CoFlux: Robustly Correlating KPIs by Fluctuations for Service Troubleshooting

Ya Su[1], Youjian Zhao[1], Wentao Xia[1], Rong Liu[2], Jiahao Bu[1], Jing Zhu[1],

Yuanpu Cao[3], Haibin Li[1,4], Chenhao Niu[1], Yiyin Zhang[5], Zhaogang Wang[5], Dan Pei[1]

1 清華大學 Tsinghua University

2 STEVENS INSTITUTE of TECHNOLOGY THE INNOVATION UNIVERSITY

3 南開大學 Nankai University

4 国防科学技術大学 National University of Defense Technology

5 阿里巴巴 Alibaba.com

Outline

Background

Algorithm

Evaluation

Case Studies

Outline

Background

Algorithm

Evaluation

Case Studies

Internet-based Services

- Internet-based services are everywhere.

- Service interruptions are inevitable.

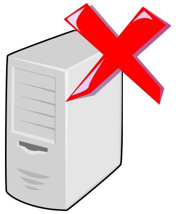

- Service Troubleshooting is necessary but challenging because of the interweaved anomalies.

KPIs and Fluctuations

- **KPIs (Key Performance Indicators):** A set of performance metrics that monitor the service.
- **Fluctuations (or Flux-features):** Anomalous changes in KPIs which could be indicated by prediction errors.

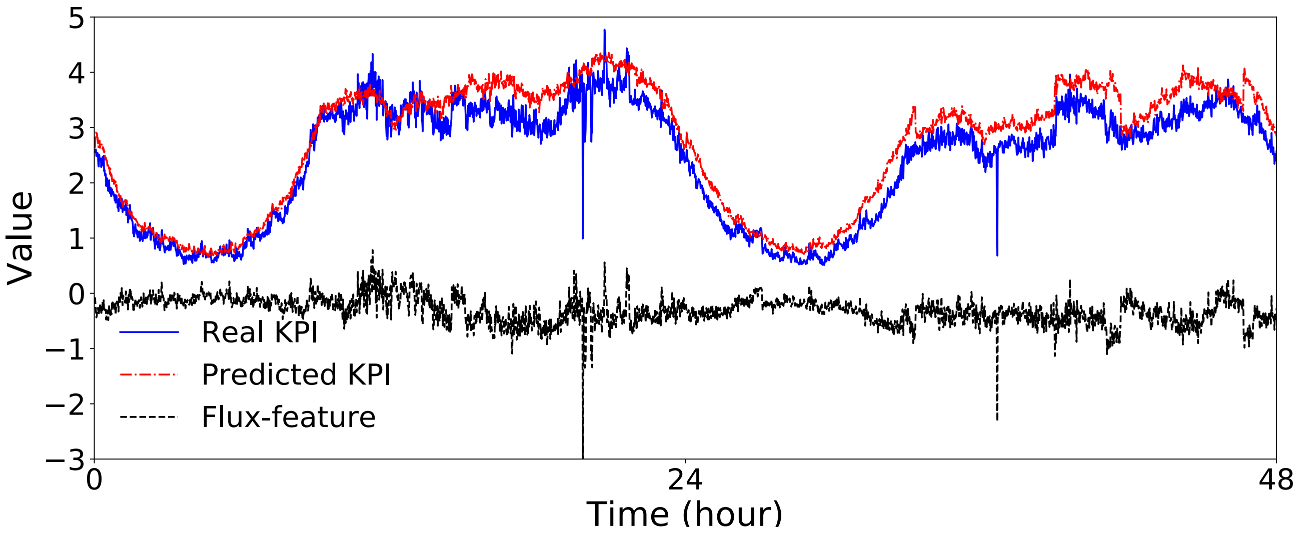

Flux-correlation

- For two KPIs X and Y, we want to answer three questions:
  - Q1: Existence of flux-correlation ($X \sim Y$ or $X \nsim Y$). If yes, then:

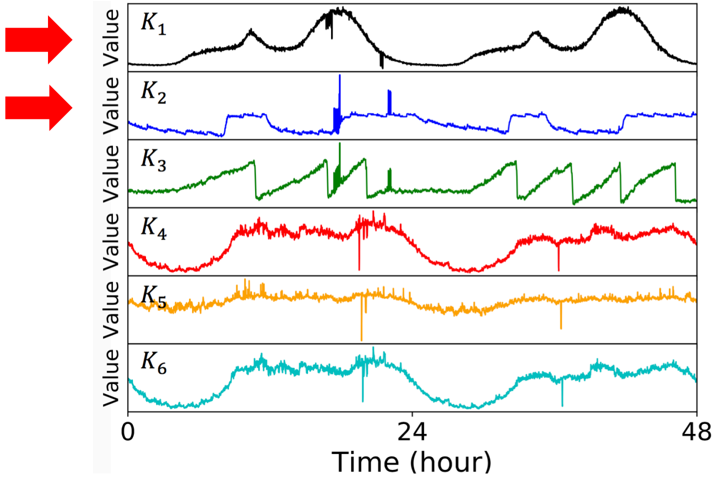

(a) Six KPIs

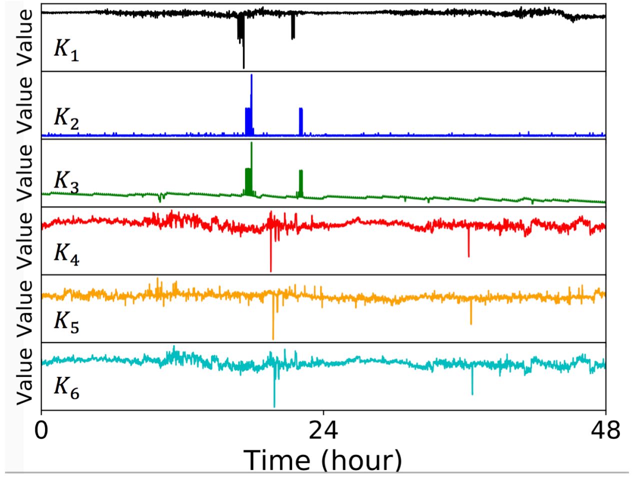

(b) Fluctuations (Flux-feature) of the KPIs in (a)

Flux-correlation

- For two KPIs X and Y, we want to answer three questions:
  - Q1: Existence of flux-correlation ($X \sim Y$ or $X \nsim Y$). If yes, then:
  - Q2: Temporal order, $X \longrightarrow Y$ or $X \longleftrightarrow Y$.

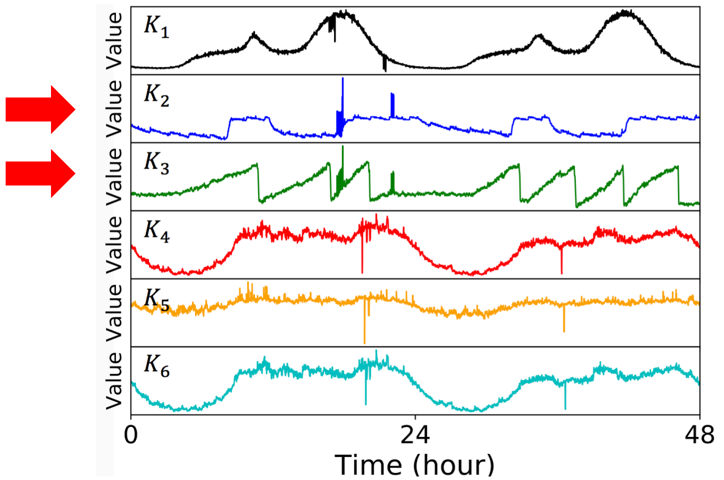

(a) Six KPIs

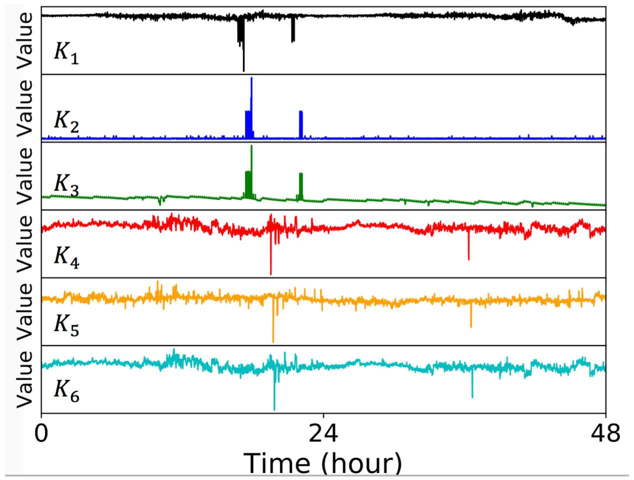

(b) Fluctuations (Flux-feature) of the KPIs in (a)

Flux-correlation

- For two KPIs X and Y, we want to answer three questions:
  - Q1: Existence of flux-correlation ($X \sim Y$ or $X \nsim Y$). If yes, then:
  - Q2: Temporal order, $X \longrightarrow Y$ or $X \longleftrightarrow Y$.
  - Q3: Direction, Positive or Negative.

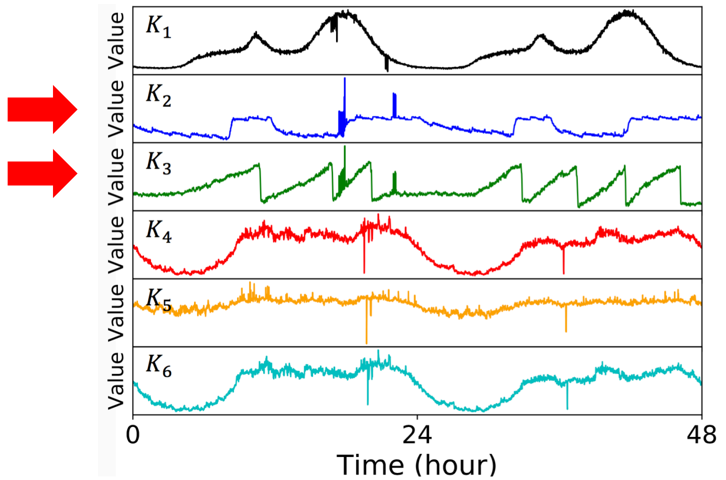

(a) Six KPIs

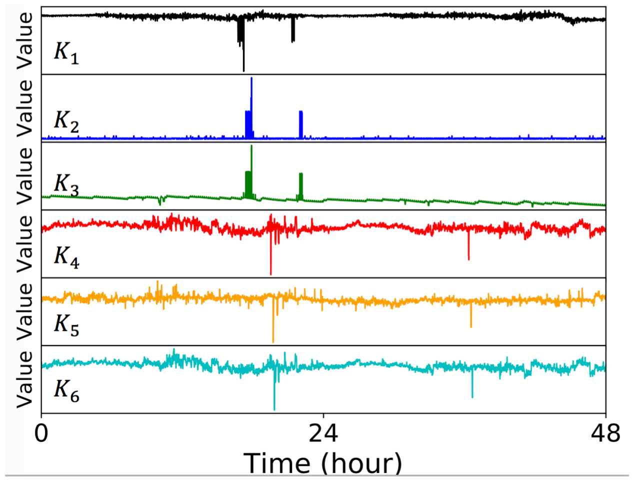

(b) Fluctuations (Flux-feature) of the KPIs in (a)

Goal

Flux-correlations of KPIs

Service Troubleshooting

- Alert Compression
- Recommending TOP N causes
- Constructing Fluctuation Propagation Chains

Related Work

	Traditional Correlation methods					Statistical models from other fields		
	Pearson Correlation	Spearman Correlation	Granger causality [ICDM 2012]	Cross Correlation	J-measure [SIGKDD 2014]	SIG [DSN 2010]	VARMA	Co-Integration
Fluctuation analysis	✘	✘	✘	✘	✔	✔	✘	✘
Temporal order	✘	✘	✔	✔	✘	✔	✘	✘
Direction	✔	✔	✘	✔	✘	✘	✘	✘

Challenges

- **Challenge 1:** To the best of our knowledge, there is no generic mechanism for fluctuation extraction.

- **Challenge 2:** Flux-correlation should not be based on anomaly detection of because of its difficulty.[IMC 2015]

- **Challenge 3:** Two flux-correlated KPIs may present different patterns.

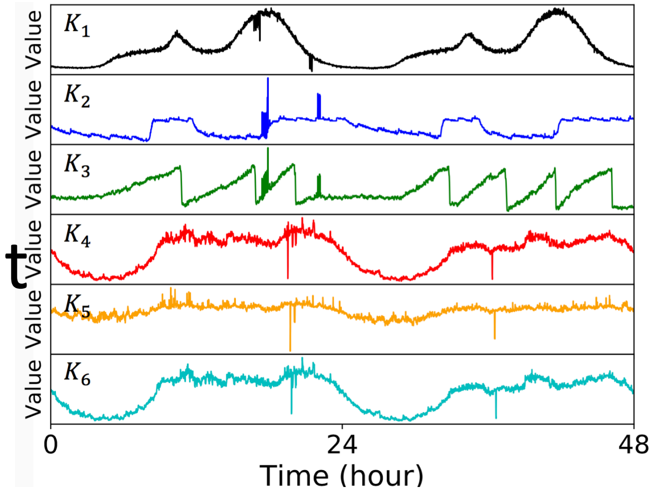

Outline

Background

Algorithm

Evaluation

Case Studies

Model Architecture

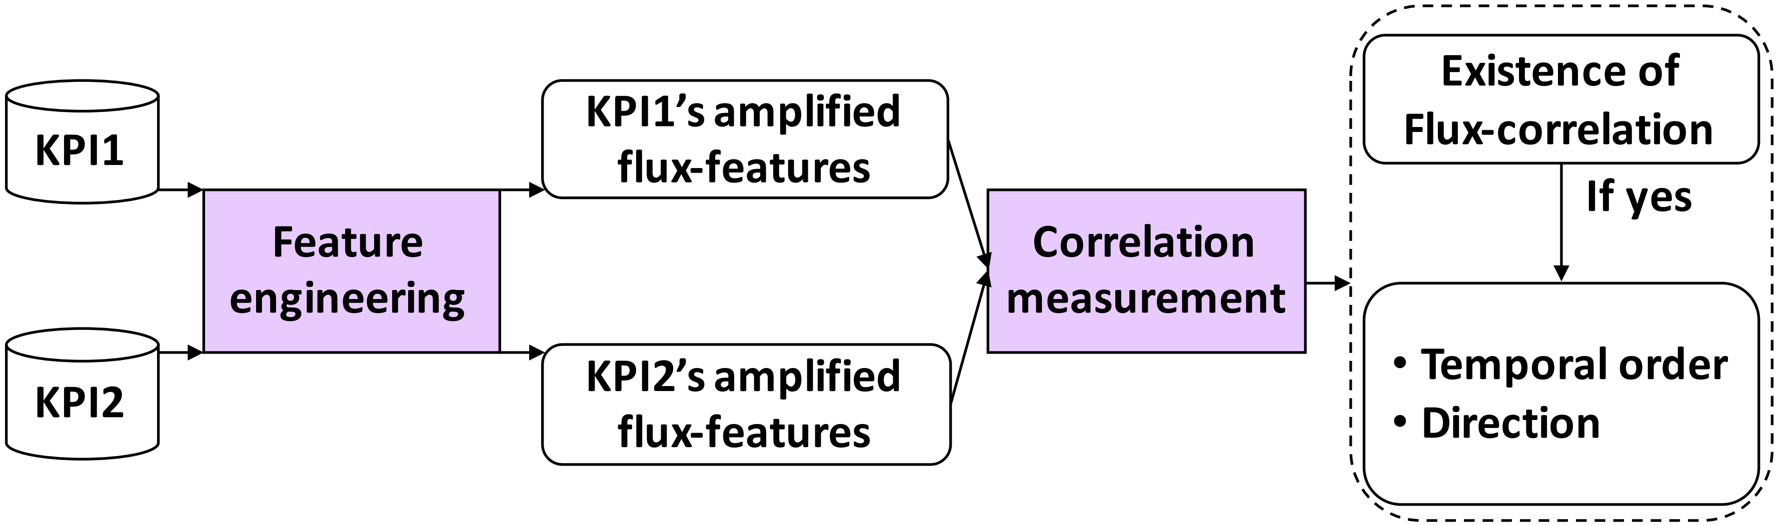

CoFlux

Feature engineering

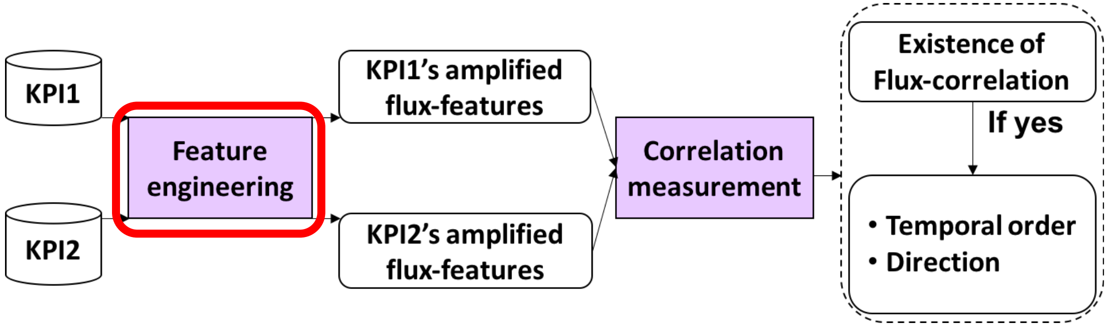

- Feature extraction: Apply time series prediction models with parameters as flux-feature detectors.

Prediction models and detectors

<table>
<tr><th>Prediction models/ # of detectors</th><th>Parameter Configurations</th></tr>
<tr><td>Diff / 2</td><td>Last-day, last-week</td></tr>
<tr><td>Holt-Winters / 64</td><td>$\alpha, \beta, \gamma = \{0.2, 0.4, 0.6, 0.8\}$</td></tr>
<tr><td>Historical average / 4</td><td rowspan="4">Win = 1, 2, 3, 4 weeks</td></tr>
<tr><td>Historical median / 4</td></tr>
<tr><td>TSD / 4</td></tr>
<tr><td>TSD median / 4</td></tr>
<tr><td>Wavelet / 4</td><td>Win = 1, 3, 5, 7 days</td></tr>
<tr><td colspan="2">In total： 7 prediction models / 86 detectors</td></tr>
</table>

Challenge 1 & 2

Feature engineering

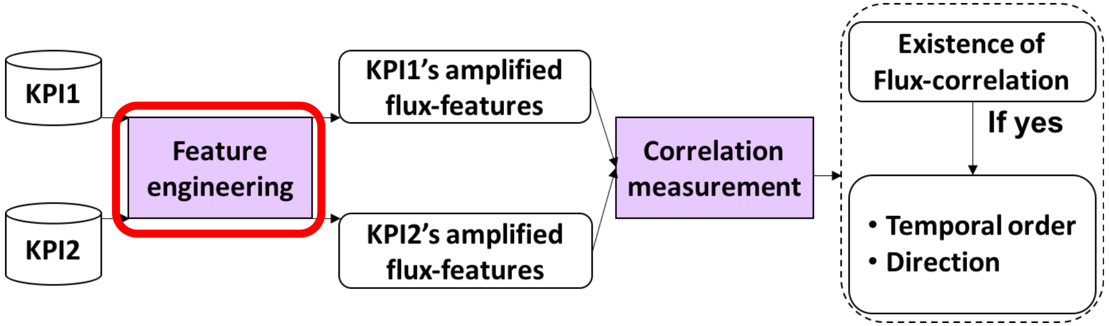

- Feature amplification:
  - Apply z-score to normalize the flux-feature.
  - To reduce the influence of noises, we use modified exponential activation to strengthen large fluctuations.

Challenge 1 & 2

$$f(\alpha,\beta,x)=\begin{cases} e^{min(x,\beta)\times\alpha}-1\ ,for\ x\geqslant 0\\ -e^{min(|x|,\beta)\times\alpha}+1\ ,for\ x<0\end{cases}$$

Correlation measurement

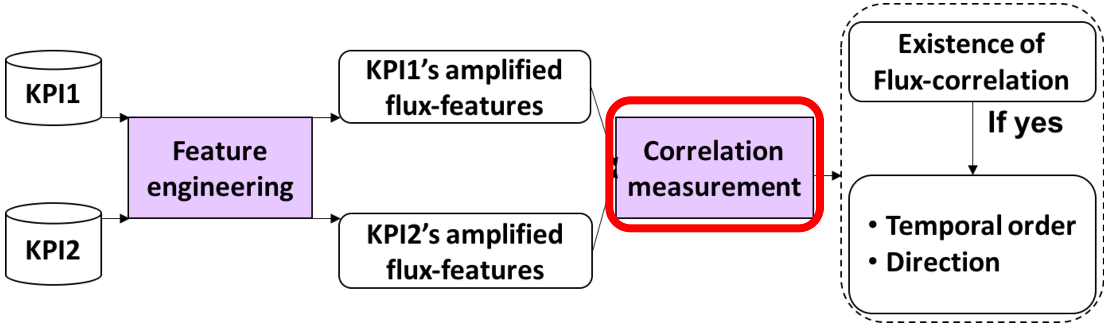

Algorithm 1: Correlation measurement

```
Input: afxSet: Set of amplified flux-features of KPI X
       afySet: Set of amplified flux-features of KPI Y
       coTHR: Threshold of existence of flux-correlation
resultSet ← [ ]
// Set of candidate flux-correlation results
for afx in afxSet do
    for afy in afySet do
        resultSet ← FCC(afx, afy)    // Eq. 4
if abs(max(resultSet[:,0])) > abs(min(resultSet[:,0])) then
    [ccV, shiftV] = max(resultSet) /* ccV: correlation value
        about the existence of flux-correlation;
        shiftV: shifted value of X when get ccV */
else [ccV, shiftV] = min(resultSet);
if abs(ccV) ⩾ coTHR then
    if shiftV = 0 then
        if ccV ⩾ 0 then X ←+→ Y;
        else X ←−→ Y;
    if shiftV < 0 then
        if ccV ⩾ 0 then X −+→ Y;
        else X −−→ Y;
    if shiftV > 0 then
        if ccV ⩾ 0 then Y −+→ X;
        else Y −−→ X;
else X ↚ Y;
```

- We apply the Cross-correlation to measure the correlation results of flux-features.

Challenge 3

$$R(G_s, H) = \sum_{i=-l+1}^{l-1} G_s[i] \times H[i]$$

$$CC(G_s, H) = \frac{R(G_s, H)}{\sqrt{R(G, G) \times R(H, H)}}$$

$$minCC = \min_s \left(CC(G_s, H)\right),\ s1 = \arg\min_s \left(CC(G_s, H)\right)$$

$$maxCC = \max_s \left(CC(G_s, H)\right),\ s2 = \arg\max_s \left(CC(G_s, H)\right)$$

$$FCC(G, H) = \begin{cases} [minCC,\ s1],\ for\ |maxCC| < |minCC| \\ [maxCC,\ s2],\ for\ |maxCC| \geq |minCC| \end{cases}$$

Outline

Background

Algorithm

Evaluation

Case Studies

CoFlux VS Baseline Models

- Datasets:
  - **Dataset I:** flux-correlated KPIs with different time series characteristics.
  - **Dataset II:** flux-correlated KPIs with homogeneous time series characteristics.

Best F1-scores of eight algorithms

Data set	Algorithms	Best F1-Score		
		Existence	Temporal order	Direction
I	CoFlux	**0.8412**	**0.9608**	**0.9579**
	J-measure	**0.7213**	N/A	N/A
	SIG	0.5381	**1.0**	N/A
	Pearson①	0.3106	N/A	0.6127
	Pearson②	0.5909	N/A	0.6945
	Granger①	0.2864	0.9009	N/A
	Granger②	0.4128	0.8952	N/A
	Cross-correlation	0.3613	0.9320	**0.9814**
II	CoFlux	**0.9026**	**0.9206**	0.9987
	J-measure	**0.8462**	N/A	N/A
	SIG	0.7706	0.8012	N/A
	Pearson①	0.7193	N/A	0.9845
	Pearson②	0.7828	N/A	**1.0**
	Granger①	0.4533	0.9025	N/A
	Granger②	0.6732	**0.9141**	N/A
	Cross-correlation	0.7494	0.7781	**1.0**

Analysis about CoFlux

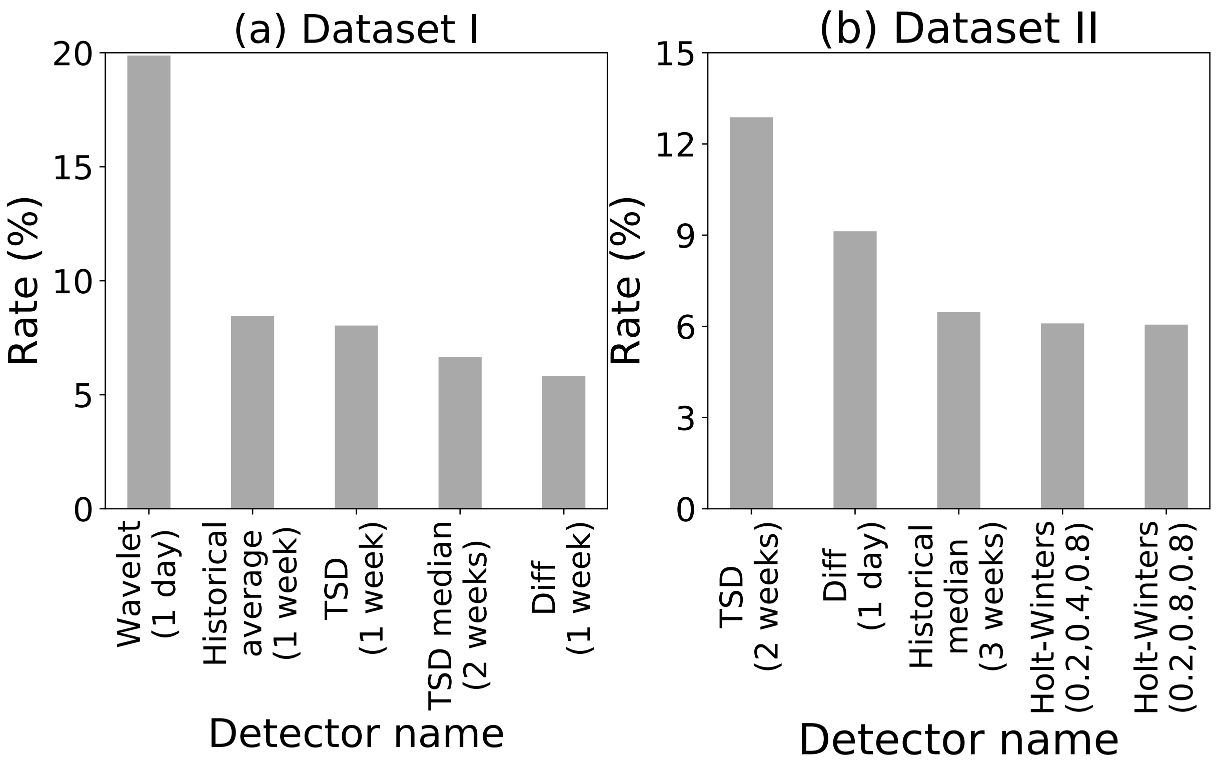

Top 5 detectors which give the flux-correlation results.

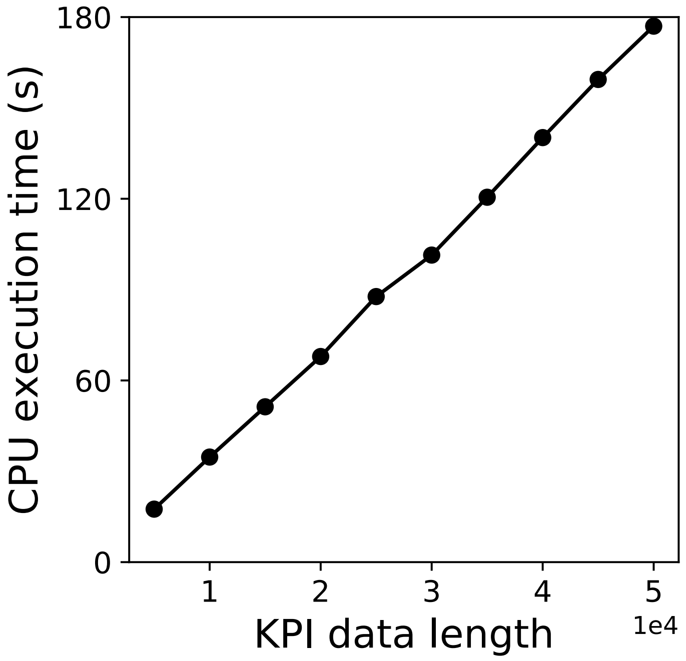

Efficiency by varying data length.

Analysis about CoFlux

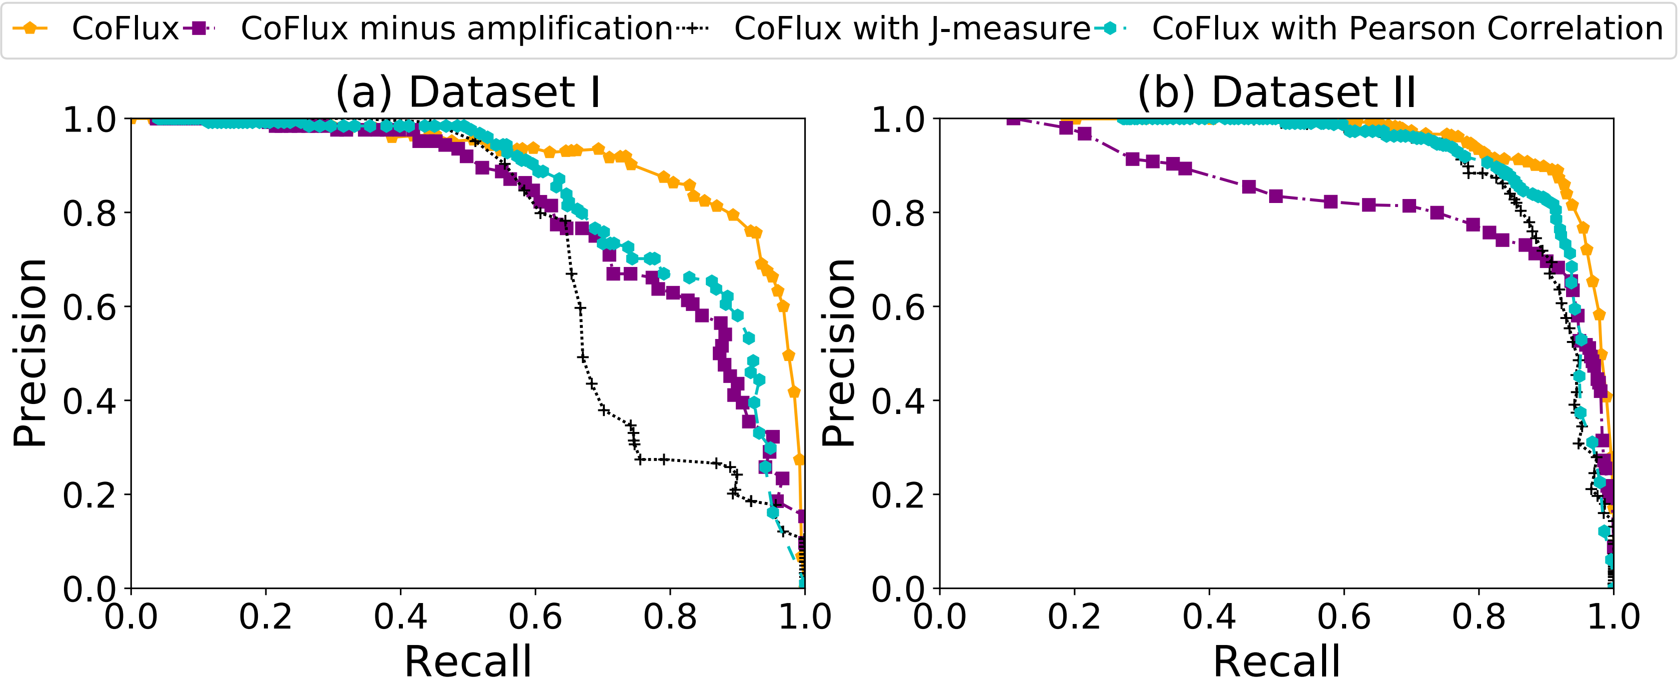

PRCs about the existence of flux-correlation among CoFlux and its variants.

Outline

Background

Algorithm

Evaluation

Case Studies

Goal

Flux-correlations of KPIs → Service Troubleshooting

- Alert Compression
- Recommending TOP N causes
- Constructing Fluctuation Propagation Chains

Clustering KPIs for alert compression

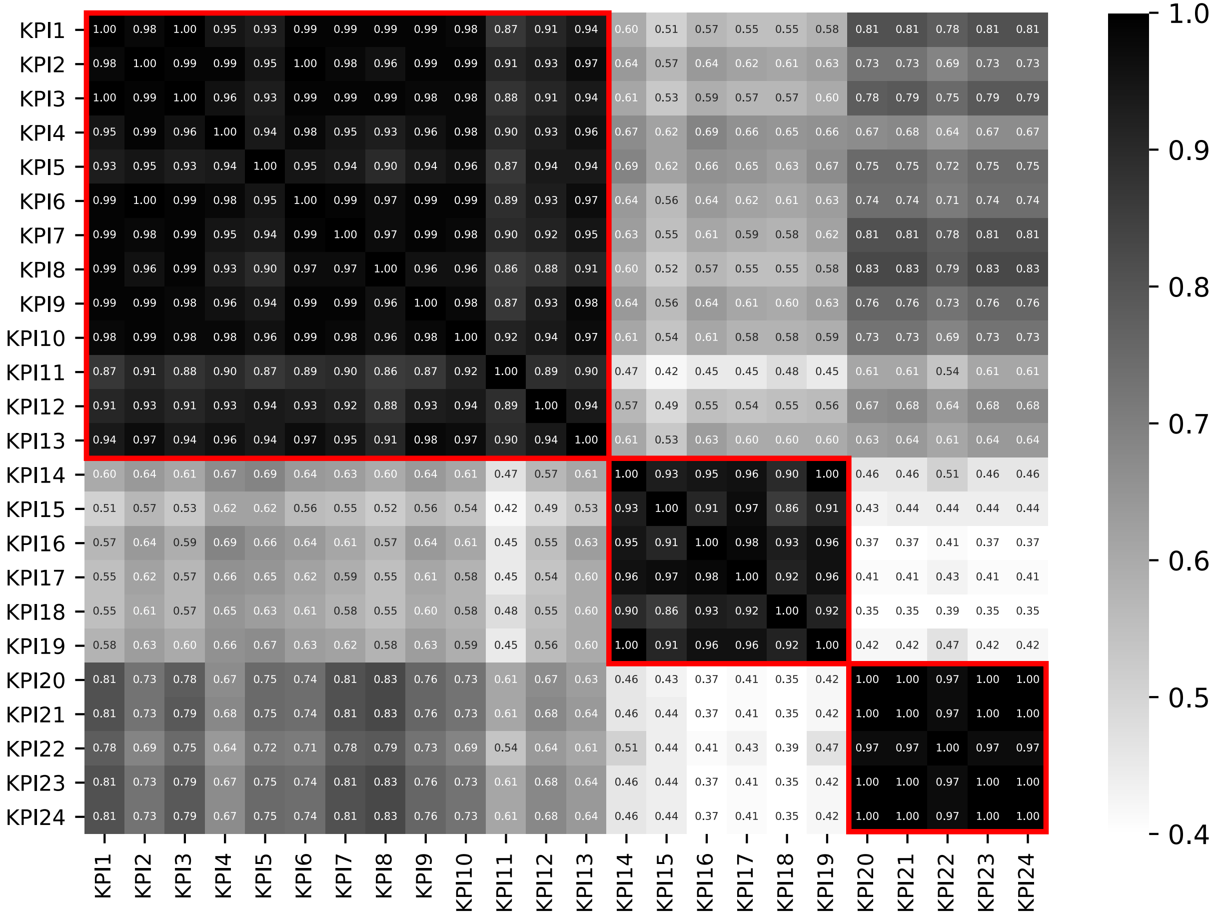

Heat map visualization for clustering results of 24 KPIs.

Recommending Top N flux-correlated KPIs

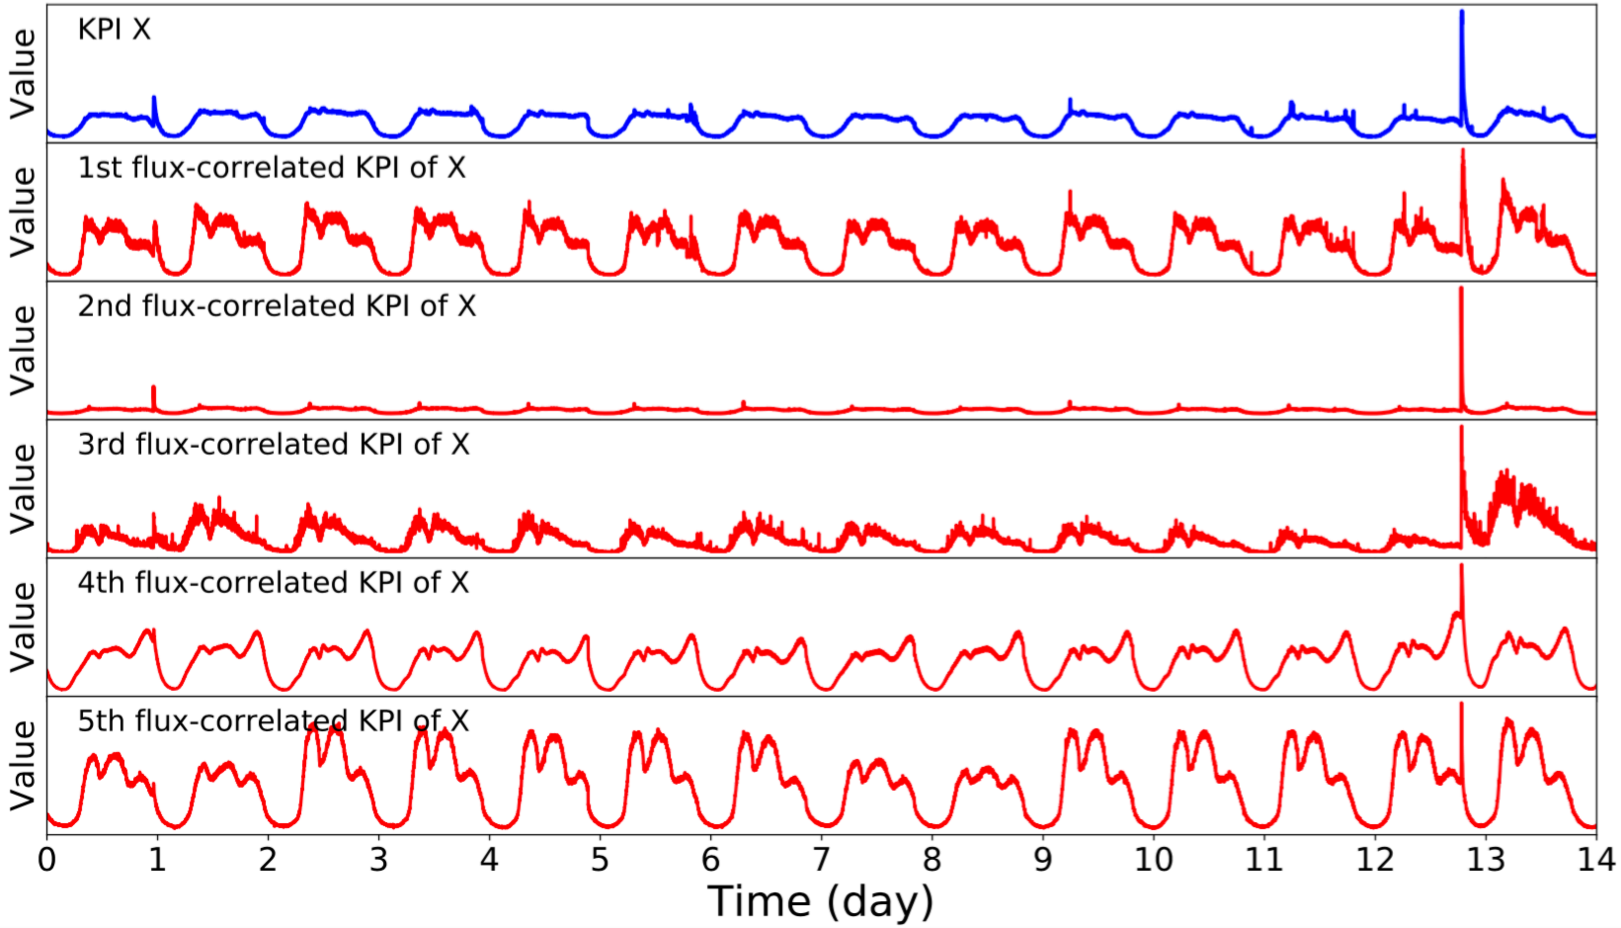

Top 5 flux-correlated KPIs for a given KPI X.

Constructing fluctuation propagation chains

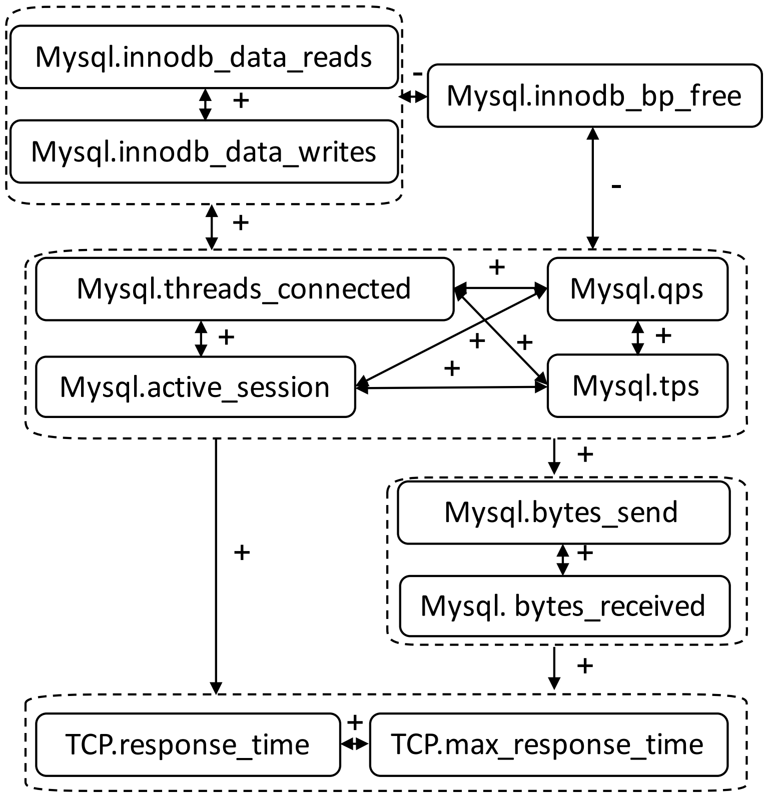

Fluctuation propagation chains of a database service constructed by the **operators**

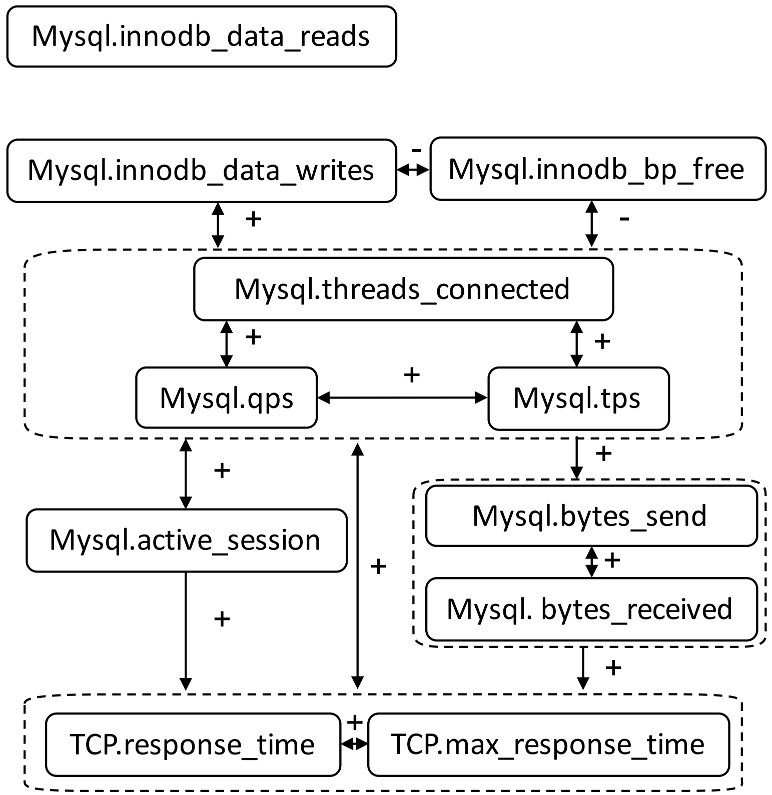

Fluctuation propagation chains of a database service constructed by **CoFlux**

Conclusion

- To the best of our knowledge, this paper is the first attempt to formulate flux-correlation and study it in detail in the domain of Internet service operations management.

- CoFlux includes a robust set of flux-features and a robust Correlation score.

- Our extensive experiments have demonstrated that CoFlux significantly outperforming the baseline algorithms and their variants.

Thank you!
Q & A

su-y16@mails.tsinghua.edu.cn

IWQOS 2019