expected starting time as the sum of the expected executing time of the leading tasks' in the ready queue. This procedure continues until the entire ready queue becomes a list ordered according to their expected utilities. We remove the ones with expected utility lower than the threshold.

On-line preemptive scheduling: The Preemptive scheduling algorithm belongs to a new family of real-time service oriented scheduling problems. As the complementarily of our previous non-preemptive algorithm (Liu *et al.*, 2010), real time tasks are scheduled preemptively with the objective of maximizing the total utility time.

The preemptive scheduling heuristics is to judiciously accept, schedule and cancel real-time services when necessary to maximize the efficiency. The new scheduling algorithm has much better performance than an earlier scheduling approach based on a similar model does.

Algorithm 3: On-line preemptive efficiency scheduling method:

- 1: Input: Let $\{T_1, T_2, \dots, T_k\}$ be the accepted tasks in the ready queue and let e_i be the expected execution time of T_i . Let current time be t and let T_0 be the task currently being executed. Let the expected utility density threshold be μ .
- 2:
- 3: if a new task, i.e., T_p arrives then
- 4: Check if T_p should preempt the current task or not;
- 5: if Preemption allowed then
- 6: T_p preempts the current task and starts being executed;
- 7: end if
- 8: if Preemption not allowed then
- 9: Accept T_p if $\frac{U_p(e_0)}{e_0} > \mu e$;
- 10: Reject T_p if $\frac{U_p(e_0)}{e_0} > \frac{\mu e_0}{e_0}$;
- 11: end if
- 12: Remove T_j in the ready queue if $\frac{U_j(e_0)}{e_j} > \mu$;
- 13: end if
- 14:
- 15: if At preemption check point then
- 16: PREEMPTION CHECKING;
- 17: end if
- 18:
- 19: if T_0 is completed then
- 20: Choose the highest expected utility density task t_i to run.
- 21: Remove T_j in the ready queue if $\frac{U_j(e_i)}{e_j} > \mu$;
- 22: end if
- 23:
- 24: if $t =$ the critical time of τ_0 then
- 25: Abort τ_0 immediately;
- 26: Choose the highest expected utility density task τ_i to run.
- 27: Remove τ_j in the ready queue if $\frac{\bar{U}_j(\bar{C}_i)}{C_j} \leq \delta$;
- 28: end if

The details of our scheduling are described in algorithm 3. There are five main parts in the scheduling. They are the preemption checking, feasibility checking, task selecting, scheduling point checking and critical point checking. When new tasks are added in to ready queue, no matter whether there is preemption or not, the feasibility checking will work to check if the new ready queue is feasible or not. If any task cannot meet the requirement, it will be removed from the ready queue. Scheduling point checking makes sure all the left tasks in the expected accrued utility density task to run when the server is idle. The critical point checking will always monitor the current running task's state to prevent the server wasting time on the non-profitable running task. The preemption checking works when there is a prosperous task wants to preempt the current task. The combination of these parts guarantees to judiciously schedule the tasks for achieving high accumulated total utilities. It is worthy to talk more about the preemption checking part in details, because improper aggressive preemption will worsen the scheduling performance. From Algorithm.4 we can see that if a task can be finished successfully before its deadline even in its worst case, the scheduling will protect the current running task from being preempted by any other tasks. Otherwise, if a prosperous task has an expected accrued utility density which is larger than the current running task's conditional expected utility density by at least a value equals to the pre-set preemption threshold, the preemption is permitted.

Algorithm 4: Verification of Preemptive method

- 1: Input: Let T_0 be the task currently being executed and T_p be the task wants to preempt T_0 , current time be t , $U(T_0, t)$ be the conditional expected utility density of T_0 at time t , e_0 is the remaining expected time of T_0 . $\bar{U}_p(t)$ be the expected utility density of T_p ;
- 2: if the expected density is greater, then
- 3: Check what is T_0 's worst case finish time;
- 4: if T_0 can be finished before its deadline even in the worst case then
- 5: Preemption is not allowed;
- 6: end if
- 7: if T_0 's worst case will miss as its deadline then
- 8: Preemption allowed;
- 9: end if
- 10: end if

The feasibility check is one more part deserves detail description. In this part, scheduling simulates the real execution sequence for the left tasks in ready queue and check following this sequence, if all of them can satisfy the requirement or not. The thing needs to be

discussed is how to determine the sequence of the left tasks. From equation (1), (2) and (3), we can clearly see that the expected utility of running a task depends heavily on variable T , i.e., the time when the task can start. If we know the execution order and thus the expected starting time for tasks in the ready queue, we will be able to quantify the expected utility density of each task more accurately. In algorithm.5, we show our utility metric based on a speculated execution order of the tasks in the ready queue.

The general idea to generate the speculated execution order is as follows. We first calculate the expected utility density for each task in the ready queue based on the expected finishing time to the current running task. Then the task with the largest one is assumed to be the first task that will be executed after the current task is finished. Based on this assumption, we then calculate the expected utilities for the rest of the tasks in the ready queue and select the next task. This process continues until all tasks in the ready queue are put in order. When completed, we essentially generate a speculated execution order for the tasks in the ready queue and, at the same time, calculate the corresponding expected utility density for each task.

Algorithm 5: Verification of Preemptive method

- 1: Input: Let $T = \{t_1, t_2, \dots, t_k\}$ be the accepted tasks in the ready queue, let r_i, e_i represent the arrival time and expected execution time of t_i . Let the current time be t .
- 2: Output: The new list $T' = \{t'_1, t'_2, \dots, t'_k\}$ with the speculate execution order and their corresponding expected utility density $\hat{U}(T'_j)$ for $\tau'_j, 1 \leq j \leq k$.
- 3: if a task T_0 is being executed then
- 4: $T = r_0 + e_0$;
- 5: else
- 6: $T = t$;
- 7: end if
- 8: While T is not empty list do
- 9: for Each task e in T do the list
- 10: Calculate gain based on Eq. 1-3;
- 11: end for
- 12: Select t_j with the highest priority;
- 13: Add t_j to the end T' the list;
- 14: Execution time = Execution time(t)+executing time for task.
- 15: Remove t_j from list;
- 16: end

The investigations and the comparison of the performance of the algorithms have been studied and simulated under a variety of conditions.

RESULTS

To evaluate the performance of the algorithms, certain investigations were done using different experimental setup.

Experiment set up: The test cases in our experiments were randomly generated. Specifically, B, W and D were randomly generated such that they are uniformly distributed within interval of $[1, 10], [30, 50]$ and $[40, 50]$, respectively. The execution time of a task is assumed to be evenly distributed between interval of

$$[B, W], \text{ i.e. } f(t) = \frac{1}{W-B}$$

G, L were assumed to be linear functions, i.e., $G(t) = -a_g(t - D)$ in the range of $[0, D]$ and $L(t) = a_l t$. The gradient for $G(t)$ and $L(t)$, i.e., a_g and a_l were randomly picked from the interval of $[4, 10]$ and $[1, 5]$, respectively. Task release times' intervals follow the exponential distribution with $\mu = 2$. The utility threshold μ is set to 0. We conducted three different groups of experiments to study and compare the performance of different approaches under different conditions. The results are reported as follows.

DISCUSSION

We first constructed 1000 task sets, each of which consists of 20 tasks. Figure 1-3 plot the accrued utility, accrued profit, as well as the accrued penalty for three different approaches: Non Pre-emptive, Preemptive and Nephela. For ease of presentation, we only show 50 sets of results in the figures. The horizontal axis is the index of the experiment sets.

The graph in Fig. 3 shows that the variation between the total utility and the experiment sets. From that we can know, how the values of total utility will increased to the corresponding values of experiment sets. And this graph shows that, preemptive results are having higher total utility than the Non preemptive and execution graph of the Nephela.

The graph in Fig. 4 shows that the variation between the total profit and the experiment sets. From that we can know, how the values of total profit will increased to the corresponding values of experiment sets. And this graph shows that, preemptive results are having higher profit than the Non preemptive and execution graph of the Nephela.

This graph in Fig. 5 shows that the variation between the total penalty and the experiment sets. From that we can know, how the values of total penalty will increased to the corresponding values of experiment sets. And this graph shows that, Non preemptive results are having higher profit than the preemptive and execution graph of the Nephela.




Fig. 3: Utility gain




Fig. 4: Profit gain




Fig. 5: Penalty cause

CONCLUSION

The popularity of the Internet has grown enormously, which has presented a great opportunity for providing real-time services over the Internet. We have discussed the challenges and opportunities for efficient parallel data processing (Chaiken *et al.*, 2008) in cloud environments and presented Nephela, the first data processing framework to exploit the dynamic resource provisioning offered by today's IaaS clouds. We have described Nephela's basic architecture and presented a performance comparison to the well-established data processing framework Hadoop. The performance evaluation gives a first impression on how the ability to assign specific virtual machine types to specific tasks of a processing job, as well as the

possibility to automatically allocate/deallocate virtual machines in the course of a job execution, can help to improve the overall resource utilization and, consequently, reduce the processing cost. The on-line real-time service system should be compatible with preemption in respect that it is necessary and befitting for nowadays' service requests. Our experimental results clearly show that our proposed preemptive scheduling algorithm is effective in this regard.

In this study, we present a novel Turnaround time utility scheduling approach which focuses on both the high priority and the low priority tasks that arrive for scheduling. This study can be viewed as the extended version of Nephela (Warneke and Kao, 2011). It is also a significant improvement compared to non-preemptive scheduling (Liu *et al.*, 2010) in which, the preemptive approaches better than the non-preemptive counterpart. Our extensive experimental results clearly show that our proposed preemptive method can outperform the non-preemptive approach.

REFERENCES

- AEC2, 2011. Amazon Elastic Compute Cloud. Amazon EC2.
- Armburst, M., A. Fox, R. Griffith, A.D. Joseph and R. Katz *et al.*, 2009. Above the clouds: A Berkeley view of cloud computing. Miscellaneous.
- AWSLLC 2011a. Amazon elastic MapReduce. Amazon Web Services LLC.
- AWSLLC, 2011b. Amazon simple storage service. Amazon Web Services LLC.
- Casati, F and Shan, M., 2001. Definition, Execution, Analysis and Optimization of Composite E-Services. IEEE Data Eng. Bull.
- Chaiken, R., B. Jenkins, P. Larson, B. Ramsey and D. Shakib *et al.*, 2008. SCOPE: Easy and efficient parallel processing of massive data sets. Proc. VLDB Endowment, 1: 1265-1276. DOI: 10.1145/1454159.1454166
- Dornemann, T., E. Juhnke and B. Freisleben, 2009. On-demand resource provisioning for BPEL workflows using amazon's elastic compute cloud. Proceedings of the 9th IEEE/ACM International Symposium Cluster Computing and the Grid, May 18-21, IEEE Xplore Press, Shanghai, pp: 140-147. DOI: 10.1109/CCGRID.2009.30
- Goldberg, D.E., 1989. Genetic Algorithms in Search, Optimization and Machine Learning. 1st Edn., Addison-Wesley, Reading, MA., ISBN: 0201157675, pp: 412.

- Li, P, H. Wu, B. Ravindran and E. D. Jensen, April 2006. A utility accrual scheduling algorithm for real-time activities with mutual exclusion resource constraints. *IEEE Trans. Comput.*, 55: 454-469. DOI: 10.1109/TC.2006.47
- Liu, S., G. Quan and S. Ren, 2010. On-line scheduling of real-time services for cloud computing. *Proceedings of the 6th World Congress on Services*, Jul. 5-10, IEEE Xplore Press, Miami, FL., pp: 459-464. DOI: 10.1109/SERVICES.2010.109
- Ravindran, B., E.D. Jensen and P. Li, 2005. On recent advances in time/utility function real-time scheduling and resource management. *Proceedings of the 8th IEEE International Symposium on Object-Oriented Real-Time Distributed Computing*, May 18-20, IEEE Xplore Press, USA., pp: 55-60. DOI: 10.1109/ISORC.2005.39
- Warneke, D and O. Kao, 2009. Nephele: Efficient parallel data processing in the cloud. *Proceedings of the 2nd Workshop Many-Task Computing on Grids and Supercomputers*, Nov. 14-20, ACM, Portland, OR, USA., pp: 1-10. ISBN: 978-1-60558-714-1 DOI: 10.1145/1646468.1646476
- Warneke, D and O. Kao, 2011. Exploiting Dynamic resource allocation for efficient parallel data processing in the cloud. *IEEE Trans. Parallel Distributed Syst.*, 22: 985-997. DOI: 10.1109/TPDS.2011.65
- White, T., 2010. *Hadoop: The Definitive Guide*. 2nd Edn., O'Reilly Media, Beijing, ISBN: 1449389732, pp: 600.
- Yu, Y., S. Ren, N. Chen and X. Wang, 2010. Profit and penalty aware (PP-aware) scheduling for tasks with variable task execution time. *Proceedings of the 2010 ACM Symposium on Applied Computing*, Mar. 22-26, ACM, Sierre, Switzerland, pp: 334-339. DOI: 10.1145/1774088.1774159
- Zaharia, M., D. Borthakur, J.S. Sarma, K. Elmeleegy and S. Shenker *et al.*, 2009. Job scheduling for multi-user mapreduce clusters. EECS Department, University of California, Berkeley.