



THE UNIVERSITY OF
SYDNEY

SCHOOL OF INFORMATION TECHNOLOGIES

**PRELIMINARY RESULTS ON USING MATCHING ALGORITHMS IN MAP-REDUCE
APPLICATIONS**

TECHNICAL REPORT 672

NIKZAD BABAI RIZVANDI, JAVID TAHERI AND ALBERT Y. ZOMAYA

MARCH, 2011

Preliminary Results on Using Matching Algorithms in Map-Reduce Applications

Nikzad Babaii Rizvandi^{1,2}, Javid Taheri¹, Albert Y. Zomaya¹

¹ Center for Distributed and High Performance Computing, School of Information Technologies, University of Sydney

² National ICT Australia (NICTA), Australian Technology Park

Sydney, Australia

nikzad@it.usyd.edu.au

Abstract—In this paper, we study CPU utilization time patterns of several Map-Reduce applications. After extracting running patterns of several applications, the patterns with their statistical information are saved in a reference database to be later used to tweak system parameters to efficiently execute unknown applications in future. To achieve this goal, CPU utilization patterns of new applications along with its statistical information are compared with the already known ones in the reference database to find/predict their most probable execution patterns. Because of different patterns lengths, the Dynamic Time Warping (DTW) is utilized for such comparison; a statistical analysis is then applied to DTWs' outcomes to select the most suitable candidates. Moreover, under a hypothesis, another algorithm is proposed to classify applications under similar CPU utilization patterns. Three standard applications (WordCount, Exim Mainlog parsing and Terasort) are used to evaluate our hypothesis in tweaking system parameters in executing similar applications. Results were very promising and showed effectiveness of our approach on pseudo-distributed Map-Reduce platforms

Index Terms—Map-Reduce, Pattern Matching, Configuration parameters, statistical analysis.

I. INTRODUCTION

Recently, businesses have started using Map-Reduce as a popular computation framework for processing large-scaled data in both public and private clouds; e.g., many Internet endeavors are already deploying Map-Reduce platforms to analyze their core businesses by mining their produced data. Therefore, there is a significant benefit to application developers in understanding performance trade-offs in Map-Reduce-style computations in order to better utilize their computational resources [1].

Map-Reduce users typically run a few number of applications for a long time. For example, Facebook, which is based on Hadoop (Apache implementation of Map-Reduce in Java), is using Map-Reduce to read its daily produced log files and filter database information depending on the incoming queries. Such applications are repeated million times per day in Facebook. Another example is Yahoo where around 80-90% of their jobs is based on Hadoop[2]. The typical applications

here are searching among large quantities of data, indexing the documents and returning appropriate information to incoming queries. Similar to Facebook, these applications are run million times per day for different purposes.

One of the major problems with direct influence on Map-Reduce performance is tweaking/tuning the effective configuration parameters [3] (e.g., number of mappers, number of reducers, input file size and so on) for efficient execution of an application. These optimal values not only are very hard to properly set, but also can significantly change from one application to another. Furthermore, obtaining these optimal values usually needs running an application for several times with different configuration parameters values: a very time consuming and costly procedure. Therefore, it becomes more important to find the optimal values for these parameters before actual running of such application on Map-Reduce platforms.

Our approach, in this work, is an attempt toward solving this problem by predicting uncertain CPU utilization pattern of new applications based on the already known ones in a database. More specifically, we propose a two-phase approach to extract patterns and find statistical similarity in uncertain CPU utilization patterns of Map-Reduce applications. In the first phase, profiling, few applications are run with different sets of Map-Reduce configuration parameters for several times to collect their execution/utilization profiles in a Linux environment. Upon obtaining such information –the CPU utilization time series of these applications– their statistical information at each point are obtained. Then, these uncertain CPU utilization values are stored in a reference database to be later used in the second phase, matching. In the matching phase, a pattern-matching algorithm is deployed to find similarity between stored CPU utilization profiles and the new application.

To demonstrate our approach, section 2 highlights the related works in this area. Section 3 provides some theoretical background for pattern matching in uncertain time series. Section 4 explains our approach in which pattern matching is used to predict behavior of unknown applications. Section 5 details our experimental setup to gauge efficiency of our approach and introduces a hypothesis to classify applications.

Discussion and analysis is presented in section 6, followed by conclusion in section 7.

II. RELATED WORKS

Early works on analyzing/improving Map-Reduce performance started almost since 2005; such as an approach by Zaharia et al [4] that addressed problem of improving the performance of Hadoop for heterogeneous environments. Their approach was based on the critical assumption in Hadoop that works on homogeneous cluster nodes where tasks progress linearly. Hadoop utilizes these assumptions to efficiently schedule tasks and (re)execute the stragglers. Their work introduced a new scheduling policy to overcome these assumptions. Besides their work, there are many other approaches to enhance or analysis the performance of different parts of Map-Reduce frameworks, particularly in scheduling [5], energy efficiency [1, 6-7] and workload optimization[8]. A statistics-driven workload modeling was introduced in [7] to effectively evaluate design decisions in scaling, configuration and scheduling. The framework in this work was utilized to make appropriate suggestions to improve the energy efficiency of Map-Reduce. A modeling method was proposed in [9] for finding the total execution time of a Map-Reduce application. It used Kernel Canonical Correlation Analysis to obtain the correlation between the performance feature vectors extracted from Map-Reduce job logs, and map time, reduce time, and total execution time. These features were acknowledged as critical characteristics for establishing any scheduling decisions. Recent works in [9-10] reported a basic model for Map-Reduce computation utilizations. Here, at first, the map and reduce phases were modeled using dynamic linear programming independently; then, these phases were combined to build a global optimal strategy for Map-Reduce scheduling and resource allocation.

The other part of our approach in this work is inspired by another discipline (Speaker recognition) in which similarity of objects is also the center of attention and therefore very important. In speaker recognition (or signature verification) applications, it has been already validated that if two voices (or signatures) are significantly similar – based on a same set of parameters as well as their combinations –; then, they are most probably produced by a unique person [11]. Inspired by this well proved fact, our proposed technique in this paper hypothesizes the same logic with the idea of pattern feature extraction and matching, an area which is widely used in pattern recognition, sequence matching in bio-informatics and machine vision. Here, we extract the CPU utilization pattern of unknown/new Map-Reduce applications for a small amount of data (not the whole data) and compare its results with already known patterns in a reference database to find similarity. Such similarity will show how much an application is similar to another application. As a result, the optimal values of configuration parameters for unknown/new applications can be set based on the already calculated optimal values for known similar application in the database.

III. THEORITICAL BACKGROUND

Pattern matching is a well-known approach – particularly in pattern recognition – to transform a time series pattern into a mathematical space. Such transformation is essential to extract the most suitable running features of an application before comparing it with reference application in a database to find its similar pairs. Such approaches have two general phases: (1) profiling phase, and (2) matching phase. In the profiling phase, the time series patterns of several applications are extracted. After applying some mathematical operations on these patterns, they are stored in a database as references during the matching phase. In matching phase, the same procedure is repeated for an unknown/new application first; and then, the time series of this application are compared with the saved time series of applications in database by using a pattern matching algorithm to find the most similar ones.

A. Uncertain time series

A time series $\varphi_c(\cdot)$ is called certain time series when its data value are fixed/certain: $\varphi_c(\cdot) = [\varphi_c[1], \dots, \varphi_c[N]]$ where $\varphi_c[i]$ in the value of time series at time i . on the contrary, a time series $\varphi_u(\cdot) = [\varphi_u[1], \dots, \varphi_u[N]]$ is called uncertain when there is uncertainty in its data [12]. Therefore, it can be formulated as:

$$\varphi_u[i] = \varphi_c[i] + e_\varphi[i]$$

Where $e_u[i]$ is the amount of error/uncertainty in i^{th} point. Due to this uncertainty, the value of each point is considered as independent random variable with statistical mean ($\mu_\varphi[i]$) and standard deviation ($\sigma_\varphi[i]$). These values are calculated during analyzing time series in profiling phase.

In Map-Reduce application, the length of CPU utilization time series and their values of each point change (due to temporal changes in system such as CPU operation state and device response delay) for several execution of an application even with the same set of configuration parameters. Therefore in our experiments, we consider the CPU utilization time series as an uncertain time series and try to use its statistical information in our similarity measurements.

B. Pattern matching

One of the important problems in data mining is to measure similarity between two data series; similarity measurement algorithms have been frequently used in pattern matching, classification and sequence alignment in bio-informatics. The measurement of similarity between two uncertain time series means to find a function: $SIM(\varphi_u, \varphi_u)$ where $\varphi_u(\cdot)$ and $\varphi_u(\cdot)$ are two time series without the same length. This function is typically designed as $0 \leq SIM(\varphi_u, \varphi_u) \leq 1$, where greater values means higher similarities. In this case, $SIM(\varphi_u, \varphi_u) = 1$ should be obtained for identical series only, and, $SIM(\varphi_u, \varphi_u) = 0$ should reflect no similarity at all. Toward this end, similarity between two uncertain time series is represented by defining a specific distance between them called the “similarity distance”. General approaches like

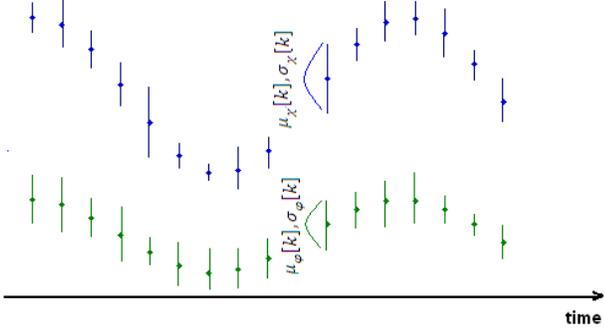


Figure 1. The distance between two uncertain time series and normal distribution of uncertainty in k^{th} points.

Dynamic Time Warping (DTW) cannot be directly utilized to find the similarity due to uncertainty of these time series.

1. Dynamic Time Warping (DTW)

DTW is generally used to calculate the similarity distance between two certain time series of different lengths. A simple method to overcome unevenness of the series is to resample one series to match the other before comparison. This method, however, usually results in unacceptable outcomes as the time series usually would not be logically/correctly aligned. DTW uses a nonlinear search to overcome this problem and map corresponding samples to each other. This method uses the following mathematic recursive formulas to obtain similarity between two certain time series $\varphi_c(\cdot) = [\varphi_c[1], \dots, \varphi_c[N]]$ and $\chi_c(\cdot) = [\chi_c[1], \dots, \chi_c[M]]$ where $N \geq M$.

$$D(i, j) = \begin{cases} D(i, j-1) \\ D(i-1, j) \\ D(i-1, j-1) \end{cases} + d(\varphi_c[i], \chi_c[j]) \quad (1)$$

where $d(\cdot, \cdot)$ is the Euclidean distance between corresponding points in both series as:

$$d(\varphi_c[i], \chi_c[j]) = \|CPU(\varphi_c[i]) - CPU(\chi_c[j])\|$$

Here, $CPU(\varphi_c[i])$ is the value of CPU utilization at time i in φ_c . Results of these formulation is the $D[\varphi_c, \chi_c]$ matrix in which each of its elements - $D(i, j)$ - reflects the minimum distance between $[\varphi_c[1], \chi_c[1]]$ to $[\varphi_c[i], \chi_c[j]]$. As a result, $D(N, M)$ would reflect the similarity distance between φ_c and χ_c . In this case, φ'_c and χ'_c can be made from φ_c and χ_c , respectively, with the same length so that $\chi'_c[i]$ is aligned with $\varphi'_c[i]$; φ'_c and χ'_c are always made from φ_c and χ_c , respectively, by repeating some of their elements based on $D[\varphi_c, \chi_c]$.

Although DTW gives a powerful way to find the similarity between two certain time series with different lengths, it is not directly useful for comparison of two uncertain time series. However as DTW produces another time series (φ'_c, χ'_c) with the same lengths, in this paper we use this ability to make the uncertain time series to be the same length.

After applying DTW on the certain time series ($\varphi_c = \text{mean}(\varphi_u)$, $\chi_c = \text{mean}(\chi_u)$) parts of the two uncertain time series (φ_u, χ_u), the comparison between two uncertain time series with different lengths:

$$\varphi_u[i] = \frac{\text{mean}(\varphi_u[i])}{\varphi_c[i]} + \frac{\text{var}(\varphi_u[i])}{e_{\varphi}[i]}, 1 \leq i \leq N \text{ and}$$

$$\chi_u[j] = \frac{\text{mean}(\chi_u[j])}{\chi_c[j]} + \frac{\text{var}(\chi_u[j])}{e_{\chi}[j]}, 1 \leq j \leq M, \quad \text{where}$$

$N \geq M$, changes to calculate the Euclidian distance between two uncertain but same length time series

$$\varphi'_u[i] = \frac{\text{mean}(\varphi'_u[i])}{\varphi'_c[i]} + \frac{\text{var}(\varphi'_u[i])}{e_{\varphi'}[i]}, 1 \leq i \leq R \text{ and}$$

$$\chi'_u[j] = \frac{\text{mean}(\chi'_u[j])}{\chi'_c[j]} + \frac{\text{var}(\chi'_u[j])}{e_{\chi'}[j]}, 1 \leq j \leq R$$

where $R \geq N, M$. In another word,

$$[\varphi'_c, \chi'_c] = DTW(\varphi_c, \chi_c) \quad (2)$$

In this paper, DTW is utilized to provide same length data series for φ_u and χ_u . It is also worth nothing that DTW does not effect statistical mean and variance of each point in these two uncertain time series. In another word, if DTW maps $\varphi'_u[i]$ to $\varphi_c[j]$, then

$$\begin{cases} \text{mean}(\varphi'_u[i]) = \text{mean}(\varphi_c[j]) \\ \text{var}(\varphi'_u[i]) = \text{var}(\varphi_c[j]) \end{cases}$$

2. Similarity measurement

After DTW, two certain time series φ'_c and χ'_c are similar when the square Euclidian distance between them is less than a distance threshold (τ):

$$DST(\varphi'_c, \chi'_c) = \sum_{i=1}^N (\varphi'_c[i] - \chi'_c[i])^2 \leq \tau$$

However, in uncertain time series φ'_u and χ'_u , the problem is not straightforward as before. In this case, the problem of similarity becomes a probability problem as [12]:

$$\Pr \left(DST(\varphi'_u, \chi'_u) = \sum_{i=1}^N D^2[i] \leq \tau \right) \geq \tau \quad (3)$$

where $D[i]$ is a random variable equal to $\varphi_u[i] - \chi_u[i]$. This means two uncertain time series are similar when the probability of their Euclidian distance is more than a pre-defined threshold ($0 < \tau \leq 1$).

Because $\varphi'_u[i]$ and $\chi'_u[i]$ are independent random variables (figure 1), both $D[i]$ and $DST(\varphi'_c, \chi'_c)$ are also independent random variables. Therefore, if $\langle \mu_{\varphi'}[i], \sigma_{\varphi'}[i] \rangle$ and $\langle \mu_{\chi'}[i], \sigma_{\chi'}[i] \rangle$ are <statistical mean, standard derivation> of $\varphi'_u[i]$ and $\chi'_u[i]$, respectively, according to [12], $DST(\varphi'_u, \chi'_u)$ has normal distribution as:

$$DST(\varphi'_u, \chi'_u) \sim \mathbb{N} \left(\sum_{i=1}^N E(D^2[i]), \sum_{i=1}^N \text{Var}(D^2[i]) \right) \quad (4)$$

where

$$\sum_{i=1}^N E(D^2[i]) = \sum_{i=1}^N \left[\mu_{\varphi'}^2[i] + \sigma_{\varphi'}^2[i] - 2\mu_{\varphi'}[i]\mu_{\chi'}[i] \right] + \sum_{i=1}^N \left[\mu_{\chi'}^2[i] + \sigma_{\chi'}^2[i] \right] \quad (5)$$

and

$$\sum_{i=1}^N Var(D^2[i]) = 4 \sum_{i=1}^N \left[(\sigma_{\varphi'}^2[i] + \sigma_{\chi'}^2[i]) (\mu_{\varphi'}[i] - \mu_{\chi'}[i])^2 \right] \quad (6)$$

Then the standard normal distribution function of $DST(\varphi'_u, \chi'_u)$ can be calculated as:

$$DST_{norm}(\varphi'_u, \chi'_u) \sim \mathbb{N}(0,1) = \frac{DST(\varphi'_u, \chi'_u) - \sum_{i=1}^N E(D^2[i])}{\sqrt{\sum_{i=1}^N Var(D^2[i])}} \quad (7)$$

Therefore the problem in Eqn.(3) changes to:

$$\Pr(DST_{norm}(\varphi'_u, \chi'_u) \leq r_{norm}(\varphi'_u, \chi'_u)) \geq \tau \quad (8)$$

Definition 1: $r_{boundary, norm}$ is a minimum distance bound value that finds the lower bound for the standard normal probability in Eqn.(8). In another word[12]:

$$\Pr(DST_{norm}(\varphi'_u, \chi'_u) \leq r_{boundary, norm}) = \tau \quad (9)$$

Where $r_{boundary, norm} = \sqrt{2} \times \text{erf}^{-1}(2\tau - 1)$ for standard normal distribution and $\text{erf}(\cdot)$ is an error function obtained from statistics tables[13]. When working on $DST(\varphi'_u, \chi'_u)$ instead of $DST_{norm}(\varphi'_u, \chi'_u)$:

$$r_{boundary} = \frac{r_{boundary, norm}^2 - \sum_{i=1}^N E(D^2[i])}{\sqrt{\sum_{i=1}^N Var(D^2[i])}} \quad (10)$$

Definition 2:

Two uncertain time series φ'_u and χ'_u are similar with probability more than τ [12]:

$$\Pr(DST(\varphi'_u, \chi'_u) \leq r) \geq \tau \quad (11)$$

when

$$r \geq r_{boundary} \quad (12)$$

Obviously, $r_{boundary}$ defines the minimum distance between two uncertain series with probability τ . In another word,

$$\Pr(DST(\varphi'_u, \chi'_u) = r_{boundary}) = \tau \quad (13)$$

Based on these equations, we only assume using uncertain time series for the rest of this paper; thus, we will use φ' , χ' , $\mu_{\varphi'}$, $\sigma_{\varphi'}$, $\mu_{\chi'}$ and $\sigma_{\chi'}$ instead of φ'_u , χ'_u , $\mu_{\varphi'_u}$, $\sigma_{\varphi'_u}$, $\mu_{\chi'_u}$ and $\sigma_{\chi'_u}$, respectively.

C. Problem definition

Map-Reduce, introduced by Google in 2004 [14], is a framework for processing large quantities of data on distributed systems. The computation of this framework has two major phases: Map and Reduce.

In the Map phase, after copying the input file to the Map-Reduce file system and split the file into smaller files, data inside the split files are converted into $\langle key, value \rangle$ format (e.g. *key* can be a line number and *value* can be a word in an essay). These $\langle key, value \rangle$ pairs are entered to the mappers and the first part of processing are applied on them. In fact, as mappers in such framework are designed independent, Map-Reduce applications are always naturally ready for parallelization. This parallelization, however, can be bounded sometimes because of other issues such as the nature of the data source and/or the numbers of CPUs have access to the data.

In the Reduce phase, after finishing the Map phase, a network intensive job starts to collect intermediate produced $\langle key, value \rangle$ pairs by mappers to reducers. Here, depending on the Map-Reduce configuration, a sort/shuffle stage may also be applied to expedite the whole process. After that, map operations with the same intermediate *key* will be presented to the same reducers. The result is concurrently produced and written in output files (typically one output file) in the file system.

The process of converting an algorithm into independent mappers and reducers causes Map-Reduce to be inefficient for algorithms with sequential nature. In fact, Map-Reduce is designed for computing on significantly large quantities of data instead of making complicated computation on a small amount of data [15]. Due to its simple structure, Map-Reduce is suffering from several serious issues, particularly in scheduling, energy efficiency and resource allocation.

In distributed computing systems, Map-Reduce has been known as a large-scale data processing or CPU intensive job [3, 15-16]. It is also well known that CPU utilization is the most important part of running an application on Map-Reduce. Therefore, optimizing the amount of CPU an application needs becomes important for customers to hire enough CPU resources from cloud providers as well as for cloud providers to schedule incoming jobs properly.

In this paper, we will study the similarity between uncertain CPU utilization time series of an incoming application with the analyzed applications in a reference database for different sets of configuration parameter values. If the uncertain CPU utilization time series of an unknown/new application is found to be adequately similar to uncertain CPU utilization time series of another application in database—for fairly the same and all sets of configuration parameters values—, then, it can be assumed that the CPU utilization behavior of both applications would be the same for other sets of configuration parameters values as well. This fact can be used in two ways: firstly, if the optimal values of the configuration parameters are obtained for one application, these optimal values lead us to optimal configuration values of other similar applications

Profiling phase

1. For i^{th} application in database (φ_i):
2. For j^{th} set of configuration parameters values (M_j, R_j, FS_j, I_j):
3. Counter=1
4. Do
5. Run application with the j^{th} set of parameters on a small set of data
6. Capture CPU utilization Time Series with SysStat (φ_{ij})
7. Counter = Counter + 1
8. While (Counter <= 10)
9. For k^{th} point in φ_{ij}
10. Calculate $\langle \mu_{\varphi_i}[k], \sigma_{\varphi_i}[k] \rangle$
11. $\mu_{\varphi_i} = [\mu_{\varphi_i}[1], \dots, \mu_{\varphi_i}[N]]$
12. $\sigma_{\varphi_i} = [\sigma_{\varphi_i}[1], \dots, \sigma_{\varphi_i}[N]]$
13. End
14. Save $\{\varphi_i, (M_j, R_j, FS_j, I_j), \mu_{\varphi_i}, \sigma_{\varphi_i}\}$ in Reference database
15. End
16. End

(a)

Matching phase

For a new unknown application (χ):

“Extract Statistical Information”

1. For i^{th} application in database (φ_i):
2. For j^{th} set of configuration parameters values (M_j, R_j, FS_j, I_j):
3. Counter=1
4. Do
5. Run χ with the j^{th} set of parameters values on a small set of data
6. Capture CPU utilization Time Series with SysStat (χ_j)
7. Counter = Counter + 1
8. While (Counter <= 10)
9. Calculate μ_{φ_i} and μ_{χ} under j^{th} set of parameters values
10. $[\mu_{\varphi_i}', \mu_{\chi}'] = \text{DTW}(\mu_{\varphi_i}, \mu_{\chi})$: align mean times series of χ_j to mean time series of φ_i and form new mean time series φ_i' and χ_j'
11. For k^{th} point in both new uncertain time series φ_i' and χ_j'
12. Calculate $\langle \mu_{\varphi_i'}[k], \sigma_{\varphi_i'}[k] \rangle$ and $\langle \mu_{\chi_j'}[k], \sigma_{\chi_j'}[k] \rangle$
13. $\mu_{\chi_j'} = [\mu_{\chi_j'}[1], \dots, \mu_{\chi_j'}[R]]$ and $\sigma_{\chi_j'} = [\sigma_{\chi_j'}[1], \dots, \sigma_{\chi_j'}[R]]$
14. $\mu_{\varphi_i'} = [\mu_{\varphi_i'}[1], \dots, \mu_{\varphi_i'}[R]]$ and $\sigma_{\varphi_i'} = [\sigma_{\varphi_i'}[1], \dots, \sigma_{\varphi_i'}[R]]$
15. End
16. Form $\{\chi_j', (M_j, R_j, FS_j, I_j), \mu_{\chi_j'}, \sigma_{\chi_j'}\}$
17. End
18. End

“Candidate Selection”

19. Set pre-defined Probability threshold ($\tau = 0.95$)
20. For j^{th} set of configuration parameters values (M_j, R_j, FS_j, I_j):
21. For i^{th} application in database (φ_i):
22. Calculate joint mean and variance of distance between φ_i' and χ_j' From Eqn. (5-6)
23. Calculate r_{boundary} from Eqn. (10)
24. $\langle \varphi_i, r_{\text{boundary}} \rangle$ is added to candidature pool of χ_j
25. End
26. End

In candidature pool of χ_j , the application with lowest r_{boundary} is chosen as the highest similar application to χ_j

(b)

Figure 2. the detailed algorithms of profiling and matching phases.

too; secondly, this approach allows us to properly categorize applications in several classes with the same CPU utilization behavioral patterns.

IV. PATTERN MATCHING IN MAP-REDUCE APPLICATIONS

In this section, we describe our technique to find the similarity between uncertain CPU utilization time series of different

Map-Reduce applications. Our approach is consisted of two phases: profiling and matching.

A. Profiling phase

In the profiling phase, CPU utilization time series of several Map-Reduce applications in database along with their statistical information is extracted. For each application, we generate a set of experiments with different values of four Map-Reduce configuration parameters on a given platform. These parameters are: number of mappers, number of reducers, size of split file systems and size of the input file.

The algorithm related to this phase has been shown in Figure 2-a. While running each experiment, the CPU utilization time series of the experiment is gathered to build a trace to be later used as the training data –this statistic can be gathered easily in Linux with the SysStat monitoring package [17]. Within the system, we sample the CPU usage of the experiment in native system from starting mappers till finishing reducers with time interval of one second (Figure 3). Because of the temporal changes, several running of an experiment with the same set of configuration parameters values results in different values in each point of the extracted CPU utilization time series resulting in uncertain CPU utilization time series. Therefore, each experiment with the same set of configuration parameters values is repeated ten times to extract the statistical $\langle \text{mean}, \text{variance} \rangle$ of each point of the time series. Then, the time series with its related set of configuration parameters values as well as its statistical features are stored in the reference database. The algorithm indicates that for the first application, the application is run for the first set of configuration parameters values on a small set of data and repeated 10 times. At the same time as its CPU Utilization Time Series (CUTS) is captured by SysStat package. This application is then re-run for the second set of configuration parameters values and its CUTS is also captured. This procedure is continued for all applications in database to profile different applications with several sets of configuration parameters values (lines 2-13 in Figure 2-a).

B. Matching phase

In the matching phase, the profiling procedure for gathering time series of an unknown/new application is repeated and then followed by the several steps to find its similarity with already known applications. As shown in Figure 2-b, the matching phase consists of two stages: “Statistical information extraction” and “Candidate selection”. In the former, CPU utilization time series of a new unknown application (χ) is captured by SysStat package. Then statistical $\langle \text{mean}, \text{variance} \rangle$ at each point $\langle \mu_{\chi}[k], \sigma_{\chi}[k] \rangle$ of the time series are extracted under several sets of configuration parameters values. To extract this statistical information, the new application is re-run ten times with the same set of configuration parameters values and the CPU utilization time series is captured in each run. As the length of the new application time series is different from the time series of

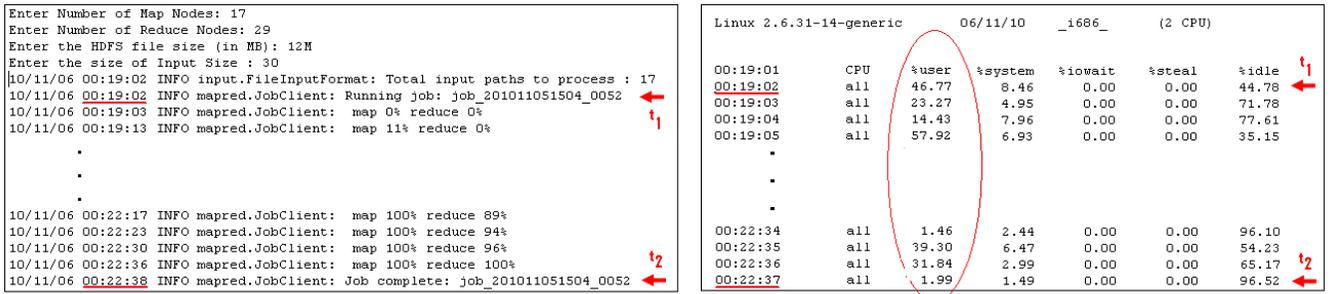


Figure 3. The procedure of capturing CPU Utilization Time Series of a Map-Reduce application

applications in reference database(φ_i), so it is mandatory to make them with the same length. DTW is used here to twist both time series. Then two new uncertain time series are built for each application (φ'_i and χ') which are then analyzed to extract their statistical information at each point $\langle \mu_{\varphi'}[k], \sigma_{\varphi'}[k] \rangle$ and $\langle \mu_{\chi'}[k], \sigma_{\chi'}[k] \rangle$.

In the later stage, Candidate selection, the mathematical analysis described in Section B-2 is applied to calculate the similarity between twisted version of uncertain time series in database (φ'_i) and the new unknown application (χ'). Consequently, based on Eqn.(13) the time series in database which gives the minimum $r_{boundary}$ for predefined Euclidian distance probability (τ) are chosen as the most similar application to the new application in the candidature pool. Raising the value of probability threshold (τ) will reduce the number of applications in candidature pool; and consequently, increases the similarity selection accuracy.

V. EXPERIMENTAL RESULTS

A. Experimental setting

Three standard applications are used to evaluate the effectiveness of our method. Our method has been implemented and evaluated on a pseudo-distributed Map-Reduce framework. In such framework, all five Hadoop daemons (namenode, jobtracker, secondary namenode, datanode and task tracker) are distributed over cores/processors of a single laptop PC. Hadoop writes all files to the Hadoop Distributed File System (HDFS), and all services and daemons communicate over local TCP sockets for inter-process communication. In our evaluation, the system runs Hadoop version 0.20.2 that is Apache implementation of Map-Reduce developed in Java [2]; at the same time, the

SysStat package is executed in another terminal to monitor/extract the CPU utilization time series of applications (in the native system) [17]. For an experiment with a specific set of Map-Reduce configuration parameters values, statistics are gathered from “running job” stage to the “job completion” stage (arrows in Figure 3-left) with sampling time interval of one second. All CPU usages samples are then combined to form CPU utilization time series of an experiment.

We have tested our experiments on a Dell Latitude E4300 laptop with two processors: Intel Centrino model 2.26GHz, 64-bit; 2 x 2GB memory; 80GB Disk. For each application in both profiling and matching phases there are 15 sets of configuration parameters values where the number of mappers and reducers are chosen between 1 to 40 and the size of file system and the size of input file vary between 1Mbyte to 50Mbyte and 10MB to 1GB, respectively.

Our benchmark applications are WordCount, TeraSort, and Exim_mainlog parsing.

- **WordCount[18-19]:** This application reads data from a text file and counts the frequency of each word. Results are written in another text file; each line of the output file contains a word and the number of its occurrence, separated by a TAB. In running a WordCount application on Map-Reduce, each mapper picks a line as input and breaks it into words $\langle key, value \rangle$. Then it assign a $\langle key, value \rangle$ pair to each word as $\langle word, 1 \rangle$. In the reduce stage, each reducer counts the values of pairs with the same key and returns occurrence frequency (the number occurrence) for each word,
- **TeraSort:** This application is a standard map/reduce sorting algorithm – except for a custom practitioner that uses a sorted list of $N - 1$ sampled $keys$ with predefined

ranges for each reducer. In particular, all *keys* with $Sample[i - 1] \leq key \leq Sample[i]$ are sent to i^{th} reducer. This guarantees that the output of the i^{th} reduce are always less than outputs of the $(i + 1)^{th}$ reducer.

- **Exim mainlog parsing [20]:** Exim is a message transfer agent (MTA) for logging information of sent/received emails on Unix systems. This information that is saved in `exim_mainlog` files usually results in producing extremely large files in mailservers. To organize such massive amount of information, a Map-Reduce application is used to parse the data – in an `exim_mainlog` file – into individual transactions; each separated and arranged by a unique transaction ID.

B. Results

Each application is executed on different amounts of input data for different values of the four sets of configuration parameters to form a CPU utilization time series related to these parameters.

1. Application similarity

Table 1 indicates the minimum distance ($r_{boundary}$) between CPU utilization patterns of the pairs of the three applications in our experiments for Euclidian distance probability of %95 ($\tau = 0.95$). The blue numbers indicates the lowest minimum distance between two instances of applications while the red numbers refer to the second lowest minimum distance between the applications. As can be seen, each element in the diagonal line includes one of the blue or red numbers. If the difference between a red and blue number in a column is low (which it is refer to the Table), our hypothesis is that the two applications are most similar when they are run with the same set of configuration parameters values. Based on this hypothesis, the Candidate selection of the matching algorithm in figure 2-b will be changed to that in Figure 4. This hypothesis can lead us to find an optimal set of configuration parameters values for a new unknown applications or a way to categorize the applications with the same CPU utilization pattern in the same class. Assume we have N applications in our database $\varphi = \{\varphi_1, \dots, \varphi_N\}$ and we know their optimal configuration parameters, which this optimally may be on the optimal number of mappers or reducers or optimal usage of CPU resources. For a new unknown application χ , we execute this application and other applications with the same sets of configuration parameters. Then the most similar one –defined as the one which has lowest minimum distance for almost all sets of configuration parameters– is chosen. Therefore, because these two applications can be categorized in the same class, the optimal set of configuration values in the database can also be applicable for optimal running of the new unknown application.

		Exim MainLog Parsing				
WordCount		Set-1	Set-2	Set-3	Set-4	
		Set-1	24044	117017	94472	228071
		Set-2	80648	64063	58351	138222
		Set-3	79431	61232	56114	104255
		Set-4	147014	83655	81434	70427

(a)

		Exim MainLog Parsing				
Terasort		Set-1	Set-2	Set-3	Set-4	
		Set-1	25400	63102	65606	132799
		Set-2	155038	66293	68455	61927
		Set-3	123668	76859	49876	76589
		Set-4	166234	77829	81751	72693

(b)

		WordCount				
Terasort		Set-1	Set-2	Set-3	Set-4	
		Set-1	23184	99242	75076	204094
		Set-2	146568	68586	82173	96227
		Set-3	119112	61228	53843	107029
		Set-4	174182	80311	72822	101707

(c)

TABLE 1. The minimum distance ($r_{boundary}$) between the three applications for $\tau = 0.95$: (a) Exim_mainlog parsing and WordCount, (b) Exim mainlog parsing and TeraSort, and (c) WordCount and TeraSort for different sets of configuration parameters values.

		WordCount				
WordCount		Set-1	Set-2	Set-3	Set-4	
		Set-1	1079	116850	89382	216474
		Set-2	116850	2293	61541	119652
		Set-3	89382	61541	2524	118166
		Set-4	216474	119652	118166	2980

(a)

		Exim				
Exim		Set-1	Set-2	Set-3	Set-4	
		Set-1	325	63003	69015	138062
		Set-2	63003	987	42078	71601
		Set-3	69015	42078	1136	65632
		Set-4	138062	71601	65632	1706

(b)

		Terasort				
Terasort		Set-1	Set-2	Set-3	Set-4	
		Set-1	361	135619	105092	150639
		Set-2	135619	1046	60107	74998
		Set-3	105092	60107	1102	69114
		Set-4	150639	74998	69114	1588

(c)

TABLE 2. The minimum distance ($r_{boundary}$) between each application for $\tau = 0.95$ and different sets of configuration parameters values.

2. Auto-application similarity

A side question is the similarity between different CPU utilization time series/pattern of an application under the same/different sets of configuration parameters values. In another word, if there is no dependency between these parameters and CPU utilization pattern, then, $r_{boundary}$ must be close to the case when an instance of application with a set of configuration parameters is compared with itself. Table 2 refers to this question. The diagonal line is the comparison

between $r_{boundary}$ of an instance of application with itself while the rest are the comparison between two instances of an application. Here, the diagonal values must be the lowest Euclidian distance among the values in Tables 1 and 2. Nevertheless, comparing the first column of Table 1-a and Table 2-a, except first row, shows that $r_{boundary}$ between WordCount and Exim is sometimes much lower than $r_{boundary}$ between two instances of WordCount for set-2, set-3 and set-4. As a result, the CPU behavior of an application changes from one set of parameters to another set.

3. τ and minimum distance ($r_{boundary}$) relation

One of the parameters influencing minimum distance between CPU utilization time series of applications ($r_{boundary}$) is the value of Euclidian distance probability (τ). Euclidian distance probability has a direct relation on the level of similarity. As can be seen in figure 5, in all experiments, increasing τ results in raising the value of $r_{boundary}$. From a mathematical point of view, this fact is correct as refer to Eqn.(9-10), higher value of (τ) should result in higher value for $erf^{-1}(2\tau - 1)$ and therefore higher value for minimum distance ($r_{boundary}$).

C. Future word

One issue that has not been addressed in this paper is to utilize all resources (CPU, Disk and Memory); this requires three uncertain time series to be extracted for one application on our pseudo-distributed platform. Therefore, to find the similarity between two applications, three uncertain time series of the first application should be compared with their related uncertain time series from the second application—this will significantly increase the computation complexity. However, using the other resources patterns may not increase the similarity accuracy. Most of the disk utilization in Map-Reduce comes from copying data to Map-Reduce DFS and vice-versa that are done before/after starting/finishing the application execution. Also, disk used between map and reduce phases to keep the temporary data generated from map phase only become important when mappers produce a large amount of intermediate data for delivering to reducers. Utilizing memory may increase the similarity accuracy as well, but its influence may not be as much as CPU utilization pattern. In fact, memory is generally used for keeping the temporal data during computation that is tightly related to the CPU pattern of an application. Therefore, it may not give more information than CPU pattern. Also, as it was mentioned, Map-Reduce is generally designed to execute CPU intensive jobs; and therefore, it is expected that its CPU pattern be more important than other patterns.

Another issue is about running the application on a cluster instead of pseudo-distributed mode. We expect extra complexity when applications are run on real clusters. Toward this end, for an N-node cluster, three uncertain time series will be extracted from each node of cluster (CPU, Disk and Memory). Therefore, $3N$ time series is obtained for an

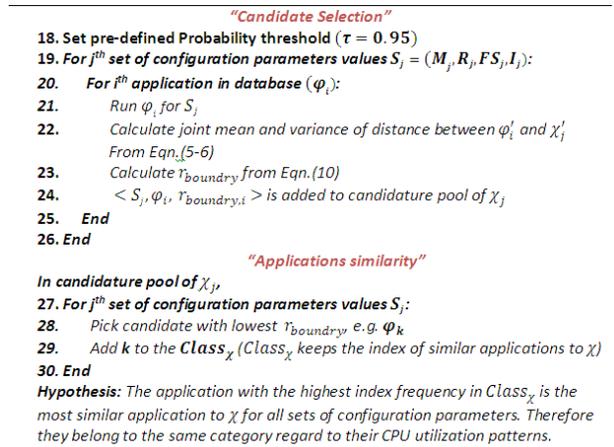


Figure 4. New Candidate selection accompany with applications matching algorithm based on our hypothesis

application. As a result, the similarity problem will turn into comparing $3N$ uncertain time series from one application to its corresponding uncertain time series from another application; a computationally expensive attempt.

Our idea to solve this problem in future is to extract fourier/wavelet coefficients of an uncertain time series and use them instead of original time series; this will result in a much shorter uncertain series than the original series. If all uncertain time series are transformed to Fourier/Wavelet domain by M coefficients; then, the problem of similarity between two applications for pre-defined value of $\tau = 0.95$ changes to obtain minimum distance between corresponding $3N$ Fourier/wavelet coefficients uncertain series, with the same length (M), of two applications. As two new uncertain series have the same length, then simple Euclidian distance calculation instead of DTW can be utilized to find the similarity. However, using Fourier/Wavelet coefficients need to solve some challenging problems such as choosing an appropriate number of coefficients (M) or the suitable wavelet family. Also, as such transformation captures the most important data with high frequency, so it is too likely to lose some important low frequency data when the Fourier/Wavelet family or their coefficient numbers are not chosen properly.

VI. CONCLUSION

This paper presents a new statistical approach to find the similarity among uncertain CPU utilization time series of applications on Map-Reduce clusters. After applying DTW on the time series of two applications, the statistical minimum distance between two uncertain time series/patterns is calculated by assuming independent normal distribution for each point at both time series. Our experiments on three applications (WordCount, Exim Mainlog parsing and TeraSort) show that applications follow different CPU utilization pattern when their data is uncertain.

VII. ACKNOWLEDGMENT

The work reported in this paper is in part supported by National ICT Australia (NICTA). Professor A.Y. Zomaya's work is supported by an Australian Research Council Grant LP0884070.

REFERENCES

[1] Y. Chen, *et al.*, "Towards Understanding Cloud Performance Tradeoffs Using Statistical Workload Analysis and Replay," University of California at Berkeley, Technical Report No. UCB/EECS-2010-81, 2010.

[2] *Hadoop-0.20.2*. Available: <http://www.apache.org/dyn/closer.cgi/hadoop/core>

[3] S. Babu, "Towards automatic optimization of Map-Reduce programs," presented at the 1st ACM symposium on Cloud computing, Indianapolis, Indiana, USA, 2010.

[4] M. Zaharia, *et al.*, "Improving Map-Reduce Performance in Heterogeneous Environments," *8th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2008)*, pp. 29-42, 18 December 2008.

[5] M. Zaharia, *et al.*, "Job Scheduling for Multi-User Map-Reduce Clusters," University of California at Berkeley, Technical Report No. UCB/EECS-2009-55, 2009.

[6] J. Leverich and C. Kozyrakis, "On the Energy (In)efficiency of Hadoop Clusters," *SIGOPS Oper. Syst. Rev.*, vol. 44, pp. 61-65, 2010.

[7] Y. Chen, *et al.*, "Statistical Workloads for Energy Efficient Map-Reduce," University of California at Berkeley, Technical Report No. UCB/EECS-2010-6, 2010.

[8] T. Sandholm and K. Lai, "Map-Reduce optimization using regulated dynamic prioritization," presented at the the eleventh international joint conference on Measurement and modeling of computer systems, Seattle, WA, USA, 2009.

[9] A. Wieder, *et al.*, "Brief Announcement: Modelling Map-Reduce for Optimal Execution in the Cloud," presented at the Proceeding of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing, Zurich, Switzerland, 2010.

[10] A. Wieder, *et al.*, "Conductor: orchestrating the clouds," presented at the 4th International Workshop on Large Scale Distributed Systems and Middleware, Zurich, Switzerland, 2010.

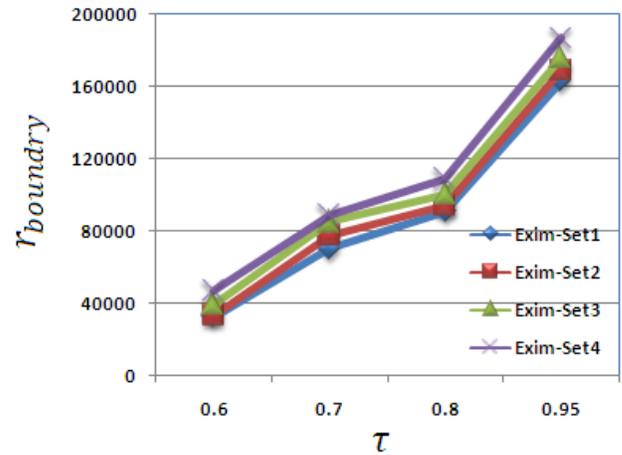
[11] I. I. Shahin and N. Botros, "Speaker identification using dynamic time warping with stress compensation technique" presented at the IEEE Southeastcon '98, 1998.

[12] M.-Y. Yeh, *et al.*, "PROUD: a probabilistic approach to processing similarity queries over uncertain data streams," presented at the 12th International Conference on Extending Database Technology: Advances in Database Technology, Saint Petersburg, Russia, 2009.

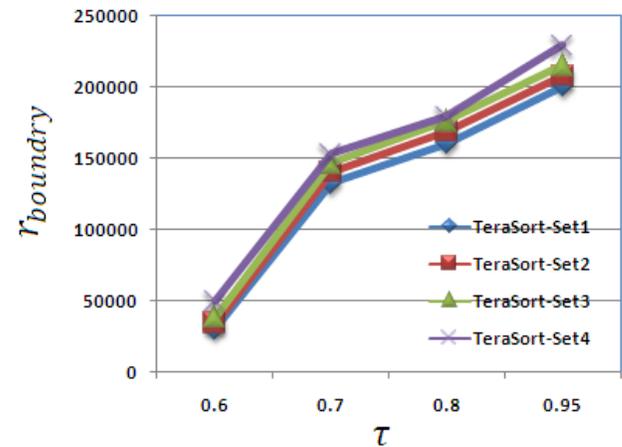
[13] *Error function table*. Available: http://www.geophysik.uni-muenchen.de/~malservi/GlobaleGeophysik2/erf_tables.pdf

[14] J. Dean and S. Ghemawat, "Map-Reduce: simplified data processing on large clusters," *Commun. ACM*, vol. 51, pp. 107-113, 2008.

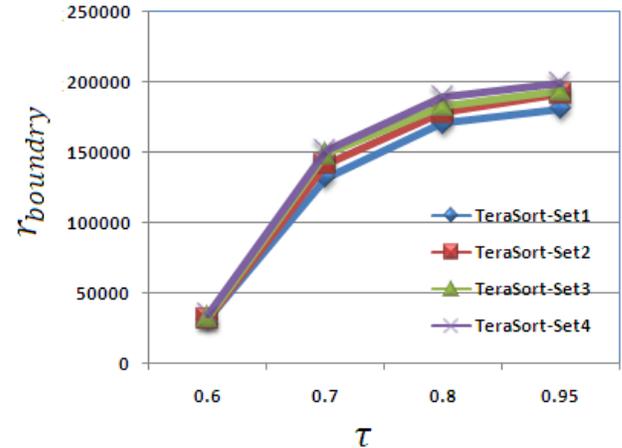
[15] *Hadoop Developer Training*. Available: <http://www.cloudera.com/wp-content/uploads/2010/01/1-ThinkingAtScale.pdf>



(a)



(b)



(c)

Figure 5. The variation of $r_{boundary}$ on the value of Euclidian distance probability (τ). (a) between WordCount-Set1 and Exim-Set 1,2,3,4, (b) between WordCount-Set1 and TeraSort-Set1,2,3,4, and (c) between Exim-Set1 and TeraSort-Set1,2,3,4

[16] S. Groot and M. Kitsuregawa, "Jumbo: Beyond Map-Reduce for Workload Balancing," presented at the 36th

- International Conference on Very Large Data Bases, Singapore 2010.
- [17] *Sysstat-9.1.6*. Available: <http://perso.orange.fr/sebastien.godard/>
- [18] *Hadoop* *wiki*. Available: <http://wiki.apache.org/hadoop/WordCount>
- [19] *Running Hadoop On Ubuntu Linux (Single-Node Cluster)*. Available: [http://www.michael-noll.com/wiki/Running_Hadoop_On_Ubuntu_Linux_\(Single-Node_Cluster\)](http://www.michael-noll.com/wiki/Running_Hadoop_On_Ubuntu_Linux_(Single-Node_Cluster))
- [20] *Hadoop example for Exim logs with Python*. Available: <http://blog.gnucom.cc/2010/hadoop-example-for-exim-logs-with-python/>

School of Information Technologies
Faculty of Engineering & Information
Technologies
Level 2, SIT Building, J12
The University of Sydney
NSW 2006 Australia

T +61 2 9351 3423
F +61 2 9351 3838
E sit.information@sydney.edu.au
sydney.edu.au/it

ABN 15 211 513 464
CRICOS 00026A