

Visualizing Multidimensional Health Status of Data Centers

Tommy Dang
Dept. of Computer Science
Texas Tech University
Lubbock, TX, USA
tommy.dang@ttu.edu

Abstract—Monitoring data centers is challenging due to their size, complexity, and dynamic nature. This project proposes a visual approach for situational awareness and health monitoring of high-performance computing systems. The visualization requirements are expanded on the following dimensions: 1) High performance computing spatial layout, 2) Temporal domain (historical vs. real-time tracking), and 3) System health services such as temperature, CPU load, memory usage, fan speed, and power consumption. To show the effectiveness of our design, we demonstrate the developed prototype on a medium-scale high-performance computing system of 10 racks and 467 hosts. The work was developed using feedback from both industrial and academic domain experts.

Index Terms—Data center visualization, high-performance computing systems, Redfish RESTful API, Nagios Core, base-board management controller

I. INTRODUCTION

Data centers are increasingly complex and hence monitoring such systems is a daunting task for system administrators. In 2013, the Distributed Management Task Force (DMTF) released *Redfish* [1], an open industry standard specification for server configuration that aims to supersede IPMI over the network (IPMI over LAN). The web servers are expected to provide end users with simple, secure management of scalable platform hardware by enhancing security and reliability to Baseboard Management Controller (BMC). Redfish uses industry standard RESTful API architecture over Hyper Text Transfer Protocol Secure (HTTPS) using JavaScript Object Notation (JSON) format based on Open Data Protocol (OData) as depicted in Figure 1. Redfish-enabled technologies enable IT systems to operate in harmony. Most importantly, a Redfish-enabled system delivers, and single API for everything [2]. One of the critical advantages of REST APIs is that it introduces a great deal of flexibility, allowing any software clients with RESTful capabilities to interact with *Redfish* firmware.

Building on top of Redfish RESTful interface, this paper introduces a visual analytic prototype for monitoring high-performance computing (HPC) system events. The goals of this prototype are: 1) to monitor the multidimensional health status of multiple hosts and racks in real-time, 2) to support system administrators in detecting unusual correlation of these services in a complex and highly dynamic system, and 3) to help in performing system troubleshooting and debugging. We demonstrate our prototype at a data center at a university.




Fig. 1. Redfish reference architecture [2].

The rest of this paper is organized as follows: We first summarize existing techniques for visualizing HPC in the next Section. Then we provide an overview of visualization tasks and describe the design and supported interactive features of our prototype in detail. In Section IV, we present the results of an informal study with HPC experts from a university and an industrial company. We argue that using visual features of pairwise projection we can capture abnormal correlations between various health dimensions in the data center. More importantly, these visual features can be computed automatically and notify the system administrator significant events. Finally, we conclude our paper with future plans for visual monitoring and predicting through the integration with Nagios and Redfish API.

II. EXISTING APPROACHES

Although not fully explored, there are many existing tools that support monitoring high-performance computing systems. For example, LLView [3] is a client-server based application which allows monitoring the utilization of clusters controlled by batch systems like IBM LoadLeveler, PBSpro, Torque, or IBM Blue Gene system database. However, multidimensional analysis proposed in this work is not in focus of LLview.




Fig. 2. Main window of LLview [3]: Graphical monitoring of batch system controlled cluster.

Ganglia [4] is an open source PHP-based web front-end interface that allows users to gather information via real-time dynamic web pages. This information, including CPU usage, memory usage, disk usage, network statistics, the number of running processes, is plotted on similar graphs. It leverages the widely used XML technologies for data representation. The advantage of real-time responsive on the web front-end, however, leads to high latency due to the size of the XML tree. Thus, Ganglia is not practical on the less powerful machines when the amount of data is large.

Nagios Core [5] is another open-source monitoring system that is capable of handling a variety of servers and operating systems with the industry standard. It provides a primary web interface for the core monitoring engine as depicted in Figure 3(b). To trace a problem, however, an administrator cannot capture a holistic monitoring view by using the Nagios web interface. Even though basic filtering operations are provided, system status overview, which is useful to correlate the isolated temporal/spatial issues, can be lost [6]. Figure 3(a) shows an example of the JSON query on CPU temperatures (on the left) and the corresponding results (on the right) returned from the server.

These existing tools inspect system status independently. This paper takes into account the correlations between various dimensions such as CPU temperature, memory usage, fan speed, power consumption, and I/O bandwidth. Analyzing the relationships of multidimensional data is essential to understand system behaviors as a whole [7].




Fig. 3. Nagios Core [5] interfaces: (a) System summary in a simple listing format and (b) Json query on CPU temperatures and the corresponding results.

III. VISUALIZATION COMPONENTS

Through in-depth discussions with domain experts, we have identified the set of design goals: 1) Provides spatial and temporal overview across hosts and racks, 2) Allows system administrators to quickly narrow down to the event of interest for system debugging, and 3) Inspects the correlation of system health services in a single view. To meet these goals, this paper proposes several visualization tasks:

- **Overview Display (T1).** Display an overview of real-time system status on the corresponding spatial layout [8].
- **Details-On-Demand (T2).** Users can inspect multidimensional historical data of a host in the system via a simple click [9].
- **Filtering (T3).** Highlight critical events on a host [10] and the associated time stamps [11].
- **Multidimensional analysis (T4)** Explore the correlations between dimensions [12], [13] such as fan speeds, memory usage, and power consumption.

We leverage the Nagios Core engine for data retrieval through a RESTful API web interface. In other words, we iteratively request health status of every host in the system. For each host, we obtain a set of updated status as shown on the right panel of Figure 3(a). In the next section, we present our approach for displaying status updates of hundreds of hosts




Fig. 4. Visualizing power consumption on Wednesday, October 17, 2018, of the 467-node Quanah cluster at Texas Tech University.

within a single view. In the next section, we first introduce the HPC system at Texas Tech University and then visualization components and the supported interactions.

A. HPC system spatial layout

Figure 4 shows a snapshot power consumption of the 467-node Quanah cluster at Texas Tech University at noon on Sunday, May 13, 2018 (the visualization task **T1**). Within a rack, hosts are listed top down. Rows are power consumption time series (the temporal distance between two consecutive cells is 3 minutes, the updating period set up on the cluster. Rectangles on each row are the computed power consumption, which is colored: red is high power consumption nodes while blue is low power consumption nodes.

B. Multidimensional analysis of health status

Through our discussions with domain experts, it is essential to have a holistic view of the entire data center on multiple dimensions for system monitoring and especially diagnosis. Therefore, our system supports a multi-axis visualization of hosts on demand (visualization task **T2**). In particular, when users mouse over a host, a pop-up window is displayed to unveil the details of host information as depicted in Figure 5. Historical temperatures of each CPU in the selected host are presented in line graphs: The vertical axis represents the CPU

temperature (from 0 to 100 degree of Fahrenheit) while the horizontal axis represents the time.

The spider or cobweb chart [14] in the lower panel of the pop-up window shows health dimensions of the selected host (visualization task **T4**) which are grouped by category: CPU temperature (degrees Fahrenheit), job load (number of jobs assigned on the the given host), memory usage (from 0% to 100%), fan speeds, and power consumption (in watts). The purpose of this chart is to enable users to detect similar multidimensional patterns and how they change over time.

We decide to use the spider chart to display multidimensional health status of data center since the mental images of spider charts allow users to capture and compare historical data between different hosts quickly. Figure 6 shows four examples of spider charts from four random hosts in the Quanah cluster at Texas Tech University. In particular, each closed curve on the spider charts represents a multidimensional observation. The observed period is from 11 am to 7 pm on Wednesday, October 17, 2018.

Users can request to show multidimensional spider chart across hosts in the system at a time stamp. Figure 7 shows our multidimensional visualization of 467 hosts at Texas Tech University at 2 pm on Wednesday, October 17, 2018. Brushing and linking between the summary view and HPC spatial layout




Fig. 5. Visualizing CPU temperature on Wednesday, October 17, 2018 of the 467-node Quanah cluster. Hosts experiencing a sudden temperature increase of 25 degrees in 3 minutes.




Fig. 6. Visualizing multidimensional health status of four different hosts in the data center.

can be done via *mouse over*. Our prototype also supports a range of interactive features, such as zooming and filtering (visualization task T3). In Figure 5, we only display hosts which experienced a sudden CPU heat up of over 25 degrees Fahrenheit within the two consecutive time stamps.

C. Implementation

Our prototype is implemented in D3.js [15]. The application, source code, sample data, and demo video are provided via our GitHub project repository, located at <https://github.com/iDataVisualizationLab/HPCC>.

IV. DISCUSSION AND FUTURE WORK

Our multidimensional analysis approach received positive feedback from both industrial and academic domain experts and encouraged to explore further. One potential direction is inspecting pairwise correlations within a 2D projection [16]. Figure 8 show an example of the 467-node Quanah cluster projected on fan speed (x axis) vs. CPU temperature (y axis). The plot reveals two clusters along the x axis and an outlier (the *compute 4-17*) at the lower left corner which has low fan speed and temperature.

Instead of having a human looking into every single scatter-plot, visual features [17] can help to simplify this process by




Fig. 7. Visualizing multidimensional health status at 2 pm on Wednesday, October 17, 2018, of the 467-node QuanaH cluster at Texas Tech University.




Fig. 8. Visualizing multidimensional health status of four different hosts in the data center.

highlighting only unusual correlations. The following examples in Figure 9 summarize possible 2D projections that can be captured using these visual features [18].

V. CONCLUSION

This paper presents a graphical tool for situational awareness and multidimensional health status of data centers using real-time data gathered through the industry-standard Redfish protocol and API (v.s. the aging IMPI protocol). The system has two components: visual feature extraction for pairwise projection and spider chart for visualizing high-dimensional health status. Our prototype supports a wide range of interactive features such as brushing and linking, zooming, and filtering.

In future work, we want to incorporate machine learning framework, such as TensorFlow [20] to predict the system




Fig. 9. Some example scatterplots and their measures [19]. In each row, the scatterplots with a low score on the associated measures are on the left while the scatterplots with a high score on the associated visual features are on the right.

health status. The model will be trained on historical data, make real-time predictions, and raise the alarm to the system administrator for timely actions. Google recently released tensorflow.js (<https://js.tensorflow.org/>) which can be naturally adapted in our project. Virtual Reality and Augmented Reality (VR and AR) are also interesting future extensions. VR and AR interfaces enable real-time monitoring beyond traditional displays by experiencing interactive visualization techniques on mobile devices as well as within immersive environments.

REFERENCES

- [1] S. D. Organization, "Distributed management task force," 2013, <https://www.dmtf.org/standards/redfish>.
- [2] V. of the Industry, "How redfish specifications can improve simplicity, visibility and control of it," 2016, <https://datacenterfrontier.com/redfish-specifications/> [Accessed date: October, 2018].
- [3] C. Karbach and J. Valder, "System Monitoring: LLview," in *Computational Science and Mathematical Methods*. Introduction to the programming and usage of the supercomputer resources at Jlich November 2015 (Course no. 78a/2015 in the training programme of Forschungszentrum Jlich), Jlich (Germany), 26 Nov 2015 - 27 Nov 2015, Nov 2015. [Online]. Available: <https://juser.fz-juelich.de/record/279901>
- [4] M. L. Massie, B. N. Chun, and D. E. Culler, "The ganglia distributed monitoring system: design, implementation, and experience," *Parallel Computing*, vol. 30, no. 7, pp. 817–840, 2004.
- [5] W. Barth, *Nagios: System and network monitoring*. No Starch Press, 2008.

- [6] N. Andrienko, G. Andrienko, and P. Gatala, "Exploratory spatio-temporal visualization: an analytical review," *Journal of Visual Languages & Computing*, vol. 14, no. 6, pp. 503–541, 2003.
- [7] W. Allcock, E. Felix, M. Lowe, R. Rheinheimer, and J. Fullop, "Challenges of hpc monitoring," in *SC '11: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, Nov 2011, pp. 1–6.
- [8] D. A. Keim, C. Panse, and M. Sips, "Information visualization : Scope, techniques and opportunities for geovisualization," in *Exploring Geovisualization*, J. Dykes, Ed. Oxford: Elsevier, 2004, pp. 1–17.
- [9] B. Shneiderman, "The eyes have it: A task by data type taxonomy for information visualizations," in *Proceedings of the 1996 IEEE Symposium on Visual Languages*, ser. VL '96. Washington, DC, USA: IEEE Computer Society, 1996, pp. 336–. [Online]. Available: <http://dl.acm.org/citation.cfm?id=832277.834354>
- [10] R. Amar, J. Eagan, and J. Stasko, "Low-level components of analytic activity in information visualization," in *Proc. of the IEEE Symposium on Information Visualization*, 2005, pp. 15–24.
- [11] T. N. Dang and L. Wilkinson, "Timeexplorer: Similarity search time series by their signatures," in *Proceedings, Part I, of the 9th International Symposium on Advances in Visual Computing - Volume 8033*. Berlin, Heidelberg: Springer-Verlag, 2013, pp. 280–289.
- [12] A. Anand, L. Wilkinson, and T. N. Dang, "Visual pattern discovery using random projections," in *Visual Analytics Science and Technology (VAST), 2012 IEEE Conference on*, 2012, pp. 43–52.
- [13] T. N. Dang, A. Anand, and L. Wilkinson, "TimeSeer: Scagnostics for high-dimensional time series," *IEEE Trans. Vis. Comput. Graph.*, vol. 19, no. 3, pp. 470–483, March 2013.
- [14] N. Bremer, "A different look for the d3.js radar chart," 2015, <https://www.visualcinnamon.com/2015/10/different-look-d3-radar-chart>[Accessed date: October, 2018].
- [15] M. Bostock, V. Ogievetsky, and J. Heer, "D3 data-driven documents," *IEEE Trans. Vis. Comput. Graph.*, vol. 17, no. 12, pp. 2301–2309, 2011.
- [16] L. Wilkinson, A. Anand, and R. Grossman, "Graph-theoretic scagnostics," in *Proceedings of the IEEE Information Visualization 2005*. IEEE Computer Society Press, 2005, pp. 157–164.
- [17] L. Wilkinson, A. Anand, and Grossman, "D3 data-driven documents," *IEEE Trans. Vis. Comput. Graph.*, vol. 17, no. 12, pp. 2301–2309, 2011.
- [18] T. Dang and L. Wilkinson, "Transforming scagnostics to reveal hidden features," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 1624–1632, Dec 2014.
- [19] T. Dang and Wilkinson, "Scagexplorer: Exploring scatterplots by their scagnostics," in *Proceedings of the 2014 IEEE Pacific Visualization Symposium*, ser. PACIFICVIS '14. Washington, DC, USA: IEEE Computer Society, 2014, pp. 73–80. [Online]. Available: <http://dx.doi.org/10.1109/PacificVis.2014.42>
- [20] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>