VLDB2020Tokyo

Diagnosing Root Causes of Intermittent Slow Queries in Cloud Databases

Minghua Ma, Zheng Yin, Shenglin Zhang, Sheng Wang, Christopher Zheng, Xinhao Jiang, Hanwen Hu, Cheng Luo, Yilin Li, Nengjun Qiu, Feifei Li, Changcheng Chen, Dan Pei

清華大學 Tsinghua University

Alibaba.com

南開大學 Nankai University

Cloud Databases

- Stability -> Constant Monitoring -> Diagnosis Slow Queries

Slow Queries

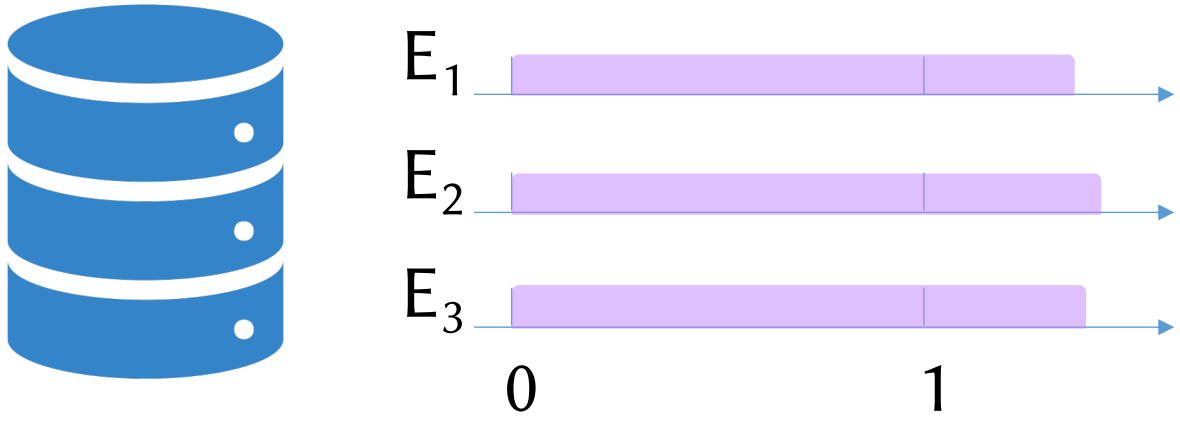

SQL A Execution Time

Slow queries result from **internal reasons**:

- complexity of SQL
- lack of indexes
- poorly-written SQL statements...

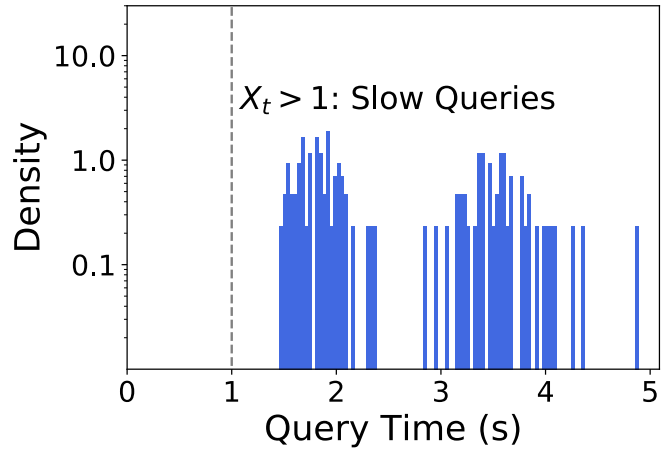

Intermittent Slow Queries (iSQs)

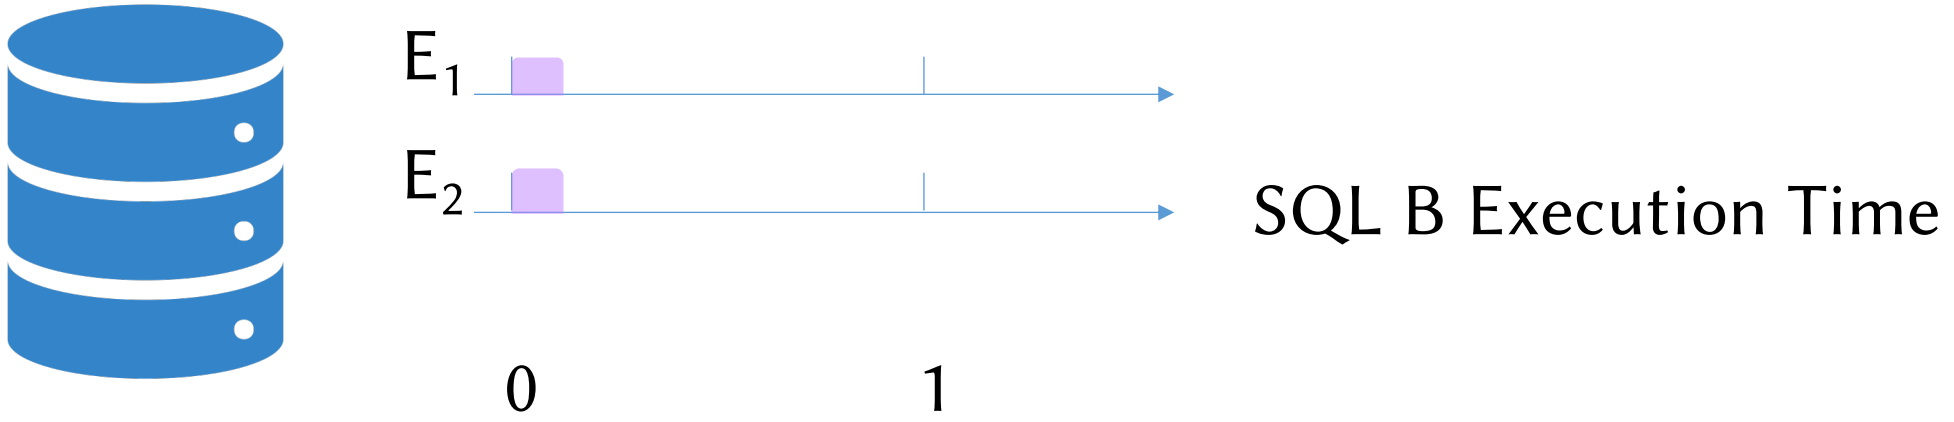

Intermittent Slow Queries (iSQs)

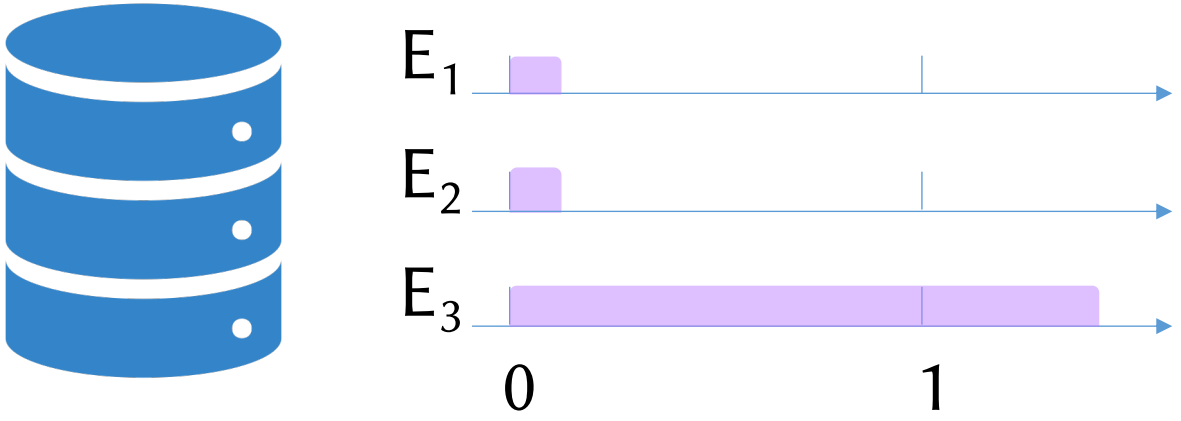

SQL B Execution Time

Slow queries result from **external reasons**:

- Intensive workload
- Host I/O bottleneck
- Accompanying slow SQL...

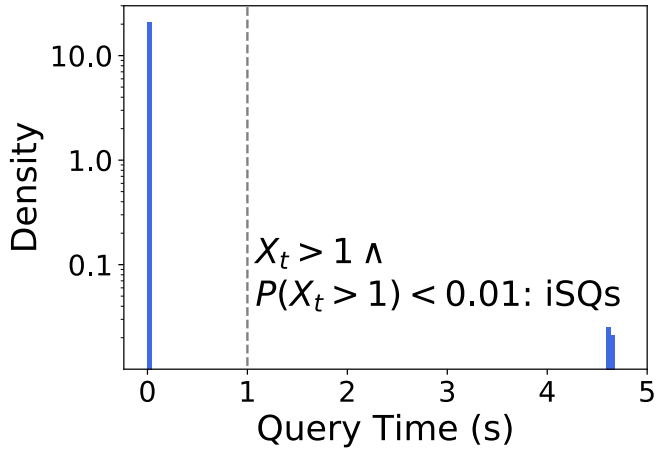

Definition of iSQs

X_t: one particular query execution time

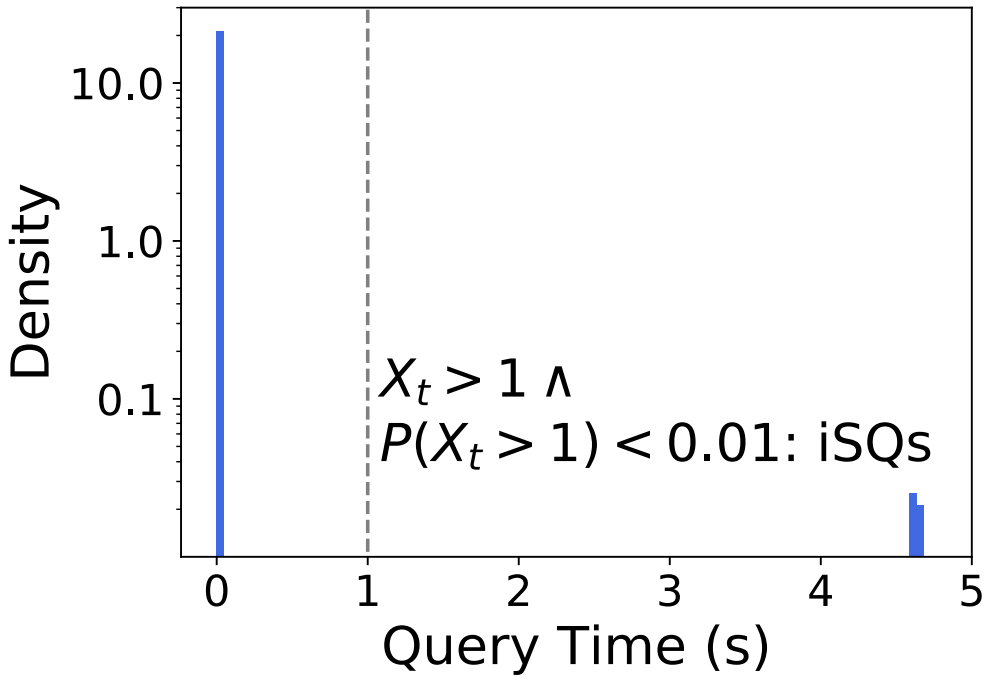

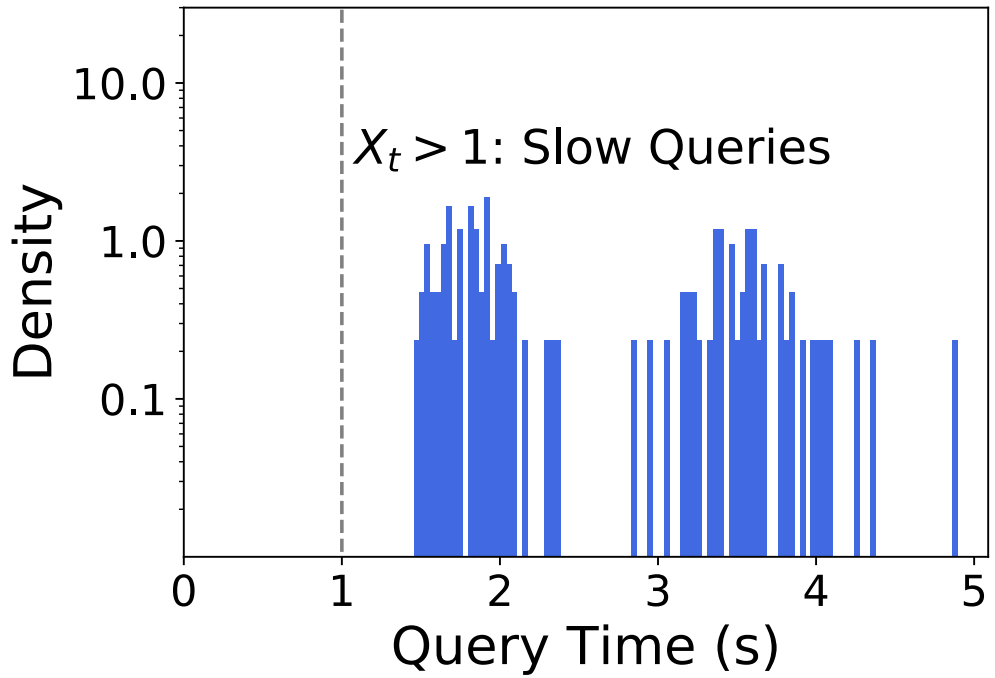

Thresholds are set empirically on Alibaba Database

Impact of iSQs

Most of iSQs are interactive queries

iSQs -> Poor user experience -> Revenue loss

Diagnosing Root Causes of iSQs in the Cloud

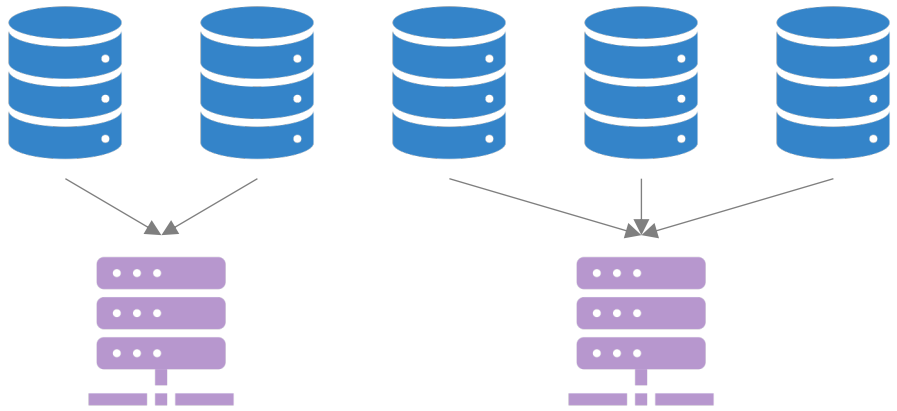

DB Instance

Physical Machine

Multiple database instances may reside on the same physical machines, which can cause resource contentions.

Diagnosing Root Causes of iSQs in the Cloud

Cloud Features

- Instance Migrations
- Database Expansions
- Storage Decoupling...

Resource Type

- CPU
- Network
- I/O...

Complexity infrastructures of cloud databases make it harder for DBAs to diagnose root causes of iSQs.

Outline

What's iSQ?

Why it's challenging?

How to diagnose it?

Evaluation

Challenge1: Anomaly Diversity

Different types of database Key Performance Indicators (KPIs)

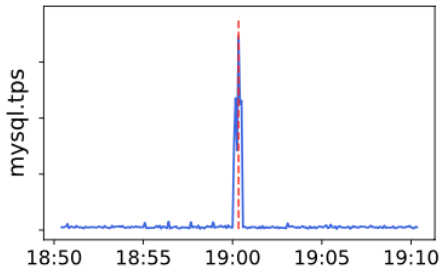

(a) Spike Up

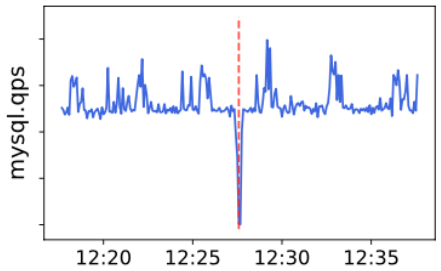

(b) Spike Down

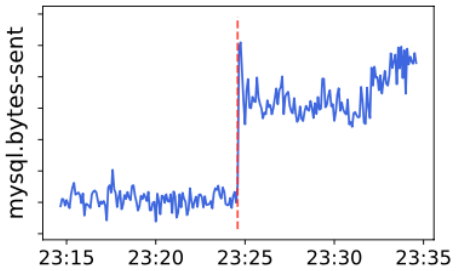

(c) Level Shift Up

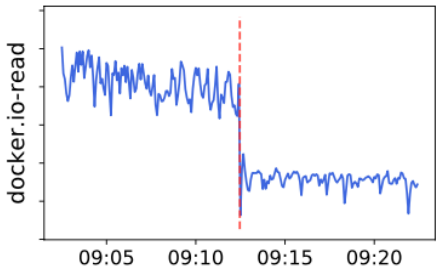

(d) Level Shift Down

Current binary anomaly detectors generally overlook and over-generalize the types of anomalies.

Challenge2: Labeling Overheads

Tens of thousands of iSQs per day in Alibaba Database

Scan hundreds of KPIs to find out the root cause of an iSQ

Manually labeling root causes is massive work;
Reproducing known root causes in a testbed is not feasible.

Challenge3: Interpretable Models

Being able to explain or narrate what causes the problem when it arises is essential in cloud databases

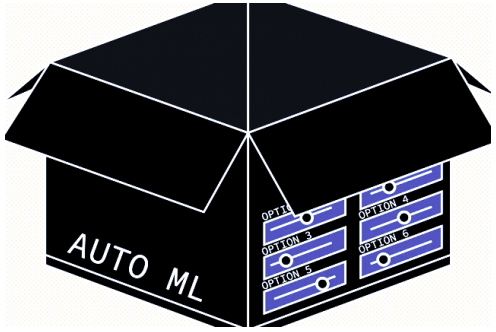

An inevitable trade-off exists between a model's accuracy and its interpretability to human.

iSQUAD

Intermittent Slow Queries
Anomaly Diagnoser

iSQUAD Overview

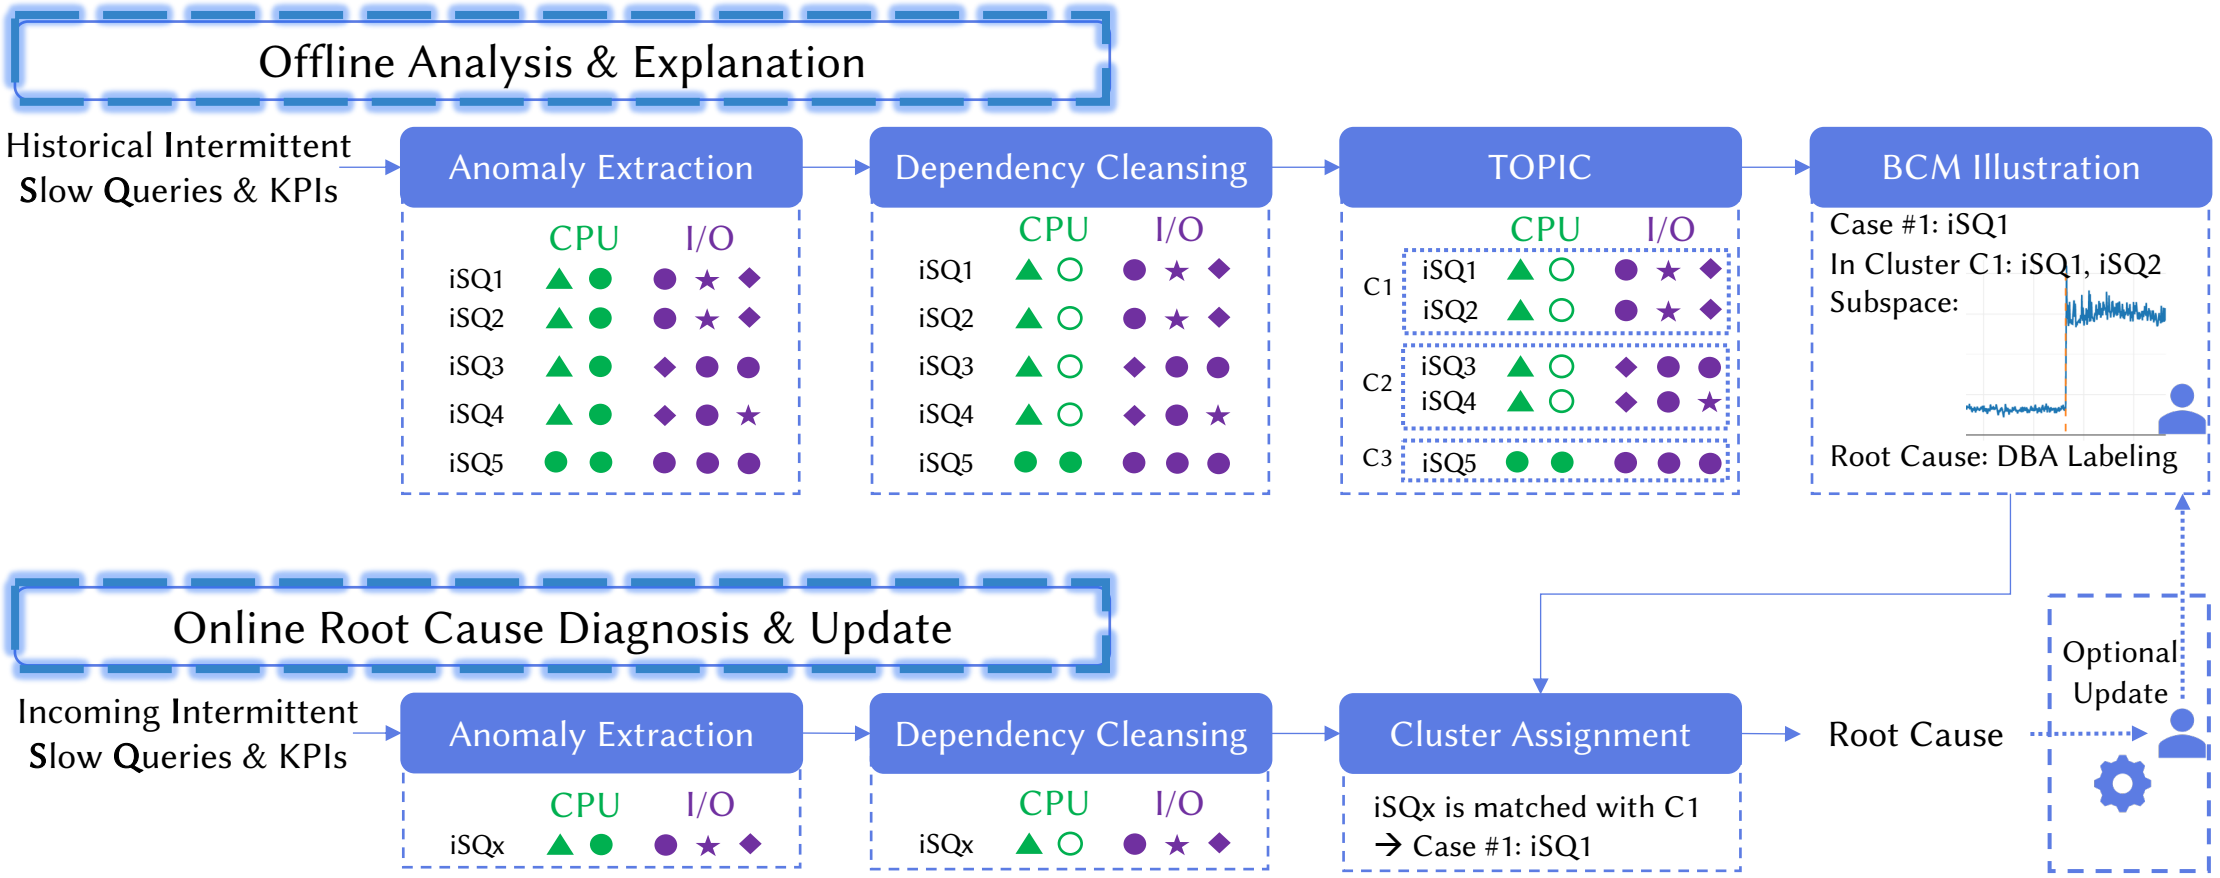

iSQUAD Overview

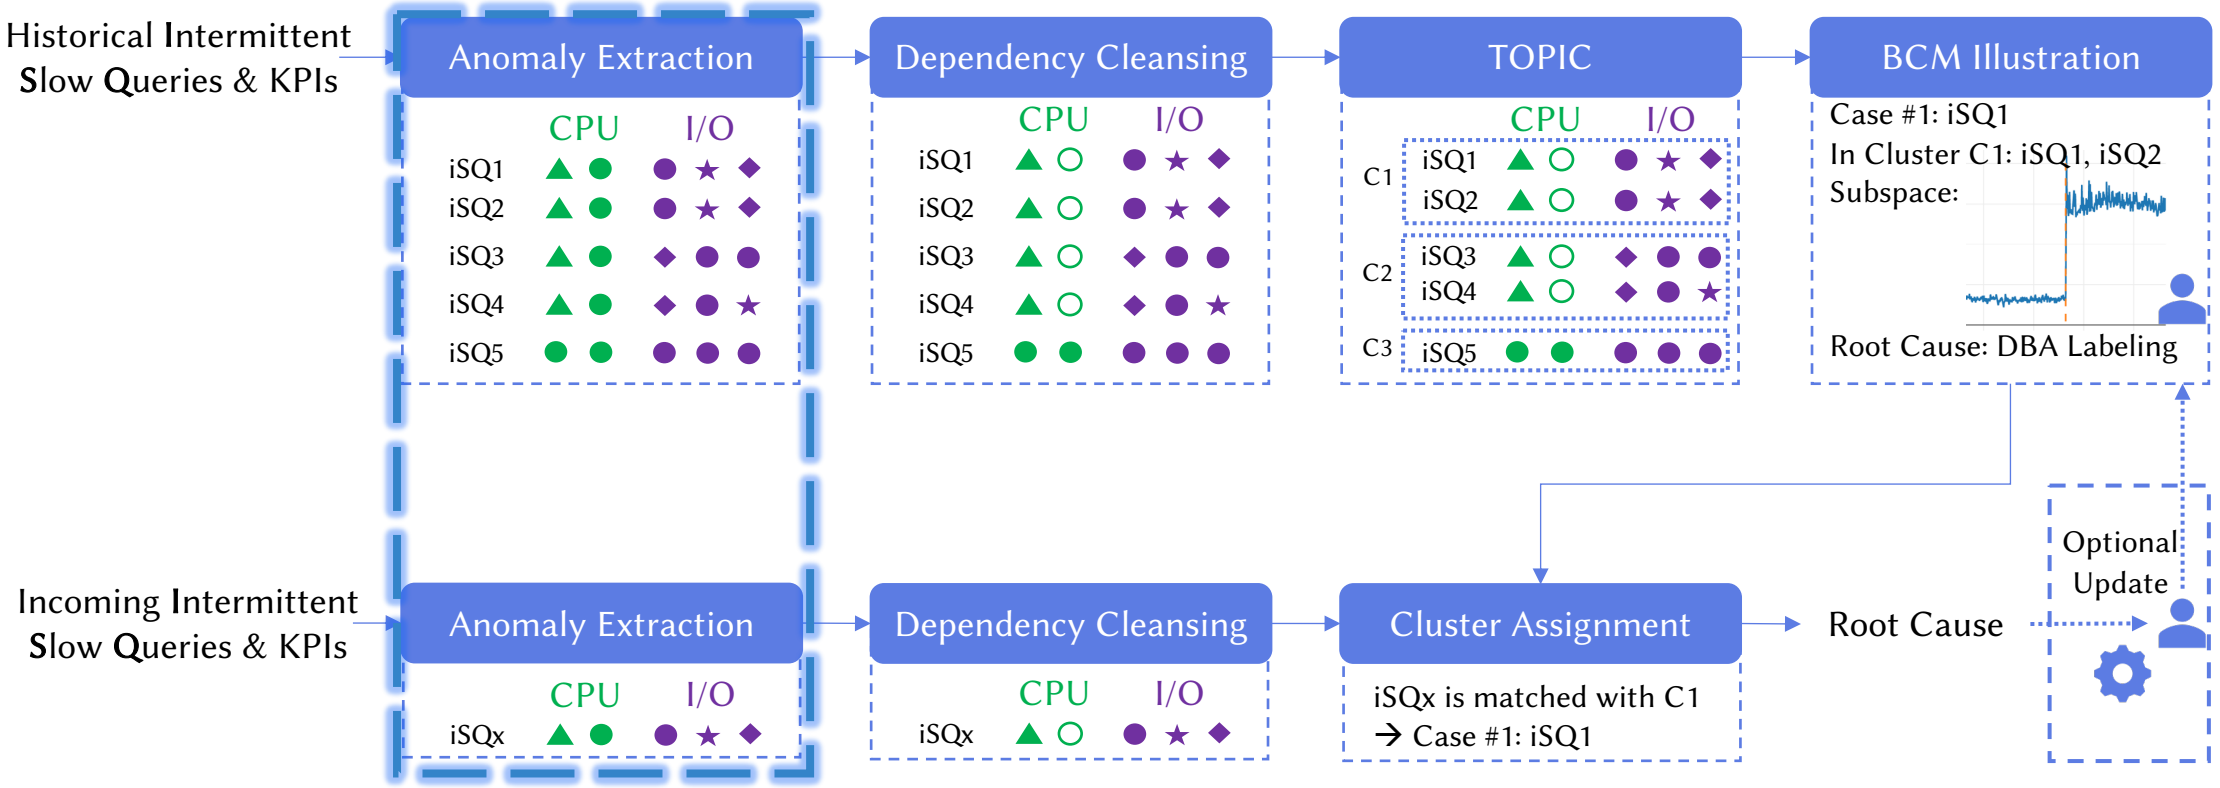

iSQUAD Overview

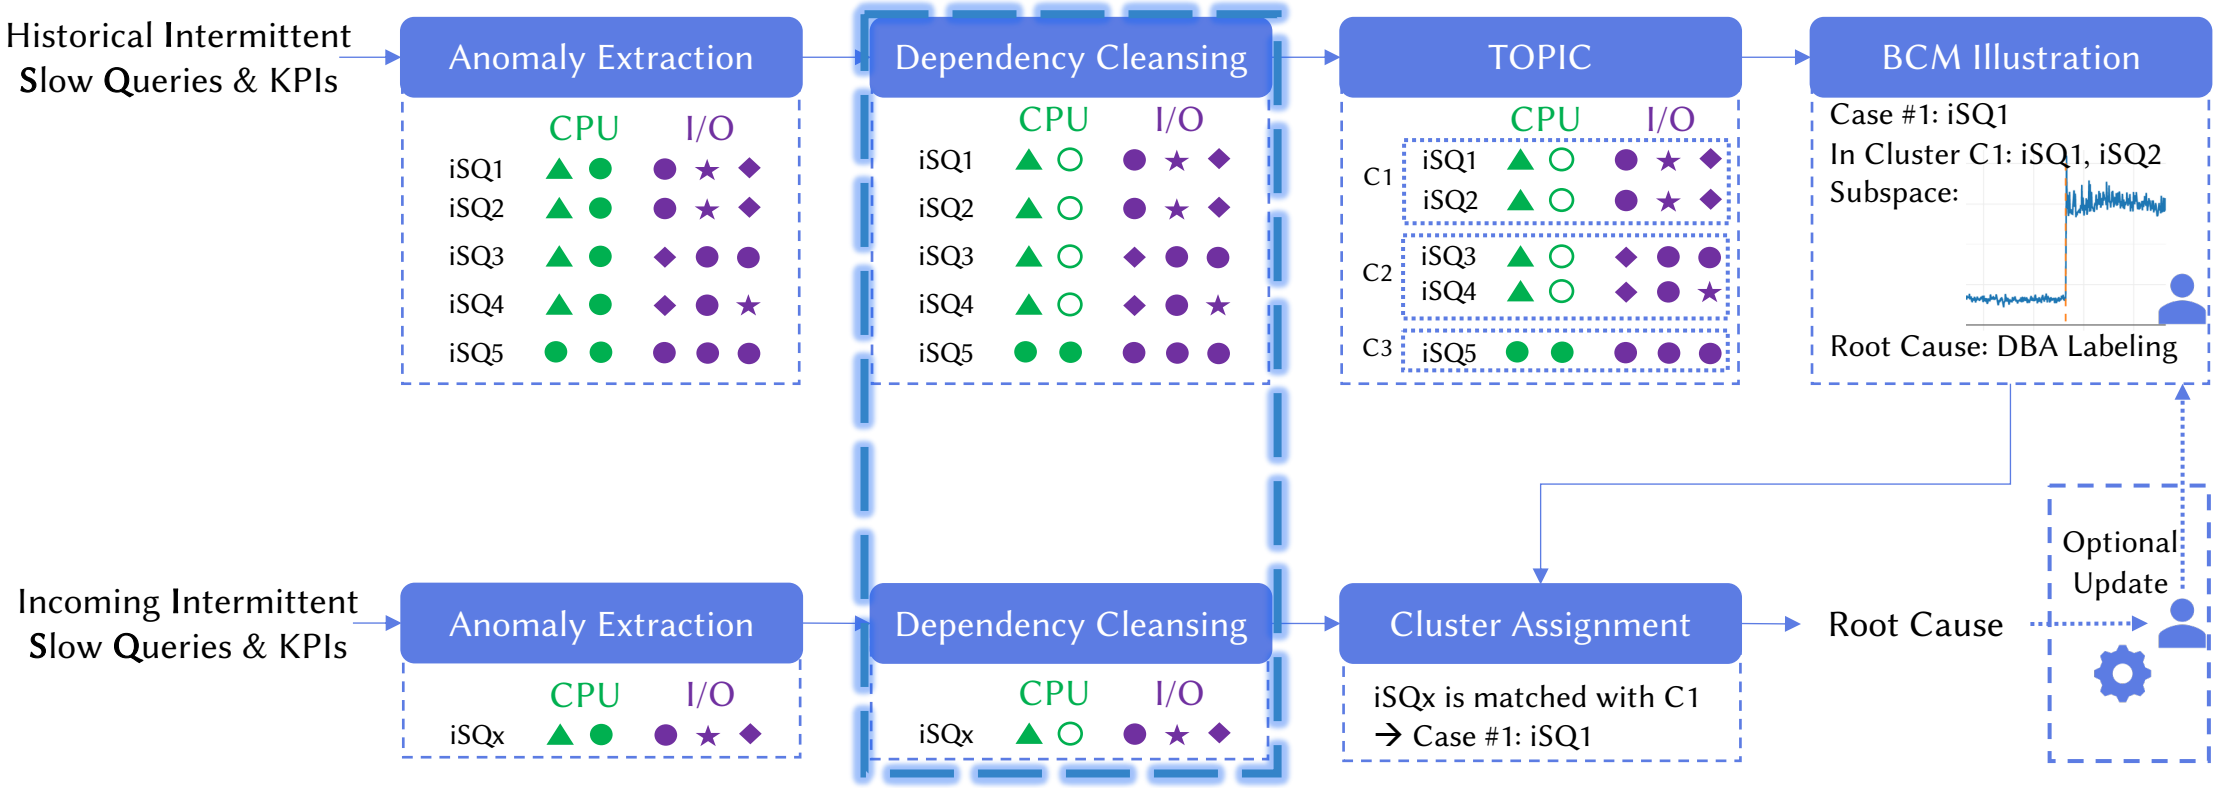

iSQUAD Overview

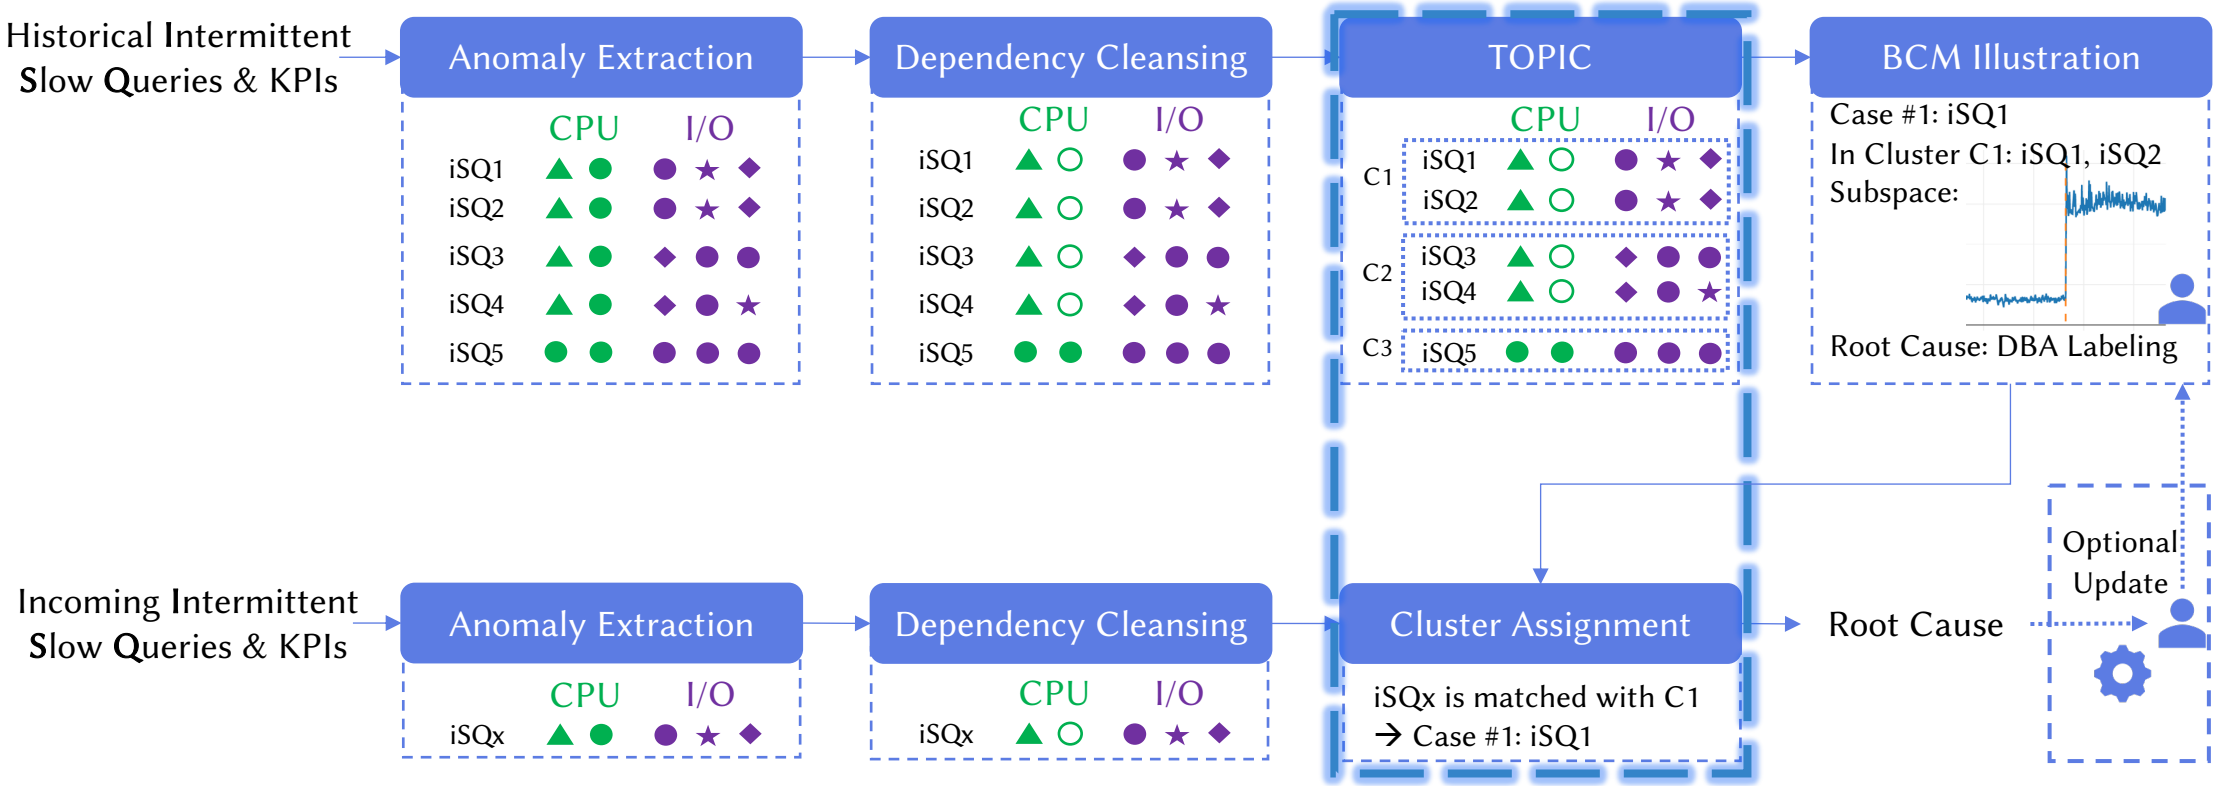

iSQUAD Overview

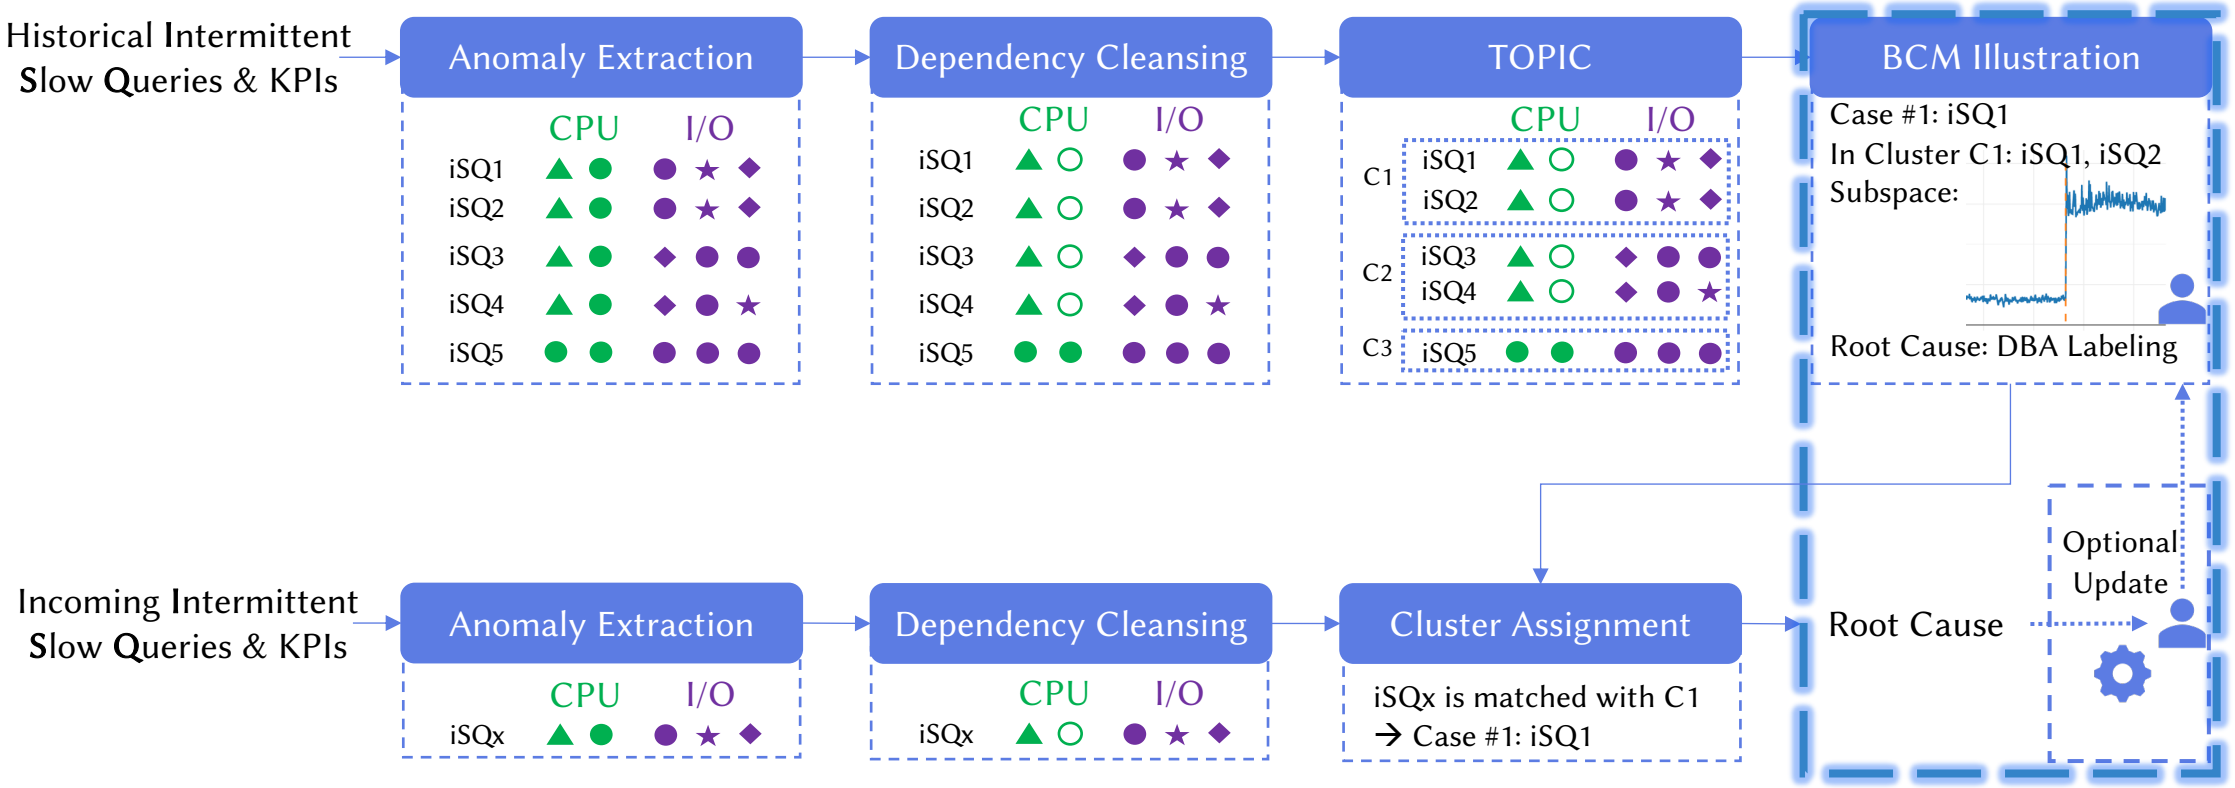

Anomaly Extraction

KPIs are important to locate iSQs' root causes.
The anomaly types of KPIs should be paid attention to.

Anomaly Extraction

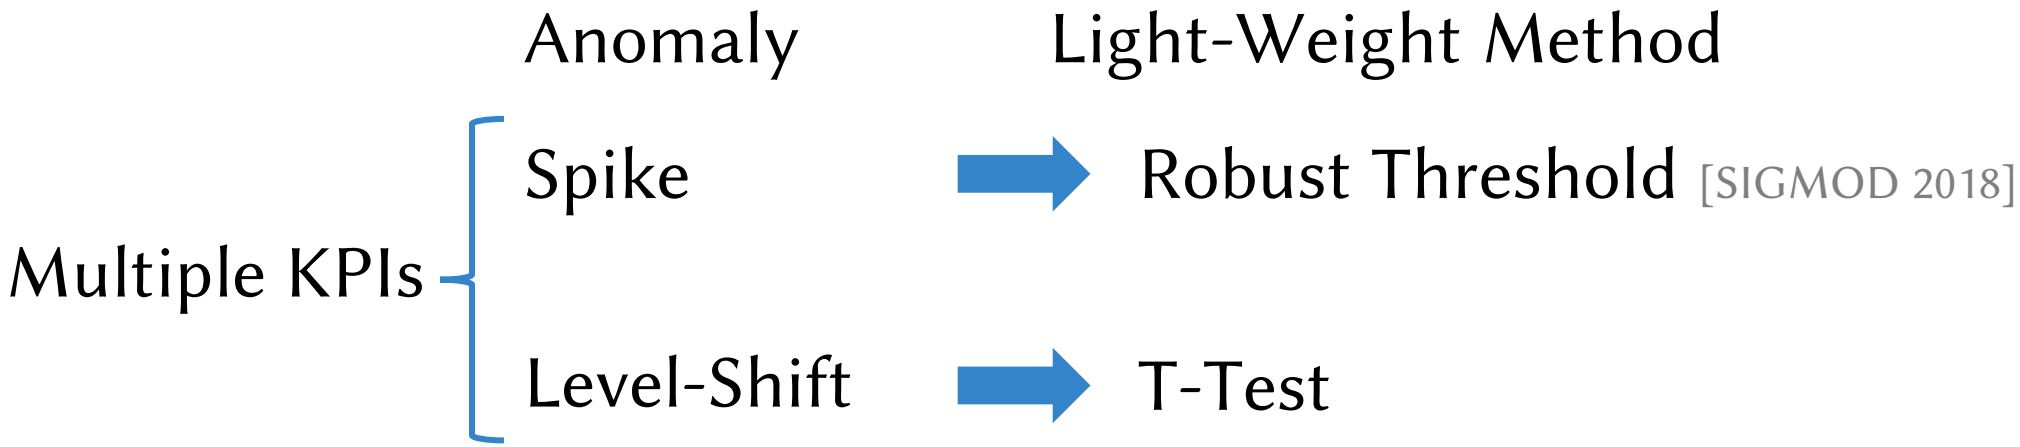

Anomaly	Method	F1-Score (%)	Time (s)
Spike	**Robust Threshold**	**98.7**	**0.19**
	dSPOT [KDD 2017]	81.0	15.11
Level-Shift	**T-Test**	**92.6**	**0.23**
	iSST [ISSRE 2018]	60.7	6.06

Dependency Cleansing

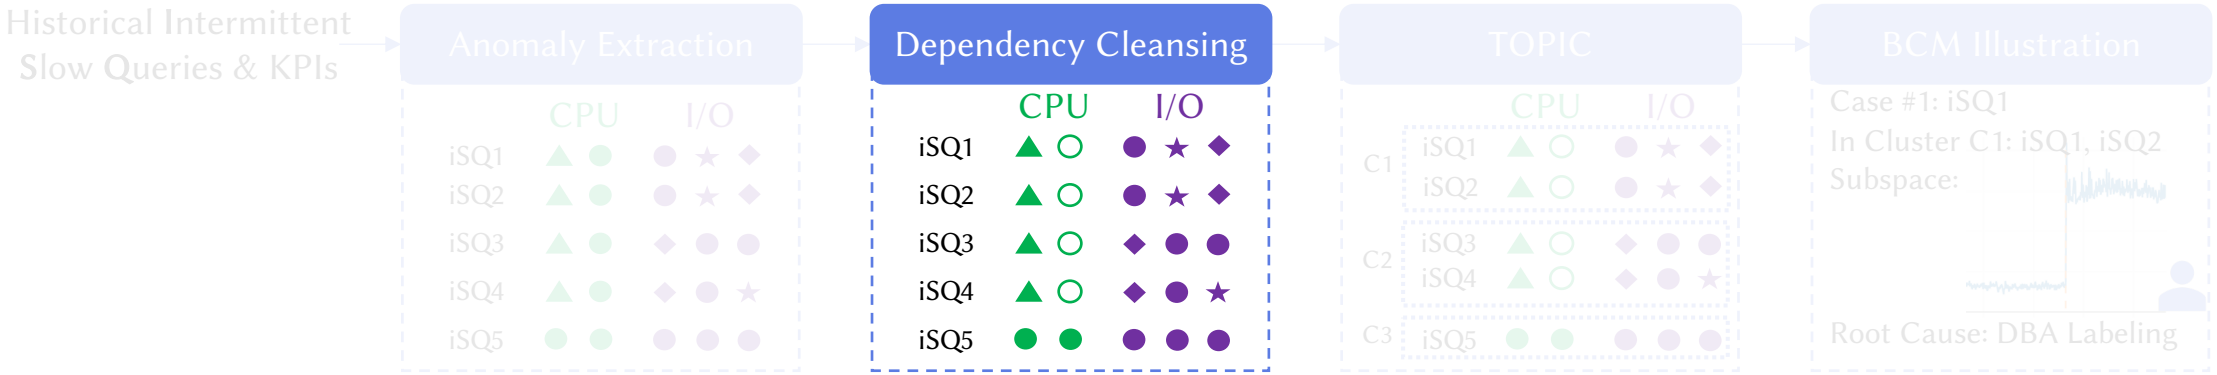

KPI anomalies are highly correlated.
e.g. docker.cpu-usage (instance) -> cpu.usage (host)

Dependency Cleansing

Based on the association rule mining between two KPIs to determine whether the two KPIs have a correlation

$$confidence(A \rightarrow B) = \frac{|A \cap B|}{|A|}$$

Method	Precision (%)	Recall (%)	F1-Score (%)
Confidence	90.9	**100**	**95.2**
MI [Sigmod 2016]	**100**	40	57.1
Gain Ratio [Infocom 2016]	87.5	70	77.8

TOPIC: Type-Oriented Pattern Integration Clustering

Similar symptoms are correlated to the same root cause.

TOPIC: Type-Oriented Pattern Integration Clustering

- KPI Type – KPIs are classified into eight types by DBAs
- Anomaly Pattern – Similarity calculate by matching coefficient
- Clustering key idea – hierarchically merge similar pattern

KPI Type	CPU	I/O	Network	Workload	...
iSQ1	▲ ●	● ★ ◆ ▲ ▲ ● ★	● ★ ★ ●	▼ ● ▼ ●	...
iSQ2	▲ ●	● ◆ ★ ▲ ▲ ◆ ★	◆ ★ ★ ●	▼ ● ▼ ●	...
Similarity	100%	57% (4/7)	75% (3/4)	100%	...

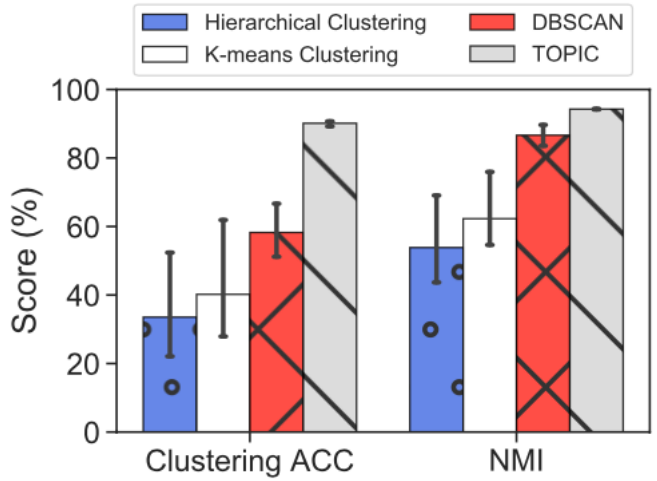

BCM Illustration

Bayesian Case Model (BCM) is a framework for extracting prototypical cases and feature subspace [NeurIPS 2014].

BCM Illustration

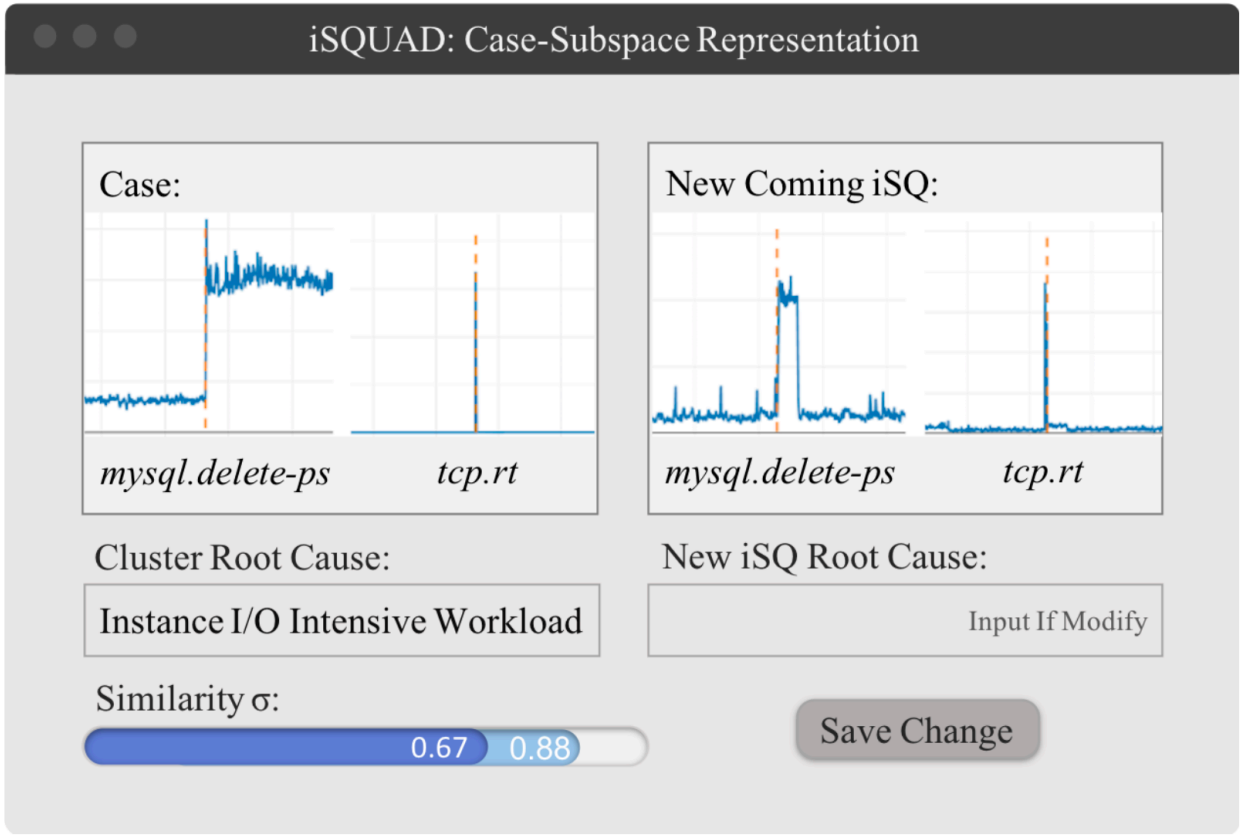

- Initial labeling root cause
- Visualization case and feature (anomaly KPI) subspace
- New coming iSQ's root cause modification
- Labeling new clusters

iSQUAD Runs in Alibaba Cloud Database

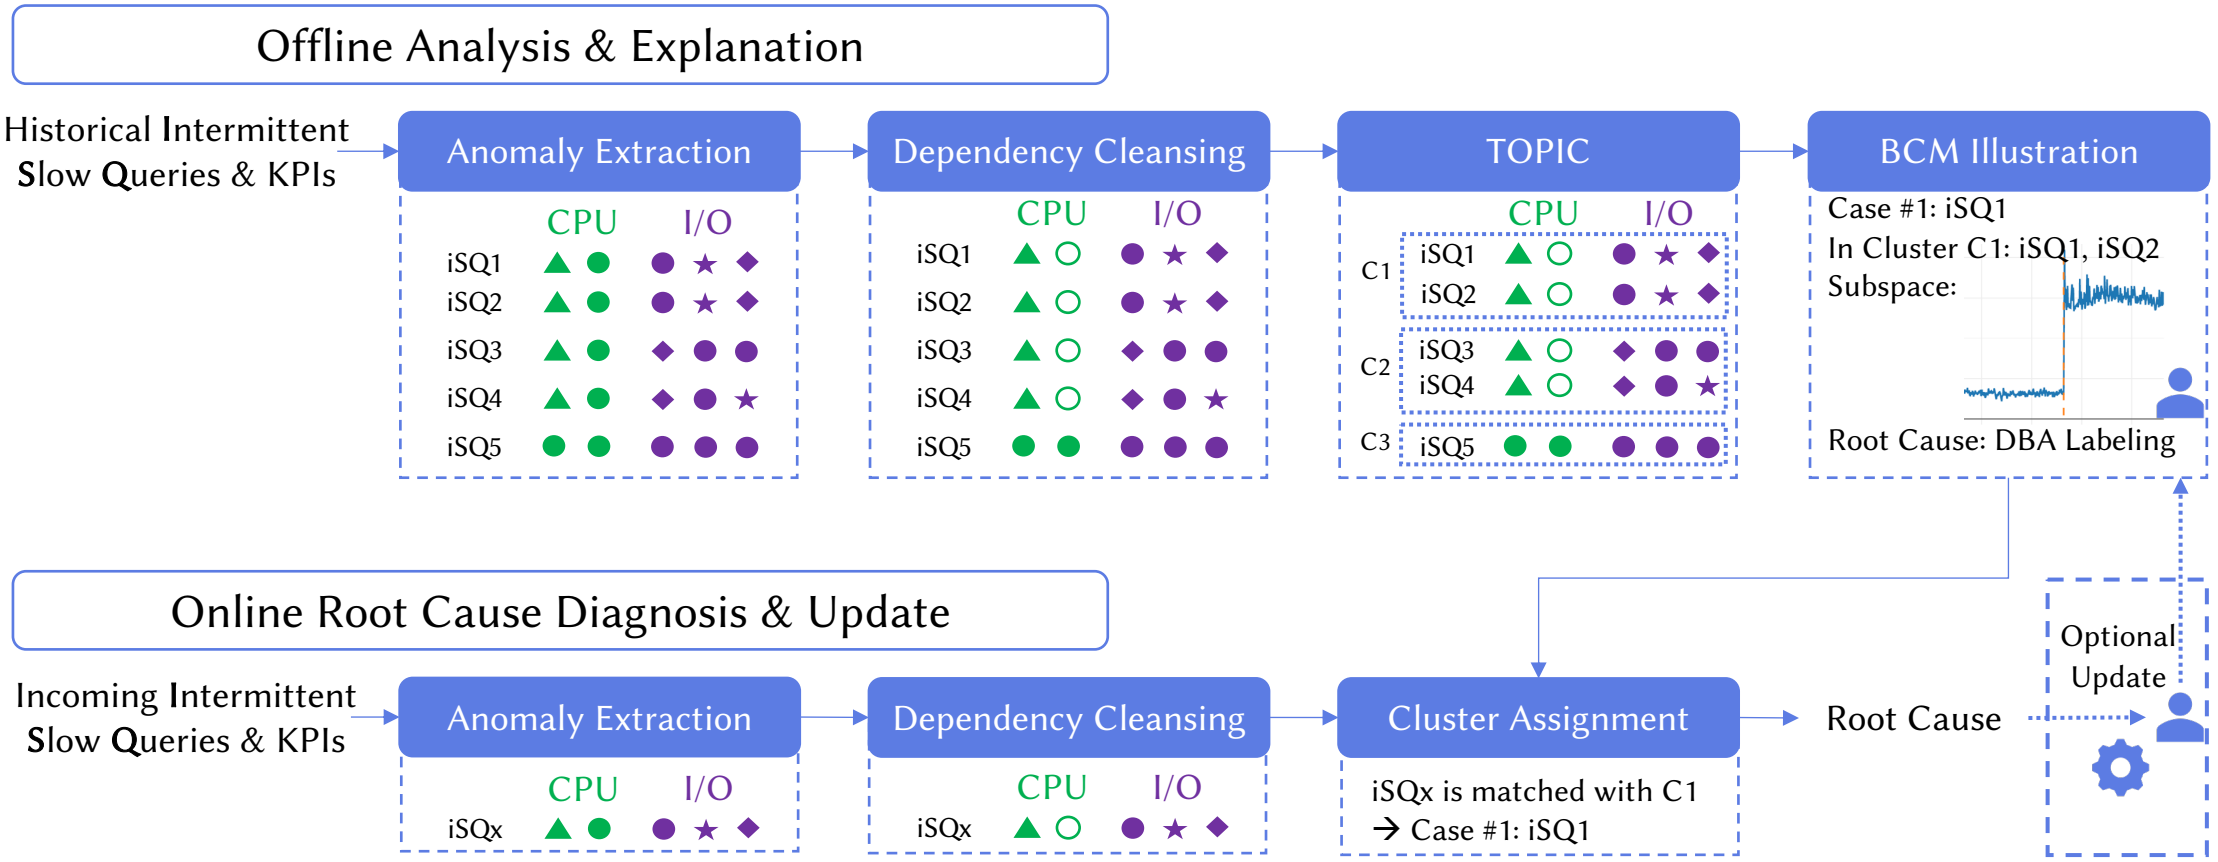

Outline

What's iSQ?

Why it's challenging?

How to diagnose it?

Evaluation

iSQUAD Accuracy

iSQ ground truth labeled by DBAs

No.	Root Cause	Offline	Online
1	Instance CPU Intensive Workload	27.6%	34.5%
2	Host I/O Bottleneck	17.2%	17.2%
3	Instance I/O Intensive Workload	0.9%	15.8%
4	Accompanying Slow SQL	8.6%	9.0%
5	Instance CPU & I/O Intensive Workload	8.1%	4.8%
6	Host CPU Bottleneck	7.5%	4.1%
7	Host Network Bottleneck	6.9%	4.1%
8	External Operations	6.9%	3.5%
9	Database Internal Problem	3.4%	3.5%
10	Unknown Problem	2.9%	3.5%

End to end performance

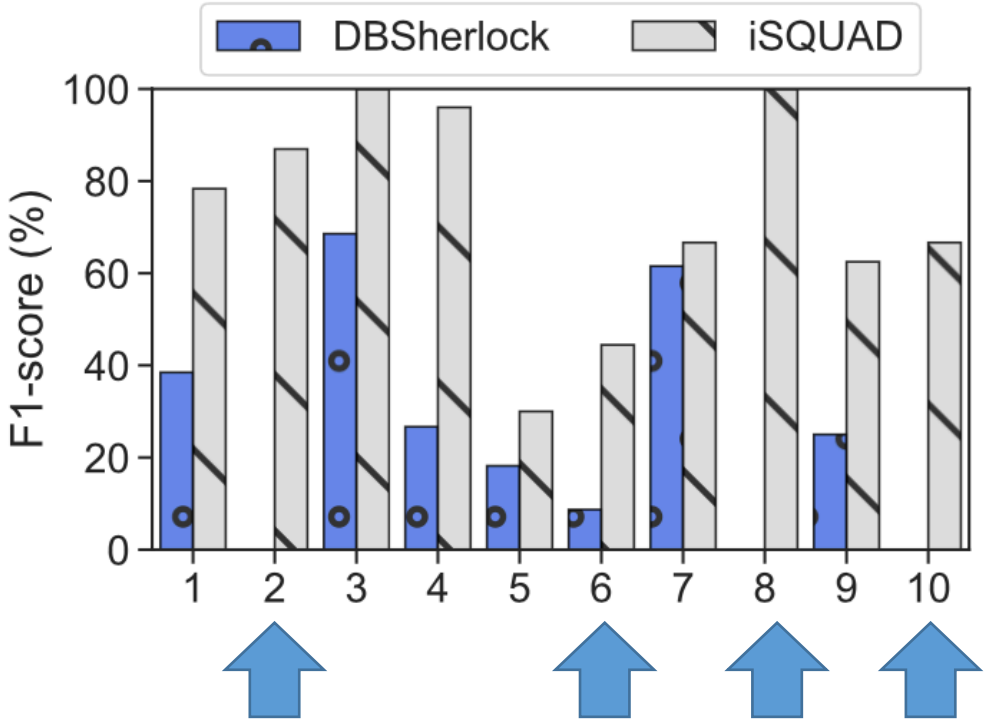

Root causes are not included in DBSherlock [Sigmod 2016]

More in Our Paper

- iSQUAD Efficiency
- BCM Effectiveness
- Parameter Sensitivity
- Contribution of Components
- Multiple Root Causes
- Generality of iSQUAD
- Root Causes to Actions

Diagnosing Root Causes of Intermittent Slow Queries in Cloud Databases

Minghua Ma[*1,2], Zheng Yin[2], Shenglin Zhang[†1], Sheng Wang[2], Christopher Zheng[1], Xinhao Jiang[1], Hanwen Hu[1], Cheng Luo[1], Yilin Li[1], Nengjun Qiu[2], Feifei Li[2], Changcheng Chen[2] and Dan Pei[1]

[1]Tsinghua University, {mmh16, luo-c18, liyilin16}@mails.tsinghua.edu.cn, peidan@tsinghua.edu.cn
[2]Alibaba Group, {yinzheng.yz, sh.wang, nengjun.qnj, lifeifei}@alibaba-inc.com
[3]Nankai University, zhangsl@nankai.edu.cn

ABSTRACT

With the growing market of cloud databases, careful detection and elimination of slow queries are of great importance to service stability. Previous studies focus on optimizing the slow queries that result from internal reasons (*e.g.*, poorly-written SQLs). In this work, we discover a different set of slow queries which might be more hazardous to database users than other slow queries. We name such queries **Intermittent Slow Queries (iSQs)**, because they usually result from *intermittent* performance issues that are external (*e.g.*, at database or machine levels). Diagnosing root causes of iSQs is a tough but very valuable task.

This paper presents **iSQUAD**, Intermittent Slow QUery Anomaly Diagnoser, a framework that can diagnose the root causes of iSQs with a loose requirement for human intervention. Due to the complexity of this issue, a machine learning approach comes to light naturally to draw the interconnection between iSQs and root causes, but it faces challenges in terms of versatility, labeling overhead and interpretability. To tackle these challenges, we design four components, *i.e.*, Anomaly Extraction, Dependency Cleansing, Type-Oriented Pattern Integration Clustering (TOPIC) and Bayesian Case Model. iSQUAD consists of an *offline clustering & explanation* stage and an *online root cause diagnosis & update* stage. DBAs need to label each iSQ cluster only once at the offline stage unless a new type of iSQs emerges at the online stage. Our evaluations on real-world datasets from Alibaba OLTP Database show that iSQUAD achieves an iSQ root cause diagnosis average F1-score of 80.4%, and outperforms existing diagnostic tools in terms of accuracy and efficiency.

PVLDB Reference Format:
M. Ma, Z. Yin, S. Zhang, S. Wang, C. Zheng, X Jiang, H. Hu, C. Luo, Y. Li, N. Qiu, F. Li, C. Chen and D. Pei. Diagnosing Root Causes of Intermittent Slow Queries in Cloud Databases. *PVLDB*, 13(8): 1176-1189, 2020.
DOI: https://doi.org/10.14778/3389133.3389136

[*]Work was done while the author was interning at Alibaba Group.
[†]Work was done while the author was a visiting scholar at Alibaba Group.

Proceedings of the VLDB Endowment, Vol. 13, No. 8
ISSN 2150-8097.
DOI: https://doi.org/10.14778/3389133.3389136

1. INTRODUCTION

The growing cloud database services, such as Amazon Relational Database Service, Azure SQL Database, Google Cloud SQL and Alibaba OLTP Database, are critical infrastructures that support daily operations and businesses of enterprises. Service interruptions or performance hiccups in databases can lead to severe revenue loss and brand damage. Therefore, databases are always under constant monitoring, where the detection and elimination of slow queries are of great importance to service stability. Most database systems, such as MySQL, Oracle, SQL Server, automatically log detailed information of those queries whose completion time is over a user-defined threshold [7, 37, 43], *i.e.*, slow queries. Some slow queries result from internal reasons, such as nature of complexity, lack of indexes and poorly-written SQL statements, which can be automatically analyzed and optimized [13, 32, 34, 42]. Many other slow queries, however, result from *intermittent* performance issues that are external (*e.g.*, at database or machine levels), and we name them **Intermittent Slow Queries (iSQs)**.

Usually, iSQs are the cardinal symptom of performance issues or even failures in cloud databases. As iSQs are intermittent, service developers and customers expect them to be responsive as normal, where sudden increases of latency have huge impacts. For example, during web browsing, an iSQ may lead to unexpected web page loading delay. It has been reported that every 0.1s of loading delay would cost Amazon 1% in sales, and every 0.5s of additional load delay for Google search results would led to a 20% drop in traffic [30]. We obtain several performance issue records carefully noted by DBAs of Alibaba OLTP Database in a year span: when a performance issue occurs, a burst of iSQs lasts for minutes. As a matter of fact, manually diagnosing root causes of iSQs takes tens of minutes, which is both time consuming and error-prone.

Diagnosing root causes of iSQs gets crucial and challenging in cloud. First, iSQ occurrences become increasingly common. Multiple database instances may reside on the same physical machines for better utilization, which in turn can cause inter-database resource contentions. Second, root causes of iSQs vary greatly. Infrastructures of cloud databases are more complex than those of on-premise databases [29], making it harder for DBAs to diagnose root causes. Precisely, this complexity can be triggered by instance migrations, expansions, storage decoupling, *etc.* Third, massive database instances in cloud make iSQs great in population. For example, tens of thousands of iSQs are generated in Alibaba OLTP Database per day. In addition, roughly 83% of enterprise workloads are forecasted to be in the cloud by 2020 [12]. This trend makes it critical to efficiently diagnose the root causes of iSQs.

In this work, we aim to diagnose root causes of iSQs in cloud databases with minimal human intervention. We learn about symp-

Conclusion

- **Motivation**: identify the problem of Intermittent Slow Queries in cloud databases
- **Challenge**: anomaly diversity, labeling overheads, interpretability
- **Solution**: iSQUAD framework for iSQs root cause diagnosis
  - Anomaly Extraction, Dependency Cleansing, TOPIC, Bayesian Case Model
- **Deployment**: iSQUAD runs in Alibaba Cloud Database

iSQUAD
VLDB 2020

Thank you

mmh16@mails.tsinghua.edu.cn