Failure Detection and Localization in Component Based Systems by Online Tracking

Haifeng Chen Guofei Jiang Cristian Ungureanu Kenji Yoshihira
{haifeng, gfj, cristian, kenji}@nec-labs.com
NEC Laboratories America, Inc.
4 Independence Way, Princeton, NJ 08540, USA

ABSTRACT

The increasing complexity of today's systems makes fast and accurate failure detection essential for their use in mission-critical applications. Various monitoring methods provide a large amount of data about system's behavior. Analyzing this data with advanced statistical methods holds the promise of not only detecting the errors faster, but also detecting errors which are difficult to catch with current monitoring tools. Two challenges to building such detection tools are: the high dimensionality of observation data, which makes the models expensive to apply, and frequent system changes, which make the models expensive to update. In this paper, we present algorithms to reduce the dimensionality of data in a way that makes it easy to adapt to system changes. We decompose the observation data into signal and noise subspaces. Two statistics, the Hotelling T^2 score and squared prediction error (SPE) are calculated to represent the data characteristics in signal and noise subspaces respectively. Instead of tracking the original data, we use a sequentially discounting expectation maximization (SDEM) algorithm to learn the distribution of the two extracted statistics. A failure event can then be detected based on the abnormal change of the distribution. Applying our technique to component interaction data in a simple e-commerce application shows better accuracy than building independent profiles for each component. Additionally, experiments on synthetic data show that the detection accuracy is high even for changing systems.

Categories and Subject Descriptors

K.6.4 [**Management of Computing and Information Systems**]: System Management; I.2.6 [**Artificial Intelligence**]: Learning

General Terms

Algorithms, Management

Keywords

Subspace decomposition, online tracking, statistics, failure detection, distributed computing, Internet services

1. INTRODUCTION

Recent increases in the complexity of software systems coupled with a demand that they function 24×7 with only minutes of downtime per year, require significant advances in management capabilities of these systems. While designers spend significant effort in building-in resilience to failures, the complexity of these systems makes anticipating all the failure modes all but impossible. A complementary way of increasing availability is to shorten the time to recover from failures when they do occur. This paper proposes a new method for online detection and localization of failures, thus contributing to a significant reduction of the recovery duration.

Detection of failures in large dynamic systems is a challenging task because failure events appear rarely and may not have fixed behavior. The high dimensionality of observation data, together with the changes in normal system behavior, due to software upgrading, web content and user behavior changes, make detection even more difficult. One main contribution of this paper is to propose an *online dynamic tracking approach for high dimensional data*, and apply it to monitor the component interactions for detecting system failures. Since the tracking strategy does not require human intervention, our algorithm can be applied to achieve automatic adaptation to system changes.

In recent years, component based architectures (J2EE, .Net) have become prevalent in developing large scale Web applications. According to the program logic, a component dynamically calls other components to fulfill service requests. We collect the frequencies of interactions between components as a feature to characterize a system's health. Typically, when a component fails, its communication pattern with others will change. By tracking their interactions over time, we can detect system failures.

One main obstacle in applying component interaction tracking is the high dimensionality of observations. For a system with l components, the dimension of interaction data is l^2. However, we observe that in most real systems, each component only interacts with a small number of other components. This allows us to decompose the original high dimensional space into signal and noise subspaces. Two important statistics, the *Hotelling T^2 score* and the *squared prediction error* (SPE) [5], are computed to represent the characteristics of data distribution in the two subspaces. An online tracking algorithm, called *sequentially discounting expectation maximization* (SDEM) algorithm, is employed to learn the distribution of these two statistics. The anomaly of a new sample is then determined by comparing the distribution model before and after that sample is learned. Meanwhile, the signal and noise subspaces are also updated along time. In both SDEM algorithm and subspace updating, an *exponentially weighted moving average* (EWMA) filter is employed to enhance the adaptivity of our algo-

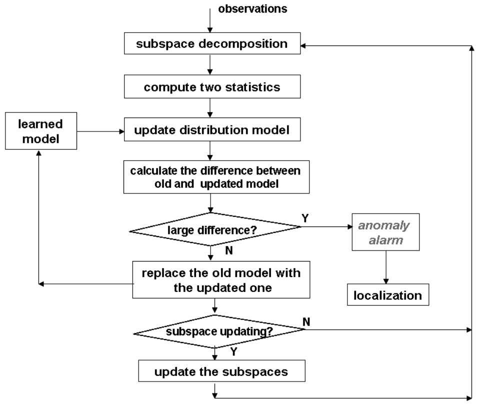

Figure 1: The building blocks of our online detector.

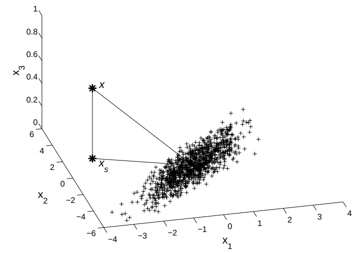

Figure 2: Geometric interpretation of Hotelling T^2 and SPE. $\boldsymbol{x}_s$ is the projection of x onto the signal space (a plane) spanned by the cluster of points. The Hotelling T^2 score indicates the Mahalanobis distance from $\boldsymbol{x}_s$ to the origin in the signal space. SPE measures the distance between x and $\boldsymbol{x}_s$.

and $\{\sigma_{r+1}, \cdots, \sigma_n\}$ belong to the diagonals of Σ_n with $\sigma_r >> \sigma_{r+1}$. The set of orthonormal vectors $U_s = [\mathbf{u}_1, \mathbf{u}_2, \cdots \mathbf{u}_r]$ forms the bases of signal space S_s. The projection matrix P_s onto the signal space would be given by $P_s = U_s U_s^\top$. Since the noise subspace is the orthogonal complement of signal subspace, $S_n = S_s^\perp$, the projection onto noise subspace S_n can be written as $P_n = I - P_s$. Any vector $\boldsymbol{x} \in R^p$ can be represented by a summation of two projection vectors from two subspaces S_s and S_n

$$\begin{aligned} \boldsymbol{x} &= \boldsymbol{x}_s + \boldsymbol{x}_n \\ &= P_s\boldsymbol{x} + (I - P_s)\boldsymbol{x} \ . \end{aligned} \tag{3}$$

The subspace decomposition can also be accomplished by eigen analysis of the correlation matrix C, which is expressed as

$$C = \frac{1}{n} X X^\top = \frac{1}{n} U \Sigma^2 U^\top \tag{4}$$

where the columns of U are actually the eigenvectors of C, and the eigenvalues of C are related to the diagonals of matrix Σ.

Once the observations have been decomposed into signal and noise subspaces, we extract some statistics to describe the data distributions in two subspaces. One is the Hotelling T^2 score, which measures the variation of each sample in signal subspace. For a new sample vector $\boldsymbol{x}$, it is expressed as

$$T^2 = \boldsymbol{x}^\top U_s \Sigma_s^{-1} U_s^\top \boldsymbol{x} \ . \tag{5}$$

Another statistic, the squared prediction error (SPE), indicates how well each sample conforms to the model, measured by the projection of sample vector on the residual space

$$SPE = \|P_n \boldsymbol{x}\|^2 = \|(I - P_s)\boldsymbol{x}\|^2 \ . \tag{6}$$

The geometric interpretation of these two statistics is shown in Figure 2. In this figure, the signal subspace is constructed by 1000 normally distributed 3D samples in a 2D plane. Given a new sample $\boldsymbol{x}$, its projection onto the signal subspace (the plane) is denoted as $\boldsymbol{x}_s$. The Hotelling T^2 score reveals the Mahalanobis distance from $\boldsymbol{x}_s$ to the origin in the plane. SPE measures the distance from the sample $\boldsymbol{x}$ to its projection in the signal space $\boldsymbol{x}_s$. Note in this example the data are centered, then the value T^2 is also related to the Mahalanobis distance from the projected sample to the mean of all samples.

If the data obey a multivariate normal distribution, the Hotelling T^2 score is χ^2 distributed. Then it is easy to set a threshold for this statistic based on a significance level. Similarly if we assume the noise $\boldsymbol{x}_n$ is normally distributed, a control limit for SPE can be obtained to distinguish outliers from inliers. The work in [8][6] follows these assumptions in the field of process control. However, in real situations the data are arbitrarily distributed and unknown. It is hard to determine those thresholds. For example, if the data are bimodally distributed in the signal subspace with large gap between two distributions, then the threshold for T^2 is meaningless. Furthermore, the data distribution is sometimes changing over time, the thresholds that are determined during training would become invalid after a certain time period. In order to avoid any prior assumptions, we present the SDEM algorithm to dynamically track the distribution of those two statistics. Note here we only employ two statistics to represent the original data. Even though such representation works well in our failure detection experiments, it is our future research to apply more statistics to sufficiently reveal the original high dimensional data distribution. For instance, we can decompose the original data space into several subspaces instead of only two, and extract statistics to represent the distribution of data projections in every subspace.

3.2 Online Detector

The sequentially dynamic expectation maximization (SDEM) is a sub-algorithm of SmartSifter, an online unsupervised outlier detector developed by Yamanishi et al.[7]. It uses a Gaussian mixture model to represent the probability density over the domain of continuous variables $\boldsymbol{z}$, $p(\boldsymbol{z}|\boldsymbol{\theta}) = \sum_{i=1}^{k} c_i p(\boldsymbol{z}|\boldsymbol{\mu}_i, \Lambda_i)$, where k is a positive integer, $c_i \geq 0$, $\sum_{i=1}^{k} c_i = 1$ and each $p(\boldsymbol{z}|\boldsymbol{\mu}_i, \Lambda_i)$ is a d-dimensional Gaussian distribution with density specified by mean $\boldsymbol{\mu}_i$, and covariance matrix Λ_i

$$p(\boldsymbol{z}|\boldsymbol{\mu}_i, \Lambda_i) = \frac{1}{(2\pi)^{d/2}|\Lambda_i|^{1/2}} exp(-\frac{1}{2}(\boldsymbol{z} - \boldsymbol{\mu}_i)^\top \Lambda_i^{-1} (\boldsymbol{z} - \boldsymbol{\mu}_i))$$

where $i = 1, \cdots, k$ and d is the dimension of each datum. In our failure detection algorithm, $\boldsymbol{z}$ is the vector of two statistics as described in Section 3.1 and $d = 2$. We set the parameter vector

$$\boldsymbol{\theta} = (c_1, \boldsymbol{\mu}_1, \Lambda_1, \cdots, c_k, \boldsymbol{\mu}_k, \Lambda_k) \ .$$

Every time the datum is input, the SDEM algorithm, as described in Figure 3, estimates the parameter $\boldsymbol{\theta}$ and hence learns the distribution model. The EWMA filter is employed in the parameter estimation in order to discount past examples. The forgetting parameter ρ is related to the degree of discounting. Intuitively, the smaller ρ is, a larger effect the SDEM algorithm has from past examples. Such mechanism makes the SDEM adaptive to non-stationary data sources, e.g., when drifting sources of time series are tackled.

Another parameter α is introduced in the SDEM algorithm in order to improve the stability of the estimates of c_i, which is set