A Comprehensive Informative Metric for Summarizing HPC System Status

| Yawei Hui | Byung-Hoon Park parkbh@ornl.gov Computer Science and Mathematics Division | | Christian Engelmann engelmannc@ornl.gov Computer Science and Mathematics Division | | |
|---|---|--|---|--|--|
| huiy@ornl.gov | | | | | |
| Computer Science and Mathematics Division | | | | | |
| Oak Ridge National Laboratory, | Oak Ridge | National Laboratory, | Oak Ridge National Laboratory, | | |
| Oak Ridge, TN, USA | Oak | Ridge, TN, USA | Oak Ridge, TN, USA | | |
| | In | troduction | | | |
| The Titan supercomputer at the Oak Ridge Leadership Computing Fa 27 Peta FLOP Cray System • 18,688 AMD 16-core Opteron CPUs • 693.5 TiB M Tesla K20X GPUs • 400 hundred cabinets with 24 blades per cabinet and 4 nodes per | acility (OLCF): emory • 18,688 NVidia er blade | The Log processing by Spark and Cassandra-based ANalytics (LogSCAN) [1] framework: Residing on the Compute and Data Environment for Science (CADES) cloud at ORNL • 256 cores, 328GB memory, and 1TB SSD storage • Titan system log events are stored using Cassandra • Data analyses are performed using Spark | | | |

Titan System Log Event Data Characteristics

| Counts F | Percentage ID Description | | |
|--|---|---|---|
| 10 | 0 1 DVS Confusion | | |
| 2,998,492 | 2 2 NVRM Xid | | |
| 5,671,348 | 4 3 Machine Check Exception (MCE) | | |
| 1,229 | 0 4 NVRM DBE | | |
| 49 | 0 5 Unknown GPU Error (UGE) | | I. J. M. M. M. MARTICALING INTERVIEW JCEPTCONCECCUTION INTERVIEW. |
| 302,969 | 0.2 6 Graphics Engine Error (GEE) | Jan 28 Jan 31 Feb 3 Feb 9 | Feb 12 Feb 15 Feb |
| 5,732 | 0 7 Kernel Panic | 100 | |
| | | | |
| 782,337 | 0.5 8 Out of Memory (OOM) | $Feb^{1/5} Nar^{1/5} Apr^{1/5} Mal^{1/5} Jun^{1/5} Jun^{1/5} Jun^{1/5} Sep^{1/5} Oct^{1/5} Nov^{1/5} Dec^{1/5} Jan^{1/6}$ | [lan 20 2015 to 5ab 10 2015] |
| 782,337 16,938,194 | 0.58Out of Memory (OOM)11.69HWERR | $Fe^{b^{1/5}} N^{ar^{1/5}} P^{r^{1/5}} N^{ar^{1/5}} P^{r^{1/5}} N^{a^{1/5}} D^{n^{1/5}} D^{n^{1/5}} P^{n^{1/5}} O^{d^{1/5}} P^{n^{1/5}} D^{e^{c^{1/5}}} D^{e^{c^{1/5}}} D^{an^{1/6}} P^{n^{1/6}}$ Figure 2: Event counts for different event types (in different colors) durin | g [Jan. 28, 2015 to Feb. 18, 2015]. |
| 782,337 16,938,194 1,215,780 | 0.58Out of Memory (OOM)11.69HWERR0.810Seg. Fault | $\frac{Fe^{b'15}}{10^6} Ne^{t'15}} Ne^{t'15} Ne^{$ | g [Jan. 28, 2015 to Feb. 18, 2015]. |
| 782,337 16,938,194 1,215,780 43,268,141 | 0.5 8 Out of Memory (OOM) 11.6 9 HWERR 0.8 10 Seg. Fault 29.5 11 Lustre | $\frac{Feb}{Nat} \sqrt{5} + \frac{Feb}{Nat} \sqrt{5} + \frac{Feb}{Nat}$ | g [Jan. 28, 2015 to Feb. 18, 2015]. |
| 782,337 16,938,194 1,215,780 43,268,141 31,498,746 | 0.5 8 Out of Memory (OOM) 11.6 9 HWERR 0.8 10 Seg. Fault 29.5 11 Lustre 21.5 12 LNet | Feb ¹⁵ Ne ¹⁵ Ap ¹⁵ Ap ¹⁵ De ¹⁵ D | g [Jan. 28, 2015 to Feb. 18, 2015]. |
| 782,337 16,938,194 1,215,780 43,268,141 31,498,746 992,997 | 0.5 8 Out of Memory (OOM) 11.6 9 HWERR 0.8 10 Seg. Fault 29.5 11 Lustre 21.5 12 LNet 0.7 13 LNet Error | Feb ^{1/5} Na ^{1/5} Na ^{1/5} Na ^{1/5} Na ^{1/5} Na ^{1/5} Ju ^{1/5} Se ^p ^{1/5} Cd ^{1/5} No ^{1/5} De ^{c^{1/5}} Ja ^{1/6} 10 ⁶ Interpret dependence of the second d | g [Jan. 28, 2015 to Feb. 18, 2015]. |
| 782,337 16,938,194 1,215,780 43,268,141 31,498,746 992,997 42,809,426 | 0.5 8 Out of Memory (OOM) 11.6 9 HWERR 0.8 10 Seg. Fault 29.5 11 Lustre 21.5 12 LNet 0.7 13 LNet Error 29.2 22 Lustre Error | Feb ¹⁵ Na ¹⁵ PP ¹⁵ Na ¹⁵ Ju ¹⁵ Ju ¹⁵ Ju ¹⁵ Se ^{p15} Od ¹⁵ P ⁰¹ De ^{c15} Ja ¹⁶ De ^{c15} Ja ¹⁶ Figure 2: Event counts for different event types (in different colors) durin | g [Jan. 28, 2015 to Feb. 18, 2015]. |
| 782,337 16,938,194 1,215,780 43,268,141 31,498,746 992,997 42,809,426 146,485,450 | 0.5 8 Out of Memory (OOM) 11.6 9 HWERR 0.8 10 Seg. Fault 29.5 11 Lustre 21.5 12 LNet 0.7 13 LNet Error 29.2 22 Lustre Error | Feb ⁻¹⁵ N ³ N ³ D ¹⁵ N ³ D ¹⁵ | g [Jan. 28, 2015 to Feb. 18, 2015]. |

Table 1: Summary of event types and their occurrences (146.5 million

in total) in Titan's logs during [Jan. 2015, Mar. 2018].

| Event Type | Source |
|------------|--|
| 3 | c10-4c2s6n0 |
| 22 | c7-2c1s6n0 |
| 12 | c20-0c1s4n3 |
| 11 | c2-1c0s1n2 |
| 9 | c3-5c0s4n2 |
| 11 | c18-7c0s2n1 |
| | Event Type 3 22 12 11 9 11 |

Table 2: An example of the event table. The event type takes the ID value from Table 1.



Figure 1: Total event counts in calendar year 2015, 2016 and 2017 from Titan. Resolution in time is by hour.



Figure 3: The nodal layouts for every event type and the combined total for one-hour window prior to "2015-02-09 01:49:54". Each layout has dimensions of [300, 64] in pixel and each pixel represents a unique Titan node with its coordinates [X, Y] (definition in following section).

It is impossible to efficiently inspect system status by relying on the counts either of total system events or any specific event types!

| | Solution | | | | | | | | |
|---------------------------------|---------------------|-------|---|--|------------------------------------|------------------|------------------------------------|-----|---|
| General Form of Data Table Prin | | Princ | Principal Components [2] in Feature Space | | Shannon Entropy [3] | | | | |
| | Feature 1 Feature 2 | ••• | Feature N | | $SVD \Longrightarrow \sigma_i$ | | | | System Information Entropy (SIE) [1] |
| Record 1 | | | | σ_i : <i>i</i> -th varia | ance out of k eigenvalues of the S | VD decomposition | b | | |
| Record 2 | | | | | | | $\sum_{n} n$ | | $\mathbf{U}(t)$ |
| | | | | Variand | ce Distribution of Principal C | components | $H = - \lambda \xi_i log_h(\xi_i)$ | (2) | $W(t) = b^{H(t)}$ (3) |
| | | | | | | | | | |
| | | | | | σ_i | | | | |
| | | | | | $\xi_i = \frac{1}{\nabla k}$ | (1) | Entropy: in a general "b-ary" form | | b: the logarithmic base used in calculating H. In |
| Record M | | | | | $\sum_{1}^{n} \sigma_{i}$ | | | | our analysis, b = 10. |
| | | | | | | | | | |
| Event Table Layout | | | | Application of System Information Entrop | y (SIE) | | | | |
| Source vs. Type | | | F | | | | | | |

| Source | Machine Check Exception | Out of Memory | Seg. Fault | Lustre | (H0 |
|-------------|----------------------------|---------------|------------|--------|--------|
| c0-2c2s1n2 | 1 | 0 | 0 | 0 | ц Ц |
| c10-4c0s0n0 | 0 | 0 | 0 | 4 | |
| c10-4c0s3n0 | 0 | 1 | 0 | 3 | > |
| c10-4c1s6n1 | 0 | 0 | 3 | 1 | > |

Table 3: An example of the time-windowed event table for one hour starting from "2016 Jan. 27, 05:00:00" and ending at "2016 Jan. 27, 06:00:00". For an illustrative purpose, only 4 rows are shown out of a total of 400.



Nodal Map

 $\begin{array}{ll} X = 12 \times Cabinet_{Row} + 4 \times Chassis + Node \\ Y = 8 \times Cabinet_{Column} + Slot \end{array}$

Figure 4: The nodal map for total events in an hour prior to "2015-02-09 01:49:54". The layout has dimensions of [300, 64] in pixel and each pixel represents a unique Titan node with its coordinates [X, Y] translated from its source name, like those shown in Table 3. For detailed definition of the transformation, please see [1].





Figure 5: System Information Entropies are plotted for two event data layouts - "Source Type" (TOP) and "Nodal Map" (BOTTOM).



50 ← -



Figure 6: Close-ups for time series of "Nodal Map" SIE (TOP) and "Total Count" (BOTTOM), from 12:00 Feb. 13 to 06:00 Feb. 14, 2015.

Figure 7: Titan nodal layouts showing event aggregation patterns at representative moments from 18:30 to 22:00 (interval at 30 mins) on Feb. 13, 2015.

Q.,

2

2

1.5

Conclusion and References

With no assumption about the relative significances among system properties, we created a general-purpose metric which reflects system status by taking a concise form of a time series of system information entropy (SIE). We demonstrated that this system metric can comprehensively and sensitively summarize overall system characteristics without compromising on computational efficiency.

- B. H. Park, Y. Hui, S. Boehm, R. A. Ashraf, C. Layton, and C. Engelmann. "A Big Data Analytics Framework for HPC Log Data: Three Case Studies Using the Titan Supercomputer Log", HPCMASPA, IEEE International Conference on Cluster Computing (CLUSTER), Belfast, Northern Ireland, 2018
- 2. K. Pearson F.R.S. LIII. "On lines and planes of closest fit to systems of points in space". The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 2(11):559–572, 1901
- C. E. Shannon. "A mathematical theory of communication". The Bell System Technical Journal, 27(3):379– 423, 1948



