Detecting Anomalies for High Performance Computing Resilience

Shamir J. Quiñones Dueño, Emmanuel Avilés Sáez,
Yael M. Camacho Bonaparte, and Houssain Kettani,
Electrical and Computer Engineering and Computer Science Department
Polytechnic University of Puerto Rico
San Juan, PR 00919

George Ostrouchov
Computer Science and Mathematics Division
Oak Ridge National Laboratory
Oak Ridge, TN 37831

***Abstract*—Supercomputers are being used increasingly by scientists and engineers to process data intensive applications. System restore points (checkpoints) are used to restore a process at a previously saved state if a system component fails. When failures become too frequent, it is no longer feasible to make application progress with checkpoints. Proactive measures that migrate application components to healthy nodes can increase time between failures and enable application progress. However, proactive measures require timely system information and an ability to predict where failures are likely to occur. This project uses data collected from system nodes to identify anomalous node behavior. Detecting anomalies is the first step to identifying failures and eventually developing a failure prediction capability. The main result of this project is a number of analysis tools for anomaly identification that are based on the R open source software environment for statistical computing and graphics and on the [1]GGobi open source visualization program for exploring high-dimensional data. The tools do not assume any specific set of system attributes. Given a large collection of system attributes recorded at some time intervals, the tools use only those attributes that contain information. Considering the informative attributes in a high-dimensional space, the tools identify anomalies and automatically find the attributes that are most responsible for the identified anomalies. Further exploration of the anomalies and their attributes is enabled by the GGobi visualization program.**

Keywords-High Performance, Failures, Nodes.

I. INTRODUCTION

As high performance computing systems get larger and more complex, time between node failures decreases. Proactive measures that migrate application components to healthy nodes can increase time between failures and enable application progress. However, proactive measures require timely system information on hardware and software failures and an ability to predict where failures are likely to occur. Detecting anomalies is the first step to identifying failures and eventually developing a failure prediction capability. The main result of this project is a number of analysis tools for anomaly identification that are based on the R open source software environment for statistical computing and graphics and on the GGobi open source visualization program for exploring high-dimensional data. The tools developed are for numerical data, they do not assume any specific set of system attributes, and they use only those attributes that display variability.

II. CLUSTERING

Given a large collection of system attributes recorded at regular time intervals, the data are divided into node-periods and one set of derived attributes is constructed for each node-period. These node-periods are then clustered. Clustering is a method of unsupervised learning in which high-dimensional data sets are grouped together based on the relationships among them. A distance metric is used to measure the relationships among the data. We use the Euclidean distance on data that is scaled to have the same variability in each dimension so that results are invariant to units of attribute measurement. A single linkage hierarchical clustering algorithm constructs a dendrogram from the data. Single linkage clustering can be viewed as a recursive process of pairwise merging of the closest pairs. It is this merging process that defines the dendrogram tree structure.

We use the R system for statistical computing and graphics [2] and its single-linkage clustering facilities to process and cluster the data. The result is a cluster dendrogram that identifies the clusters for a given linkage value. Single linkage clustering can also be viewed as the computation of the minimum spanning tree of the data. The dendrogram of our data is shown in Figure 1. The large amount of data results in overplotting in the detailed portion of the dendrogram. However, we are interested in the larger structures that are shown in red by cutting the dendrogram at 20 clusters, as shown in Figure 2. After excluding the largest cluster with about 80% of the data, the remaining 19 much smaller clusters can be studied further as potential anomalies.

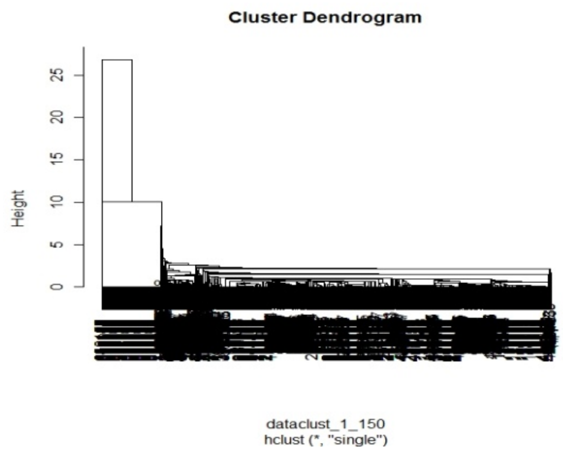

Figure 1. Hierarchical cluster dendrogram of data collected from xtorc.

III. ANOMALY DETECTION AND EXPLORATION

Starting with a set of clusters defined by the dendrogram in Figure 2, we compute an anomaly measure that is inversely proportional to cluster size. Node-periods are then displayed by GGobi as two-dimensional scatter plot projections from the high-dimensional attribute space. Using R to control GGobi[3], we color node-periods by their anomaly measure ranging from blue as not anomalous, through yellow to red as highly anomalous.

We accomplish this by writing R scripts that transfer the cluster colors computed in R to the GGobi display.

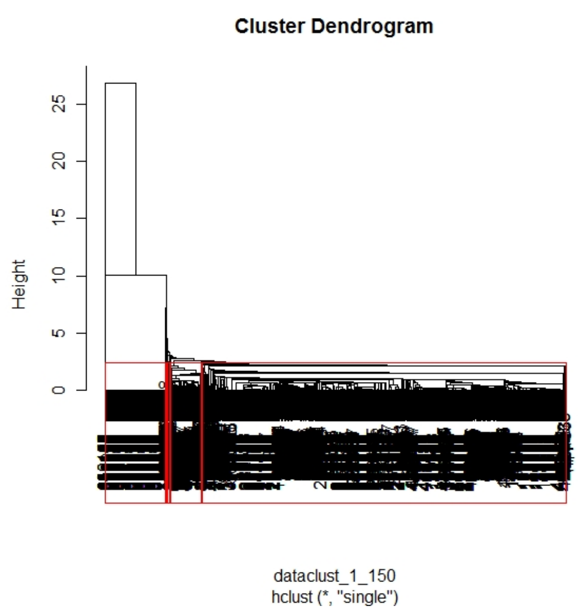

Figure 2. The same dendrogram cut at 20 clusters.

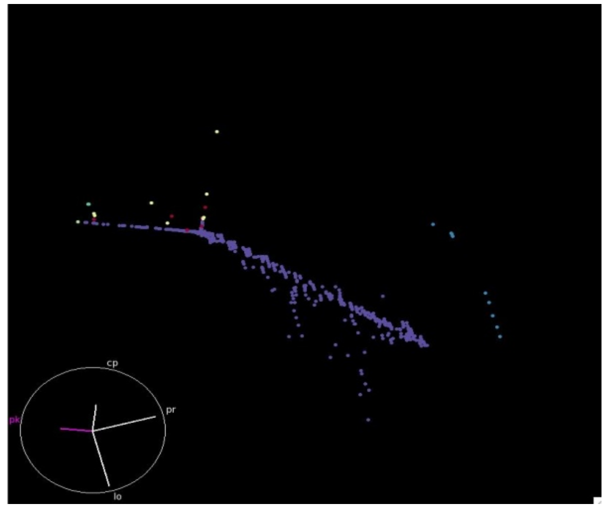

Figure 3. A GGobi projection of node-periods in a combination of four attributes emphasizing anomalies.

GGobi can then be used to explore the anomalies in any projection of the multivariate attribute space. The projection in Figure 3 was selected for its visual separation of the anomalies from the main body of the data. It also identifies the four attributes that are responsible for the separation. Figure 4, shows a rectangular arrangement of the node-periodes, where periods are on the vertical axis increasing from the bottom to the top and nodes are on the horizontal axis. Same colors are used in both figures. Note that some nodes are missing because they were down.

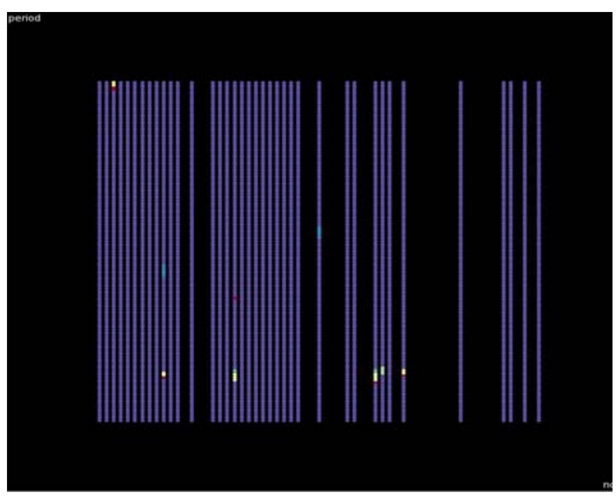

Figure 4. A GGobi projection of node-periods in a node by period arrangement. Faults, which show up as anomalies, were injected during data collection.

IV. DISPLAYING HIGH-DIMENSIONAL ATTRIBUTE RELATIONSHIPS

Given a set of anomalous clusters, it is interesting to know which attributes are most responsible for each cluster's separation from the main body of the data. To explore how to do this in an automated way, we compute the contribution of each attribute to the distances between a given small cluster and the largest cluster containing the majority of the data. The four attributes with the largest contribution are used to display the node-periods in a series of conditional scatter plots [2]. The conditional scatter plots show a four-dimensional relationship by partitioning the data on two attribute domains and displaying a scatter plot of two other attributes for each partition. Typically, the four-dimensional relationship of the displayed attributes provides an explanation of the anomalies and clues to why they are different from the rest of the node-periods.

For example, Figure 5 shows scatter plots for cpu_wio by bytes_in conditioned on partitions of pkts_in and pkts_out. Similarly, Figure 6 shows scatter plots of load_one by cpu_wio conditioned on partitions of pkts_in and proc_run. In both cases, the range of the conditioning variable is shown by the boldly colored section at the top of each scatter plot. While this technique shows promise, better automated attribute selection is needed. Selecting just on maximum contribution to distances does not provide sufficient separation of anomalous data.

V. CONCLUSION

We have developed a set of analysis tools that can be used in high performance computing operations to detect and analyze anomalous behavior. Identifying anomalies is the first step toward timely failure detection. The development of these tools required learning many new tools and concepts including the Unix operating system, the R environment for statistical computing and graphics, the GGobi visualization program, clustering techniques, high-dimensional data, and statistics. Overall, the experience was mainly educative.

VI. FUTURE WORK

Detecting anomalies leads to identifying failures and eventually to the development of a failure prediction capability. Future work includes classification of anomalies into those that require attention and those that are only informative. A set of discovered anomalies can be used to collect data segments that contain predictive information about those anomalies. With appropriate data, failure predictive capabilities can be built.

VII. ACKNOWLEDGEMENTS

This work was funded in part by thank the Department of Energy (DOE) under the Faculty and Student Teams (FaST) program at Oak Ridge National Laboratory (ORNL). It was also funded in part by the Department of Defense and the Extreme Scale Software Center at the Oak Ridge National Laboratory.

REFERENCES

[1] GGobi data visualization system, www.ggobi.org

[2] The R project for statistical computing, www.r-project.org

[3] D. Cook and D. Swayne, "Interactive and Dynamic Graphics for Data Analysis, with R and Ggobi," Springer XVIII, 190 p., 2007.

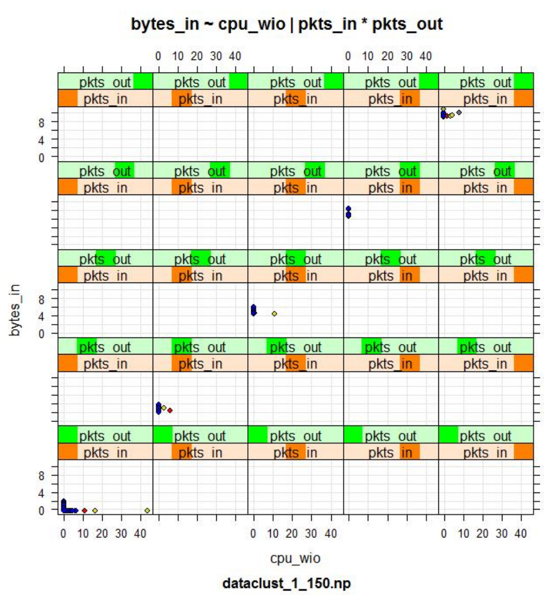

Figure 5. Conditional scatter plot series showing a four-dimensional explanation of the anomalies.

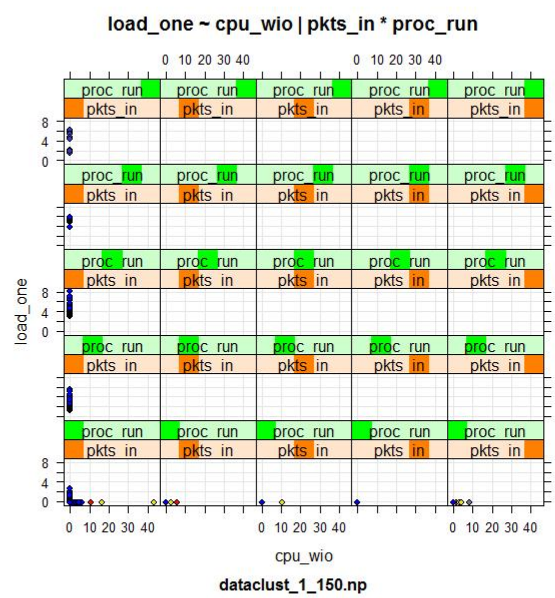

Figure 6. The same data showing another four-dimensional explanation discovered in GGobi.