

# AID: Efficient Prediction of Aggregated Intensity of Dependency in Large-scale Cloud Systems

Tianyi Yang\*, Jiacheng Shen\*, Yuxin Su\*, Xiao Ling†, Yongqiang Yang†, and Michael R. Lyu\*

\*Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, China.

Email: {tyyang, jcshen, yxsu, lyu}@cse.cuhk.edu.hk

†Computing and Networking Innovation Lab, Cloud BU, Huawei

Email: {lingxiao1, yangyongqiang}@huawei.com

**Abstract**—Service reliability is one of the key challenges that cloud providers have to deal with. In cloud systems, unplanned service failures may cause severe cascading impacts on their dependent services, deteriorating customer satisfaction. Predicting the cascading impacts accurately and efficiently is critical to the operation and maintenance of cloud systems. Existing approaches identify whether one service depends on another via distributed tracing but no prior work focused on discriminating to what extent the dependency between cloud services is. In this paper, we survey the outages and the procedure for failure diagnosis in two cloud providers to motivate the definition of the intensity of dependency. We define the intensity of dependency between two services as how much the status of the callee service influences the caller service. Then we propose AID, the first approach to predict the intensity of dependencies between cloud services. AID first generates a set of candidate dependency pairs from the spans. AID then represents the status of each cloud service with a multivariate time series aggregated from the spans. With the representation of services, AID calculates the similarities between the statuses of the caller and the callee of each candidate pair. Finally, AID aggregates the similarities to produce a unified value as the intensity of the dependency. We evaluate AID on the data collected from an open-source microservice benchmark and a cloud system in production. The experimental results show that AID can efficiently and accurately predict the intensity of dependencies. We further demonstrate the usefulness of our method in a large-scale commercial cloud system.

**Index Terms**—cloud computing, software reliability, AIOps, service dependency

## I. INTRODUCTION

Service reliability is one of the key challenges that cloud providers have to deal with. The common practice nowadays is developing and deploying small, independent, and loosely coupled cloud microservices that collectively serve users' requests. The microservices that serve the same purpose are called cloud services<sup>1</sup>. The microservices communicate with each other through well-defined APIs. Such an architecture is called microservice architecture [1]. The microservice architecture has been widely adopted in cloud systems because of its reliability and flexibility. Under this architecture, microservice management frameworks like Kubernetes will be responsible for managing the life cycles of microservices.

Yuxin Su is the corresponding author.

<sup>1</sup>For simplicity, in this paper, "cloud service" and "cloud microservice" are interchangeable when they are used alone.

Developers can focus on the application logic instead of the bothering tasks of resource management and failure recovery.

Although microservice management frameworks provide automatic mechanisms for failure recovery, unplanned service failures may still cause severe cascading effects. For example, failures of critical services that provide basic request routing functions will impact the invocation of cloud services, slow down request processing, and deteriorate customer satisfaction. Therefore, evaluating the impact of service failures rapidly and accurately is critical to the operation and maintenance of cloud systems. Knowing the scope of the impact, reliability engineers can put more emphasis on services that have greater impacts on others.

A failed service will only affect services that will invoke it. In other words, service invocations cause dependencies between services. Many recent approaches [2], [3] propose to use the dependencies of services to approximate their failure impact. All the services and dependencies in a cloud system collectively construct a directed graph of services, which is also called a dependency graph. Identifying whether one service depends on another in cloud systems can be well solved by industrial tracing frameworks like Dapper and Jaeger. By using these frameworks, all the invocations between the caller and callee services can be recorded as traces that are composed of spans. The attributes about each invocation, like duration, status, invoked service name, timestamp, etc., are recorded in each span. Based on the spans, current dependency detection methods treat the dependency as a binary value indicating whether one service invokes another or not.

However, modeling the relations of services solely with binary dependencies is not precise enough. To show the insufficiency of existing methods, we first conduct an empirical study on the outages of Amazon Web Service and Huawei Cloud. We point out that it is inefficient to conduct failure diagnosis and recovery based on binary dependencies. This is because the different dependencies of a cloud service impact the cloud service in different ways. Manual examination of different dependencies without any priority is inefficient, especially in cloud systems where the number of dependencies could be large. Based on this observation, we argue that it will be helpful if the dependency can be measured as a continuous value that indicates the intensity of this dependency. Specifically, by checking services that are dependent on the failed service

with large intensity values, on-call engineers (OCEs) can find the root cause of a system failure with a higher probability. By recovering the services that are strongly dependent on the failed one, the whole system could be restored faster.

To improve the reliability of cloud systems, in this paper, we propose AID, an end-to-end approach to predict the intensity of dependencies between cloud microservices for cascading failure prediction. We first generate a set of candidate dependency pairs from the spans. Then we distribute each span into different fixed-length bins according to their timestamp and service name. We calculate the statistics of all spans in each bin as the Key Performance Indicators (KPIs) for the bin. The KPIs of one service form a multivariate time series that will be treated as the representation of the service's status. For each candidate dependency pair, we calculate the similarities between the statuses of the two services in the pair. Finally, we aggregate the similarities to produce a unified value as the intensity of the pair.

To show the effectiveness of AID, we evaluate AID on two datasets. One is a simulated dataset, and the other is an industrial dataset. For the simulated dataset, we deploy train-ticket, an open-source microservice benchmark system, simulate users' requests, and collect the traces. For the industrial dataset, we collect the traces from a production cloud system. Then we evaluate AID on the datasets and compare its performance with several baselines. The experimental results show that our proposed method can accurately measure the intensity of dependencies and outperform the baselines. Furthermore, we showcase the successful usage of our method in a large-scale production cloud system. In addition, we release both datasets to facilitate future studies.

The main contributions of this work are highlighted as follows:

- We propose AID, the first method to quantify the intensity of dependencies between different services.
- The evaluation results show the effectiveness and efficiency of the proposed method.
- We release a simulated dataset and an industrial dataset from a production cloud system to facilitate future studies.

**Organization.** The remainder of this paper is organized as follows. Section II provides motivation and background knowledge that underpin our approach. We describe our survey and empirical study on real outages that motivate the proposed method in Section III. Section IV elaborates on the method in detail. Section V introduces the datasets, baselines and shows the experimental results. Successful use cases of the proposed method in a production cloud system are demonstrated in Section VI. We discuss the practical usage, the perceived limitations, and the possible threats to validity in Section VII. Section VIII introduces related works. The last section, Section IX, concludes this paper and lists directions for future exploration.

## II. BACKGROUND

In this section, we briefly describe the service-oriented architecture of cloud systems and the distributed tracing tools in cloud systems. Then we present the main techniques, i.e., time series similarity analysis, that underpin our approach.

### A. The Architecture of Cloud Systems

Modern cloud systems are often constructed from a complex and large-scale hierarchy of distributed software modules [4]. The common practice nowadays is to develop and deploy these software modules as cloud microservices that collectively comprise multiple large cloud services [5]. Microservices are small, independent, and loosely coupled software modules that can be deployed independently [1]. Different microservices serve different responsibilities [6] like user authentication, resource allocation, virtual network management, billing, etc. When an external request arrives at the cloud system, the request will be routed through the system and served by dozens of different cloud services and microservices. The microservices communicate with each other through well-defined APIs and, therefore, can be refactored and scaled independently and dynamically to adapt to incidents like surges of requests and service failures [7]. Such an architecture is called microservice architecture [1].

The microservice architecture becomes increasingly popular due to its high flexibility, reusability, and scalability [8]. It enables agile development and supports polyglot programming, i.e., microservices developed under different technical stacks can work together smoothly. However, the loosely coupled nature of microservices makes it difficult for engineers to conduct system maintenance. Different microservices in a large cloud system are usually developed and managed by separate teams. Each team only has access to their own services as well as services that are closely related, which means they only have a local view of the whole system [9].

As a result, the failure diagnosis, fault localization, and performance debugging in a large cloud system become more complex than ever [10]–[12]. Despite various fault tolerance mechanisms introduced by modern cloud systems, it is still possible for minor anomalies to magnify their impact and escalate into system outages. As exemplified in Section III-A, when a cloud service enters an anomalous state and does not return results in a timely manner, other services that depend on it will also suffer from the increased request latency. Such anomalous states can propagate through the service-calling structure and eventually affect the entire system, resulting in a degraded user experience or even a service outage.

### B. Distributed Tracing

For commercial cloud providers, it is crucial to troubleshoot and fix the failures in a timely manner because massive user applications may be affected even by a small service failure [13]. Distributed tracing is a crucial technique for gaining insight and observability to cloud systems.

In large-scale cloud systems, a request is usually handled by multiple chained service invocations. As clues to defective

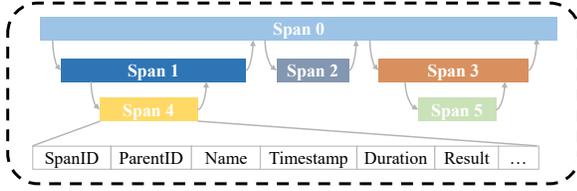


Fig. 1. A trace log with six spans.

services are hidden in the intricate network of services, it is difficult for even knowledgeable OCEs to keep track of how a request is processed in the cloud system. Distributed tracing provides an approach to monitor the execution path of each request. For chained service invocations, e.g., service A invokes service B, and service B invokes service C, it is important to know the status of each service invocation, including the result, the duration of execution, etc. By adding hooks to the services and microservices of the cloud system, a distributed tracing system [14]–[16] can record the contextual information of each service invocation. Such records are called span logs, abbreviated as spans. A span represents a logical unit of execution that is handled by a microservice in a cloud system. All the spans that serve for the same request collectively form a directed graph of spans, as illustrated in Figure 1. Such a directed graph of spans generated by a request is called a piece of trace log, abbreviated as a trace. A trace represents an execution path through the cloud system. With a trace, engineers can track how the request propagates through the cloud system. Collectively analyzing the traces of the entire cloud system can help engineers obtain in-depth latency reports that could assist failure diagnosis, fault localization, and surface performance degradation in the cloud system.

Span ID	e22f30bdbfd09134
Parent Span ID	b42a04bf18997d5d
Name	ts-preserve-service
Timestamp ( $\mu s$ )	1618589098705000
Duration ( $\mu s$ )	1126
Result	SUCCESS
Trace ID	c0d17d481f47bdd9
Additional Logs	...

Fig. 2. A span generated by the train-ticket benchmark.

Although the actual implementation of distributed tracing systems varies a lot, the types of information they record are similar. For clarity, we formally describe the attributes of spans as follows. Suppose we have a trace  $T$  composed of spans  $\{s_1, s_2, \dots, s_n\}$ , a span  $s_i \in T$  contains the following attributes<sup>2</sup>:

- $s_i^{id}$ : The ID of span  $s_i$ ,
- $s_i^{pid}$ : The ID of the parent span of  $s_i$ ,
- $s_i^{tid}$ : The ID of the trace that  $s_i$  belongs to,

<sup>2</sup>Other additional contextual information [17] is omitted as we do not use them in our method.

- $s_i^{name}$ : The name of service/microservice corresponding to  $s_i$ ,
- $s_i^{ts}$ : The time stamp of  $s_i$ ,
- $s_i^d$ : The duration of execution of  $s_i$ , and
- $s_i^r$ : The result of execution of  $s_i$ .

Figure 2 illustrates a span generated by the train-ticket benchmark [18]. It means that service `ts-preserve-service` was invoked at 04:58 on April 17, 2020. The duration of execution is 1126  $\mu s$ , and the execution result is SUCCESS.

### C. Time Series Similarity Analysis

Time series data are ubiquitous. One important task in time series data mining is to measure the similarity between two time series. Similar to human intuition, the similarity measure is usually based on the similarity between the shapes of two time series [19].

Dynamic time warping (DTW) [20] is a widely-used similarity measure when two time series have the same overall component shapes but are not aligned on the timeline. It attempts to align two time series along a timeline by distorting the timeline for one time series so that its converted form is better aligned with the second time series. DTW was initially used in speech recognition applications [20] and extended and optimized by many works [21]–[23].

## III. MOTIVATIONS

The research described in this paper is motivated by the maintenance of a real-world cloud system in production. In this section, we first survey thirteen publicly known service outages that severely affected Amazon Web Services (AWS) from 2011 to 2020. Among the thirteen outages, we identify five that are related to service dependency and summarized the consequences of inappropriate management of service dependency. Second, we empirically study the diagnosis records of five real outages in the cloud system of Huawei Cloud that are related to inappropriate management of service dependency. Our study indicates that the information in the traces has not been used efficiently and current practice heavily relies on the engineers’ familiarity with the dependencies in the system. Lastly, we propose to measure the intensity of dependency in terms of status propagation between dependent cloud services. We demonstrate the usefulness of the intensity by motivating examples in real cloud systems.

### A. A Survey of the Outages in AWS

Service outages are inevitable in the cloud [24]. In this section, we empirically analyzed over 1000 incidents of Huawei Cloud in 2019 and thirteen publicly known major outages<sup>3</sup> of AWS from 2011 to 2020. Among the incidents of Huawei Cloud, we found that improper service dependency is the most frequent reason for failures in Huawei Cloud. Among the outage summaries of AWS, we also identified that five

<sup>3</sup><https://aws.amazon.com/premiumsupport/technology/pes/>

TABLE I  
SUMMARY OF AWS OUTAGES RELATED TO SERVICE DEPENDENCY.

Date	Consequences	
	Cascading Failure	Slow Recovery
Apr 21, 2011	✓	
June 29, 2012		✓
Oct 22, 2012	✓	
Aug 7, 2014		✓
Nov, 25 2020	✓	✓

of the outages (38%) are related<sup>4</sup> to service dependency. As shown in Table I, among the five outages that are related to service dependency, three of them are due to cascading failures triggered by erroneous upgrades of services. During the failure recovery, the inappropriate dependencies lead to slow failure recovery in three outages.

AWS is the worldwide leading cloud provider. It operates in many regions, each consisting of multiple Availability Zones (AZs). Each AZ uses separate physical facilities and independently provides various cloud services [5], including Steam Data Processing (Kinesis), API Usage Analysis (Cognito), Customer Dashboard (Cloudwatch), Elastic Compute Cloud (EC2), Relational Database Service (RDS), Elastic Load Balancing (ELB), and Low-level Block Storage (EBS), etc. For brevity's sake, we simplify the dependencies as 1) EC2, RDS, and ELB all depend on EBS, and 2) Cognito and Cloudwatch depend on Kinesis<sup>5</sup>.

The outages on April 21, 2011, and October 22, 2012, are both caused by erroneous upgrades of EBS. When EBS failed, the services that depend on EBS, i.e., EC2, ELB, and RDS, are all affected. The cascading failures resulted in service disruptions of over 48 hours in the US-East-1 Region of AWS.

The outages on June 29, 2012, and August 7, 2014, are both triggered by the blackouts. After the blackout, the RDS and ELB services restarted quickly as expected, but they are still unable to fully recover because they both depend on EBS service which, at that time, can not recover simultaneously. The slow failure recovery incurred by service dependencies affected the service availability for days in the US-East-1 Region and the EU West-1 Region of AWS. As a follow-up optimization, ELB service reduced the dependency on EBS after the outage in 2014.

On November, 25 2020, the erroneous upgrade of Kinesis lead to its failure, cascadingly causing the failure of Cognito and Cloudwatch. More severely, during the recovery, AWS could not notify the customers via the normal way because the normal customer notification service also relied on Cognito. Due to the inner mechanism of Kinesis, the recovery of Kinesis took more than ten hours. Thus the recoveries of Cognito

<sup>4</sup>The outages are usually caused by various reasons that mutually affect each other. Service dependency is one of the reasons, so we use the word "related".

<sup>5</sup>The actual dependency relations between these services are complicated. We omit the details here.

and Cloudwatch were also slowed down. As a follow-up optimization, Cognito and Cloudwatch services reduced the dependency on Kinesis after the severe outage.

### B. Drawbacks of Current Failure Diagnosis Methods

To gain more knowledge about the procedure of failure diagnosis in industrial circumstances, we first interviewed engineers in Huawei Cloud<sup>6</sup>. Then we summarize the procedure of failure diagnosis, and point out the drawbacks of current practice in Huawei Cloud.

In Huawei Cloud, the failure diagnosis can be triggered by two systems, i.e., the customer support system and the monitoring system. When a customer experiences a service disruption, the customer can submit a support ticket in the customer support system. The on-call engineers will distribute the support ticket to the corresponding engineers responsible for the service. The monitoring system, on the other hand, monitors the Key Performance Indicators (KPIs) and the logs of each service in the cloud system. If the KPIs or the number of erroneous logs of one service increased abnormally or reached predefined thresholds, the monitoring system will send an alert to the corresponding engineers. Upon receiving the support ticket or alert, engineers start diagnosing the failures.

We summarize the common practice of failure diagnosis in Huawei Cloud as follows. Suppose the anomalous service is A, OCEs will first check whether the failure is caused by the faults of service A (e.g., an erroneous upgrade). If so, the development team of service A will handle the failure. If service A is in good condition, OCEs will analyze the status of all services that A depends on. The status includes the number of calls, the error rate, etc. If they found the failure of a service B is likely to cause the failure of service A, then engineers will continue to investigate service B. Recall that all the services construct a directed graph where each node represents a service. The failure diagnosis procedure can be viewed as a recursive search on the service dependency graph.

The practice works well in small cloud systems that contain tens of cloud services. However, the dependencies in large-scale cloud systems are much more complicated [10], making manual failure diagnosis inefficient and difficult for engineers. Engineers may have trouble identifying the cause of the failure. In this case, the development teams of all cloud services have to check whether the failure is caused by their corresponding services. Sometimes engineers may infer the possible causes of a failure, but it heavily relies on the engineer's familiarity with the dependencies in the system. In summary, the complex dependency relations in large-scale cloud systems make failure diagnosis difficult, and current practice is inefficient and dependent on the human experience.

### C. Intensity of Service Dependency

A cloud system is composed of many services. The dependency between two services is caused by one service invoking the other via predefined APIs. Existing tools [2], [25], [26]

<sup>6</sup>AWS does not disclose the detailed procedures of failure diagnosis related to the five outages, so we cannot analyze the aforementioned outages in depth.

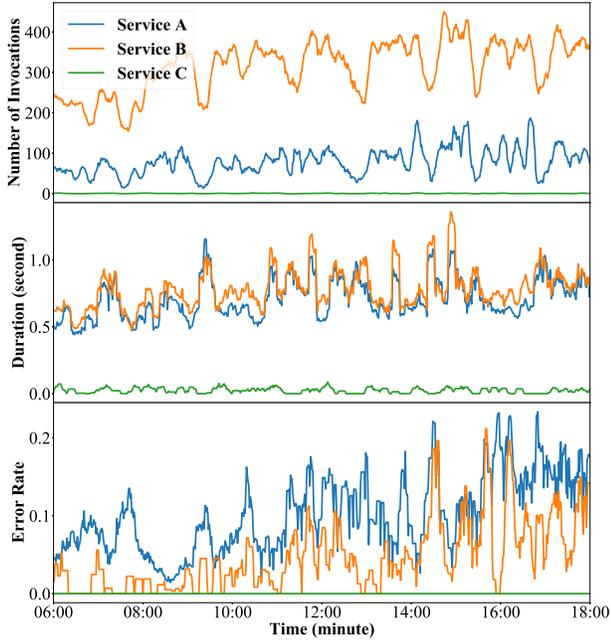


Fig. 3. The statuses of service A, B and C. A invokes B and C but B has a greater effect on A.

treat the dependency as a binary relation, i.e., if the caller service invokes the callee service, then the caller is dependent on the callee. We suggest that this binary dependency metric is not fine-grained enough for cloud maintenance. Figure 3 shows the statuses of three services<sup>7</sup> A, B, and C in Huawei Cloud. Service A invokes both service B and service C. Service B encountered failures. The x-axis represents time in minute. The y-axes represent the number of invocations per minute, the average duration of invocations per minute, and the error rate per minute of A, B, and C. Although service A invokes service B and service C, it is obvious that the statuses of B and C influence the status of A in different degrees. The reason is that the functionalities provided by service A and B are creating virtual machines, and allocating block storage, respectively. Creating a virtual machine requires allocating one or more block storage. Thus, the failure of service B inevitably affects service A. On the contrary, due to the fault tolerance mechanism of service A, the failure of service C will not affect service A a lot. Thus, it is more accurate to say that the intensity of dependency between service A and service B is higher than the intensity of dependency between service A and service C. As can be seen in Figure 3, the similarity of the statuses reflect the difference in the intensities.

Ideally, if the development team of every cloud microservice accurately provides the intensity of dependencies for every dependent services, the failure diagnosis could be accelerated. OCEs can prioritize the services that exhibit higher intensity of dependency instead of inspecting all the dependent services (Section III-B) if they have accurate intensity information.

<sup>7</sup>For confidentiality reasons, we cannot reveal the names of related services.

However, due to the complexity and the fast-evolving nature of cloud systems [27], manually maintaining the dependency relations with intensity is very difficult. As a result, OCEs often struggle in diagnosing failures due to the lack of intensities. In order to relieve the pressure on OCEs, we propose to predict the intensity of dependency from the statuses of services.

#### IV. APPROACH

In this section, we present AID, a framework for predicting the Aggregated Intensity of service Dependency in large-scale cloud systems. We first present the overall workflow of AID. Then we elaborate on each step in detail, i.e., candidate selection, service status generation, and intensity prediction.

##### A. Overview

The overall workflow of AID is illustrated in Figure 4. AID consists of three steps: candidate selection, status generation, and intensity prediction. Given the raw traces, AID first generates a set of candidate service pairs  $(P, C)$  where service  $P$  directly invokes service  $C$  (Section IV-B). The intuition is that direct service invocation incurs direct dependency to some degree. Indirect dependencies through the transitivity of service invocation will be discussed in Section VII-A. For status generation, we generate the status of all services (Section IV-C). The status of one service is composed of three aspects of dependency, i.e., number of invocations, duration of invocations, error of invocations. Each aspect of the service's status contains one or more Key Performance Indicators (KPIs), depending on the actual implementation of the distributed tracing system. A KPI is an aggregated value of a service status of all the spans of a service in a fixed time interval, e.g., 1 minute. We use the statistical indicators of each aggregation as the values of the KPIs. Motivated by the experience of engineers introduced in Section III-B, we propose to predict the intensity of service dependencies from the similarity of the statuses of dependent services. The intuition behind using the similarity of time series is to evaluate the propagation of service statuses. The intensity prediction step (Section IV-D) predicts the intensity of dependency by measuring the similarity between two service's statuses. The similarity between two service's statuses is a normalized and weighted average of the similarity of all the KPIs of the two services. We calculate the similarity between two KPIs by a dynamic status warping algorithm. Finally, AID produces the dependency graph with intensity.

##### B. Candidate Selection

In general, direct service invocations can be divided into two categories, i.e., synchronous invocations and asynchronous invocations. Modern tracing mechanisms can keep track of both synchronous and asynchronous invocations [28]. Given all the raw traces of the cloud system, in this step, we generate a candidate dependency set  $Cand$ . The candidate dependency set  $Cand$  contains service invocation pairs  $(P_1, C_1), (P_2, C_2), \dots, (P_n, C_n)$ . Each pair  $(P_i, C_i)$  in the candidate dependency set denotes that the service named

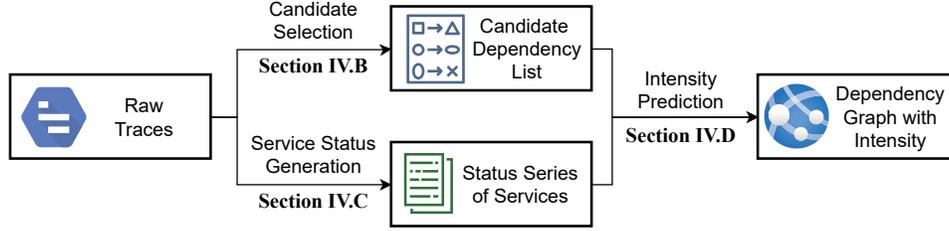


Fig. 4. The overall workflow of AID.

$P_i$  invokes the service named  $C_i$  at least once. Therefore, service  $P_i$  depends on service  $C_i$ . This step is to shrink the search space of possible dependent pairs because the service invocations indicate direct dependencies.

To generate the candidate dependency set, we need to know the name of the caller service and the callee service. The name of callee service is clearly recorded in the span, but the name of the caller service is not. Hence, we first augment each span  $s$  by adding another attribute  $s^{pname}$  which denotes the service name of the parent span. Specifically, the augmentation of attribute  $s^{pname}$  is achieved by 1) looking for another span  $s'$  whose  $id$  is the same as  $s^{pid}$ , and 2) set the  $name$  of  $s'$  as  $s^{pname}$ . Then we iterate over all the spans and add  $(s^{pname}, s^{name})$  to the candidate dependency set by the union operation.

For example, assuming the name of services are the same as the index of spans, the six spans in Figure 1 will result in a candidate set of  $\{(Service_0, Service_1), (Service_1, Service_4), (Service_0, Service_2), (Service_0, Service_3), (Service_3, Service_5)\}$ .

### C. Service Status Generation

In this step, we generate the status of all cloud services from the traces. We start by defining the status of a cloud service (i.e., service status) and then describe the procedure of service status generation.

*Definition of Service Status:* A service invocation is composed of three logical components, i.e., the caller service, the callee service, and the network communication. In particular, the caller service initiates an invocation to the callee service via the network. The callee service then processes the invocation, during which it may invoke other services. After the processing is finished, the callee service will send the result, e.g., status, to the caller service via the network. Hence, we could derive three aspects of service invocations: *initiation of invocation, processing, result*. As service invocations occur repeatedly, the three aspects of service invocations can derive three aspects of service dependency:

- Number of Invocations: The number of invocations from the caller to the callee,
- Duration of Invocations: The duration of invocations,
- Error of Invocations: The number of successful invocations from the caller to the callee.

*Representation of Service Status:* In a cloud system, the spans record information about every invocation. Intuitively, the status of a cloud service can be easily obtained from the spans of that service. Inspired by the common practice in cloud monitoring [29], we distribute the spans of one service into many bins according to the spans' timestamps. Each bin accepts spans whose timestamp is in a short, fixed-length period. We denote the length of the short period as  $\tau$ . For example, the span shown in Figure 2 will be put in the bin of `ts-preserve-service` at time 04:58, April 17 2020. We can then represent the status of a cloud service in a short period by the statistical indicators of all the spans in the corresponding bin.

Formally, given all the spans in the cloud system over a long period  $T$ , we first initiate  $S \times N$  empty bins of the predefined size  $\tau$ .  $S$  is the number of microservices.  $N$ , determined by  $\frac{T}{\tau}$ , is the number of bins. Then we distribute all spans into different bins according to their timestamp  $s^{ts}$  and service name  $s^{name}$ . After that, we calculate the following three types of indicators as the KPIs for each bin.

- $invo_t^M$ : Total number of invocations (spans) in the bin;
- $err_t^M$ : Error rate of the bin, i.e., the number of errors divided by the number of invocations;
- $dur_t^M$ : Averaged duration of all spans in the bin;

where  $t$  is the time of the bin and  $M$  is the microservice name of the bin. If a service is not invoked in a particular bin (i.e., the corresponding bin is empty), all the KPIs will be zero. In the end, we get the KPIs of every service  $M$  at every period  $t$ . Ordering the bins by  $t$ , we get three time series of KPIs for each cloud service, denoted as  $invo^M$ ,  $err^M$ , and  $dur^M$ . We name the time series of server KPIs as status series.

### D. Intensity Prediction

In this paper, we define the intensity of dependency between two services as *how much the status of the callee service influences the status of the caller service*. The step of intensity prediction quantitatively predicts the intensity of dependency by measuring the similarity between two services' status series. Specifically, we calculate the similarity of two different status series with dynamic status warping and aggregate all the similarities to get the overall similarity.

1) *Dynamic Status Warping:* Inspired by the dynamic time warping algorithm (DTW) [30], we propose the dynamic status warping (DSW) algorithm (Algorithm 1) to calculate the

---

**Algorithm 1:** Dynamic Status Warping

---

**Input:** The status series of caller service and callee service  $status^P, status^C$ ; duration series of callee  $dur^C$ , estimated round trip time  $\delta_{rtt}$ , max time drift  $\delta_d$

**Output:** The similarity between two status series

```
1 Set the warping window  $w = \max(dur^C) + \delta_{rtt}$ 
2  $K = \text{length}(status^C)$ 
3  $N = \text{length}(status^P)$ 
4 Initialize the cost matrix  $\mathbf{C} \in \mathbb{R}^{K \times N}$ , set the initial
  values as  $+\infty$ 
5  $\mathbf{C}_{1,1} = (status_1^P - status_1^C)^2$ 
6 for  $i = 2 \dots \min(\delta_d, K)$  do // Initialize the first
  column
7    $\mathbf{C}_{i,1} = \mathbf{C}_{i-1,1} + (status_1^P - status_i^C)^2$ 
8 end
9 for  $j = 2 \dots \min(w + \delta_d, N)$  do // Initialize the first
  row
10   $\mathbf{C}_{1,j} = \mathbf{C}_{1,j-1} + (status_j^P - status_1^C)^2$ 
11 end
12 for  $i = 2 \dots K$  do
13   for  $j = \max(2, i - \delta_d) \dots \min(N, i + w + \delta_d)$  do
14      $\mathbf{C}_{i,j} = \min(\mathbf{C}_{i-1,j-1}, \mathbf{C}_{i-1,j}, \mathbf{C}_{i,j-1}) +$ 
       $(status_j^P - status_i^C)^2$ 
15   end
16 end
17 return  $\mathbf{C}_{K,N}$ 
```

---

distance between two status series. DSW automatically warps the time in chronological order to make the two status series as similar as possible and get the similarity by summing the cost of warping. It utilizes dynamic programming to calculate an optimal matching between two status series. Given two services  $P, C$ , and their status series  $invo^P, invo^C, err^P, err^C, dur^P$ , and  $dur^C$ , the warping from the callee  $C$  to the caller  $P$  is specially designed for the cloud environment. The design considerations include:

*Directed warping:* Due to the latency of the network and the time of processing, it takes some time for the status of the callee service to affect the status of the caller service. Therefore, different from dynamic time warping, the time warping of DSW is directed, meaning that the matching from the callee to the caller can only happen in chronological order.

*Adaptive propagation window:* In cloud systems, after the round trip time ( $\delta_{rtt}$ ) plus the duration of request processing, the caller can receive the result of an invocation. Thus, the size of the directed warping window  $w$  is automatically set as the maximum duration of the callee's spans plus  $\delta_{rtt}$ .

*Time drift:* The machine time may drift due to issues with time synchronization in cloud systems, so we add an undirected time drift  $\delta_d$  to the warping window.

In summary,  $status_i^C$  can only be matched with one of  $[status_{i-\delta_d}^P, status_{i+w+\delta_d}^P]$ . The DSW returns the warping cost  $\mathbf{C}_{M,N}$  as the measure of similarity.

TABLE II  
DATASET STATISTICS.

Dataset	TT	Industry <sup>9</sup>
# Microservices	25	192
# Spans	17,471,024	About 1.0e10
# Strong	18	67
# Weak	1	8

2) *Similarity Aggregation:* For all  $(P_i, C_i) \in Cand$ , we calculate similarities between their status series, denoted as  $d_{invo}^{(P_i, C_i)}$ ,  $d_{err}^{(P_i, C_i)}$ , and  $d_{dur}^{(P_i, C_i)}$ . We normalize the similarity across the whole candidate set with a min-max normalization with Equation 1, where  $status \in \{invo, err, dur\}$ .

$$d_{status}^{(P_i, C_i)} = \frac{d_{status}^{(P_i, C_i)} - \min(d_{status}^{(P, C)})}{\max(d_{status}^{(P, C)}) - \min(d_{status}^{(P, C)})} \quad (1)$$

The intensity of dependency between  $P_i$  and  $C_i$  is the average similarity of all three similarities between their status series.

$$I^{(P_i, C_i)} = \frac{1}{3} \sum_{status \in S} d_{status}^{(P_i, C_i)}, S = \{invo, err, dur\} \quad (2)$$

Finally, we can build the dependency graph with intensity from the candidate set and the corresponding intensity values.

## V. EXPERIMENTS

In this section, we evaluate AID on both a simulated dataset and an industrial dataset. Particularly, we aim to answer the following research questions (RQs):

- **RQ1.** How effective is AID in predicting the intensity of dependency?
- **RQ2.** What is the impact of different parameter settings?
- **RQ3.** What is the impact of different similarity measures?
- **RQ4.** How efficient is AID?

### A. Experimental Setup

1) *Dataset:* To show the practical effectiveness of AID, we further conduct experiments on the simulated dataset and an industrial dataset from the cloud system of Huawei Cloud. Since there are no existing datasets of trace logs, we deploy a benchmark microservice system to simulate a real cloud system. We simulate user requests and collect the generated trace logs to construct the simulated dataset. We release both datasets with the paper to facilitate future studies in this field<sup>8</sup>.

**Simulated dataset:** For the simulated dataset, we deploy train-ticket [18], an open-source microservice benchmark, for data collection. Train-ticket is a web-based ticketing system with 25 microservices, through which users can search for

<sup>8</sup><https://github.com/OpsPAI/aid>

<sup>9</sup>We only labeled 75 dependencies that the engineers are familiar with.

tickets, reserve tickets, and pay for the reserved tickets. An open-source tracing framework, Jaeger, is used to trace all the API calls. To generate traces, we develop a request simulator that simulates normal users’ access to the ticketing system. The simulator will log in to the system, search for tickets, reserve a ticket according to the results of the search, and pay for the ticket. Then we collect the traces from Jaeger and transform the traces into 17,471,024 spans. The dataset is termed as “TT” in Table II.

**Industrial dataset:** Apart from the simulated dataset, we also collected traces from a region of Huawei Cloud to evaluate AID. To support tens of millions of users worldwide, the cloud system of Huawei Cloud contains numerous cloud services and microservices. The service invocations are monitored and recorded by an independently developed distributed tracing system. The complex dependency relations in the cloud system increase the burden of OCEs. The OCEs can diagnose problematic microservices timely if the intensity of dependencies can be automatically detected in real-time. To evaluate the practical effectiveness of our method, we collected a 7-day-long trace dataset with 192 microservices in April 2021. The dataset is termed as “Industry” in Table II.

**Manual labeling:** Since our method is unsupervised, labels are only for evaluation. Neither of the datasets has labels about the intensity of dependency, so manual labeling is needed. We set two candidate labels for the intensity of dependency, i.e., “strong” and “weak”. Given a candidate dependency pair  $(P, C)$ , if the failure of service  $C$  will cause the failure of service  $P$ , the intensity between  $(P, C)$  should be labeled “strong”; otherwise it should be labeled “weak”. For the simulated dataset, two Ph.D. students inspect the source code of all microservices and label every service dependency independently. For the industrial dataset, several senior engineers are invited to manually label the intensity of dependency. In both processes, disagreement on labels will be discussed until consensus is reached. Finally, we convert the “strong” labels to 1 and the “weak” labels to 0 so that they can be effectively compared with the computed intensities.

The statistics of the datasets are listed in Table II. “# Microservices” denotes the number of microservices in the dataset. “# Spans” denotes the number of spans in the dataset. “# Strong” and “# Weak” denote the number of dependencies that are labeled with “strong” or “weak” respectively.

2) *Baselines:* Since there is no existing work that measures the intensity of service dependency, we use Pearson correlation coefficient, Spearman correlation coefficient, and Kendall Rank correlation coefficient as the baseline. Particularly, we calculate correlation on the status series of a candidate dependency pair  $(P, C)$ , denoted as  $corr_{status}^{(P,C)}$  and  $corr_{s_{status}}^{(P,C)}$ . For the baselines, we directly use the implementation from the Python package `scipy`. We map the correlation to  $[0, 1]$  with the function  $f(x) = (x+1)/2$ . The intensities of dependencies are then produced in the same way as Equation 2.

3) *Evaluation Metrics:* We employ Cross Entropy (CE), Mean Absolute Error (MAE), and Root Mean Squared Error

TABLE III  
PERFORMANCE COMPARISON OF DIFFERENT METHODS ON TWO DATASETS

Dataset	Method	Metric		
		CE	MAE	RMSE
TT	Pearson	0.6872	<b>0.3305</b>	0.4388
	Spearman	0.7512	0.3735	0.4697
	Kendall	0.6464	0.3749	0.4577
	AID	<b>0.4562</b>	0.3435	<b>0.3859</b>
Industry	Pearson	0.6076	0.4524	0.4563
	Spearman	0.6030	0.4501	0.4537
	Kendall	0.6258	0.4636	0.4656
	AID	<b>0.3270</b>	<b>0.1751</b>	<b>0.3044</b>

(RMSE), as calculated in Equation 3, to evaluate the effectiveness of AID in predicting the intensity of dependency.

$$CE = \frac{1}{N} \sum_{i=1}^N -[y_i \cdot \log(p_i) + (1 - y_i) \cdot \log(1 - p_i)]$$

$$MAE = \frac{\sum_{i=1}^N |y_i - p_i|}{n} \quad (3)$$

$$RMSE = \sqrt{\frac{\sum_{i=1}^N (y_i - p_i)^2}{N}}$$

Specifically, cross entropy calculates the difference between the probability distributions of the label and the prediction. Mean absolute error and root mean squared error measures the absolute and squared error. Lower CE, MAE, and RMSE values indicate a better prediction.

4) *Experimental Environments:* We run the experiments on the simulated dataset on a Linux server with Intel Xeon E5-2670 CPU @ 2.40GHZ and 128 GB RAM. The experiments on the industrial dataset run on a Laptop with Intel Core i7 CPU @ 2.60 GHz and 16 GB RAM.

*B. RQ1: How effective is AID in predicting the intensity of dependency?*

To study the effectiveness of AID, we compare its performance with the baseline models on both the simulated dataset and the industrial dataset collected from Huawei Cloud. For the parameters of AID, we set the bin size  $\tau = 1 \text{ minute}$ , the estimated round trip time  $\delta_{rtt} = 0$ . Specially, we set the max time drift  $\delta_d = 1 \text{ minute}$  for the industrial dataset and set  $\delta_d = 0$  for the simulated dataset. We do this because the simulated dataset is deployed in a single server, so the time drift will not be a problem. In addition, we use moving average to smoothen the status series for the baselines and our method. The outputs are scalar values ranging from 0 to 1. A larger value indicates higher intensity. The overall performance is shown in Table III, where we mark the smallest loss for each loss metric and dataset.

AID achieves the best performance on the industrial dataset and reduces the loss by 45.8%, 61.1%, and 33.2% in terms of cross entropy, mean absolute error, and root mean squared error. On the simulated dataset, AID achieves the best performance in terms of cross entropy and root mean squared error. Pearson correlation coefficient marginally outperforms AID on the simulated dataset. The improvement of AID on the simulated dataset is smaller than that on the industrial dataset. This is because the benchmark for simulation did incorporate very few fault tolerance mechanisms, making most of the dependencies strong. Moreover, since the service invocations of the TT benchmark are very fast, the statuses of TT's services are relatively similar, making simple baselines and our approach perform similarly.

C. RQ2: What is the impact of different parameter settings?

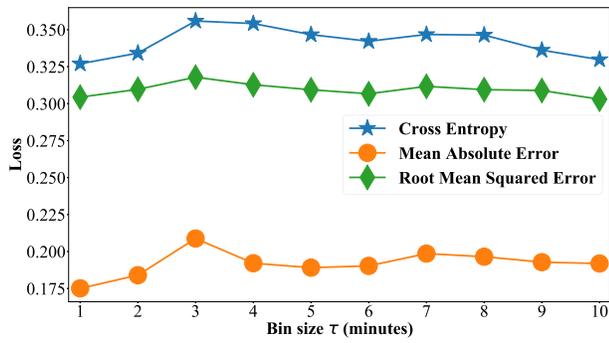


Fig. 5. Prediction loss under different bin size  $\tau$ .

Since the estimated round trip time  $\delta_{rtt}$  and the max time drift  $\delta_d$  are minuscule, we only study the impact of the bin size  $\tau$ . As the range of time of the simulated dataset is small, we only study the impact of the bin size  $\tau$  in the industrial dataset. In particular, we conduct experiments on with the bin size  $\tau \in [1, 10](minutes)$ , and keep  $\delta_{rtt} = 0$  and  $\delta_d = 1 minute$ . We did not set larger bin sizes because larger bin sizes result in more coarse-grained sampling of the service status, which will add difficulty to the similarity calculation in the subsequent DSW algorithm.

Figure 5 shows the prediction loss under different bin size  $\tau$ . The x-axis denotes the bin size and the y-axis shows the three loss metrics. The results indicate that the impact of different bin sizes in a reasonable range is small, but  $\tau = 1 minute$  gives the best performance on the industrial dataset.

D. RQ3: What is the impact of different similarity measures?

We further study the impact of different similarity measures on both datasets.  $AID_{DSW}$  denotes AID that uses the proposed DSW to measure the similarity between status series.  $AID_{DTW}$  denotes AID that uses the DTW [30] to measure the similarity. We keep the bin size  $\tau = 1 minute$  and the estimated round trip time  $\delta_{rtt} = 0$  as usual. Similar to previous experiments, we set the max time drift  $\delta_d = 1 minute$  for the industrial dataset and set  $\delta_d = 0$  for the simulated dataset.

TABLE IV  
THE IMPACT OF DIFFERENT SIMILARITY MEASURES

Dataset / Bin size	Method	Metric		
		CE	MAE	RMSE
TT / 1min	$AID_{DSW}$	0.4562	0.3435	0.3859
	$AID_{DTW}$	0.4494	0.3467	0.3832
Industry / 1min	$AID_{DSW}$	0.3270	0.1751	0.3044
	$AID_{DTW}$	0.3584	0.1996	0.3169

Table IV shows the performance of  $AID_{DSW}$  and  $AID_{DTW}$  on both datasets. On the industrial dataset, the proposed DSW algorithm improves the performance, but on the simulated dataset, the performance is almost the same. This is probably because the duration of spans in the simulated dataset is too small so that the effect of directed warping is weak. The results imply that the proposed DSW algorithm works better in real-world cloud environments.

E. RQ4: How efficient is AID?

The most time-consuming operations are the candidate selection and service status generation steps because we have to iterate over all the spans in the cloud system. Theoretically, the time complexities of the candidate selection and service status generation steps are  $O(S)$ , where  $S$  is the number of spans to process in the cloud system. In practice, the industrial dataset contains about  $1.0 \times 10^{10}$  spans, so we process it with a distributed computing service in Huawei Cloud. Since the preprocessing is dynamically scheduled and mixed with other teams' tasks, we do not count the time spent on it. For the intensity prediction step, the time complexity is  $O(kN^2)$ , where  $N = \frac{T}{\tau}$  is the number of bins and  $k$  is proportional to the warping window  $w$ . In practice, the intensity prediction step takes 155 seconds on average to process two status series both with 1440 bins on a laptop. Since the similarity calculation of different  $(P, C)$  pairs are independent, we could easily parallelize the intensity prediction step to further improve the time efficiency.

## VI. CASE STUDY

In Huawei Cloud, AID has been successfully incorporated into the dependency management system that serves hundreds to thousands of cloud services. Figure 6 illustrates the conceptual workflow. AID processes trace logs and continuously updates the aggregated intensity in the dependency management system. The reliability engineers will categorize the intensity into different levels by referring to both the output of AID and their domain expertise. Then the dependency management system will provide reference to the engineers in optimizing dependencies and mitigating cascading failures.

A. Optimization of Dependencies

In a cloud system, service failures are inevitable, but we can prevent the failures from affecting other services by optimizing

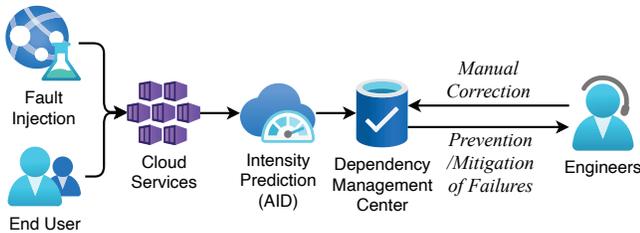


Fig. 6. The use case of AID.

improper dependencies. AID assists in the discovery of unnecessary strong dependency on critical cloud services. If a critical cloud service depends on another service with high intensity, the dependency management system will remind the engineers to check whether the dependency is necessary. If the dependency is unnecessary, the development team has to reduce the intensity of the dependency to improve the robustness of the critical cloud service. Since AID’s deployment, more than ten unnecessary dependencies of critical cloud services have been discovered by AID and optimized by the development team.

### B. Mitigation of Cascading Failures

AID also assists in the mitigation of cascading failures. During a cascading failure, AID can provide the latest intensity of dependency to OCEs, so that they can diagnose service failures efficiently. In addition, when a cascading failure occurs, OCEs can limit the traffic to critical cloud services and recover the dependencies marked as “strong” first. By doing so, the service disruption can get under control. Once a critical failure occurs, the manually confirmed “strong” dependencies will be treated with high priority. We conduct field interviews with OCEs to collect feedback. Based on the feedback, we have seen our method shedding light on reducing the impact of critical failures.

## VII. DISCUSSION

### A. Practical Usage and Perceived Limitations

1) *Indirect Dependencies*: In this work, we mainly considered direct dependencies, which is caused by direct service invocations. The proposed approach does not explicitly consider indirect dependencies through transitivity of service invocation because the intensity of indirect dependencies can be easily inferred from direct dependencies. In practice, the intensity of indirect dependencies can be inferred by a “cascading conduction mechanism” that if A intensively depends on B and B intensively depends on C then A intensively depends on C. The proposed approach also works well on dependencies caused by circuit breakers as long as the circuit breakers work transparently.

2) *Extension of Service Status*: In this paper, we only derive three aspects of service invocations, i.e., *number of invocations*, *duration of invocations*, *error of invocations*. We utilized them because they are part of the state-of-the-art tracing system. Other aspects like the content of invocation responses can also be important to determine status. In practice, cloud

providers can incorporate additional information to extend the representation of service status in their own implementation of AID.

3) *Limitations on Asynchronous Invocations*: Although modern tracing mechanisms can keep track of asynchronous invocations, AID may suffer from inaccuracies when dealing with asynchronous invocations. This is because the max time drift  $\delta_d$  in Algorithm 1 is hard to estimate for asynchronous service invocations. Furthermore, if the traces of synchronous and asynchronous invocations are mixed, AID may not work well since the time drift of synchronous and asynchronous invocations usually differs a lot. We leave this problem as future work.

### B. Threat to Validity

In this work, we identified the following major threats to validity.

1) *Labeling accuracy*: In this paper, we propose to measure the intensity of service dependency with AID. To evaluate the practical usage of AID, we conduct experiments on a simulated dataset and an industrial dataset. As it is a new relation between cloud services, manual labeling is needed for the evaluation. The evaluation on the industrial dataset requires engineers to manually inspect the dependencies and label the intensity of dependencies. Limited by the experience of engineers, the label may not be 100% accurate. The fast evolution of cloud services may also change their fault tolerance mechanism, resulting in inaccurate labels. However, the engineers we invited have rich domain knowledge and are in charge of the architecture design of the cloud system of Huawei Cloud. They also discuss with each other when there are disagreements. Moreover, the labeled dependencies are the core cloud services in Huawei Cloud, so the intensity of dependencies are stable during the data collection period. We believe the amount of inaccurate labels is small (if exists). Most importantly, our method is unsupervised, so inaccurate labels will not affect the prediction results of the proposed method.

2) *Insufficiency of the simulation*: For the evaluation purpose, we deploy an open-source microservice benchmark to simulate a real cloud system. The benchmark only contains 25 microservices, which is far below the number of cloud microservices in a real cloud. Additionally, the implementation of the open-source benchmark did not fully consider the fault tolerance, resulting in only one weak dependency in the simulation. Hence, the simulated dataset may not exhibit some common attributes of a real cloud system. For example, the proportion of “strong” dependency in the simulated dataset is twice the proportion of “strong” dependency in the industrial dataset. However, the insufficiency of the simulation will not hinder the practical usefulness of AID in the real cloud system. On the contrary, as we show in Section V, the proposed method works better on the industrial dataset. The experimental results on the simulated dataset only confirm the insufficiency of the simulation.

## VIII. RELATED WORK

### A. Cloud Monitoring

Monitoring cloud services properly with low overhead is the key to provide reliable services. Distributed Tracing, as a means of monitoring distributed cloud services, has been widely studied in the literature. All the distributed tracing approaches can be classified as intrusive tracing and non-intrusive tracing. Intrusive tracing requires modification to application code either in run time or at compile time. Google proposes Dapper [14] to help engineers understand system behavior and reasoning about performance issues. It reduces the tracing overhead by sampling and restricting the instrumentation number. X-Trace [16] monitors and reconstructs the whole request path from a client by modifying all the network protocols and embedding the tracing data to the package header.

Non-intrusive tracing approaches do not require code modification and usually have a lower overhead. Normally, these approaches leverage information like the system runtime logs and the source code to reconstruct the real event traces. Zhao et. al. [31] propose lprof to reconstruct the execution flow of distributed systems using the runtime log of these systems. lprof conducts static analysis on the call-graph of request processing code of the system to attribute a log output to a client request. Chow et. al. [32] also leverage system runtime logs to conduct performance monitoring and analysis. They propose *ÜberTrace* to reconstruct traces from the existing logs, then use *The Mystery Machine* to construct a causal model and conduct analyses. Stitch uses pattern matching on logs to reconstruct the hierarchical relationship of events in a system. Pensieve [33] automatically reconstructs a chain of causally dependent events that leads to a system failure exploiting the log files and system bytecode.

### B. Dependency Mining

Automatically discovering service dependencies is critical to cloud system administration and maintenance. There are two major types of dependency mining approaches, i.e., passive dependency mining and active ones. Passive dependency mining generates service dependency based purely on the runtime logs or KPIs. Shah et. al. [34] propose to use Recurrent Neural Networks (RNNs) to analyze and extract dependencies in KPIs and use the discovered dependencies to identify early indicators for a given performance metric, analyze lagged and temporal dependencies, and to improve forecast accuracy. EIDefrawy et. al. [35] use Transfer Entropy to passively mine the dependencies. Luo et. al. [25] apply log parsing and Bayesian decision theory to estimate the direction of dependencies among services. They employ time delay consistency to reduce false dependencies. Zand et. al. [36] construct a service correlation graph based on network measures and extract dependencies using hypothesis-testing. They further compute an importance metric for network's components to facilitate administration. CloudScout [3] employs Pearson Product-moment Correction Coefficient over machine-level

KPIs such as TCP/UDP connection numbers and CPU utilization to calculate the similarity between different services. The similarity measure is used to cluster different services together and to conduct VM consolidation based on the service clusters. Unlike all these approaches that mostly use physical machine metrics to infer service dependencies, our method is designed for the emerging microservice architecture and utilizes the trace logs that directly record service invocations.

Active dependency mining requires modification to services. Ma et. al. propose GMAT [2], which generates service dependencies in the microservice architecture leveraging the reflection feature of Java and visualizes the dependencies to engineers. Rippler [26] extracts the dependencies by randomly injecting temporal perturbation patterns in request arrival timings for different services and investigates the propagation of the patterns. Wang et. al. [37] constructs a service knowledge graph using real-time measures, operational metadata, and business features. They propose new metrics to measure the popularity of services based on their dependencies. Novotny et. al. [38] focus on mining dependencies on the highly dynamic mobile networks. They use local monitors to collect local views of dependencies and generate a global view of dependency on demand.

## IX. CONCLUSION

In this paper, we first conduct a comprehensive empirical study on the maintenance of AWS and Huawei Cloud. We identify the inefficiency in failure diagnosis and recovery with the binary-valued dependencies and define the intensity of dependency for the first time. To facilitate cloud maintenance, we propose AID, the first approach to predict the intensity of dependencies between cloud microservices. AID first generates a set of candidate dependency pairs from the spans. AID then represents the status of each cloud service with a multivariate time series aggregated from the spans and calculates the similarities between the statuses of the caller and callee of each candidate pair. Finally, AID aggregate the similarities to produce a unified value as the intensity of the dependency. For the evaluation, we collect and manually label a new dataset from an open-source microservice benchmark and evaluate AID on it. Furthermore, we evaluate AID using the data of Huawei Cloud and showcase the practical usage of AID. Both the evaluation results and case studies show the efficiency and effectiveness of AID. In the future, we plan to incorporate more information from the traces and other service logs for more accurate predictions.

## ACKNOWLEDGMENT

The work was supported by Key-Area Research and Development Program of Guangdong Province (No. 2020B010165002) and the Research Grants Council of the Hong Kong Special Administrative Region, China (CUHK 14210920).

## REFERENCES

- [1] Microsoft, "Microservices architecture style," 2019. [Online]. Available: <https://docs.microsoft.com/en-us/azure/architecture/guide/architecture-styles/microservices>
- [2] S.-P. Ma, C.-Y. Fan, Y. Chuang, W.-T. Lee, S.-J. Lee, and N.-L. Hsueh, "Using service dependency graph to analyze and test microservices," in *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, vol. 02, 2018, pp. 81–86.
- [3] J. Yin, X. Zhao, Y. Tang, C. Zhi, Z. Chen, and Z. Wu, "Cloudscout: A non-intrusive approach to service dependency discovery," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 5, pp. 1271–1284, 2016.
- [4] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "Above the clouds: A Berkeley view of cloud computing," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-28, Feb 2009.
- [5] R. DeFauw, A. Chigani, and N. Harris, "It resilience within aws cloud, part ii: Architecture and patterns," 2021. [Online]. Available: <https://aws.amazon.com/blogs/architecture/it-resilience-within-aws-cloud-part-ii-architecture-and-patterns/>
- [6] D. L. Oppenheimer and D. A. Patterson, "Architecture and dependability of large-scale internet services," *IEEE Internet Comput.*, vol. 6, no. 5, pp. 41–49, 2002.
- [7] M. Villamizar, O. Garcés, H. Castro, M. Verano, L. Salamanca, R. Casallas, and S. Gil, "Evaluating the monolithic and the microservice architecture pattern to deploy web applications in the cloud," in *2015 10th Computing Colombian Conference (10CCC)*. IEEE, 2015, pp. 583–590.
- [8] A. Balalaie, A. Heydarnoori, and P. Jamshidi, "Microservices architecture enables devops: Migration to a cloud-native architecture," *IEEE Softw.*, vol. 33, no. 3, pp. 42–52, 2016.
- [9] Y. Wang, G. Li, Z. Wang, Y. Kang, Y. Zhou, H. Zhang, F. Gao, J. Sun, L. Yang, P. Lee, Z. Xu, P. Zhao, B. Qiao, L. Li, X. Zhang, and Q. Lin, "Fast outage analysis of large-scale production clouds with service correlation mining," in *43rd IEEE/ACM International Conference on Software Engineering, ICSE 2021, Madrid, Spain, 22-30 May 2021*. IEEE, 2021, pp. 885–896.
- [10] Y. Gan, Y. Zhang, K. Hu, D. Cheng, Y. He, M. Pancholi, and C. Delimitrou, "Seer: Leveraging big data to navigate the complexity of performance debugging in cloud microservices," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2019, Providence, RI, USA, April 13-17, 2019*. ACM, 2019, pp. 19–33.
- [11] P. Wang, J. Xu, M. Ma, W. Lin, D. Pan, Y. Wang, and P. Chen, "Cloudranger: Root cause identification for cloud native systems," in *18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGRID 2018, Washington, DC, USA, May 1-4, 2018*. IEEE Computer Society, 2018, pp. 492–502.
- [12] P. Chen, Y. Qi, P. Zheng, and D. Hou, "Causeinfer: Automatic and distributed performance diagnosis with hierarchical causality graph in large distributed systems," in *2014 IEEE Conference on Computer Communications, INFOCOM 2014, Toronto, Canada, April 27 - May 2, 2014*. IEEE, 2014, pp. 1887–1895.
- [13] J. Chen, X. He, Q. Lin, Y. Xu, H. Zhang, D. Hao, F. Gao, Z. Xu, Y. Dang, and D. Zhang, "An empirical investigation of incident triage for online service systems," in *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice, ICSE (SEIP) 2019, Montreal, QC, Canada, May 25-31, 2019*. IEEE / ACM, 2019, pp. 111–120.
- [14] B. H. Sigelman, L. A. Barroso, M. Burrows, P. Stephenson, M. Plakal, D. Beaver, S. Jaspán, and C. Shanbhag, "Dapper, a large-scale distributed systems tracing infrastructure," Google, Inc., Tech. Rep., 2010. [Online]. Available: <https://research.google.com/archive/papers/dapper-2010-1.pdf>
- [15] P. Barham, R. Isaacs, R. Mortier, and D. Narayanan, "Magpie: Online modelling and performance-aware systems," in *9th Workshop on Hot Topics in Operating Systems (HotOS IX)*. USENIX Association, May 2003.
- [16] R. Fonseca, G. Porter, R. H. Katz, S. Shenker, and I. Stoica, "X-trace: A pervasive network tracing framework," in *4th Symposium on Networked Systems Design and Implementation (NSDI 2007), April 11-13, 2007, Cambridge, Massachusetts, USA, Proceedings*. USENIX, 2007.
- [17] OpenTracing, "The opentracing semantic specification," 2021. [Online]. Available: <https://opentracing.io/docs/overview/spans/>
- [18] X. Zhou, X. Peng, T. Xie, J. Sun, C. Ji, W. Li, and D. Ding, "Fault analysis and debugging of microservice systems: Industrial survey, benchmark system, and empirical study," *IEEE Trans. Software Eng.*, vol. 47, no. 2, pp. 243–260, 2021.
- [19] A. Fakhrazari and H. Vakilzadian, "A survey on time series data mining," in *IEEE International Conference on Electro Information Technology, EIT 2017, Lincoln, NE, USA, May 14-17, 2017*. IEEE, 2017, pp. 476–481.
- [20] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE transactions on acoustics, speech, and signal processing*, vol. 26, no. 1, pp. 43–49, 1978.
- [21] D. J. Berndt and J. Clifford, "Using dynamic time warping to find patterns in time series," in *Knowledge Discovery in Databases: Papers from the 1994 AAAI Workshop, Seattle, Washington, USA, July 1994. Technical Report WS-94-03*. AAAI Press, 1994, pp. 359–370.
- [22] V. Niennattrakul and C. A. Ratanamahatana, "On clustering multimedia time series data using k-means and dynamic time warping," in *2007 International Conference on Multimedia and Ubiquitous Engineering (MUE 2007), 26-28 April 2007, Seoul, Korea*. IEEE Computer Society, 2007, pp. 733–738.
- [23] A. Mueen, H. Hamooni, and T. Estrada, "Time series join on subsequence correlation," in *2014 IEEE International Conference on Data Mining, ICDM 2014, Shenzhen, China, December 14-17, 2014*. IEEE Computer Society, 2014, pp. 450–459.
- [24] Z. Chen, Y. Kang, L. Li, X. Zhang, H. Zhang, H. Xu, Y. Zhou, L. Yang, J. Sun, Z. Xu, Y. Dang, F. Gao, P. Zhao, B. Qiao, Q. Lin, D. Zhang, and M. R. Lyu, "Towards intelligent incident management: why we need it and how we make it," in *ESEC/FSE '20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020*. ACM, 2020, pp. 1487–1497.
- [25] J.-G. Lou, Q. Fu, Y. Wang, and J. Li, "Mining dependency in distributed systems through unstructured logs analysis," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 1, pp. 91–96, 2010.
- [26] A. Zand, G. Vigna, R. Kemmerer, and C. Kruegel, "Rippler: Delay injection for service dependency detection," in *IEEE INFOCOM 2014-IEEE Conference on Computer Communications*. IEEE, 2014, pp. 2157–2165.
- [27] A. B. M. B. Alam, A. Haque, and M. Zulkermine, "CREM: A cloud reliability evaluation model," in *IEEE Global Communications Conference, GLOBECOM 2018, Abu Dhabi, United Arab Emirates, December 9-13, 2018*. IEEE, 2018, pp. 1–6.
- [28] OpenTracing, "Opentracing spring cloud," 2021. [Online]. Available: <https://github.com/opentracing-contrib/java-spring-cloud>
- [29] G. Aceto, A. Botta, W. De Donato, and A. Pescapè, "Cloud monitoring: A survey," *Computer Networks*, vol. 57, no. 9, pp. 2093–2115, 2013.
- [30] E. J. Keogh, "Exact indexing of dynamic time warping," in *Proceedings of 28th International Conference on Very Large Data Bases, VLDB 2002, Hong Kong, August 20-23, 2002*. Morgan Kaufmann, 2002, pp. 406–417.
- [31] X. Zhao, Y. Zhang, D. Lion, M. F. Ullah, Y. Luo, D. Yuan, and M. Stumm, "lprof: A non-intrusive request flow profiler for distributed systems," in *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*. USENIX Association, Oct. 2014, pp. 629–644.
- [32] M. Chow, D. Meisner, J. Flinn, D. Peek, and T. F. Wenisch, "The mystery machine: End-to-end performance analysis of large-scale internet services," in *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*. USENIX Association, Oct. 2014, pp. 217–231.
- [33] Y. Zhang, S. Makarov, X. Ren, D. Lion, and D. Yuan, "Pensieve: Non-intrusive failure reproduction for distributed systems using the event chaining approach," in *Proceedings of the 26th Symposium on Operating Systems Principles*, 2017, pp. 19–33.
- [34] S. Y. Shah, Z. Yuan, S. Lu, and P. Zerfos, "Dependency analysis of cloud applications for performance monitoring using recurrent neural networks," in *2017 IEEE International Conference on Big Data (Big Data)*. IEEE, 2017, pp. 1534–1543.
- [35] K. EIDefrawy, T. Kim, and P. Sylla, "Automated inference of dependencies of network services and applications via transfer entropy," in *2016 IEEE 40th Annual Computer Software and Applications Conference (COMPSAC)*, vol. 2. IEEE, 2016, pp. 32–37.

- [36] A. Zand, A. Houmansadr, G. Vigna, R. Kemmerer, and C. Kruegel, "Know your achilles' heel: Automatic detection of network critical services," in *Proceedings of the 31st Annual Computer Security Applications Conference*, 2015, pp. 41–50.
- [37] H. Wang, C. Shah, P. Sathaye, A. Nahata, and S. Katariya, "Service application knowledge graph and dependency system," in *2019 34th IEEE/ACM International Conference on Automated Software Engineering Workshop (ASEW)*. IEEE, 2019, pp. 134–136.
- [38] P. Novotny, B. J. Ko, and A. L. Wolf, "On-demand discovery of software service dependencies in manets," *IEEE Transactions on Network and Service Management*, vol. 12, no. 2, pp. 278–292, 2015.