ϵ-Diagnosis: Unsupervised and Real-time Diagnosis of Small-window Long-tail Latency in Large-scale Microservice Platforms

Huasong Shan, Yuan Chen, Haifeng Liu, Yunpeng Zhang, Xiao Xiao,
Xiaofeng He, Min Li and Wei Ding

JD.com

Problem

Small-window long-tail latency(SWLT)

Long tail latency at extremely small timescales (e.g., 1 minute, even 1 second) for web services deployed in container based microservice platforms

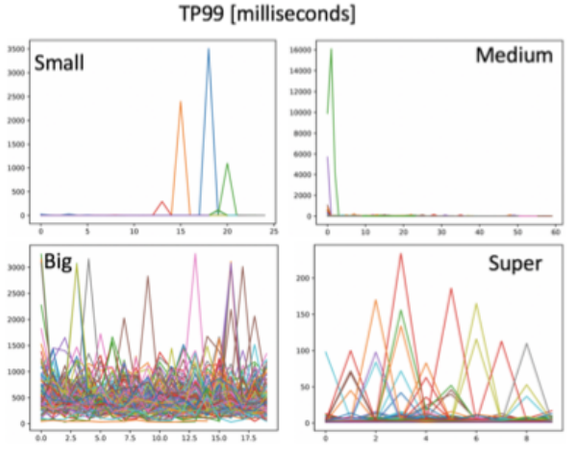

Characterization of SWLT

- High-variance
- Frequent-shift

Goal

- The algorithm and the system can quickly diagnose root-causes at runtime with low computation cost
- The algorithm can significantly reduce the problem space (metrics) while guaranteeing not to miss any actual root-cause metrics

System Architecture

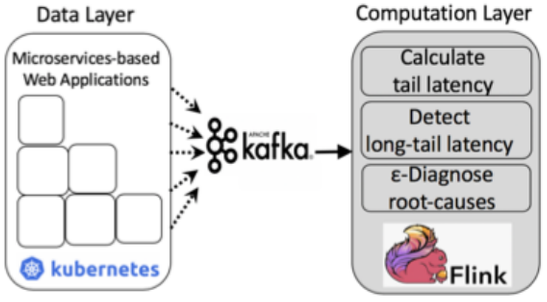

Solution

ϵ-Diagnosis

- Selecting two samples from the snapshot
- Two-sample null hypothesis test
- ϵ-Statistics (Energy distance correlation)

Algorithm 1 Pseudo-code for the ϵ-Diagnosis algorithm

Input: small-window long-tail latency, M time-series metrics of N containers, confidence threshold α, alarm window
Output: candidate root-cause metrics, problematic containers

```
procedure ϵ-DIAGNOSIS
    for ContainerN ← 1 to N do
        S_A = getAnomalySample
        S_N = getNormalSample
        for MetricM ← 1 to M do
            ( ρ (S_A, S_N), P ) = Calculate Energy distance correlation coefficient of S_A and S_N using Equation (3) with P-value
            if P < α then
                /* Reject Hypothesis S_A != S_N*/
                Add MetricM as a candidate root-cause metric
                Add ContainerN as a candidate problematic contaiter
            else
                /* Accept Hypothesis S_A = S_N */
            end if
        end for
    end for
end procedure
```

Energy distance correlation

$$\rho^2(S_A, S_N) = \begin{cases} \frac{cov^2(S_A, S_N)}{\sqrt{\sigma^2(S_A)\sigma^2(S_N)}}, & \sigma^2(S_A)\,\sigma^2(S_N) > 0\,. \\ 0, & \sigma^2(S_A)\,\sigma^2(S_N) = 0\,. \end{cases}$$

Operational Results in Real-Production

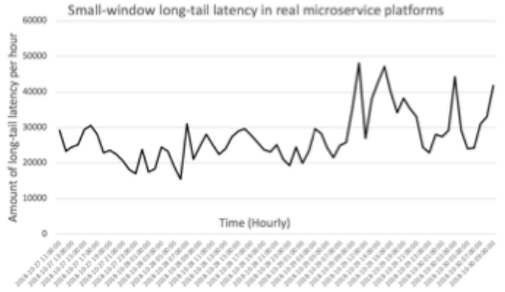

Service rate of ϵ-Diagnosis in a real-life production

Evaluations

Baseline

- Pearson's distance
- K-NN: K-Nearest Neighbor
- MST: Minimal Spanning Tree

ϵ-Diagnosis: energy distance correlation

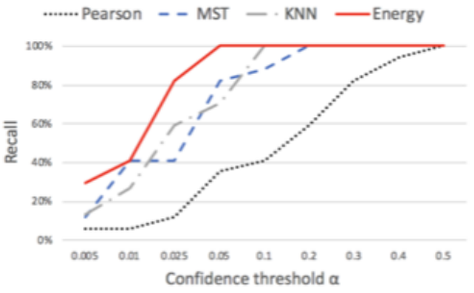

Energy can reach 100% recall quickly as α increases

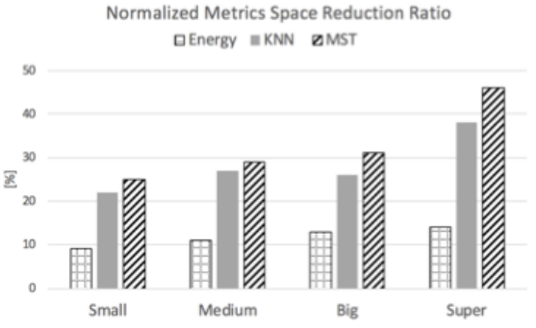

ϵ-Diagnosis(Energy) reduces metrics to approximately 10%

Conclusion and Future Works

- ϵ-Diagnosis finished to diagnose metrics anomaly of each container in large-scale microservice platforms
- Further root-cause diagnosis will extend other use cases for microservice–based web applications at JD.com

JD.COM

THE WEB CONFERENCE 2019