

# Resource Distribution Estimation for Data-Intensive Workloads: Give Me My Share & No One Gets Hurt!

Alireza Khoshkbarforoushha<sup>1,2</sup>(✉), Rajiv Ranjan<sup>1,2</sup>, and Peter Strazdins<sup>1</sup>

<sup>1</sup> Australian National University, Canberra, Australia  
`peter.strazdins@cs.anu.edu.au`

<sup>2</sup> CSIRO, Canberra, Australia  
`{alireza.khoshkbarforoushha,raj.ranjan}@csiro.au`

**Abstract.** Robust resource share estimation of data-intensive workloads is integral to efficient workload management in a (virtualized) cluster where multiple systems co-exist and share the same infrastructure. However, developing a reliable resource estimator is quite challenging due to (i) heterogeneity of workloads (e.g. stream processing, batch processing, transactional, etc.) in a multi-system shared cluster, (ii) limited (in batch processing) or complete uncertainties (in stream processing) on input data size or arrival rates, and (iii) changing configurations from run to run. To address above challenges, we propose an inclusive framework and related techniques for *workload profiling*, *similar job identification*, and *resource distribution prediction* in a cluster. Our analysis shows that the framework can successfully estimate the whole spectrum of resource usage as probability distribution functions for wide ranges of data-intensive workloads.

**Keywords:** Resource estimation · Big data workload · Multi-cluster workload management · Distribution prediction · Data-intensive systems

## 1 Introduction

Datacenter-scale computing for big data analytics workloads has seen a surge in adoption due to availability and affordability of large-scale data processing systems which transform the traditional data mining and machine learning techniques into easy to program and deploy distributed analytics applications. In such an environment, robust resource usage prediction of data-intensive jobs is integral to make efficient workload management decisions in different scenarios such as workload migration between clouds or workload scheduling in a multi-system cluster.

Cost reduction is one of the main drivers for migrating the workloads between cloud providers (e.g. Amazon AWS) [10]. For example, the Azure Cost Estimator<sup>1</sup> has been designed to assist the infrastructure manager to either assess the

---

<sup>1</sup> <http://www.microsoft.com/en-us/download/details.aspx?id=43376>.

running cost of the existing on-premises workloads or analyse how much she can save by moving the hosted servers on other cloud platforms (e.g. AWS) to Azure. This class of tools typically does analysis irrespective of the workload resource usage patterns, thereby is not able to extrapolate the analysis beyond the current conditions of the workloads. However, robust resource usage prediction paves the way for predictive analysis of workload migration costs between clouds. From another standpoint, such estimation is integral to efficient workload scheduling in a multi-system cluster.

In a multi-system cluster, data-intensive workloads are typically classified as either productions or best effort jobs [5]. Production jobs (e.g. Oozie), unlike the best-effort ad-hoc jobs, are business-critical, meaning that missing their service level agreements (SLAs) can have substantial financial impact. State of the art workload scheduling techniques [5, 8, 14] focus on guaranteeing the SLAs for this class of workload subject to fairness, capacity, priority, and throughput maximization. Majority of these techniques [5, 8] need *the cost of the workload* to be specified a priori to be able to define appropriate resource sharing policies.

However, developing a reliable resource estimator for production jobs is a hard research problem due to: (i) heterogeneity of workloads pertaining to each class of big data systems, (ii) limited (in batch processing) or complete uncertainties (in stream processing) on input data size/arrival rates and schema, and (iii) changing configurations (e.g. number of mappers/reducers or spouts/bolts respectively in Hadoop and Storm) from run to run. To address above problems, we propose an inclusive framework and related techniques for resource distribution estimation of heterogeneous big data workloads in a shared cluster.

To this end, the proposed framework first generates a set of data-intensive specific job templates (JT) by applying a clustering technique on a set of job characteristics. These templates are then used to identify similar jobs. Once the templates are generated, an statistical machine learning (ML) model is built for each of them by exploiting the past execution traces within the template. We argue that the existing single point resource estimator is not adequate for describing the whole spectrum of resource usage of big data workloads. Therefore, we introduce the novel approach of applying mixture density networks (MDN) as an underlying ML technique to approximate the probability distributions by means of finite mixture of Gaussians.

Therefore, the main contributions of the proposed resource estimation framework for big data analytics workload are: (i) Introducing appropriate techniques to profile and identify similar jobs within a heterogeneous workloads from both batch and stream data processing systems, (ii) Proposing a novel approach of estimating the full spectrum of resource usage in form of probability density functions (pdfs).

The remainder of the paper is organized as follows: The next section explores the related work. Section 3 presents the overview of the proposed framework. In this section, we propose our ideas on *how to define templates for similar jobs identification*, and *how to build resource distribution prediction models*. Section 4 presents initial results on distribution based resource modelling along with some

discussions on how the predicted pdfs can be utilized in workload management scenarios. The paper ends with some concluding remarks and the plan for future work in Sect. 5.

## 2 Related Work

In this section we present a high level discussion on the past research done in the domain of job performance estimation. Following that, discussion of related work including performance estimation of declarative (SQL-style) and procedural (MapReduce style) data-intensive workload is presented.

**Job Performance Estimation.** Job performance estimation, in general, using historical information of *similar* job has been studied in the past in domain of parallel computing [17, 18]. The key difference among the existing approaches is the way they tackle the problem of identifying similar jobs. In particular, authors in [18] focus on application characteristics such as user of jobs, jobs' submission time, jobs' arguments and on the number of physical servers on which the jobs are submitted for the definition of application similarity. Recently, authors in [17] propose the novel use of clone detection technique to determine the clone level of a newly submitted job with respect to the jobs in the execution history and to predict the resource requirement of the new job depending on its clone level.

**High-Level Data Intensive Frameworks.** There are a number of related work on runtime and resource usage estimation in the context of DBMS [1, 13]. In the majority of related work, different statistical ML techniques are applied for estimating query performance. These approaches typically build statistical models using past query executions and a representative set of query features (query plan and/or operator level features) which possibly have the high predictive power in terms of resource or performance measures.

When it comes to MapReduce ecosystem, a major fraction of big data cluster workloads are generated by a handful of high-level frameworks such as Hive, Pig, Giraph [4]. This opens up an opportunity to train ML techniques against a set of finite recurring operators to estimate workload performance. For example, authors in [7] apply the Kernel Canonical Correlation Analysis (KCCA) to the Hive execution plan operators. They conclude that *only* a set of *low level features* pertaining to Hive query execution such as the number of maps and reduces, bytes read locally, bytes read from HDFS lead to accurate performance modelling. However, the provided low level features are *not* available before actual query execution. Therefore, above technique cannot be applied for performance prediction of new incoming workloads.

Similarly, authors in [16] propose a technique that predicts the runtime performance for a fixed set of queries running over varying input data sets. Specifically, it splits each query into several segments where each segment's performance is estimated using uni-variate linear regression. Next the estimates are plugged into a global analytical model to predict the overall query runtime. Since

modelling a small and finite space of relational operators might not be adequate for all MapReduce workloads (e.g. iterative analytics), several studies focused on more fine-grained analysis of MapReduce job performance analysis as discussed next.

**Procedural Data Intensive Workload.** Herodotou et al. [9] propose a self-tuning system for Hadoop that uses performance models with the goal of workload tuning, finding the best configuration settings for a given workload and a cluster infrastructure. Along these lines, the authors in [20] first build the job profile from the job past executions or by executing the workload on a smaller data set using an automated profiling tool. Then, they apply the performance bounds of completion time of different job phases to predict the job completion time as a function of the input dataset size and allocated resources.

In terms of runtime prediction of ML algorithms executing on top of MapReduce ecosystem, authors in [15] present an experimental methodology for predicting the runtime of iterative algorithms written in Apache Giraph. To do so, they conduct sample runs for capturing the algorithm’s convergence trend and per-iteration key input features.

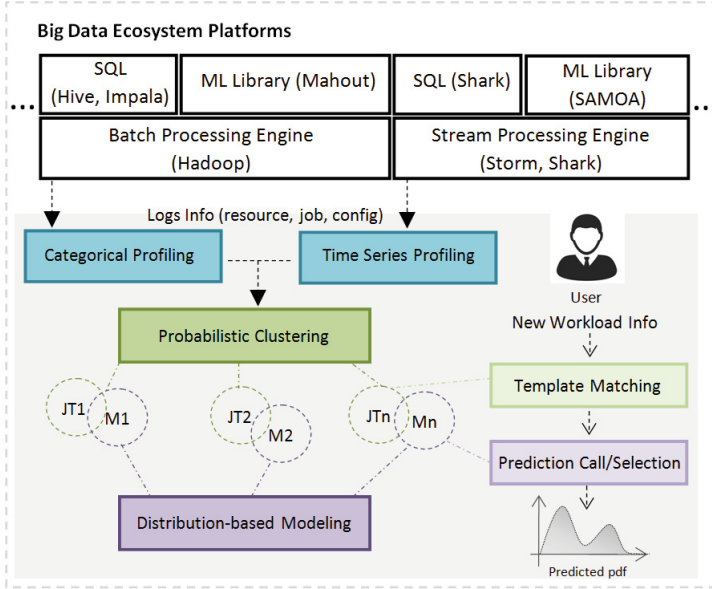
**Concluding Remarks.** In summary, all of the above studies focused on single type of workload in particular Hadoop, while in reality big data workloads are heterogeneous consisting of multiple types of systems (e.g. Apache Hive, Apache Storm, etc.) and jobs. In contrast, our framework considers heterogeneous workloads where different jobs and queries from either batch or stream data processing systems running side by side. Moreover, existing approaches estimate resource and performance as a single point value which is neither expressive enough nor does it capture the possible variances due to resource contentions and interferences from other workloads. In contrast, we use a distribution prediction technique that describes the resource usage as conditional distribution functions.

### 3 Overview of the Proposed Framework

The problem of resource requirement estimation for future job in a multi-system cluster is decomposed in (i) characterizing the similar jobs that have executed in the past, and (ii) building a prediction model based on the collected statistics. The workflow in the proposed framework, as shown in Fig. 1, is as follows. Firstly, workload profiling is conducted in order to collect required information for defining similar jobs and building resource models. Secondly, a set of job templates are generated based on the collected features and their corresponding values. Finally, a distribution ML model is trained and built for each job templates which is responsible for resource prediction of a new incoming workload.

#### 3.1 Similarity Definition and Template Generation

The routine nature of productions jobs allows us to build the resource model based on the past execution profiles of the *same* or *similar* jobs. Thus, building



**Fig. 1.** Comprehensive framework for performance distribution prediction of heterogeneous data-intensive workloads.

profiles is the first step. Our framework considers both batch and streaming workloads, thereby the profiles will be logged as either categorical or time series files respectively. We next look into *what* to profile.

Resource usage  $r$  is the function of a job  $j$  that is executed on input data (stream)  $d$  using configuration parameter settings  $c$ :

$$r = f(j, d, c) \tag{1}$$

This means that resource modelling requires logging appropriate information about job, its data flow and configurations along with the resources in use, as follows:

– **Job Profile**

- *Input Data Stats*: such as number of Input Records/Bytes (e.g. in Hive tables), Average arrival rates (in Stream processing), Data file format, etc.
- *Job Metadata*: including Job Name, User who submitted the job, Submission time, File input path, and so on.
- *Runtime Stats*: such as number of mappers and reducers per stage (as in Hive), Response time, Latency/Throughput (in stream processing).

– **Resource Profile** including CPU time and utilization, Memory usage, local/network I/O.

### – Configuration Profile

- *Compact<sup>2</sup> job configs*: such as number of reducers in Hive which can be set for a specific query using `mapred.reduce.tasks` property and override the cluster wide settings.
- *Compact cluster configs*: such as the number of virtual CPU cores for each reduce task of a job which is set by `mapreduce.reduce.cpu.vcores` property in Apache Yarn [19].

The job profiling is a recurrent process, means that the mentioned profiles will be generated for every submitted job. However, to avoid any performance degradation, two design principles need to be realized. First, statistics should be collected passively, without affecting the performance. For example, in [12] we used *dstat* (<http://dag.wiee.rs/home-made/dstat/>) as a lightweight python-based tool that collects OS and system statistics non-intrusively. Second, profiling process should be tractable via enabling a feature to on/off profiling process. This means that every job is profiled unless it is deactivated by the user of the final prototype.

Upon extracting required information from job history and big data cluster logs, we now focus on the second challenge, that is characterizing the similar jobs that have executed in the past. We need to identify a set of classes of jobs (i.e. job templates) that exhibit a similar resource usage pattern.

Modern big data clusters run a diverse mix of applications and production workloads [19], thereby characterizing similar jobs is challenging. Although difficult, we argue that appropriate clustering techniques along with the proper job execution and big data system configuration profiles lead to formation of fitting templates.

In our problem domain the clusters (i.e. candidate job templates) need not be disjoint, and the same job can be associated to several classes. Because two jobs can be compared in many ways. For example, our initial analysis on synthetic MapReduce workloads of Facebook<sup>3</sup> [4] demonstrates that the (`submit_time_seconds`, `hdfs_input_path`) is a proper candidate job template since jobs with the close submission time and same input path have roughly same map input byte size, shuffle bytes, etc. Yet another template is (`user`, `submit_time_seconds`, `hdfs_input_path`) which is more restricted. Intuitively, a certain job can be assigned to both of these templates with respect to its characteristics.

Therefore, we use a probabilistic clustering technique, Expectation Maximization (EM) algorithm [6] which is a soft clustering technique. EM finds clusters by determining a mixture of Gaussians that fit a given data set. Each Gaussian has a mean and covariance matrix. The prior probability for each Gaussian is the fraction of points in the cluster defined by that Gaussian. These parameters can be either initialized by randomly selecting means of the Gaussians or by using the output of K-means algorithms for the initial centres.

<sup>2</sup> Due to the large number of configuration parameters, only a subset of settings which have substantial impacts on resource and performance measures need to be logged.

<sup>3</sup> <https://github.com/SWIMProjectUCB/SWIM/wiki>.

EM converges on a locally optimal solution by iteratively updating means and variances.

Once the candidate job templates (i.e. clusters) becomes available, different validity measures which falls broadly into internal, relative, and external validations are used to evaluate clustering results. The output of this step is a set of final job templates through which we are able to identify who belongs to whom. To do so, for a new job, the template membership probability is calculated. The higher probability seemingly shows the more the templates match the job's characteristics.

### 3.2 Distribution Prediction

Many workload scheduling studies [5, 8] formulate and optimize the policies subject to fairness, capacity, and priority, conditioned on having the cost of a workload a priori in a multidimensional space representing different resources. In contrast, there exists criticism of the off-line predictions [14] due to the interferences from other workloads concurrently running and sharing the same resources at runtime.

As stressed in related work, the state of the art resource and performance estimation techniques for data intensive workloads only provide the conditional mean of the point of interest. However, even running the same query on the same data with constant configuration show different performance and resource behaviour. In response, we argue that the *distribution* estimation provides an expressive description of the target values (e.g. runtime, CPU time) and the possible variances due to resource contention.

We adopt a novel approach of workload resource distribution prediction using Mixture Density Networks (MDN) [3]. An MDN fuses a Gaussian mixture model (GMM) with feed-forward neural networks. In MDN, the distribution of the outputs  $t$  is described by a parametric model whose parameters are determined by the output of a neural network, which takes  $x$  as inputs. Specifically, an MDN maps a set of input features  $x$  to the parameters of a GMM including mixture weights  $\alpha_i$ , mean  $\mu_i$ , and variance  $\sigma^2$  which in turn produces the full *pdf* of an output feature  $t$ , conditioned on the input vector. Detailed discussion on the proposed approach are available in accompanying technical reports [11, 12] which discuss how to predict the resource and performance distribution for batch (Hive workloads) [11] and stream data processing (i.e. continuous queries) [12]. We will show the efficacy of distribution as opposed to existing techniques in workload resource modelling in the next section.


Note that due to the large number of profiles and historical logs, building a model could be prohibitively expensive, though we already showed [11] that the training time of the MDN linearly grows with respect to the training data size. Thus, enabling a *maximum history* feature as in [18] which indicates the maximum number of data points to be used for building/refreshing a model is inevitable.

Once the model is built and trained, it can then be invoked when a similarity measures assign a new job to one of the existing classes. Since the job templates

overlap, a single job may be associated to multiple job templates and their corresponding ML model. Therefore, the prediction with the smallest confidence interval will be selected.

## 4 Initial Results

In this section, we present some compact yet lucid results on how the distribution prediction look like and how one can utilize them for appropriate policy setting across a shared big data cluster.



**Fig. 2.** (a) a sample predicted pdf for CPU Time of an input test from Hive workload. (b) a sample predicted pdf for CPU Utilization, selected from a test dataset of linear road benchmark (Color figure online).

Figure 2(a) plots a sample predicted pdf for CPU time for one of the experiments conducted on TPC-H in [11]. The predicted pdf is corresponding to a test input from Template-7 (Q7) of TPC-H against 100 GB database size. To demonstrate the whole possible range of CPU time values under Q7, the histograms for 30 instance queries based on Q7 from test set are shown as well.

As we can see, the predicted distribution adequately estimates the CPU time distribution in which they show high probability around the target value. More importantly, they provide information about the whole spectrum of resource usage. In particular, the predicted pdf shows highly probable CPU time in ranges (0.25, 0.4) which are consistent with the actual distribution. Note that the predicted pdf is concerned with one single input, thereby the resulted uncertainty of pdf for the range (0.6, 0.9) is justifiable.

In a similar manner, Fig. 2(b) depicts a sample predicted pdf and actual CPU usage in terms of normalized histogram for one of the experiments on linear road benchmark [2] queries conducted in [12]. As the figure indicates, the estimated pdf approximates the actual resource usage closely. The predicted pdf provides a complete description of the statistical properties of the CPU utilization through which we are not only able to capture the observation point, but also the whole spectrum of the resource usage. In contrast, a best estimation from the existing resource estimation techniques [1, 7, 9, 13, 20] merely provides the point which is



visualized by solid red vertical line. Unlike pdfs, with such estimation we are not able to directly calculate any valuable statistical measures (e.g. variance, expectation) about the target data. Detailed evaluation of the proposed approach and its comparison with the state of the art single point estimator can be found in accompanying technical reports [11, 12].

Once the resource usage distribution becomes available, it can then be used to define appropriate resource pool policy for critical SLA-driven workloads. For example, Cloudera Manager<sup>4</sup> provides the ability to statically allocate resources using Linux control groups through which one can allocate services (e.g. Hive, Storm) a *percentage* of total resources. Static resource pools isolate the services in the cluster from one another, so that load on one service has a bounded impact on the others. With distribution based resource estimation, we are able to determine appropriate percentage of resource shares for a workload before actual execution, without sacrificing the cluster throughputs and utilization.

## 5 Conclusions and Future Work

This paper proposed an inclusive framework and related techniques for resource usage distribution prediction of heterogeneous big data workloads in a cluster. To this end, our framework uses the clustering techniques along with the statistical machine learning algorithm to identify similar jobs and build a distribution-based prediction models. The initial results show that the approach is capable of estimating resource usage distribution accurately, through which we are able to define more reliable resource sharing policies aiming at guaranteeing SLA subject to fairness, capacity, priority, and throughput maximization.

As an ongoing work, we plan to complete and evaluate the template generation phase using real-world workload traces collected from our private CSIRO big data cluster and possibly more traces from other companies. Following that, we also plan to accommodate the distribution predictions in cost-optimized resource provisioning of big data analytics flows in a datacenter cloud.

## References

1. Akdere, M., Çetintemel, U., Riondato, M., Upfal, E., Zdonik, S.B.: Learning-based query performance modeling and prediction. In: 2012 IEEE 28th International Conference on Data Engineering (ICDE), pp. 390–401. IEEE (2012)
2. Arasu, A., Cherniack, M., Galvez, E., Maier, D., Maskey, A.S., Ryvkina, E., Stonebraker, M., Tibbetts, R.: Linear road: a stream data management benchmark. In: Proceedings of the Thirtieth International Conference on Very Large Data Bases, vol. 30, pp. 480–491. VLDB Endowment (2004)
3. Bishop, C.M.: Mixture density networks (1994)
4. Chen, Y., Alspaugh, S., Katz, R.: Interactive analytical processing in big data systems: a cross-industry study of mapreduce workloads. VLDB 5(12), 1802–1813 (2012)

---

<sup>4</sup> <http://www.cloudera.com/content/cloudera/en/products-and-services/cloudera-enterprise/cloudera-manager.html>.

5. Curino, C., Difallah, D.E., Douglas, C., Krishnan, S., Ramakrishnan, R., Rao, S.: Reservation-based scheduling: if you're late don't blame us! In: Proceedings of the ACM Symposium on Cloud Computing, pp. 1–14. ACM (2014)
6. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the em algorithm. *J. R. Stat. Soc. Ser. B (Methodological)* **39**, 1–38 (1977)
7. Ganapathi, A., Chen, Y., Fox, A., Katz, R., Patterson, D.: Statistics-driven workload modeling for the cloud. In: 2010 IEEE 26th International Conference on Data Engineering Workshops (ICDEW), pp. 87–92. IEEE (2010)
8. Ghodsi, A., Zaharia, M., Hindman, B., Konwinski, A., Shenker, S., Stoica, I.: Dominant resource fairness: fair allocation of multiple resource types. In: NSDI, vol. 11, p. 24 (2011)
9. Herodotou, H., Babu, S.: Profiling, what-if analysis, and cost-based optimization of mapreduce programs. *VLDB* **4**(11), 1111–1122 (2011)
10. Jamshidi, P., Ahmad, A., Pahl, C.: Cloud migration research: a systematic review. *IEEE Trans. Cloud Comput.* **1**(2), 142–157 (2013)
11. Khoshkbarforousha, A., Ranjan, R.: Resource and performance distribution prediction for large scale analytics queries. TR-2015-01, ANU Technical report (2015)
12. Khoshkbarforousha, A., Ranjan, R., Gaire, R., Jayaraman, P.P., Hosking, J., Abbasnejad, E.: Resource usage estimation of data stream processing workloads in datacenter clouds. arXiv preprint [arXiv:1501.07020](https://arxiv.org/abs/1501.07020) (2015)
13. Li, J., König, A.C., Narasayya, V., Chaudhuri, S.: Robust estimation of resource consumption for sql queries using statistical techniques. *Proc. VLDB Endowment* **5**(11), 1555–1566 (2012)
14. Mace, J., Bodik, P., Fonseca, R., Musuvathi, M.: Retro: targeted resource management in multi-tenant distributed systems. In: NSDI. USENIX (2015)
15. Popescu, A.D., Balmin, A., Ercegovac, V., Ailamaki, A.: Predict: towards predicting the runtime of large scale iterative analytics. *Proc. VLDB Endowment* **6**(14), 1678–1689 (2013)
16. Popescu, A.D., Ercegovac, V., Balmin, A., Branco, M., Ailamaki, A.: Same queries, different data: Can we predict runtime performance? In: 2012 IEEE 28th International Conference on Data Engineering Workshops (ICDEW), pp. 275–280. IEEE (2012)
17. Sarkar, M., Mondal, T., Roy, S., Mukherjee, N.: Resource requirement prediction using clone detection technique. *Future Gener. Comput. Syst.* **29**(4), 936–952 (2013)
18. Smith, W., Foster, I., Taylor, V.: Predicting application run times using historical information. In: Feitelson, D.G., Rudolph, L. (eds.) IPPS-WS 1998, SPDP-WS 1998, and JSSPP 1998. LNCS, vol. 1459, pp. 122–142. Springer, Heidelberg (1998)
19. Vavilapalli, V.K., Murthy, A.C., Douglas, C., Agarwal, S., Konar, M., Evans, R., Graves, T., Lowe, J., Shah, H., Seth, S., et al.: Apache hadoop yarn: yet another resource negotiator. In: Proceedings of the 4th annual Symposium on Cloud Computing, p. 5. ACM (2013)
20. Verma, A., Cherkasova, L., Campbell, R.H.: Aria: automatic resource inference and allocation for mapreduce environments. In: Proceedings of the 8th ACM International Conference on Autonomic Computing, pp. 235–244. ACM (2011)