

Technical Disclosure Commons

Defensive Publications Series

November 2021

Alert Generation for Cyclically Varying Metrics

Joseph Turner

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Turner, Joseph, "Alert Generation for Cyclically Varying Metrics", Technical Disclosure Commons, (November 15, 2021)

https://www.tdcommons.org/dpubs_series/4708



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

Alert Generation for Cyclically Varying Metrics

ABSTRACT

This disclosure describes techniques for monitoring computer system metrics that exhibit cyclical (periodic) behavior. Per techniques of this disclosure, function-based thresholds are utilized for metric monitoring and generation of system alerts. The function-based thresholds are time-varying and are aligned with the troughs and peaks of the underlying metric being monitored. Metric measurements are obtained for a particular period of time. A function such as a Discrete Cosine Transform (DCT) is fitted to the obtained metric data. A threshold is set for the metric based on the fitted function at a fixed offset from the function derived value. Function-based thresholding can enable monitoring metrics that exhibit cyclical properties with greater sensitivity when compared to fixed thresholding. Deviations of metric values can be accurately detected even at troughs in the metric data while also mitigating false positives at peaks in the metric data.

KEYWORDS

- Site Reliability Engineering (SRE)
- Crash rate
- System Monitoring
- Periodic function
- Performance metric
- Queries per second (QPS)
- Fast Fourier Transform (FFT)
- Discrete Cosine Transform (DCT)

BACKGROUND

Metrics associated with computer systems, e.g., distributed or cloud computing systems, are monitored for indications of abnormal behavior. For example, metrics such as application crash rates, memory usage, etc., are routinely monitored for any early indications of downtime, software errors, infrastructure issues, etc. In many computer systems, such metrics are automatically monitored, and alerts are generated and transmitted based on the observed metrics meeting predetermined thresholds.

However, some observed client (or server) metrics have a diurnal or other time-dependent periodicity. For example, at certain times of a day, metric values may be substantially higher than at other times of the day.




Fig. 1: Time-varying metrics

Fig. 1 is an example graph that illustrates time-varying characteristics of a metric that exhibits periodically varying (e.g., cyclical) behavior. As depicted in Fig. 1, the graph of the

metric indicates a pattern of troughs and peaks. Fig. 1 additionally depicts an example fixed threshold (the dashed line) determined based on an average value (e.g., mean) of the metric over a time-period. A monitoring and/or alerting system may generate an alert based on a comparison of a current value of the metric with the fixed threshold value. However, when a metric exhibits cyclical behavior, the fixed threshold value can be too high for the times when the metric values are in their trough (lower part of the cycle) zones. Consequently, alerting systems may not detect metric abnormalities until the time that the metric values are no longer in their trough zones.

General monitoring and metrics are based on thresholds or service-level objectives (SLOs), which do not take into account the typical value for the metric based on the exact time; rather, these are typically based on constant thresholds).

DESCRIPTION

This disclosure describes techniques for monitoring computer system metrics that exhibit cyclical (periodic) behavior. Per techniques of this disclosure, function-based thresholds are utilized for metric monitoring and generation of system alerts. The function-based thresholds are time-varying and are aligned with the patterns (e.g., troughs and peaks) of the underlying metric being monitored. Historical metric data is utilized to determine a function that fits (approximates) the metric data, and the determined function is subsequently utilized to determine a current threshold for comparison with current metric data. Utilization of function-based thresholds enables comparisons that provide greater sensitivity to the time varying behavior of the metric, and for accurate detection of issues at peaks as well as at troughs in the metric data.




Fig. 2: Metric exhibiting cyclical properties

Fig. 2 is a graph of an example metric that is characterized by cyclical behavior. As depicted in Fig. 2, which is a plot of metric values as a function of time, the metric values follow a diurnal cycle, with repeated peaks and troughs. The different frequency components in the metric values can be determined via frequency analysis by applying a suitable transform, e.g., a Discrete Cosine Transform (DCT), a Fast Fourier Transform (FFT), etc.




Fig. 3: A transform function is utilized to determine dominant frequencies

Fig. 3 depicts results obtained by applying a transform function, e.g., a Discrete Cosine Transform (DCT) to the metric data. A DCT is utilized to represent an input signal, e.g., metric data, as a sum of cosine functions. The x-axis represents frequency, and the y-axis represents an amplitude of the metric data at the corresponding frequency. Dominant frequency components are indicated by their corresponding peaks, and are commonly observed at frequency corresponding to a day (24 hours), a week, etc.




Fig. 4: A fitted function can approximate the time varying properties of a metric

Fig. 4 depicts a comparison of metric measurements with an example fitted function obtained by applying a transform to the metric measurements. As can be seen in Fig 4, the fitted function approximates the metric measurements, and tracks the peaks and troughs observed in the metric measurements.

Adjustments may be made to the fitted function to additionally track non-cyclical changes in the metric measurements, e.g., if the metric measurements are observed to drift upwards or downwards in magnitude.




Fig. 5: A function-based threshold can be set at a fixed offset from the function

Fig. 5 depicts a function-based threshold for a metric, per techniques of this disclosure. As depicted in Fig. 5, a threshold may be set at a fixed offset from a value determined by the transform function. This enables a greater sensitivity in server and client process monitoring and alerting. Issues that may only have been detected and surfaced once the metric got closer to its peak value can be detected even when the metric is close to its trough value(s).




Fig. 6: Example workflow for using function-based thresholds for metric monitoring

Fig. 6 depicts an example workflow for utilization of function-based thresholds for metric monitoring, per techniques of this disclosure. Metric measurements are obtained (610) for a particular period of time, e.g., 1 week, 2 weeks, etc. A function, e.g., a Discrete Cosine Transform (DCT) is fitted (620) to the obtained metric data. The function can be utilized to determine a predicted value for the metric based on a current timestamp.

A threshold is set (630) for the metric based on the fitted function at a fixed offset from the function derived value. For example, a threshold may be set to a level that is greater than the function derived value, e.g., 25%, 50%, etc. During monitoring of the metric, a current metric value is compared to a threshold value based on the fitted function for a corresponding timestamp. Additional metric data is obtained (640) to determine an average value for the metric. The additional metric data is utilized to adjust (650) the threshold and/or fitted function to capture any updated trends in the metric data.

Function-based thresholding can enable monitoring metrics that exhibit cyclical properties with greater sensitivity when compared to fixed thresholding. Deviations (regressions) of metric data values can be accurately detected even at troughs in the metric data while also mitigating false positives at peaks in the metric data.

CONCLUSION

This disclosure describes techniques for monitoring computer system metrics that exhibit cyclical (periodic) behavior. Per techniques of this disclosure, function-based thresholds are utilized for metric monitoring and generation of system alerts. The function-based thresholds are time-varying and are aligned with the troughs and peaks of the underlying metric being monitored. Metric measurements are obtained for a particular period of time. A function such as a Discrete Cosine Transform (DCT) is fitted to the obtained metric data. A threshold is set for the metric based on the fitted function at a fixed offset from the function derived value. Function-based thresholding can enable monitoring metrics that exhibit cyclical properties with greater sensitivity when compared to fixed thresholding. Deviations of metric values can be accurately detected even at troughs in the metric data while also mitigating false positives at peaks in the metric data.