

Executing Bag of Distributed Tasks on the Cloud: Investigating the Trade-offs Between Performance and Cost

Long Thai, Blesson Varghese and Adam Barker
School of Computer Science, University of St Andrews, Fife, UK
Email: {lth2, varghese, adam.barker }@st-andrews.ac.uk

Abstract—Bag of Distributed Tasks (BoDT) can benefit from decentralised execution on the Cloud. However, there is a trade-off between the performance that can be achieved by employing a large number of Cloud VMs for the tasks and the monetary constraints that are often placed by a user. The research reported in this paper is motivated towards investigating this trade-off so that an optimal plan for deploying BoDT applications on the cloud can be generated. A heuristic algorithm, which considers the user’s preference of performance and cost is proposed and implemented. The feasibility of the algorithm is demonstrated by generating execution plans for a sample application. The key result is that the algorithm generates optimal execution plans for the application over 91% of the time.

I. INTRODUCTION

Bag of Tasks (BoT) refers to a collection of independent and identical tasks, which can be executed in sequence or in parallel. There is a subset of BoT, named **Bag of Distributed Tasks (BoDT)**, in which each task additionally requires data placed at different geographical locations for execution.




Figure 1: Difference between BoT and BoDT in executing a task

A manager and an executor are required by both BoT and BoDT for executing tasks. The key difference is illustrated in Figure 1. In BoT, a task is executed as soon as it is received by an executor. In BoDT, an executor has to request the data required by the task from a globally located source before it can be executed. This step makes scheduling the execution more complicated as the location of data sources must be included into the scheduling algorithm, which incurs an additional cost since the time required for transferring

data utilised by the task has to be taken into account in addition to the time taken for processing the task.

Centralised approaches suffer high communication costs between the executor and the data sources as all tasks are executed at one fixed location. Since a BoDT consists of independent tasks, it is possible to divide it into smaller sub-sets based on the tasks’ geographical distribution and execute each of them on an application located near their data sources. This is a decentralised approach, which can potentially reduce the communication cost.

Cloud computing offers the infrastructure to decentralise the execution of BoDT as its providers maintain multiple data centres, which are geographically distributed across the world. It is possible to rapidly deploy an application on a data centre close (in terms of network distance) to its data sources in order to reduce data transfer times.

Consider for example BoDT applications, such as Feedly¹ and Flipboard², that aggregates news items for a user by collecting them from globally distributed sources. The application can benefit from being executed on multiple Cloud data centres in a manner that will reduce the network distance between it and the data sources.

While data transfer times can be reduced in BoDT on the cloud, a different kind of a problem emerges since a user pays for the cloud resources required by the tasks. Often more resources will boost the performance of the task, but this comes at a monetary cost. Hence, there is a trade-off between the performance gain by acquiring more resources on the cloud for executing a task and the budget available for executing the task on the cloud. For example, the performance of the BoDT applications that aggregate news considered above can be maximised by gathering more resources for increasing parallelism. However, this is more expensive and may not be cost-effective.

In this paper, we set out to investigate the trade-off between performance gain and monetary cost of executing BoDT on the Cloud. To facilitate the investigation a heuristic algorithm is proposed and developed, which takes into account a user’s preference of performance over cost specific to an application and/or the user’s desired Service Level

¹<http://feedly.com/>

²<https://flipboard.com/>

Agreement (SLA). The algorithm generates a execution plan for BoDT applications based on the user preference. The feasibility of the algorithm is demonstrated using a sample user application and the key result is that for 10 out of 11 times an accurate execution plan is generated.

The contributions of the research reported in this paper are (i) providing a mathematical model of executing BoDT on the Cloud, (ii) introducing a heuristic algorithm to find the execution plan based on a user's preference, and (iii) investigating the trade-off between performance and cost in executing BoDT on the Cloud.

The remainder of this paper is organised as follows. Section II introduces the mathematical models of executing BoDT. Section III, presents a heuristic algorithm for generating an execution plan. Section IV shows the experimental evaluation of the framework. Section V highlights the related work. Section VI concludes this paper.

II. PROBLEM MODELLING

In this section, we present mathematical models for centralised and decentralised execution of BoDT, and a discussion on the trade-off between performance and cost for the decentralised execution model.

A. Centralised Execution

In the centralised execution model of BoDT on the cloud only one VM is used. The cost of this VM depends on its running time. Hence, an optimal execution model is obtained when a location to deploy a BoDT application is found such that the running time is at its minimum.

Let $T = \{t_1, t_2, \dots\}$ be the collection of tasks and $L = \{l_1, l_2, \dots\}$ be the locations of the tasks. Each task can be represented as a pair that describes its location and the size of the data, i.e. $t = (l_t, s_t)$ for $t \in T$, $l_t \in L$ and $s_t > 0$.

Let $C = \{c_1, c_2, \dots\}$ be the collection of Cloud VMs for deploying the BoDT. For simplicity, we assume that each region can only host the same type of VM. Deploying different VM types simultaneously in a region will be investigated in the future and reported elsewhere.

For the task location $l \in L$ and the VM $c \in C$, the time taken for transferring one unit of data is denoted as $trans_{l,c}$. Hence, the time taken to transfer the data of a task $t \in T$ to the VM $c \in C$ is:

$$trans_{t,c} = s_t \times trans_{l_t,c}$$

The performance of each VM $c \in C$ in the model is indicative of how long it takes to process one unit of data and is represented as $comp_c$. Since this model supports VMs of the same type in a region the performance of the VMs is identical and is denoted as $comp$. The time taken to execute a task $t \in T$ on the VM $c \in C$ is:

$$comp_{t,c} = s_t \times comp$$

The execution time of task t on VM c is:

$$exec_{t,c} = trans_{t,c} + comp_{t,c}$$

The total time taken to execute all tasks $t \in T$ on VM c is:

$$exec_{T,c} = \sum_{t \in T} exec_{t,c}$$

In order to maximise the performance, the VM should be selected so that the total execution time $exec_{T,c}$ is minimum. As the computational performance between VMs is identical, the total execution time is minimised when the data transferring time is minimised. The optimal Cloud VM c_o is selection as:

$$c_o = \arg \min_{c \in C} (exec_{T,c})$$

The running time of a Cloud VM is the sum of the total execution time and the deployment time dt , which is the time taken to start up and deploy the application on a VM. Notably, the Cloud VM is only started if and only if there is tasks assigned to it, i.e. its execution time is larger than 0. Otherwise, the Cloud VM will not be started, which means its running time is 0. Hence, the running time is calculated as:

$$running_time_{T,c} = \begin{cases} dt + exec_{T,c} & \text{if } exec_{T,c} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Cloud VMs are charged for the time blocks used and is usually by the hour (equivalent to 3600 seconds). A time block running the BoDT T on VM c is represented as:

$$tb_{T,c} = \lceil \frac{running_time_{T,c}}{3600} \rceil \quad (2)$$

The cost of each time block is identical and denoted by $unit_cost$. Therefore the total cost is

$$cost_{T,c} = tb_{T,c} \times unit_cost \quad (3)$$

This cost can be minimised by selecting a Cloud VM which takes the lowest time for executing the task.

B. Decentralised Execution

In the decentralised execution model the BoDT is divided into groups of tasks, each of which is executed by one VM.

For $c \in C$, $T_c \subset T$ denotes all tasks executed on c , we have two constraints. Firstly, the constraint that ensures that all tasks are executed, which is:

$$\cup_{c \in C} T_c = T \quad (4)$$

Secondly, the constraint that ensures that each task is executed only on one VM, which is:

$$T_i \cap T_j = \emptyset \quad \text{for } i, j \in C \text{ and } i \neq j \quad (5)$$

The execution time for each VM is:

$$exec_c = exec_{T_c} = \sum_{t \in T_c} exec_{t,c}$$

The *running_time* of each VM is calculated similar to Equation 1. Hence, the total number of used time blocks is:

$$tb_{T,C} = \sum_{c \in C} \lceil \frac{running_time_c}{3600} \rceil \quad (6)$$

Since all VMs run in parallel, the overall time taken for executing all tasks is:

$$overall_exec = \max_{c \in C} exec_c \quad (7)$$

As stated in Equation 7, the *overall_exec* depends on the VM with **highest execution time**. In other words, it does not matter if one VM manages to finish its execution very quickly, as long as the other are still running. Hence, it is beneficial to balance the tasks between VMs so that they can finish approximately at the same time.

C. Discussion

Performance of the execution of BoDT on the Cloud can be maximised by minimising the overall execution time presented in Equation 7. This can be achieved by using multiple Cloud VMs so that many tasks can be executed in parallel and the load of task execution on each VM is reduced. Cloud VMs are charged by the hour and therefore using a large number of VMs may not be cost effective if they are all required only for a small fraction of the hour. Hence, tasks must be scheduled such that the total number of charged time blocks as shown in Equation 6 are minimised. Moreover, data transfer times can be reduced by assigning task to Cloud VMs closest to the data source. However, data can be distributed all over the world which may result in a VM being assigned more tasks than others. This problem can be circumvented with the help of a balancing mechanism which will be presented in the next section.

There is a trade-off in minimising the total cost and the overall execution times shown in Equation 6 and Equation 7. A user must decide whether the BoDT application is executed for maximum performance, or for minimum costs, for optimality by taking both reasonable performance and cost effectiveness into account. If β is a factor that represents the importance of minimising overall execution time, and $0 \leq \beta \leq 1$, then the optimal plan satisfies the unity constraints shown in Equations 4, 5 and minimises:

$$score = \beta \times \overline{overall_exec} + (1 - \beta) \times \overline{tb_{T,C}} \quad (8)$$

The overall execution times and number of time blocks of each instance are normalised so that both of them have the same scale.

III. HEURISTIC ALGORITHMS

Minimising Equation 8 while satisfying all the constraints is a constraint programming problem. Although solving this problem is able to produce an optimal solution, it is expensive and can take a considerable amount of time until a final answer is found. Hence, in this section, a heuristic algorithm is proposed in order to find a plan based on given β value in a more efficient way.

A. Nearest Plan

Due to the homogeneity of cloud resources, the time to process the same amount of data between all locations should be identical. The difference in execution time comes from transferring data from its source to a VM. Intuitively, each task should be executed at the location nearest to its data source in order to minimise the communication time, thus minimising the total execution time. The plan in which all tasks are assigned to locations nearest to their data sources is **nearest plan** and presented in Algorithm 1.

Algorithm 1 Get Nearest Plan

```

1: function GET_NEAREST_PLAN( $T, C, trans$ )
2:   // create an empty list of tasks for each cloud VMs
3:   for  $c \in C$  do
4:      $T_c \leftarrow \emptyset$ 
5:   end for
6:   for  $t \in T$  do
7:      $c = \arg \min_{c \in C} trans_{t,c}$  // find the nearest VM
8:      $T_c \leftarrow T_c \cup \{t\}$ 
9:   end for
10:   $T_n \leftarrow \{T_c \text{ for } c \in C\}$ 
11:  return  $T_n$ 
12: end function

```

B. Cost Effective Plan

As mentioned in Section II-C, using as many VMs as possible is not cost effective, hence we develop an algorithm which reduces the number of charged time blocks of the *nearest_plan*. Algorithm 2 aims to find the **cost effective plan** which uses the minimum number of time blocks. The number of time blocks are minimised by filling each of them with as many tasks as possible.

Besides the distance between VMs and tasks, Algorithm 2 takes as input the nearest plan produced by Algorithm 1 and the minimum number of time blocks allowed (Line 1). The most cost-ineffective VM whose execution time has the smallest fraction of an hour is selected (Line 3). The selected VM has the highest charged but unused time.

Algorithm 2 Reduce Time Block

```
1: function REDUCE_TB( $trans, T_n, min\_tb$ )
2:    $T_{ce} \leftarrow T_n$ 
3:    $T_{c_0} \leftarrow \arg \min_{T_c \in T_{ce}} (running\_time_{T_c} \bmod 3600)$ 
4:    $t_0 \leftarrow \arg \max_{t \in T_{c_0}} exec_{t, c_0}$ 
5:    $C_1 \leftarrow (C - \{c_0\})$  ordered by  $trans_{t_0, c}$ 
6:   for  $c_1 \in C_1$  do
7:     if  $t_0$  is never moved from  $c_0$  to  $c_1$  then
8:        $tb'_{c_1} := \lceil \frac{running\_time_{c_1} + exec_{t_0, c_1}}{3600} \rceil$ 
9:       if  $tb'_{c_1} = tb_{c_1}$  then
10:         $T'_{c_0} \leftarrow T_{c_0} - t_0$ 
11:         $T'_{c_1} \leftarrow T_{c_1} \cup \{t_0\}$ 
12:         $T_{ce} \leftarrow (T_{ce} - \{T_{c_0}, T_{c_1}\}) \cup \{T'_{c_0}, T'_{c_1}\}$ 
13:        if  $tb_T > min\_tb$  then
14:          go to 3
15:        end if
16:      end if
17:    end if
18:  end for
19:  return  $T_{ce}$ 
20: end function
```

The task with the highest execution time is selected from the chosen VM (Line 4). Next, the remaining cloud VMs are sorted in descending order based on their distance to the selected task's data (Line 5). Sorting is performed to assign tasks to the closest possible VM, thus minimising the execution time.

Then, the remaining sorted VMs are iterated through in order (Line 6). A constraint in line 7 is added in order to avoid an infinite loop in which a task is moved back and forth between 2 VMs.

As the algorithm aims to lower the number of time blocks by moving tasks from one VM to another, the number of time blocks used by the VM receiving a new task must not increase. In order to ensure this, the new number of time blocks if a task is assigned to the VM (Line 8) is compared against the previous number of time blocks (Line 9).

If a VM satisfying all the constraints is found, the task is removed from the cost-ineffective VM (Line 10) and added to the selected VM (Line 11). Finally, the current plan, including the list of VMs and their assigned tasks, is updated (Line 12).

The aim is to find a plan that reflects the trade-off between performance and cost instead of a plan that has the lowest number of time blocks. The value min_tb is used to limit the number of time blocks to accommodate the cost that a user is willing to pay. The number of time blocks of the nearest plan is the upper bound to min_tb with a lower bound of 1, assuming there are always tasks to be executed.

By de-assigning a task from its current VM, the execution time is reduced, and consequentially its used time blocks. It

is possible to remove all tasks from a certain VM, i.e the VM is not used any more. So the algorithm reduces not only the number of charged time blocks but also the number of VMs. This is beneficial as each VM has an initial deployment time that is charged but cannot be used for executing tasks. This means that if a user makes use of more VMs, there is likely to be more unused time.

C. Balancing Algorithm

Algorithm 3 Balancing Algorithm

```
1: function BALANCE( $trans, T_{ce}$ )
2:    $T_b \leftarrow T_{ce}$ 
3:    $T_{c_0} \leftarrow \arg \min_{T_c \in T_b} exec_{T_c}$ 
4:    $t_0 \leftarrow \arg \max_{t \in T_{c_0}} exec_{t, c_0}$ 
5:    $C_1 \leftarrow (C - \{c_0\})$  ordered by  $trans_{t_0, c}$ 
6:   for  $c_1 \in C_1$  do
7:     if  $t_0$  is never moved from  $c_0$  to  $c_1$  then
8:        $running\_time'_{c_1} := running\_time_{c_1} +$ 
9:          $exec_{t_0, c_1}$ 
10:      if  $running\_time'_{c_1} < running\_time_{c_0}$ 
11:        then
12:           $T'_{c_0} \leftarrow T_{c_0} - t_0$ 
13:           $T'_{c_1} \leftarrow T_{c_1} \cup \{t_0\}$ 
14:           $T_b \leftarrow (T_b - \{T_{c_0}, T_{c_1}\}) \cup \{T'_{c_0}, T'_{c_1}\}$ 
15:          go to 3
16:        end if
17:      end if
18:    end for
19:  return  $T_b$ 
20: end function
```

Since Algorithm 1 and Algorithm 2 selects a VM based on the network distance, it is possible that the tasks are not evenly distributed on the VMs.

Hence, the *balancing algorithm* presented in Algorithm 3 is employed to balance tasks on all VMs and make their running times almost similar. It is quite similar to Algorithm 2 which reduces the number of time blocks, but the cloud VMs are selected differently. VMs with most wasted charged time is selected for reducing the number of time blocks, while the VMs with highest running time is selected in the balancing algorithm (Line 3).

Furthermore, the algorithm has to ensure that moving tasks does not make the receiving VM's execution time higher than a giving one's (Line 9). In other words, the balancing process aims to reduce the overall running time by moving tasks from VMs with high execution time to ones with lower execution time. Moreover, it tries to move tasks to nearest VMs possible in order to further minimise overall running time.

D. Finding Execution Plan Based on Given Trade-off

A complete solution based on the previous algorithms are presented in Algorithm 4, which finds the plan taking into account a cost versus performance trade-off provided by a user.

Algorithm 4 Find Plan Algorithm

```

1: function FIND_OPTIMAL_PLAN( $T, C, trans, \beta$ )
2:    $P \leftarrow \emptyset$ 
3:    $T_n \leftarrow GET\_NEAREST\_PLAN(T, C, trans)$ 
4:   for  $min\_tb$  from 1 to  $tb_{T_n}$  do
5:      $T_{ce} \leftarrow REDUCE\_TB(trans, T_n, min\_tb)$ 
6:      $T_b \leftarrow BALANCE(trans, T_{ce})$ 
7:     if  $T_b \notin P$  then
8:        $P \leftarrow P \cup \{T_b\}$ 
9:     end if
10:  end for
11:   $max\_tb \leftarrow \max_{T_b \in P} tb_{T_b}$ 
12:   $max\_overall\_exec \leftarrow \max_{T_b \in P} overall\_exec_{T_b}$ 
13:   $T_o = \arg \min_{T_b \in P} (\beta \times \frac{overall\_exec_{T_b}}{max\_overall\_exec}$ 
     $+ (1 - \beta) \times \frac{tb_{T_b}}{max\_tb})$ 
14:  return  $T_o$ 
15: end function

```

The inputs of the algorithm are the list of tasks (i.e. their location and data size), cloud VMs, the distances between nodes and the β values. First of all, the nearest plan is created (Line 3). Then, it is used to find the list of plans with difference min_tb values running from 1 to the *nearest_plan*'s number of time blocks (Lines 4 to 10). The score value is calculated based on β and Equation 8. Finally, the plan with lowest score is returned as an optimal plan (Line 13).

Algorithm 4 can be modified in order to find an execution plan based on given budget constraints, which is converted into the min_tb value. Hence, instead of using different min_tb values, the algorithm needs to only one min_tb value, which is derived from given budget constraint, to find an optimal plan.

E. Optimal Centralised Plan

Algorithm 5 Find Optimal Centralised Plan Algorithm

```

1: function FIND_CENTRALISED_PLAN( $T, C, trans$ )
2:   for  $c \in C$  do
3:      $T_c \leftarrow T$ 
4:   end for
5:    $T_{centralised} \leftarrow \arg \min_{c \in C} exec_{T_c}$ 
6:   return  $T_{centralised}$ 
7: end function

```

Algorithm 5 compares the decentralised and centralised approaches and finds an optimal decentralised execution

plan. The idea is to assign the tasks to all VMs and select the VM with lowest execution time.

IV. EXPERIMENTAL STUDIES

A. Setup

To evaluate our approach, we developed a word count application which downloaded text files from globally distributed hosts and counted the number of words in each file without aggregating the result. There were 3290 counting tasks, each of which consisted a text file distributed in one of 47 different nodes on PlanetLab (PL), a global research network which offers over a thousand of globally distributed machines for experimentation [1]. We used Amazon Web Service (AWS)³ which offers eight different regions for cloud VMs; thus there are 8 Cloud VMs (instance). Given the list of tasks, the communication cost and the computational cost, the BoDT is divided into smaller subsets, each of which is sent to an application deployed on the AWS instance. Since AWS instance is charged by hour, each *time_block* equals to 3600 seconds.

Prior to the experiment, we measured the communication cost (i.e. $trans_{l,c}$) and the computational cost (i.e. $comp$) by running a small test in which the applications were deployed in all AWS regions and retrieved and processed data from all PlanetLab nodes.

In order to investigate the trade-off between performance and cost, instead of providing the β value and obtaining the corresponding *optimal_plan*, we executed all possible execution plans, each of which had a different number of time blocks. The *centralised_plan* was executed in order to compare the centralised and decentralised approaches.

In order to automate the experiments, we implemented a Scala framework which can find solutions based on a given input, deploy and monitor the execution.

Our framework was able to find 6 different decentralised solutions represented as $plan_i$ for $3 \geq i \geq 8$, where i represents the number of VMs required, and predicted the number of used time blocks since no instance took more than one hour to finish its execution. Hence, it was predicted that at least 3 time blocks were required to execute the BoDT. We also found a *centralised_plan* which was predicted to use 4 time blocks. Each plan was executed three times and the average results are presented.

B. General Results

Figure 2 presents the cost of executing distributed tasks on the cloud with respect to the number of time blocks used by each plan. It shows that our framework predicted the number of used time block for all plans accurately. The decentralised approach represented by $plan_3$ is cheaper than the centralised one that requires 4 hours. In other words, the benefit of reducing the network distance can outweigh the *deploy_time* suffered by each additional VM.

³<http://aws.amazon.com/>




Figure 2: Time blocks for centralised and decentralised plans

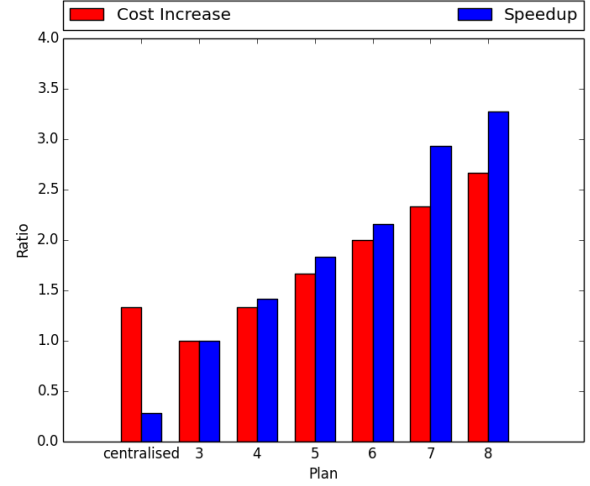


Figure 4: Increase in cost vs speedup for different plans

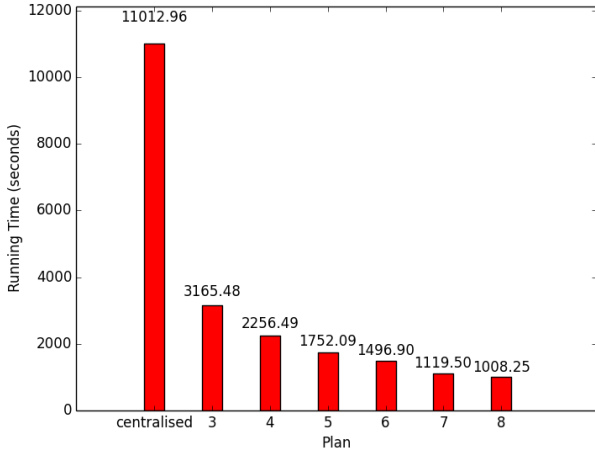


Figure 3: Running time of different plans

In order to evaluate the performance of different plans, we measured the overall running times and plotted them in Figure 3. As predicted, using more instances results lower running time. The *centralised_plan* has the highest running time, since it does not execute in parallel and most tasks are not assigned to their nearest instances.

Figure 2 and Figure 3 also highlight the trade-off between different decentralised plans; higher the cost with respect to the number of time blocks better performance with respect to execution time is obtained. Moreover, it also reflects the decision that user has to make: whether she favours performance over monetary cost or vice versa.

C. Cost vs Performance Gain

We further investigate the relationship between the increase in cost and performance, for example, whether doubling the cost would provide a twofold performance. For

this, *plan_3* is selected as the base value, the cost increase is calculated as $\frac{tb_{plan}}{tb_{plan_3}}$. The performance gain is calculated as the speedup when using *plan_3* over other plans, i.e. $speedup = \frac{running_time_{plan_3}}{running_time_{plan}}$.

Figure 4 presents the comparison between the increase in cost and performance gain. As predicted, *centralised_plan* performs worst as it does not have any speedup in spite of the increase in cost. On the other hand, for the decentralised plans, it is interesting to see that the relationship between cost increment and performance gain is not proportional. In other words, the performance gain is higher than the additional cost.




Figure 5: Average data transfer speeds for different plans

As we used cloud instances of the same type, the computational capacity is identical. However, in BoDT applications, the location of an additional instance also affects the performance gain. While moving tasks from one cloud

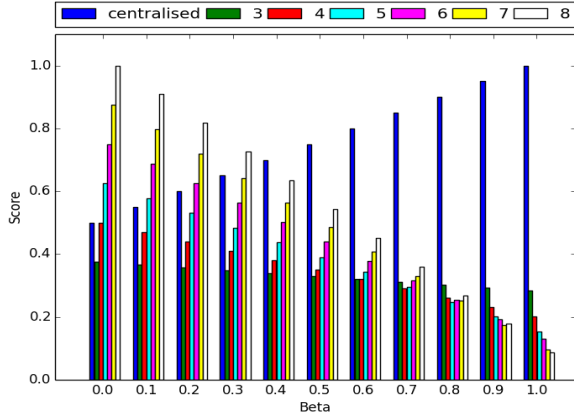


Figure 6: Scores of different β values for different plans

instance to another, the algorithm always tries to preserve data proximity, i.e. keeping the data transfer time of each task as low as possible. Figure 5 presents the average data transfer speeds for all execution plans. It shows that expensive plans (with more cloud instances) have lower data transfer speed (high data proximity). So, by employing more instances the overall performance is boosted by not only increasing computational capacity but also reducing the data transfer time.

D. Effect of β Value

Figure 6 presents the scores calculated based on Equation 8 with different values of β for all plans. It shows how β , i.e. the importance of performance over cost, affects the scores and the decision of selecting a plan based on user's preference. When $\beta = 0$, (performance is completely ignored in order to minimise cost), *plan_3* has the lowest score while *plan_8* has the highest score. As β is increased the score of *plan_8* is seen to decrease and is the lowest when $\beta = 1.0$. In all cases there is at least one solution which is better than a *centralised_plan*, i.e. decentralised approaches results in better performance and are cost-saving.

Table I is an evaluation of our approach and its accuracy of prediction of the deployment plan based on the user provided β value. The actual plan is the solution with the lowest score. The key result is that the approach makes correct predictions 10 out of 11 times.

The prediction is inaccurate in one case when $\beta = 0.6$. This is due to network instability on the PL nodes resulting in slow data transfer. Hence, VMs that download data from the PL nodes took more time to finish their execution than others. This issue will be investigated further in future work.

V. RELATED WORK

Optimising the deployment of an application on the cloud aims to select instance types, number of instances and their locations so that certain criteria such as performance

β	Prediction	Actual	Accurate
0.0	<i>plan_3</i>	<i>plan_3</i>	Yes
0.1	<i>plan_3</i>	<i>plan_3</i>	Yes
0.2	<i>plan_3</i>	<i>plan_3</i>	Yes
0.3	<i>plan_3</i>	<i>plan_3</i>	Yes
0.4	<i>plan_3</i>	<i>plan_3</i>	Yes
0.5	<i>plan_3</i>	<i>plan_3</i>	Yes
0.6	<i>plan_4</i>	<i>plan_3</i>	No
0.7	<i>plan_4</i>	<i>plan_4</i>	Yes
0.8	<i>plan_5</i>	<i>plan_5</i>	Yes
0.9	<i>plan_7</i>	<i>plan_7</i>	Yes
1.0	<i>plan_8</i>	<i>plan_8</i>	Yes

Table I: Predicted plan vs actual plan for different β values

or cost can be satisfied. The work presented in [2], [3] focused on maximising performance while minimising cost and satisfying other constraints.

Moreover, the geographical distribution of services and users are also considered while deploying an application on the cloud in order to satisfy predefined SLA [4] or minimise network distance between all nodes [5]. Similarly, [6] and [7] considered the relationship between services in order to minimise the communication costs not only from users to services but also between services. In [8], the geographical locations of web services were in into account in order to select a location to deploy workflow orchestration. Finally, in [9], the authors used edge cloud to process large amount of data based on geographical and network distances from the requesting clients, the proposed framework was able to outperform other similar systems.

Besides optimising the deployment of the application, scheduling tasks execution is also important. Data locality and proximity were investigated in [10], [11], and it was concluded that the total performance could be increased by reducing the distance between the computation and its required data. In [12], the authors used statistical methods to dynamically schedule the execution of tasks on the cloud to satisfy budget constraint and completion time goal.

Stream processing in geographically distributed environment was investigated in [13]. The authors proposed a service oriented framework to process and aggregate streaming data from globally located locations.

More recently, there are studies about the trade-off between performance and cost of executing application on the cloud. In [14], the authors investigated the usage of public cloud, which would increase the cost, to meet the application deadline. Jung et al. proposed a novel workflow scheduler which considered both the execution time and cost of cloud resources [15].

Different from existing research which use the trade-off between cost and performance as a constraint to optimise the deployment and/or execution of the application, our goal is to investigate it, i.e. comparing the relationship between the performance gain and monetary cost between different setups. Moreover, the geographical locations of data is

also taken into account as they directly affect the overall execution time. Finally, we also compare the centralised and decentralised approaches in order to emphasise the benefit of decentralising of execution on the cloud.

VI. CONCLUSIONS

Cloud computing offers the flexibility of acquiring resources without any prior commitment. Moreover, it provides different types and locations of resource for users to select from in order to maximise an application's performance. The flexibility offered by cloud computing has changed the way the industry as well as the scientific community approach their problems. However, it also introduces new questions, which is how to balance the performance gain and monetary cost of running application on the cloud.

In this paper, we investigated the trade-off between performance gain and monetary cost of executing BoDT on the cloud. Taking the location of tasks into account, we proposed a heuristic algorithm which can find where the tasks should be executed in order to satisfy user's preference of cost versus performance. The trade-off has been clearly presented in our experiment which shows the shifting between different plans when the user's preference changes. The benefit of decentralising the execution instead of running everything in the centralised location is also demonstrated.

In the future, we plan to investigate deploying multiple instances in the same region by taking a variety of cloud VMs that have different performance capabilities and costs into account. Dynamic scheduling will be investigated in order to handle undesired events, for example, network instability during execution.

ACKNOWLEDGMENTS

This research is supported by the EPSRC grant 'Working Together: Constraint Programming and Cloud Computing' (EP/K015745/1), a Royal Society Industry Fellowship 'Bringing Science to the Cloud', an EPSRC Impact Acceleration Grant (IAA) and an Amazon Web Services (AWS) Education Research Grant.

REFERENCES

- [1] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "Planetlab: An overlay testbed for broad-coverage services," *SIGCOMM Comput. Commun. Rev.*, vol. 33, pp. 3–12, July 2003.
- [2] J. Tordsson, R. S. Montero, R. Moreno-Vozmediano, and I. M. Llorente, "Cloud brokering mechanisms for optimized placement of virtual machines across multiple providers," *Future Gener. Comput. Syst.*, vol. 28, pp. 358–367, Feb. 2012.
- [3] J. L. Lucas-Simarro, R. Moreno-Vozmediano, R. S. Montero, and I. M. Llorente, "Scheduling strategies for optimal service deployment across multiple clouds," *Future Gener. Comput. Syst.*, vol. 29, pp. 1431–1441, Aug. 2013.
- [4] Q. Zhang, Q. Zhu, M. Zhani, and R. Boutaba, "Dynamic service placement in geographically distributed clouds," in *Distributed Computing Systems (ICDCS), 2012 IEEE 32nd International Conference on*, pp. 526–535, June 2012.
- [5] H. Qian and Q. Wang, "Towards proximity-aware application deployment in geo-distributed clouds," *Advances in Computer Science and its Applications*, vol. 2, no. 3, 2013.
- [6] Y. Kang, Z. Zheng, and M. Lyu, "A latency-aware co-deployment mechanism for cloud-based services," in *Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on*, pp. 630–637, June 2012.
- [7] J. Zhu, Z. Zheng, Y. Zhou, and M. Lyu, "Scaling service-oriented applications into geo-distributed clouds," in *Service Oriented System Engineering (SOSE), 2013 IEEE 7th International Symposium on*, pp. 335–340, March 2013.
- [8] M. Luckeneder and A. Barker, "Location, location, location: Data-intensive distributed computing in the cloud," in *In Proceedings of IEEE CloudCom 2013*, pp. 647–653, Dec 2013.
- [9] M. Ryden, K. Oh, A. Chandra, and J. B. Weissman, "Nebula: Distributed edge cloud for data intensive computing," 2014.
- [10] K. Ranganathan and I. Foster, "Decoupling computation and data scheduling in distributed data-intensive applications," in *Proceedings of the 11th IEEE International Symposium on High Performance Distributed Computing, HPDC '02*, (Washington, DC, USA), pp. 352–, IEEE Computer Society, 2002.
- [11] M. Zaharia, D. Borthakur, J. Sen Sarma, K. Elmeleegy, S. Shenker, and I. Stoica, "Delay scheduling: A simple technique for achieving locality and fairness in cluster scheduling," in *Proceedings of the 5th European Conference on Computer Systems, EuroSys '10*, (New York, NY, USA), pp. 265–278, ACM, 2010.
- [12] A. Opreescu and T. Kielmann, "Bag-of-tasks scheduling under budget constraints," in *Cloud Computing Technology and Science (CloudCom), 2010 IEEE Second International Conference on*, pp. 351–359, Nov 2010.
- [13] R. Tudoran, G. Antoniu, and L. Bouge, "Sage: Geo-distributed streaming data analysis in clouds," in *Proceedings of the 2013 IEEE 27th International Symposium on Parallel and Distributed Processing Workshops and PhD Forum, IPDPSW '13*, (Washington, DC, USA), pp. 2278–2281, IEEE Computer Society, 2013.
- [14] T. A. Genez, L. F. Bittencourt, and E. R. Madeira, "On the performance-cost tradeoff for workflow scheduling in hybrid clouds," in *Utility and Cloud Computing (UCC), 2013 IEEE/ACM 6th International Conference on*, pp. 411–416, Dec 2013.
- [15] G. Jung and H. Kim, "Optimal time-cost tradeoff of parallel service workflow in federated heterogeneous clouds," in *Web Services (ICWS), 2013 IEEE 20th International Conference on*, pp. 499–506, June 2013.