

CloudDet: Interactive Visual Analysis of Anomalous Performances in Cloud Computing Systems

Ke Xu, Yun Wang, Leni Yang, Yifang Wang, Bo Qiao, Si Qin, Yong Xu, Haidong Zhang, Huamin Qu
IEEE Transactions on Visualization and Computer Graphics, 2019



Amirhossein Abbasi
Nov 2019

CloudDet



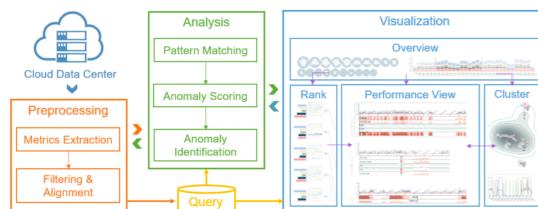
1

2

3

4

System Overview



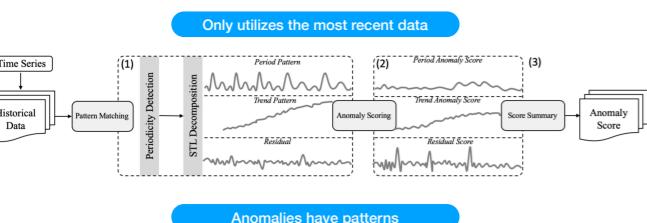
5

6

7

8

Algorithm Flow



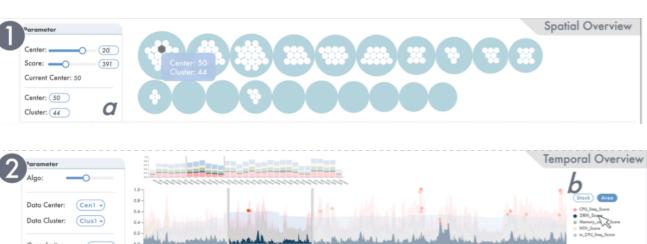
9

10

11

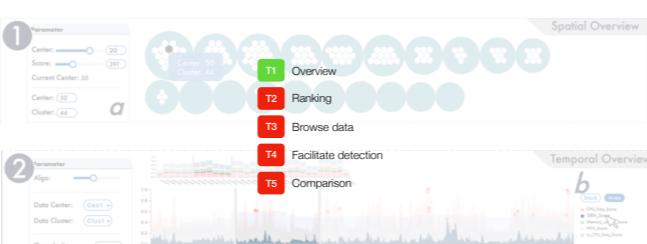
12

Spatial and Temporal Views



13

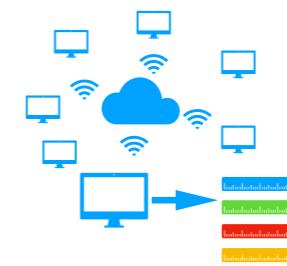
Spatial and Temporal Views



14

Motivation

- Monitoring nodes instead of monitoring applications
- Too many **false positives**, scale problem.
- Visualization of anomalies: **Intuitiveness, interaction.**
- Research Contribution:** Detection system, Visualization, Evaluation



3

Mathematics!

$$P(f_{k/N}) = ||X(f_{k/N})||^2, \quad k = 0, 1, \dots, \frac{N-1}{2}, \quad (1)$$

$$AS_{periodic} = \min\left(\left|\frac{T_n - T_{n-1}}{T_{n-1}}\right|, 1\right), \quad (4)$$

$$ACF(\tau) = \frac{1}{N} \sum_{n=0}^{N-1} d(\tau) \cdot d(n+\tau), \quad \tau = 0, 1, \dots, \frac{N-1}{2}. \quad (2)$$

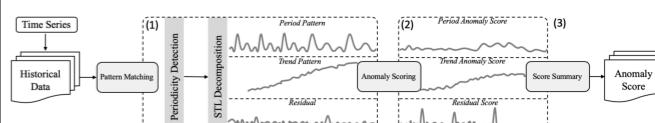
$$AS_{trend} = \min\left(\left|\frac{K_n - K_{n-1}}{K_{n-1}}\right|, 1\right), \quad (5)$$

$$d_n = S_n + Tr_n + R_n, \quad n = 1, 2, \dots, N, \quad (3)$$

$$AS_{spike} = \min\left(\left|\frac{R_n - \mu_{n-1} - 3\sigma_{n-1}}{3\sigma_{n-1}}\right|, 1\right), \quad (6)$$

$$AS = f(AS_{periodic}, AS_{trend}, AS_{spike}), \quad (7)$$

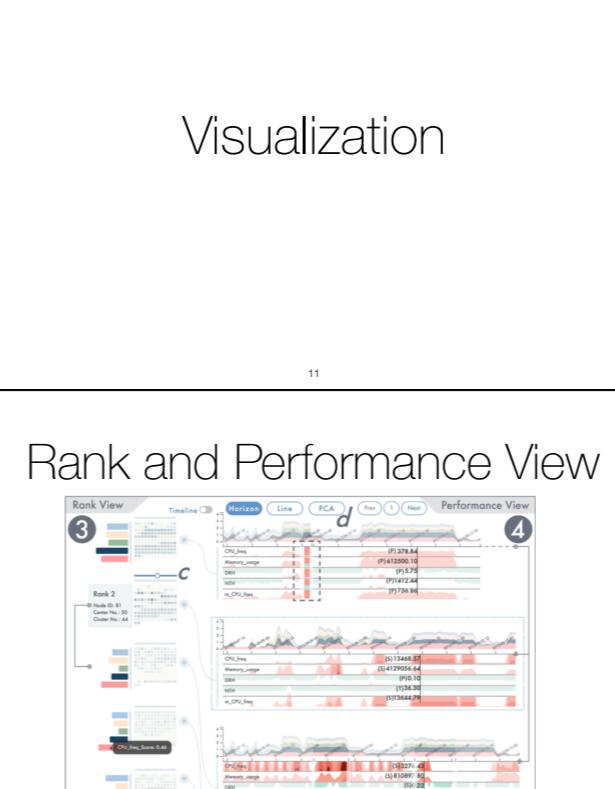
Algorithm Flow



Design Tasks

- T1 Overview of anomalies for data query
- T2 Ranking suspicious nodes dynamically
- T3 Browse data flexibly
- T4 Facilitate anomaly detection
- T5 Similarities of nodes

Visualization



15

Encoding Protocol

Global Categorical Colors: performance metrics (CPU Frequency, Memory Usage,...)



Linear Color Scheme: Anomaly score



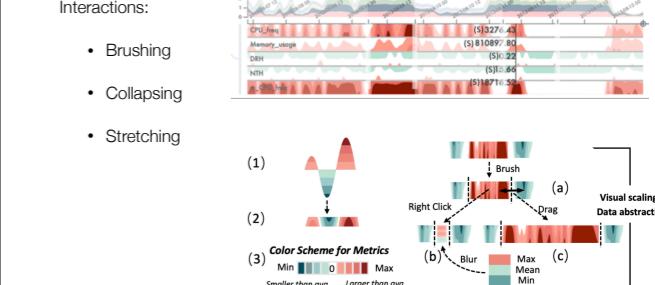
Diverging Color Scheme: Difference of performance metrics to average



Horizon Chart

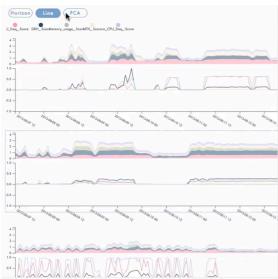
Interactions:

- Brushing
- Collapsing
- Stretching



16

Line mode



- Each line for one metric
- More conventional
- Normalize data to [-1, 1]

17

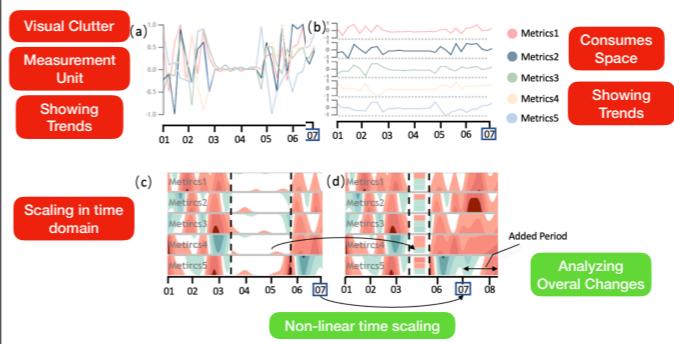
PCA mode



- Project a multivariate data to a one-dimensional time-series data
- Major Trend

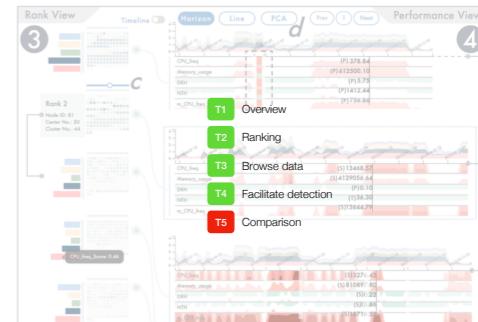
18

Alternative Designs



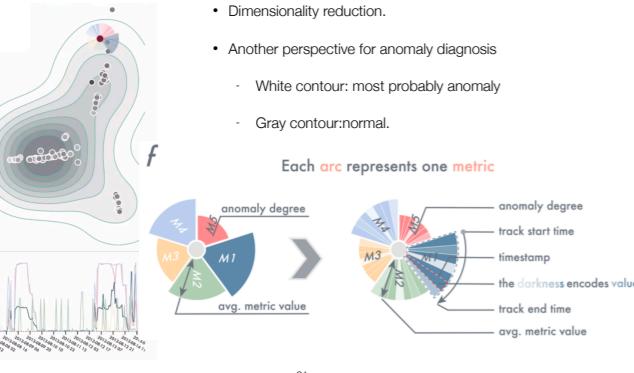
19

Rank and Performance View



20

The Cluster View



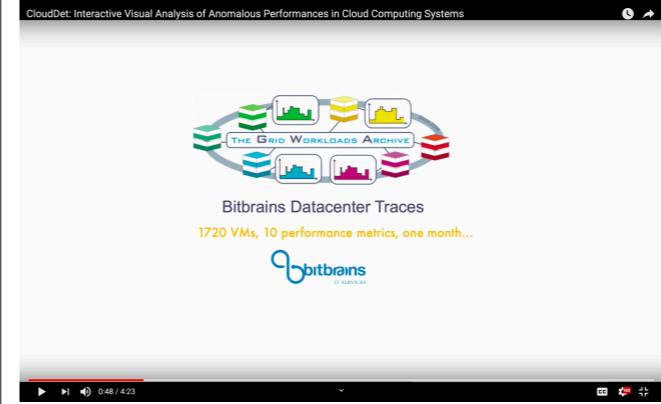
21

The Cluster View



22

Official Video (1:05 min)



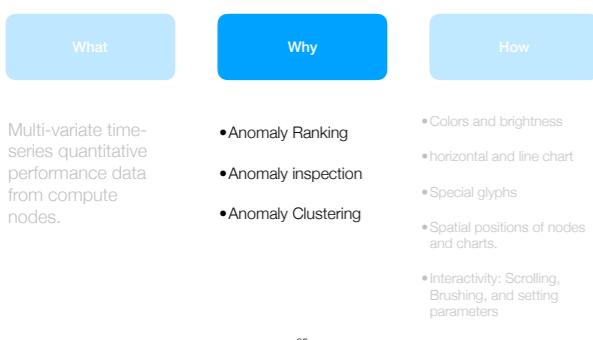
23

What-Why-How Summary



24

What-Why-How Summary



25

What-Why-How Summary



26

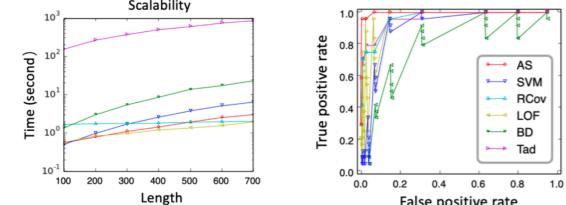
Scalability



27

Evaluation

Quantitative Evaluation



29

Case Studies

Case Study 1: Bitbrains Datacenter Traces

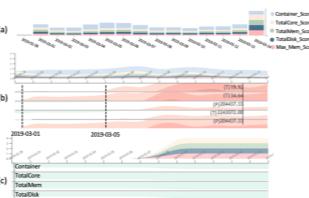
500 VMs, One month



Case Study 2: Live Cloud System Data

1,000,000 nodes, Two weeks

[100 data centers with 20 data clusters with 500 nodes each]



30

User Feedback

- Automated Anomaly Detection: **Trust** in algorithm,
- System: **Useful** and **User-friendly**, **Consistent**, too **comprehensive** and **Overwhelming**, **Need Tutorial**
- Visualization and Interaction: **Helpful**, **new perspective** for overall trend, **clear comparison**, Confess that they use chaotic line charts before.

31

Critique

Positive

- 😊 Alternative designs
- 😊 Super-scalable
- 😊 Perfect evaluation
- 😊 Very Accurate
- 😊 Special Glyphs

Negative

- 😢 Better to use non-diverging colors for horizon charts.
- 😢 Minor occlusion in the clustering.
- 😢 Make use of global colors in horizontal chart.
- 😢 Bad way for Assigning the ranks to performance.
- 😢 Empty clusters in spatial overview.
- 😢 Limitation: Just consider recent data and one metric.
- 😢 Limitation: Don't discuss why using those performance metrics for anomaly.

Question?