

Predicting the Available Bandwidth on Intra Cloud Network Links for Deadline Constrained Workflow Scheduling in Public Clouds

Rachael Shaw^(✉), Enda Howley, and Enda Barrett

College of Engineering and Informatics,
National University of Ireland, Galway, Ireland
{r.shaw4,ehowley,enda.barrett}@nuigalway.ie

Abstract. Cloud computing infrastructure has in recent times gained significant popularity for addressing the ever growing processing, storage and network requirements of scientific applications. In public cloud infrastructure predicting bandwidth availability on intra cloud network links play a pivotal role in efficiently scheduling and executing large scale data intensive workflows requiring vast amounts of network bandwidth. However, the majority of existing research focuses solely on scheduling approaches which reduce cost and makespan without considering the impact of bandwidth variability and network delays on execution performance. This work presents a time series network-aware scheduling approach to predict network conditions over time in order to improve performance by avoiding data transfers at network congested times for a more efficient execution.

Keywords: Cloud computing · Workflow scheduling · Public cloud · ARIMA modelling

1 Introduction

Data-intensive applications often modelled as workflows are routinely used throughout many fields of scientific research. Workflows play a key role in assisting scientists to orchestrate complex multi-step computational analysis on extensively large data sets. Modern day scientific workflows have advanced considerably and are becoming increasingly large, generating terabytes of data which is expected to soar over the next decade [1]. As a result, workflows require vast amounts of rich and diverse resources accessible across distributed platforms in order to address their ever growing processing, storage and network requirements.

Recently, cloud computing has emerged as a new service provisioning model which offers an alternative and more scalable solution to traditional infrastructure such as computational grids and clusters. Cloud computing services deliver

R. Shaw—This work is supported by the Irish Research Council through the Government of Ireland Postgraduate Scholarship Scheme.

on demand access to vast amounts of compute resources charged on a pay per use basis. The use of virtualisation technology enables users to dynamically procure virtual machines and release resources on demand with varying combinations of CPU, memory and storage available to meet both performance requirements and budget constraints. Through the advancement of cloud based services and in particular High Performance Computing (HPC) platforms scientists have immediate access to large scale distributed infrastructure and customized execution environments to meet their growing needs. As a result, cloud computing infrastructure is fast evolving as the target platform for executing large scale scientific applications requiring high throughput and data analysis [2].

Despite the recent introduction of HPC resources such as Amazons Cluster Compute (CC) platform which offers improved networking capabilities research has shown that the overall performance of HPC applications in public clouds remains limited by poor network throughput [3]. Large data transfers across distributed cloud resources often hinge on unstable bandwidth availability on network links due to the shared nature of the resource. Consequently, a decrease in bandwidth causes an increase in data transfer times, thus increasing total execution time and associated rental costs of cloud resources. This highlights a fundamental workflow scheduling issue which is the impact of bandwidth variability and network congestion on data transfer times between workflow tasks.

As the adoption of cloud computing services continues to grow to facilitate a new generation of scientific users with high computational and data transfer requirements, estimating network resources gains significant importance in the development of efficient and reliable schedulers capable of anticipating dynamically changing network conditions in order to generate more efficient scheduling decisions. To address this issue we propose a network-aware scheduling approach which employs a time series Autoregressive Integrated Moving Average (ARIMA) forecasting algorithm. Our approach promotes efficient utilization of limited network resources by scheduling workflow tasks agnostic to underlying network conditions in order to avoid peak hours of network congestion while also meeting a hard deadline constraint and reducing overall costs.

The remainder of the paper is structured as follows. In Sect. 2 related work in the field is discussed. Section 3 formulates the problem. Section 4 introduces time series ARIMA modelling. Section 5 presents our preliminary results. Lastly, Sect. 6 concludes the paper and discusses future work.

2 Related Research

Much of the existing work in this area has focused on scheduling algorithms which aim to satisfy user Quality of Service (QoS) requirements, namely execution time and cost [4–6]. However, these works naively assume access to unlimited network bandwidth at run time and fail to consider the implications of network constraints on the execution performance of scientific workflows. However, research has also been proposed to target issues concerning the availability of resources on workflow schedules primarily in the form of reactive techniques [7–9].

While these approaches consider changes in resource availability at run time the available bandwidth is assumed to be entirely certain. However network resources are volatile and a function of system load at discrete time intervals [10]. Additionally, some of the suggested approaches involve task rescheduling, this additional overhead may not be a feasible solution in public clouds due to the associated rental costs of computational resources.

We propose an alternative approach which we expect to impact the state-of-the-art in two regards. Firstly, we address a key limitation of current approaches by considering the implications of dynamically changing network behaviour on execution time and associated costs of leasing resources from public clouds. More concretely, we develop a network-aware scheduling algorithm capable of predicting bandwidth availability in public cloud infrastructure where the competition for network resources is far greater. Secondly, we demonstrate the use of a statistical time series ARIMA modelling approach to the workflow scheduling problem for learning and predicting bandwidth availability. In particular, we demonstrate how our approach has the capacity to better align workflow scheduling requirements with the dynamic nature of network resources in order to generate a more reliable schedule and improve execution efficiency.

3 Problem Formulation

A scientific application (workflow) is often modelled as a Directed Acyclic Graph (DAG) denoted as $W = \{V, E\}$ where $V = \{T_1, T_2, \dots, T_n\}$ is a set of tasks in the workflow which are assigned to a specific resource and E is the set of directed edges representing the data dependencies between tasks. A directed edge $E_{i,j}$ signifies that task T_i is the parent task of task T_j . Child tasks can only execute once all parent tasks have been processed and the data has been transferred, while nodes on the same level can be executed in parallel using multiple resources in order to speed up execution time. In addition, each workflow W has an assigned deadline constraint. Figure 1(a) illustrates a workflow example consisting of 7 nodes in which the edges between tasks denote the file inputs and outputs. Figure 1(b) depicts a valid schedule for the adjacent workflow. The objective of our scheduler is to find a mapping of tasks to resources at more optimal times for large data transfers that meet the specified deadline and reduces the overall execution cost.

The cloud infrastructure network model used in this work was generated based on measurements of Amazons EC2 network performance [11]. This benchmark study provides a model of the bandwidth within Amazons EU region. In order to keep the focus of the performance around the network we consider a finite set of homogeneous resources $R = \{r_1, r_2 \dots r_n\}$ which we assume have sufficient CPU, memory and disk to execute each task T_i which has a fixed processing time of 1 h. We consider file sizes that are fixed but vary across individual tasks in the workflow. The execution of task T_i on resource r_j incurs an execution cost. Generally, there is no charge for the transfer of data between tasks in the same region. Total cost is calculated as a function of processing and

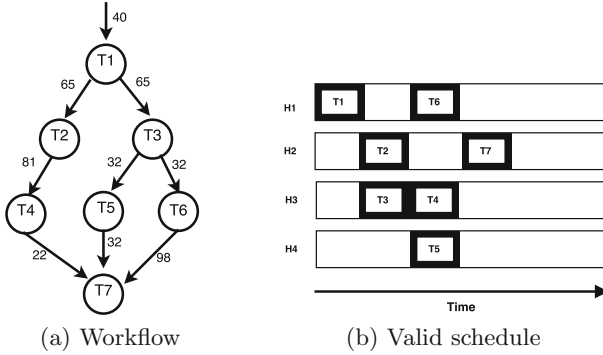


Fig. 1. Sample workflow application with valid schedule

data transfer time. In addition, total execution time can be defined as the latest completion time of all tasks executed on all cloud resources.

4 Time Series Forecasting: ARIMA Modelling

ARIMA models have become a widely popular methodology for time series forecasting [12]. An ARIMA model consists of three fundamental components denoted as (p, d, q) . Identifying a valid model is the process of finding suitable values for (p, d, q) which capture the systematic patterns in the data. The autoregressive (AR) component (p) represents the influence of past values on current values in the series. For example An $AR(1)$ model predicts future values based on the value of the preceding observation defined as:

$$y_t = \phi(y_{t-1}) + \varepsilon_t . \tag{1}$$

Where ϕ is a parameter of the model and ε_t is random variation at time t . The moving average (MA) term (q) models the random variation of the model as a combination of previous error terms. For example An $MA(1)$ model forecasts future values based on a combination of the current random variation and previous error as defined in Eq. 2:

$$y_t = \theta(\varepsilon_{t-1}) + \varepsilon_t . \tag{2}$$

Where ε_{t-1} is the value of the previous random shock and θ is a parameter of the model. The integrated component of the model (d) is the order of differencing applied to the series in order to render the series stationary. A stationary series is one whose statistical properties such as mean and variance are constant over time. A non-stationary time series is often unreliable and can result in false autocorrelations in the series. The combined model assuming differenced data is defined in Eq. 3:

$$y_t = c + \phi_1(y_{t-1}) + \dots + \phi_p(y_{p-1}) + \theta_1(\varepsilon_{t-1}) + \dots + \theta_q(\varepsilon_{t-q}) + \varepsilon_t . \tag{3}$$

In addition, Seasonal ARIMA models can be used to model highly seasonal data formed by including additional ARIMA terms $(P, D, Q)_m$, where m signifies the number of periods per season. To generate the proposed model the Box-Jenkins methodology was employed which is composed of several steps outlined below [13].

4.1 Model Identification

The bandwidth data used in this work showed a significant seasonal periodic component which occurs within each 24 h period. The Autocorrelation Function (ACF) plot also revealed strong periodic oscillations indicating the implementation of a seasonal ARIMA model as defined in Eq. 4 to capture the strong seasonality present. It also confirmed the series was non stationary.

$$\phi_p(B)\Phi_P(B^s)W_t = \theta_q(B)\Theta_Q(B^s)Z_t . \tag{4}$$

Where B denotes the lag operator, $\phi_p, \Phi_P, \theta_q, \Theta_Q$ are parameters of the seasonal and non-seasonal model components $(p, q)(P, Q)$ respectively and Z_t represents the error. In addition, a first order seasonal difference was applied to the data which proved sufficient in transforming the non stationary series into a stationary series.

To select the appropriate orders of both the non-seasonal components p, q and seasonal components P, Q of the model the ACF and Partial Autocorrelation Function (PACF) plots of the differenced data W_t were examined which identified several alternative values to select from.

4.2 Model Estimation and Diagnostics

To estimate the model parameters $\phi'_p s, \Phi'_P s, \theta'_q s, \Theta'_Q s$ in the forecast equation above the Maximum Likelihood Estimation (MLE) was adopted using R software. In order to select the best model to fit the data a common criterion known as Akaike Information Criterion (AIC) was used. This statistic as defined in Eq. 5 is a fundamental measurement of the quality of a statistical model for a series. ARIMA(0, 0, 0)(0, 1, 1)[144] was deemed the best model as it produced the lowest AIC value.

$$AIC = -2\log(L) + 2(p + q + k + 1) . \tag{5}$$

Where L is the maximum likelihood of the data, $k = 1$ if $c \neq 0$ while $k = 0$ if $c = 0$ and the final term represents the number of parameters in the model. A formal Ljung-Box test was also conducted on the residuals which generated a p-value of 0.735, which concluded that the residuals are independent.

4.3 Forecasting and Validation

To assess the accuracy of the forecasts a test set which equated to 1 week of bandwidth values over 10 min intervals was used as a comparison measure. The results

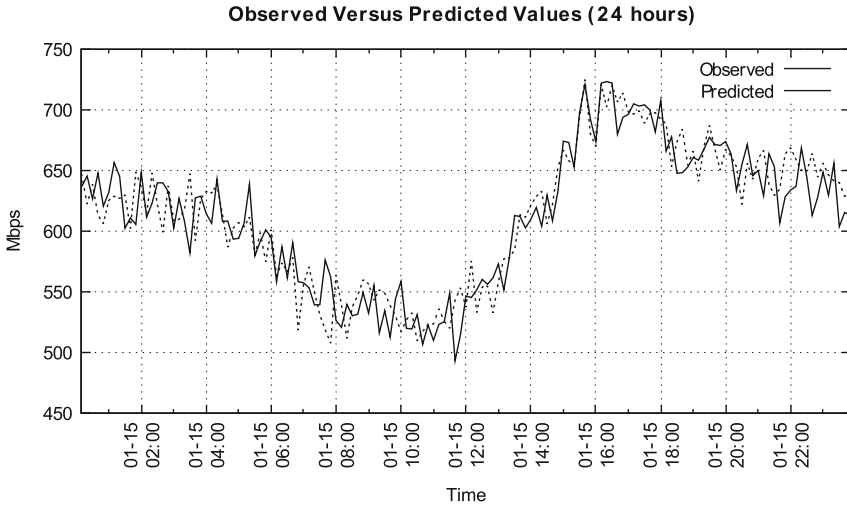


Fig. 2. Predicted and observed values for the subsequent day based on the previous weeks bandwidth values used to fit the model

showed that the forecasts generated were all within the 95% confidence limit. Additionally the Mean Absolute Percentage Error (MAPE) as given by Eq. 6 was calculated where O_t and P_t are the observed and predicted values of the time series. The MAPE from the resulting forecasts was 2.81% which is deemed highly accurate [14].

$$MAPE = \frac{100}{n} \sum_{t=1}^n \left| \frac{O_t - P_t}{O_t} \right|. \tag{6}$$

5 Preliminary Results

A cloud simulator was developed to evaluate the proposed network-aware scheduling procedure. As an initial benchmark we compare our approach to a non network-aware heuristic called *Execute-First*. To evaluate the proposed procedure using reasonable deadline constraints the *Execute-First* heuristic was run over 30 iterations calculating the makespan of each 10 min interval in a single day. The average earliest and latest finish times were computed in order to define deadline D denoted in Eq. 7, where parameter m is defined as 1, 30 and 60 to evaluate the performance of the algorithms over low, medium and high deadline constraints.

$$D = EFT_{average} + m \times (LFT_{average} - EFT_{average}). \tag{7}$$

Figure 3(a) shows the total execution time in hours for all 3 deadline categories. These results show the advantage of using the predictive capabilities of

ARIMA modelling to inform our scheduling decision, our approach selects the most opportunistic time frame within the deadline to transfer data resulting in shorter execution times. Evidently, the performance of our approach continues to increase when deadlines span over a greater number of hours. This is largely due to the visibility our ARIMA driven algorithm has over dynamically changing bandwidth availability. Figure 3(b) also shows a significant reduction in cost as our network-aware scheduler postpones execution until network conditions are more optimum. Conversely, the *Execute-First* algorithm incurs larger costs due to poor scheduling decisions resulting in longer transfer times when the network is saturated.

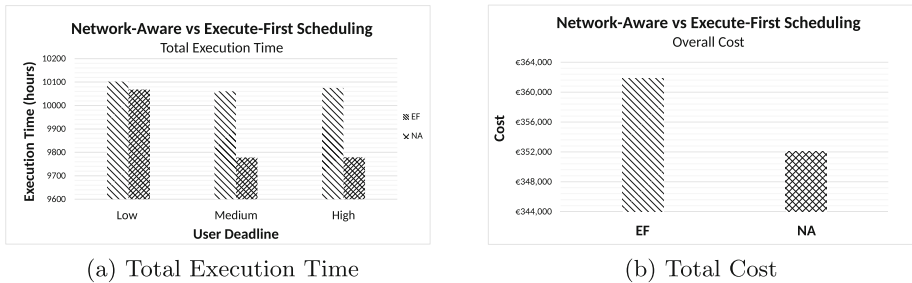


Fig. 3. Total execution time over low, medium and high deadline constraints and overall cost generated by both approaches

6 Conclusion

This work presented an efficient network-aware workflow scheduler based on time series ARIMA modelling designed to minimize total execution time and costs. Our empirical results have shown that by adopting a scheduling procedure which has the capacity to reason over the impact of dynamically changing bandwidth availability we can achieve significant cost reductions and reduce execution time. In future work we intend on extending our solution to consider heterogeneous workflow tasks and cloud resources to further optimize resource availability, while also considering the impact of additional factors such as queuing and propagation delays in order to deliver a more complete solution. Eventually, we hope to evaluate our approach using a live virtualised test bed.

Acknowledgments. The primary author would like to acknowledge the ongoing financial support provided to her by the Irish Research Council.

References

1. Yang, X., Wallom, D., Waddington, S., Wang, J., Shaon, A., Matthews, B., Wilson, M., Guo, Y., Guo, L., Blower, J.: Cloud computing in e-Science: research challenges and opportunities. *J. Supercomput.* **70**(1), 408–464 (2014). Springer

2. Lifka, D., Foster, I., Mehring, S., Parashar, M., Redfern, P., Stewart, C., Tuecke, S.: XSEDE cloud survey report. Technical report, National Science Foundation (2013)
3. Expósito, R.R., Taboada, G.L., Ramos, S., González-Domínguez, J., Touriño, J., Doallo, R.: Analysis of I/O performance on an Amazon EC2 cluster compute and high I/O platform. *J. Grid Comput.* **4**(11), 613–631 (2013). Springer
4. Abrishami, S., Naghibzadeh, M., Epema, D.: Deadline-constrained workflow scheduling algorithms for infrastructure as a service clouds. *Future Gener. Comput. Syst.* **29**(1), 158–169 (2013). Elsevier
5. Barrett, E., Howley, E., Duggan, J.: A learning architecture for scheduling workflow applications in the cloud. In: 2011 Ninth IEEE European Conference on Web Services (ECOWS), pp. 83–90. IEEE (2011)
6. Pandey, S., Wu, L., Guru, M.S., Buyya, R.: A particle swarm optimization-based heuristic for scheduling workflow applications in cloud computing environments. In: 2010 24th IEEE International Conference on Advanced Information Networking and Applications (AINA), pp. 400–407. IEEE (2010)
7. Allen, G., Angulo, D., Foster, I., Lanfermann, G., Liu, C., Radke, T., Seidel, E., Shalf, J.: The Cactus worm: experiments with dynamic resource discovery and allocation in a grid environment. *Int. J. High Perform. Comput. Appl.* **15**(4), 345–358 (2001). Sage Publications, Thousand Oaks
8. Batista, D.M., da Fonseca, N.L., Miyazawa, F.K., Granelli, F.: Self-adjustment of resource allocation for grid applications. *Comput. Netw.* **52**(9), 1762–1781 (2008). Elsevier
9. Tang, W., Jenkins, J., Meyer, F., Ross, R., Kettimuthu, R., Winkler, L., Yang, X., Lehman, T., Desai, N.: Data-aware resource scheduling for multicloud workflows: a fine-grained simulation approach. In: 2014 IEEE 6th International Conference on Cloud Computing Technology and Science (CloudCom), pp. 887–892. IEEE (2014)
10. Duggan, M., Duggan, J., Howley, E., Barrett, E.: A network aware approach for the scheduling of virtual machine migration during peak loads. *Cluster Comput.* **20**(2083), 1–12 (2017)
11. Sanghrajka, S., Mahajan, N., Sion, R.: Cloud performance benchmark series: Network performance-Amazon EC2. Technical report, Stony Brook University (2011)
12. Hyndman, R.J., Athanasopoulos, G.: *Forecasting: Principles and Practice*. OTexts, Melbourne (2013)
13. Box, G.E., Jenkins, G.M.: *Time Series Analysis, Control, and Forecasting*, vol. 3226(3228), p. 10. Holden Day, San Francisco (1976)
14. Kenneth, D.L., Ronald, K.K.: *Advances in Business and Management Forecasting*. Emerald Books, UK (1982)