

# Sieve: Actionable Insights from Monitored Metrics in Microservices

Jörg Thalheim<sup>1</sup>, Antonio Rodrigues<sup>2</sup>, Istemi Ekin Akkus<sup>3</sup>, Pramod Bhatotia<sup>1</sup>,  
Ruichuan Chen<sup>3</sup>, Bimal Viswanath<sup>4</sup>, Lei Jiao<sup>5</sup>, Christof Fetzer<sup>6</sup>

<sup>1</sup>University of Edinburgh, <sup>2</sup>Carnegie Mellon Univ., <sup>3</sup>NOKIA Bell Labs, <sup>4</sup>University of Chicago, <sup>5</sup>University of Oregon, <sup>6</sup>TU Dresden

## Abstract

Major cloud computing operators provide powerful monitoring tools to understand the current (and prior) state of the distributed systems deployed in their infrastructure. While such tools provide a detailed monitoring mechanism at scale, they also pose a significant challenge for the application developers/operators to transform the *huge space of monitored metrics* into *useful insights*. These insights are essential to build effective management tools for improving the efficiency, resiliency, and dependability of distributed systems.

This paper reports on our experience with building and deploying SIEVE—a platform to derive actionable insights from monitored metrics in distributed systems. SIEVE builds on two core components: a metrics reduction framework, and a metrics dependency extractor. More specifically, SIEVE first reduces the dimensionality of metrics by automatically filtering out unimportant metrics by observing their signal over time. Afterwards, SIEVE infers metrics dependencies between distributed components of the system using a predictive-causality model by testing for Granger Causality.

We implemented SIEVE as a generic platform and deployed it for two microservices-based distributed systems: OpenStack and Share-Latex. Our experience shows that (1) SIEVE can reduce the number of metrics by at least an order of magnitude (10 – 100×), while preserving the statistical equivalence to the total number of monitored metrics; (2) SIEVE can dramatically improve existing monitoring infrastructures by reducing the associated overheads over the entire system stack (CPU—80%, storage—90%, and network—50%); (3) Lastly, SIEVE can be effective to support a wide-range of workflows in distributed systems—we showcase two such workflows: Orchestration of autoscaling, and Root Cause Analysis (RCA). This technical report is an extended version of our conference publication [87].

## 1 Introduction

Most distributed systems are constantly monitored to understand their current (and prior) states. The main purpose of monitoring is to gain *actionable insights* that would enable a developer/operator to take appropriate actions to better manage the deployed system. Such insights are commonly used to manage the health and resource requirements as well as to investigate and recover from failures (root cause identification). For these reasons, monitoring is a crucial part of any distributed system deployment.

All major cloud computing operators provide a monitoring infrastructure for application developers (e.g., Amazon CloudWatch [2], Azure Monitor [13], Google StackDriver [5]). These platforms provide infrastructure to monitor a large number (hundreds or thousands) of various application-specific and system-level metrics associated with a cloud application. Although such systems feature scalable measurement and storage frameworks to conduct monitoring at scale, they leave the task of transforming the monitored metrics into usable knowledge to the developers. Unfortunately,

this transformation becomes difficult with the increasing size and complexity of the application.

In this paper, we share our experience on: *How can we derive actionable insights from the monitored metrics in distributed systems?* In particular, given a large number of monitored metrics across different components (or processes) in a distributed system, we want to design a platform that can derive actionable insights from the monitored metrics. This platform could be used to support a wide-range of use cases to improve the efficiency, resiliency, and reliability of distributed systems.

In this work, we focus on microservices-based distributed systems because they have become the de-facto way to design and deploy modern day large-scale web applications [51]. The microservices architecture is an ideal candidate for our study for two reasons: First, microservices-based applications have a large number of distributed components (hundreds to thousands [48, 59]) with complex communication patterns, each component usually exporting several metrics for the purposes of debugging, performance diagnosis, and application management. Second, microservices-based applications are developed at a rapid pace: new features are being continuously integrated and deployed. Every new update may fix some existing issues, introduce new features, but can also introduce a new bug. With this rapid update schedule, keeping track of the changes in the application as a whole with effects propagating to other components becomes critical for reliability, efficiency, and management purposes.

The state-of-the-art management infrastructures either rely on ad hoc techniques or custom application-specific tools. For instance, prior work in this space has mostly focused on analyzing message-level traces (instead of monitored metrics) to generate a causal model of the application to debug performance issues [33, 45]. Alternatively, developers usually create and use custom tools to address the complexity of understanding the application as a whole. For example, Netflix developed several application-specific tools for such purposes [8, 48] by instrumenting the entire application. These approaches require either complicated instrumentation or sophisticated techniques to infer happens-before relationships (for the causal model) by analyzing message trace timestamps, making them inapplicable for broader use.

This paper presents our experience with designing and building SIEVE, a system that can utilize an existing monitoring infrastructure (i.e., without changing the monitored information) to infer actionable insights for application management. SIEVE takes a data-driven approach to enable better management of microservices-based applications. At its core, SIEVE is composed of two key modules: (1) a metric reduction engine that reduces the dimensionality of the metric space by filtering out metrics that carry redundant information, (2) a metric dependency extractor that builds a causal model of the application by inferring causal relationships between metrics associated with different components.

Module (1) enables SIEVE to identify “relevant” metrics for a given application management task. For instance, it might be sufficient to monitor only a few metrics associated with error states of the application instead of the entire set when monitoring the health of the application. It is important to also note that reducing the metric space has implications for deployment costs: frameworks like Amazon CloudWatch use a per-metric charging model, and not identifying relevant metrics can significantly drive up the cost related to monitoring the application.

Module (2) is crucial for inferring actionable insights because it is able to automatically infer complex application dependencies. In a rapidly updating application, the ability to observe such complex dependencies and how they may change is important for keeping one’s understanding of the application as a whole up-to-date. Such up-to-date information can be helpful for developers to quickly react to any problem that may arise during deployment.

We implemented SIEVE as a generic platform, and deployed it with two microservices-based distributed systems: ShareLatex [23] and OpenStack [28]. Our experience shows that (1) SIEVE can reduce the number of monitored metrics by an order of magnitude (10 – 100×), while preserving the statistical equivalence to the total number of monitored metrics. In this way, the developers/operators can focus on the important metrics that actually matter. (2) SIEVE can dramatically improve the efficiency of existing metrics monitoring infrastructures by reducing the associated overheads over the entire system stack (CPU—80%, storage—90%, and network—50%). This is especially important for systems deployed in a cloud infrastructure, where the monitoring infrastructures (e.g. AWS CloudWatch) charge customers for monitoring resources. And finally, (3) SIEVE can be employed for supporting a wide-range of workflows. We showcase two such case-studies: In the first case study, we use ShareLatex [23] and show how SIEVE can help developers orchestrate autoscaling of microservices-based applications. In the second case study, we use OpenStack [28] and show how developers can take advantage of SIEVE’s ability to infer complex dependencies across various components in microservices for Root Cause Analysis (RCA). SIEVE’s source code with the full experimentation setup is publicly available: <https://sieve-microservices.github.io/>.

## 2 Overview

In this section, we first present some background on microservices-based applications and our motivation to focus on them. Afterwards, we present our goals, design overview, and its possible use cases.

### 2.1 Background and Motivation

Microservices-based applications consist of loosely-coupled distributed components (or processes) that communicate via well-defined interfaces. Designing and building applications in this way increases modularity, so that developers can work on different components and maintain them independently. These advantages make the microservices architecture the de facto design choice for large-scale web applications [51].

While increasing modularity, such an approach to developing software can also increase the application complexity: As the number of components increases, the interconnections between components also increases. Furthermore, each component usually exports several metrics for the purposes of debugging, performance diagnosis, and application management. Therefore, understanding the

**Table 1.** Number of metrics exposed by microservices.

Microservices	Number of metrics
Netflix [61]	~ 2,000,000
Quantcast [14]	~ 2,000,000
Uber [16]	~ 500,000,000
ShareLatex [23]	889
OpenStack [18, 20]	17,608

dependencies between the components and utilizing these dependencies with the exported metrics becomes a challenging task. As a result, understanding how the application performs as a whole becomes increasingly difficult.

Typical microservices-based applications are composed of hundreds of components [48, 59]. Table 1 shows real-world microservices-based applications that have tens of thousands of metrics and hundreds of components. We experimented with two such applications, ShareLatex [23] and OpenStack [19], each having several thousands of metrics and order of tens of components. The metrics in these applications come from all layers of the application like hardware counters, resource usage, business metrics or application-specific metrics.

To address this data overload issue, developers of microservices-based applications usually create ad hoc tools. For example, application programmers at Netflix developed several application-specific tools for such purposes [8, 48]. These tools, however, require the application under investigation to be instrumented, so that the communication pattern between components can be established by following requests coming into the application. This kind of instrumentation requires coordination among developers of different components, which can be a limiting factor for modularity.

Major cloud computing operators also provide monitoring tools for recording all metric data from all components. For example, Amazon CloudWatch [2], Azure Monitor [13], and Google StackDriver [5]. These monitoring tools aid in visualizing and processing metric data in real-time (i.e., for performance observation) or after an issue with the application (i.e., for debugging). These tools, however, either use a few system metrics that are hand-picked by developers based on experience, or simply record all metric data for all the components.

Relying on past experience may not always be effective due to the increasing complexity of a microservices-based application. On the other hand, recording all metric data can create significant monitoring overhead in the network and storage, or in the case of running the application in a cloud infrastructure (e.g., AWS), it can incur costs due to the provider charging the customers (e.g., CloudWatch). For these reasons, it is important to understand the dependencies between the components of a microservice-based application. Ideally, this process should not be intrusive to the application. Finally, it should help the developers to identify and minimize the critical components and metrics to monitor.

### 2.2 Design Goals

While designing SIEVE, we set the following goals.

- **Generic:** Many tools for distributed systems have specific goals, including performance debugging, root cause analysis and orchestration. Most of the time, these tools are custom-built for the application in consideration and target a certain




Figure 1. High level steps of SIEVE.

goal. Our goal is to design a generic platform that can be used for a wide-range of workflows.

- **Automatic:** The sheer number of metrics prohibits manual inspection. On the other hand, designing a generic system to help developers in many use cases might require manually adjusting some parameters for each use case. Our tool should be as automated as possible while reducing the number of metrics and extracting their relationships. However, we leave the utilization of our platform’s output to the developers, who may have different goals.
- **Efficient:** Our platform’s operation should be as efficient as possible. Minimizing analysis time becomes important when considering distributed systems, such as microservices-based applications.

**Assumptions.** While developing SIEVE, we made two assumptions.

- We assume that the developers can supply a workload generator for the application under investigation. This assumption should be reasonable and less demanding for developers than fully instrumenting each component and/or figuring out relationships across all metrics.
- It is possible for specific use cases to have additional assumptions. For example, a root cause analysis scenario would require knowing a faulty and non-faulty (correct) version of the application.

### 2.3 SIEVE Overview

The underlying intuition behind SIEVE is two-fold: Firstly, in the metric dimension, some metrics of a component may behave with similar patterns as other metrics of that component. Secondly, in the component dimension, there are dependencies between components. As a result, monitoring all metrics of all components at runtime may be unnecessary and inefficient (as components are not independent).

In this paper, we present SIEVE to reduce this complexity by systematically analyzing the application to filter collected metrics and to build a dependency graph across components. To showcase the generality of this dependency graph and its benefits, we then utilize SIEVE to orchestrate autoscaling of the ShareLatex [23] application—an online collaboration tool, and to perform Root Cause Analysis (RCA) in OpenStack [28]—a cloud management software (§4).

At a high level, SIEVE’s design follows three steps as shown in Figure 1.

**Step #1: Load the application.** SIEVE uses an application-specific load generator to stress the application under investigation. This load generator can be provided by the application developers. For example, OpenStack already uses a load generator named Rally [21]. During the load, SIEVE records the communications among components to obtain a *call graph*. This recording does not require

any modifications to the application code. In addition, SIEVE records all exposed *metrics* by all components. Note that this recording only happens during the creation of the call graph and not during runtime.

**Step #2: Reduce metrics.** After collecting the metrics, SIEVE analyzes each component and organizes its metrics into fewer groups via clustering, so that similar-behaving metrics are clustered together. After clustering, SIEVE picks a representative metric from each cluster. These representative metrics as well as their clusters in a sense characterize each component.

**Step #3: Identify dependencies.** In this step, SIEVE explores the possibilities of one component’s representative metrics affecting another component’s metrics using a pairwise comparison method: each representative metric of one component is compared with each representative metric of another component. SIEVE uses the call graph obtained in Step 1 to choose the components to be compared (i.e., components directly communicating) and the representative metrics determined in Step 2. As a result, the search space is significantly reduced compared to the naive approach of comparing all components with every other component using all metrics.

If SIEVE determines that there is a relationship between a metric of one component and another metric of another component, a dependency edge between these components is created using the corresponding metrics. The direction of the edge depends on which component is affecting the other.

### 2.4 Potential Use Cases

We envision SIEVE can be useful to the developers or operators of distributed systems to build a wide-range of management tools to improve the efficiency, reliability, and resiliency of distributed systems. Each of these cases might require some tweaks and specific knowledge about the application. Nevertheless, we think that the output of SIEVE can be a starting point. In Section 4, we showcase two of these use cases with two different applications.

**Orchestration and autoscaling of components.** The pairwise investigation of representative metrics of components produces the dependencies across components. By definition, the dependency graph shows the order of bottlenecked components. As a result of this graph, developers can have a better understanding of which components need to be scaled out first, meaning that the number of monitored components can be reduced. Furthermore, the dependencies show the metrics that are affected, meaning that one only needs to monitor a limited set of metrics rather than every metric exported by these components. In combination, these reductions are reflected in the monitoring overhead. Also, scaling decisions are based on fewer components and metrics, so that potentially conflicting scaling decisions can be avoided.

**Root cause analysis.** It is possible that updates can introduce bugs and problems into the application. Changing dependency graphs (i.e., after updates) may indicate that the problem got introduced during a particular update that caused the dependency graph to change. Identifying such changes will be useful in debugging the application by pointing to the root cause of the problem.

**Performance diagnosis.** Similar to the root cause analysis scenario, the resource-usage profile of a component may change after an update. Even if the interfaces between the components may stay the same, the update may address a previously unknown bug, introduce a new feature and/or implement the same functionality in a more efficient and/or a different way. Such changes can have effects on other parts of the application, and the complexity of the application may render these effects not easily foreseeable. The dependency graph can be utilized to understand the overall effect on the application the update can have.

**Placement decisions.** When scaling decisions are made, the resource-usage profiles of components become important, because components with similar resource-usage profiles may create contention (e.g., two components may be CPU-bound). As a result, placing a new instance of a component on a computer where other instances of the same component or instances of another component with a similar resource-usage profile run may not yield the additional benefit of the extra instance. The dependency graph lists which metrics are affected in a component under load. This information can be used to extract resource-usage profiles of components to avoid such contention during scale-up events.

### 3 Design

In this section, we detail the three steps of SIEVE.

#### 3.1 Load the Application

For our systematic analysis, we first run the application under various load conditions. This loading serves two purposes: First, the load exposes a number of metrics from the application as well as the infrastructure it runs on. These metrics are then used to identify potential relationships across components. Second, the load also enables us to obtain a call graph, so that we can identify the components that communicate with each other. The call graph is later used to reduce the amount of computation required to identify the inter-component relationships (§3.3). The load test is intended to be run in an offline step and not in production.

**Obtaining metrics.** During the load of the application, we record metrics as time series. There are two types of metrics that we can leverage for our analysis: First, there are system metrics that are obtained from the underlying operating system. These metrics report the resource usage of a microservice component, and are usually related to the hardware resources on a host. Examples include usages in CPU, memory, network and disk I/O.

Second, there are application-level metrics. Application developers often add application-specific metrics (e.g., number of active users, response time of a request in a component). Commonly-used components (e.g., databases, load balancers) and certain language runtimes (e.g., Java) may provide statistics about specific operations (e.g., query times, request counts, duration of garbage collection).

**Obtaining the call graph.** Generally speaking, applications using a microservices architecture communicate via well-defined

interfaces similar to remote procedure calls. We model these communications between the components as a directed graph, where the vertices represent the microservice components and the edges point from the caller to the callee providing the service.

By knowing which components communicate directly, we can reduce the number of component pairs we need to check to see whether they have a relation (see Section 3.3). Although it is possible to manually track this information for smaller-sized applications, this process becomes quickly difficult and error-prone with increasing number of components.

There are several ways to understand which microservice components are communicating with each other. One can instrument the application, so that each request can be traced from the point it enters the application to the point where the response is returned to the user. Dapper [84] from Google and Atlas [48, 61] from Netflix rely on instrumenting their RPC middleware to trace requests.

Another method to obtain communicating components is to monitor network traffic between hosts running those components using a tool like *tcpdump*. After obtaining the traffic, one can map the exchanged packets to the components via their source/destination addresses. This method can produce communicating component pairs by parsing all network packets, adding significant computational overhead and increasing the analysis time. Furthermore, it is possible that many microservice components are deployed onto the same host (e.g., using containers), making the packet parsing difficult due to network address translation on the host machine.

One can also observe system calls related to network operations via APIs such as *ptrace()* [31]. However, this approach adds a lot of context switches between the tracer and component under observation.

SIEVE employs *sysdig* to obtain the communicating pairs. *sysdig*[24] is a recent project providing a new method to observe system calls in a more efficient way. Utilizing a kernel module, *sysdig* provides system calls as an event stream to a user application. The event stream also contains information about the monitored processes, so that network calls can be mapped to microservice components, even if they are running in containers. Furthermore, it enables extraction of the communication peer via user-defined filters. Employing *sysdig*, we avoid the shortcomings of the above approaches: 1) We do not need to instrument the application, which makes our system more generally applicable, 2) We add little overhead to obtain the call graph of an application for our analysis (see Section 6.1.3).

#### 3.2 Reduce Metrics

The primary goal of exporting metrics is to understand the performance of applications, orchestrating them and debugging them. While the metrics exported by the application developers or commonly-used microservice components may be useful for these purposes, it is often the case that the developers have little idea regarding which ones are going to be most useful. Developers from different backgrounds may have different opinions: a developer specializing in network communications may deem network I/O as the most important metric to consider, whereas a developer with a background on algorithms may find CPU usage more valuable. As a result of these varying opinions, often times many metrics are exported.

While it may look like there is no harm in exporting as much information as possible about the application, it can create problems. Manually investigating the obtained metrics from a large number

of components becomes increasingly difficult with the increasing number of metrics and components [38]. This complexity reflects on the decisions that are needed to control and maintain the application. In addition, the overhead associated with the collection and storage of these metrics can quickly create problems. In fact, Amazon CloudWatch [2] charges its customers for the reporting of the metrics they export. As a result, the more metrics an application has to export, the bigger the cost the developers would have to bear.

One observation we make is that some metrics strongly correlate with each other and it might not be necessary to consider all of them when making decisions about the control of the application. For example, some application metrics might be strongly correlated with each other due to the redundancy in choosing which metrics to export by the developers. It is also possible that different subsystems in the same component report similar information (e.g., overall memory vs. heap usage of a process). In addition, some system metrics may offer clues regarding the application’s state: increased network I/O may indicate an increase in the number of requests.

The direct outcome of this observation is that it should be possible to reduce the dimensionality of the metrics the developers have to consider. As such, the procedure to enable this reduction should happen with minimal user effort and scale with increased numbers of metrics.

To achieve these requirements, SIEVE uses a clustering approach named *k*-Shape [76] with a pre-filtering step. While other approaches such as principal component analysis (PCA) [52] and random projections [75] can also be used for dimensionality reduction, these approaches either produce results that are not easily interpreted by developers (i.e., PCA) or sacrifice accuracy to achieve performance and have stability issues producing different results across runs (i.e., random projections). On the other hand, clustering results can be visually inspected by developers, who can also use any application-level knowledge to validate their correctness. Additionally, clustering can also uncover hidden relationships which might not have been obvious.

**Filtering unvarying metrics.** Before we use *k*-Shape, we first filter metrics with constant trend or low variance ( $var \leq 0.002$ ). These metrics cannot provide any new information regarding the relationships across components, because they are not changing according to the load applied to the application. Removing these metrics also enables us to improve the clustering results.

***k*-Shape clustering.** *k*-Shape is a recent clustering algorithm that scales linearly with the number of metrics. It uses a novel distance metric called *shape-based distance* (SBD). SBD is based on a normalized form of cross correlation (NCC) [76]. Cross correlation is calculated using Fast Fourier Transformation and normalized using the geometric mean of the autocorrelation of each individual metric’s time series. Given two time series vectors,  $\vec{x}$  and  $\vec{y}$ , SBD will take the position  $w$ , when sliding  $\vec{x}$  over  $\vec{y}$ , where the normalized cross correlation maximizes.

$$SBD(\vec{x}, \vec{y}) = 1 - \max_w (NCC_w(\vec{x}, \vec{y})) \quad (1)$$

Because *k*-Shape uses a distance metric based on the shape of the investigated time series, it can detect similarities in two time series, even if one lags the other in the time dimension. This feature is important to determine relationships across components in microservices-based applications because a change in one metric

in one component may not reflect on another component’s metrics immediately (e.g., due to the network delay of calls between components).

Additionally, *k*-Shape is robust against distortion in amplitude because data is normalized via z-normalization ( $z = \frac{x-\mu}{\sigma}$ ) before being processed. This feature is especially important because different metrics may have different units and thus, may not be directly comparable.

*k*-Shape works by initially assigning time series to clusters randomly. In every iteration, it computes new cluster centroids according to SBD with the assigned time series. These centroids are then used to update the assignment for the next iteration until the clusters converge (i.e., the assignments do not change).

We make three adjustments to employ *k*-Shape in SIEVE. First, we preprocess the collected time series to be compatible with *k*-Shape. *k*-Shape expects the observations to be equidistantly distributed in the time domain. However, during the load of the application, timeouts or lost packets can cause gaps between the measurements.

To reconstruct missing data, we use spline interpolation of the third order (cubic). A spline is defined piecewise by polynomial functions. Compared to other methods such as averages of previous values or linear interpolation, spline interpolation provides a higher degree of smoothness. It therefore introduces less distortion to the characteristics of a time-series [69]. Additionally, monitoring systems retrieve metrics at different points in time and need to be discretized to match each other. In order to increase the matching accuracy, we discretize using 500ms instead of the original 2s used in the original *k*-Shape paper [76].

Our second adjustment is to change the initial assignments of metric time series to clusters. To increase clustering performance and reduce the convergence overhead, we pre-cluster metrics according to their name similarity (e.g., Jaro distance [63]) and use these clusters as the initial assignment instead of the default random assignment. This adjustment is reasonable given that many developers use naming conventions when exporting metrics relating to the same component or resource in question (e.g., “cpu\_usage”, “cpu\_usage\_percentile”). The number of iterations to converge should decrease compared to the random assignment, because similar names indicate similar metrics. Note that this adjustment is only for performance reasons; the convergence of the *k*-Shape clustering does not require any knowledge of the variable names and would not be affected even with a random initial assignment.

During the clustering process, *k*-Shape requires the number of clusters to be previously determined. In an application with several components, each of which having various number of metrics, pre-determining the ideal number of clusters may not be straightforward. Our final adjustment is to overcome this limitation: we iteratively vary the number of clusters used by *k*-Shape and pick the number that gives the best *silhouette value* [81], which is a technique to determine the quality of the clusters. The silhouette value is  $-1$  when the assignment is wrong and  $1$  when it is a perfect assignment [32]. We use the SBD as a distance measure in the silhouette computation.

In practice, experimenting with a small number of clusters is sufficient. For our applications, seven clusters per component was sufficient, where each component had up to 300 metrics.

**Representative metrics.** After the clustering, each microservice component will have one or more clusters of metrics. The number of clusters will most likely be much smaller than the number of metrics belonging to that component. Once these clusters are obtained, SIEVE picks one representative metric from each cluster. To pick the representative metric from each cluster, SIEVE determines the SBD between each metric and the corresponding centroid of the cluster. The metric with the lowest metric is chosen as the representative metric for this cluster.

The high-level idea is that the behavior of the cluster will match this representative metric; otherwise, the rest of the metrics in the cluster would not have been in the same cluster as this metric. The set of representative metrics of a component can then be used to describe a microservice component’s behavior. These representative metrics are then used in conjunction with the call graph obtained in Section 3.1 to identify and understand the relationships across components.

### 3.3 Identify Dependencies

To better understand an application, we need to find dependencies across its components. A naïve way of accomplishing this goal would be to compare all components with each other using all possible metrics. One can clearly see that with the increasing number of components and metrics, this would not yield an effective solution.

In the previous section, we described how one can reduce the number of metrics one has to consider in this pairwise comparison by clustering and obtaining the representative metrics of each component. Still, comparing all pairs of components using this reduced set of metrics may be inefficient and redundant considering the number of components in a typical microservices-based application (e.g., tens or hundreds).

SIEVE uses the call graph obtained in Section 3.1 to reduce the number of components that need to be investigated in a pairwise fashion. For each component, we do pairwise comparisons using each representative metric of its clusters with each of its neighbouring components (i.e., callees) and their representative metrics.

SIEVE utilizes Granger Causality tests [57] in this pairwise comparison. Granger Causality tests are useful in determining whether a time series can be useful in predicting another time series: In a microservices-based application, the component interactions closely follow the path a request takes inside the application. As a result, these interactions can be predictive of the changes in the metrics of the components in the path. Granger Causality tests offer a statistical approach in understanding the relationships across these components. Informally, Granger Causality is defined as follows. If a metric  $X$  is Granger-causing another metric  $Y$ , then we can predict  $Y$  better by using the history of both  $X$  and  $Y$  compared to only using the history of  $Y$  [54].

To utilize Granger Causality tests in SIEVE, we built two linear models using the ordinary least-square method [35]. First, we compare each metric  $X_t$  with another metric  $Y_t$ . Second, we compare each metric  $X_t$  with the time-lagged version of the other metric  $Y_t$ :  $Y_{t-Lag}$ . Covering the cases with a time lag is important because the load in one component may not be reflected on another component until the second component receives API calls and starts processing them.

SIEVE utilizes short delays to build the time-lagged versions of metrics. The reason is that microservices-based applications

typically run in the same data center and their components communicate over a LAN, where typical round-trip times are in the order of milliseconds. SIEVE uses a conservative delay of 500ms for unforeseen delays.

To apply the Granger Causality tests and check whether the past values of metric  $X$  can predict the future values of metric  $Y$ , both models are compared via the F-test [70]. The null hypothesis (i.e.,  $X$  does not granger-cause  $Y$ ) is rejected if the p-value is below a critical value.

However, one has to consider various properties of the time series. For example, the F-test requires the time series to be normally distributed. The load generation used in Section 3.1 can be adjusted to accommodate this requirement. Also, the F-test might find spurious regressions when non-stationary time series are included [56]. Non-stationary time series (e.g., monotonically increasing counters for CPU and network interfaces) can be found using the Augmented Dickey-Fuller test [58]. For these time series, the first difference is taken and then used in the Granger Causality tests. Although longer trends may be lost due to the first difference, accumulating metrics such as counters do not present interesting relationships for our purposes.

After applying the Granger Causality test to each component’s representative metrics with its neighbouring component’s representative metrics, we obtain a graph. In this graph, we draw an edge between microservice components, if one metric in one component Granger-causes another metric in a neighbouring component. This edge represents the dependency between these two components and its direction is determined by Granger causality.

While Granger Causality tests are useful in determining predictive causality across microservice components, it has some limitations that we need to consider. For example, it does not cover instantaneous relationships between two variables. More importantly, it might reveal spurious relationships, if important variables are missing in the system: if both  $X$  and  $Y$  depend on a third variable  $Z$  that is not considered, any relationship found between  $X$  and  $Y$  may not be useful. Fortunately, an indicator of such a situation is that both metrics will Granger-cause each other (i.e., a bidirectional edge in the graph). SIEVE filters these edges out.

## 4 Applications

In this section, we describe two use cases to demonstrate SIEVE’s ability to handle different workflows. In particular, using SIEVE’s base design, we implemented 1) an orchestration engine for autoscaling and applied it to ShareLatex [23], and 2) a root cause analysis (RCA) engine and applied it to OpenStack [19].

### 4.1 Orchestration of Autoscaling

For the autoscaling case study, we used ShareLatex [23]—a popular collaborative LaTeX editor. ShareLatex is structured as a microservices-based application, delegating tasks to multiple well-defined components that include a KV-store, load balancer, two databases and 11 node.js based components.

SIEVE’s pairwise investigation of representative metrics of components produces the dependencies across components. By leveraging this dependency graph, our autoscaling engine helps developers to make more informed decisions regarding which components and metrics are more critical to monitor. As a result, developers can generate *scaling rules* with the goal of adjusting the number of active component instances, depending on real-time workload.




Figure 2. SIEVE’s root cause analysis methodology.

More specifically, we use SIEVE’s dependency graph and extract (1) *guiding metrics* (i.e., metrics to use in a scaling rule), (2) *scaling actions* (i.e., actions associated with reacting to varying loads by increasing/decreasing the number of instances subject to minimum/maximum thresholds), and (3) *scaling conditions* (i.e., conditions based on a guiding metric triggering the corresponding scaling action). Below, we explain how we use SIEVE to generate a scaling rule:

**#1: Metric.** We pick a metric  $m$  that appears the most in Granger Causality relations between components.

**#2: Scaling actions.** In our case study, we restrict scaling actions to scale in/out actions, with increments/decrements of a single component instance (+/−1).

**#3: Conditions.** The scale in/out thresholds are defined from the values of  $m$  according to a Service Level Agreement (SLA) condition. For ShareLatex, such an SLA condition can be to keep 90% of all request latencies below 1000ms. The thresholds for  $m$  are iteratively refined during the application loading phase.

#### 4.2 Root Cause Analysis

For the root cause analysis (RCA) case study, we used OpenStack [19, 28], a popular open-source cloud management software. OpenStack is structured as a microservices-based application with a typical deployment of ~10 (or more) individual components, each often divided into multiple sub-components [83]. Due to its scale and complexity, OpenStack is susceptible to faults and performance issues, often introduced by updates to its codebase.

In microservices-based applications such as Openstack, components can be updated quite often [62], and such updates can affect other application components. If relationships between components are complex, such effects may not be easily foreseeable, even when inter-component interfaces are unchanged (e.g., if the density of inter-component relationships is high or if the activation of relationships is selective depending on the component’s state and inputs). SIEVE’s dependency graph can be used to understand the update’s overall effect on the application: changing dependency graphs can indicate potential problems introduced by an update. By identifying such changes, SIEVE can help developers identify the root cause of the problem.

Our RCA engine leverages SIEVE to generate a list of possible root causes of an anomaly in the monitored application. More specifically, the RCA engine compares the dependency graphs of two different versions of an application: (1) a *correct* version; and (2) a *faulty* version. Similarly to [64, 66], we assume that the system anomaly (but not its *cause*) has been observed and the correct and faulty versions have been identified. The result of this comparison is a list of {*component, metric list*} pairs: the *component* item points to a component as a possible source for the issue, whereas the

Table 2. Description of dependency graph differences considered by the root cause analysis engine.

Scoping level	Differences of interest
Component metrics	Present in F version, not in C ( <i>new</i> ) Present in C version, not in F ( <i>discarded</i> )
Clusters	Cluster includes new/discarded metrics
Dep. graph edges	New/discarded edge between similar clusters Different time-lag between similar clusters Includes clusters w/ new/discarded metrics

*metric list* shows the metrics in that component potentially related to the issue, providing a more fine-grained view. With the help of this list, developers can reduce the complexity of their search for the root cause.

Figure 2 shows the five steps involved in the comparison. At each step, we extract and analyze SIEVE’s outputs at three different granularity levels: *metrics*, *clusters*, and *dependency graph edges*. The levels and corresponding differences of interest are described in Table 2. We describe the steps in more detail below.

**#1: Metric analysis.** This step analyzes the presence or absence of metrics between C and F versions. If a metric  $m$  is present in both C and F, it intuitively represents the maintenance of healthy behavior associated with  $m$ . As such, these metrics are filtered out of this step. Conversely, the appearance of a new metric (or the disappearance of a previously existing metric) between versions is likely to be related with the anomaly.

**#2: Component rankings.** In this step, we use the results of step 1 to rank components according to their *novelty score* (i.e., total number of new or discarded metrics), producing an initial group of interesting components for RCA.

**#3: Cluster analysis: novelty & similarity.** Clusters aggregate component metrics which exhibit similar behavior over time. The clusters with new or discarded metrics should be more interesting for RCA compared to the unchanged clusters of that component (with some exceptions, explained below). For a given component, we compute the novelty scores of its clusters as the sum of the number of new and discarded metrics, and produce a list of {*component, metric list*} pairs, where the *metric list* considers metrics from the clusters with higher novelty scores.

In addition, we track the similarity of a component’s clusters between C and F versions (or vice-versa). This is done to identify two events: (1) appearance (or disappearance) of edges between versions; and (2) attribute changes in relationships maintained between C and F versions (e.g., a change in Granger causality time lag). An edge between clusters  $x$  and  $y$  (belonging to components  $A$  and  $B$ , respectively) is said to be ‘maintained between versions’ if their respective metric compositions do not change






Figure 3. Pairwise adjusted mutual information (AMI) scores between 3 measurements.

significantly between C and F versions, i.e. if  $S(M_{x,C}^A) \approx S(M_{x',F}^A)$  and  $S(M_{y,C}^B) \approx S(M_{y',F}^B)$ .  $M_{x,C}^A$  and  $M_{x',F}^A$  are the metric compositions of clusters  $x$  and  $x'$  of component  $A$ , in the C and F versions, respectively.  $S$  is some measure of cluster similarity (defined below). Both events – (1) and (2) – can be an indication of an anomaly, because one would expect edges between clusters with high similarity to be maintained between versions.

We compute the *cluster similarity score*,  $S$ , according to a modified form of the Jaccard similarity coefficient

$$S = \frac{|M_{i,C}^A \cap M_{j,F}^A|}{|M_{i,C}^A|} \quad (2)$$

To eliminate the penalty imposed by new metrics added to the faulty cluster, we only consider the contents of the correct cluster in the denominator (instead of the union of  $M_{i,C}^A$  and  $M_{j,F}^A$ ).

**#4: Edge filtering.** To further reduce the list of  $\{component, metric\}$  pairs, we examine the relationships between components and clusters identified in steps 2 and 3. We identify three events:

1. Edges involving (at least) one cluster with a high novelty score
2. Appearance or disappearance of edges between clusters with high similarity
3. Changes in time lag in edges between clusters with high similarity

Event 1 isolates metrics related to edges which include at least one ‘novel’ cluster. Events 2 and 3 isolate clusters which are maintained between C and F versions, but become interesting for RCA due to a change in their relationship. Novelty and similarity scores are computed as in step 3. We define thresholds for ‘high’ novelty and similarity scores.

**#5: Final rankings.** We present a final list of  $\{component, metric\}$  pairs. The list is ordered by component, following the rank given in step 2. The *metric list* items include the metrics identified at steps 3 and 4.

## 5 Implementation

We next describe the implementation details of SIEVE. Our system implementation, including used software versions, is published at <https://sieve-microservices.github.io>. For load generation, SIEVE requires an application-specific load generator. We experimented with two microservices-based applications: ShareLatex [23] and OpenStack [19, 28]. For ShareLatex, we developed our own load

generator using Locust [11], a Python-based distributed load generation tool to simulate virtual users in the application (1,041 LoC). For OpenStack, we used Rally [21], the official benchmark suite from OpenStack.

For metric collection, SIEVE uses Telegraf [25] to collect application/system metrics and stores them in InfluxDB [7]. Telegraf seamlessly integrates with InfluxDB, supports metrics of commonly-used components (e.g., Docker, RabbitMQ, memcached) and can run custom scripts for collection of additional metrics exposed by application APIs (e.g., [20]). With this setup, SIEVE can store any time-series metrics exposed by microservice components.

For the call graph extraction, SIEVE leverages sysdig call tracer [24] to obtain which microservice components communicate with each other. We wrote custom scripts to record network system calls with source and destination IP addresses on every machine hosting the components (457 LoC). These IP addresses are then mapped to the components using the cluster manager’s service discovery mechanism.

We implemented SIEVE’s data analytics techniques in Python (2243 LoC) including metric filtering, clustering based on  $k$ -Shape, and Granger Causality. The analysis can also be distributed across multiple machines for scalability.

Lastly, we also implemented two case studies based on the SIEVE infrastructure: autoscaling in ShareLatex (720 LoC) and RCA in OpenStack (507 LoC). For our autoscaling engine, we employed Kapacitor [9] to stream metrics from InfluxDB in real-time and to install our scaling rules using its user-defined functions. For the RCA engine, we implemented two modules in Python: one module extracts metric clustering data (125 LoC) and the other module (382 LoC) compares clustering data and dependency graphs.

## 6 Evaluation

Our evaluation answers the following questions:

1. How effective is the general SIEVE framework? (§6.1)
2. How effective is SIEVE for autoscaling? (§6.2)
3. How effective is SIEVE for root cause analysis? (§6.3)

### 6.1 Sieve Evaluation

Before we evaluate SIEVE with the case studies, we evaluate SIEVE’s general properties: (a) the robustness of clustering; (b) the effectiveness of metric reduction; and (c) the monitoring overhead incurred by SIEVE’s infrastructure.






Figure 4. Average no. of metrics after SIEVE’s reduction.

**Experimental setup.** We ran our measurements on a 10 node cluster, every node with a 4-core Xeon E5405 processor, 8 GB DDR2-RAM and a 500GB HDD. For the general experiments, we loaded ShareLatex using SIEVE five times with random workloads. The random workloads also help to validate whether the model stays consistent, if no assumption about the workload is made.

### 6.1.1 Robustness

We focus on two aspects to evaluate SIEVE’s robustness. First, we investigate the *consistency* of clustering across different runs. Second, we try to *validate* whether the metrics in a cluster indeed belong together.

**Consistency.** To evaluate consistency, we compare cluster assignments produced in different measurements. A common metric to compare cluster assignments is Adjusted Mutual Information (AMI) score [88]. AMI is normalized against a random assignment and ranges from zero to one: If AMI is equal to one, both clusters match perfectly. Random assignments will be close to zero.

Figure 3 shows the AMI of cluster assignments for individual components for three independent measurements. To reduce the selection bias we apply randomized workload in a controlled environment. As a result, they should constitute a worst-case performance for the clustering. Our measurements show that the average AMI is 0.597, which is better than random assignments. Based on these measurements, we conclude the clusterings are consistent.

**Validity.** To evaluate the validity of the clusters, we choose three criteria: (1) Is there a common visible pattern between metrics in one cluster? (2) Do metrics in a cluster belong together assuming application knowledge? (3) Are the shape-based distances between metrics and their cluster centroid below a threshold (i.e., 0.3)?

We choose three clusters with different Silhouette scores (high, medium, low). According to the above criteria, we conclude the clustering algorithm can determine similar metrics. For example, application metrics such as HTTP request times and corresponding database queries are clustered together. Similar to consistency, higher Silhouette scores indicate that the clusters are more meaningful and potentially more useful for the developers. We omit the details for brevity.

### 6.1.2 Effectiveness

The purpose of clustering is to reduce the number of metrics exposed by the system without losing much information about the system behavior. To evaluate how effective our clustering is in




Figure 5. Completion time for HTTP requests when using tcpdump, sysdig or native (i.e., no monitoring).

Table 3. InfluxDB overhead before SIEVE’s reduction of metrics.

Metric	Before	After	Reduction
CPU time [s]	0.45G	0.085G	81.2 %
DB size [KB]	588.8	36.0	93.8 %
Network in [MB]	11.1	2.3	79.3 %
Network out [KB]	15.1	7.4	50.7 %

reducing the number of metrics, we compare the results of the clustering with the actual number of metrics in the application. We identified 889 unique metrics within ShareLatex, meaning that an operator would have to understand and filter these metrics. SIEVE’s clustering reduces this number to 65 (averaged across five runs). Figure 4 shows the reduction in the number of metrics for the individual components in ShareLatex. Note that this measurement is with high Silhouette scores for the clusters, which implies that the metrics reduction does not affect the quality of the clusters.

### 6.1.3 Monitoring Overhead

We evaluate SIEVE’s overhead based on two aspects. First, we compare different techniques for obtaining the call graph and show how our approach fares. Second, we investigate the overhead incurred by the application by comparing the monitoring overhead with and without using SIEVE.

**Overheads.** To measure the monitoring overhead during the loading stage, we run an experiment with 10K HTTP requests for a small static file using Apache Benchmark [12] on an Nginx web server [15]. Because the computational overhead for serving such a file is low, this experiment shows the worst-case performance for sysdig and tcpdump. Figure 5 shows the time it takes to complete the experiment. While tcpdump incurs a lower overhead than sysdig (i.e., 7% vs. 22%), it provides less context regarding the component generating the request and requires more knowledge about the network topology to obtain the call graph. sysdig provides all this information without much additional overhead.

**Gains.** To show the gains during the runtime of the application after using SIEVE, we compare the computation, storage and network usage for the metrics collected during the five measurements. We store all collected metrics in InfluxDB and measure the respective resource usage. We then repeat the same process using the metrics found by SIEVE; thus, simulating a run with the reduced metrics. Table 3 shows the relative usage of the respective resources with SIEVE. SIEVE reduces the monitoring overhead for computation, storage and network by 80%, 90% and 50%, respectively.

**Table 4.** Comparison between a traditional metric (CPU usage) and SIEVE’s selection when used as autoscaling triggers.

Metric	CPU usage	Sieve	Difference [%]
Mean CPU usage per component	5.98	9.26	+54.82
SLA violations (out of 1400 samples)	188	70	-62.77
Number of scaling actions	32	21	-34.38

## 6.2 Case-study #1: Autoscaling

We next evaluate the effectiveness of SIEVE for the orchestration of autoscaling in microservices.

**Experimental setup.** For the autoscaling case study, we used ShareLatex [23] (as described in §4.1). We used 12 t2.large VM-Instances on Amazon EC2 with 2 vCPUs, 8GB RAM and 20 GB Amazon EBS storage. This number of instances were sufficient to stress-test all components of the application. The VM instances were allocated statically during experiments as Docker containers. We created a Docker image for each ShareLatex component and used Rancher [22] as the cluster manager to deploy our containers across different hosts.

**Dataset.** We used a HTTP trace sample from soccer world cup 1998 [6] for an hour long trace. Note that the access pattern and requested resources in the world cup trace differs from the ShareLatex application. However, we used the trace to map traffic patterns for our application to generate a realistic spike workload. In particular, sessions in the HTTP trace were identified by using the client IP. Afterwards, we enqueued the sessions based on their timestamp, where a virtual user was spawned for the duration of each session and then stopped.


**Results.** We chose an SLA condition, such that 90th percentile of all request latencies should be below 1000ms. Traditional tools, such as Amazon AWS Auto Scaling [1], often use the CPU usage as the default metric to trigger autoscaling. SIEVE identified an application metric named *http-requests\_Project\_id\_GET\_mean* (Figure 6) as a better metric for autoscaling than CPU usage.

To calculate the threshold values to trigger autoscaling, we used a 5-minute sample from the peak load of our HTTP trace and iteratively refined the values to stay within the SLA condition. As a result, we found that the trigger thresholds for scaling up and down while using the CPU usage metric should be 21% and 1%, respectively. Similarly, for *http-requests\_Project\_id\_GET\_mean*, the thresholds for scaling up and down should be 1400ms and 1120ms, respectively.

After installing the scaling actions, we ran our one-hour trace. Table 4 shows the comparison when using the CPU usage and *http-requests\_Project\_id\_GET\_mean* for the scaling triggers. When SIEVE’s selection of metric was used for autoscaling triggers, the average CPU usage of each component was increased. There were also fewer SLA violations and scaling actions.

## 6.3 Case-study #2: Root Cause Analysis

To evaluate the applicability of SIEVE to root cause analysis, we reproduce two representative OpenStack anomalies, Launchpad bugs #1533942 [29] and #1590179 [30]. We selected these issues because they have well-documented root causes, providing an appropriate ground truth, and allowing for the identification of ‘correct’ and



**Figure 6.** Relations between components based on Granger Causality in ShareLatex. The dashed lines denote relationships with metric *http-requests\_Project\_id\_GET\_mean*.

‘faulty’ code versions. Moreover, these showcase SIEVE’s effectiveness while analyzing two different types of bugs: (1) a crash in the case of #1533942; and (2) performance regression (e.g. increase in latency) in the case of #1590179. We compare the documented root causes to the lists of root causes produced by our RCA engine.

### 6.3.1 Bug 1533942 : Failure to launch a VM

Bug #1533942 manifests itself as follows: when launching a new VM instance using the command line interface, one gets the error message ‘No valid host was found. There are not enough hosts available.’ despite the availability of compute nodes. Without any other directly observable output, the instance falls into ‘ERROR’ state and fails. A similar failure is used as a representative case in [55, 83].

**Root cause.** The failure is caused by the crash of an agent in the Neutron component, namely the Open vSwitch agent. The Open vSwitch agent is responsible for setting up and managing virtual networking for VM instances. The ultimate cause is traced to a configuration error in OpenStack Kolla’s deployment scripts [29].

**Experimental setup.** We deployed OpenStack components as containerized microservices using Kolla [28]. We configured Kolla to deploy 7 main OpenStack components along with 12 auxiliary components. Some components can be sub-divided in several microservices and replicated among deployment hosts, for a total of 47 microservices. This information is summarized in Table 5. We use OpenStack’s telemetry component (Ceilometer) to expose relevant OpenStack-related metrics and extract them via Telegraf.

The Openstack versions used for the correct (C) and faulty (F) versions are listed in Table 8. The configurations for the Kolla deployments of each version are publicly available <sup>1</sup>, as well as the monitoring infrastructure and evaluation scripts <sup>2</sup>.

The infrastructure consists of two m4.xlarge Amazon EC2 VM instances to run OpenStack components (16 vCPUs, 64 GB RAM and 20 GB Amazon EBS storage) and three t2.medium VM instances (2 vCPUs, 4GB RAM and 20 GB EBS storage) for the supporting components (measurement, database and deployment).

<sup>1</sup><https://github.com/sieve-microservices/kolla>

<sup>2</sup><https://github.com/sieve-microservices/rca-evaluation>



**Table 5.** Components deployed by Openstack Kolla during RCA evaluation.

Component	Purpose	# Microservices
Nova	VM computing	8
Neutron	VM networking	6
Keystone	Identity	3
Glance	VM image manag.	2
Heat	-	3
Horizon	Web UI	1
Ceilometer	Telemetry	5
Heka	Logging	3
Cron	Job scheduling	3
Open vSwitch	VM networking (aux.)	4
Elasticsearch	Search engine	1
Kibana	Data visualiz.	1
Memcached	Auth. token caching (among others)	1
Mariadb	Openstack parameter storage	1
RabbitMQ	Message broker	1
MongoDB	Ceilometer data storage	1
Telegraf	Metric collection	1
InfluxDB	Metric storage	1
PostgreSQL	Metric storage	1
<b>Totals</b>	-	<b>47</b>

is a change in the token caching strategy in-between Openstack ‘Liberty’ and ‘Mitaka’ releases.

In the context of Openstack, tokens represent the authenticated identity of a specific requester (e.g. a system user) and grants authorization for a specific Openstack action (e.g. starting a VM) [10]. Openstack supports different types of tokens, but this issue is particular to Fernet tokens [10], which do not require persistence in a database: validation is based on symmetric encryption, with secret keys kept by the Openstack identity component.

**Experimental setup.** Similarly to the bug #1533942 use case, we deployed OpenStack components as containerized microservices using Kolla [28]. We deployed 7 Openstack components, along with 12 auxiliary components, as listed in Table 5. The Openstack versions used for the correct (C) and faulty (F) versions are listed in Table 8. The configurations for the Kolla deployments of each version are publicly available <sup>3</sup>, as well as the monitoring infrastructure and evaluation scripts <sup>4</sup>.

The infrastructure consists of 3 t2.large Amazon EC2 VM instances to run OpenStack components (2 vCPUs, 8 GB RAM and 30 GB Amazon EBS storage) and 2 t2.medium VM instances (2 vCPUs, 4GB RAM and 30 GB EBS storage) for the supporting components (metric collection and storage).

**Results.** Since this bug manifests itself as a general performance degradation issue, we run the Rally tasks below to load 4 essential Openstack components, thus giving us a reasonably large ‘search space’ for RCA (details in Table 6):

- **B&D:** Same Rally task used in bug #1533942. Loads Openstack’s compute (Nova) and networking (Neutron) components.
- **AU&VT:** Authenticates a stream of user’s tokens in Keystone, Openstack’s identity component. In hindsight, since

<sup>3</sup><https://github.com/sieve-microservices/kolla>

<sup>4</sup><https://github.com/sieve-microservices/rca-evaluation>

**Table 6.** Details about Rally tasks used in RCA evaluation.

Bug #	Benchmark	# Runs	Concurr.	Details
1533942	boot and delete (B&D)	100	5	VMs up for 15-25 sec
	<b>B&amp;D</b>	25	5	VMs up for 15-25 sec
1590179	authenticate user and validate token (AU&VT)	100	5	-
	create and delete networks (C&DN)	50	5	-
	create and delete image (C&DI)	50	2	Cirros 0.35 x86_64 image

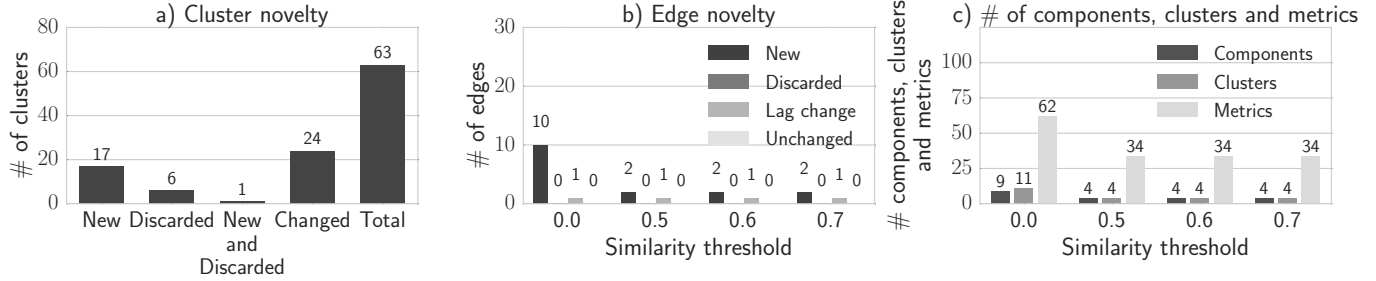
**Table 7.** OpenStack components, sorted by the number of novel metrics between correct (C) and faulty (F) versions.

Component	Changed (New/Discarded)	Total (per component)	Final ranking
Nova API	29 (7/22)	59	1
Nova libvirt	21 (0/21)	39	2
Nova scheduler	14 (7/7)	30	-
Neutron server	12 (2/10)	42	3
RabbitMQ	11 (5/6)	57	4
Neutron L3 agent	7 (0/7)	39	5
Nova novncproxy	7 (0/7)	12	-
Glance API	5 (0/5)	27	6
Neutron DHCP ag.	4 (0/4)	35	7
Nova compute	3 (0/3)	41	8
Glance registry	3 (0/3)	23	9
Haproxy	2 (1/1)	14	10
Nova conductor	2 (0/2)	29	-
Other 3 components	0 (0/0)	59	-
<b>Totals</b>	<b>113 (22/91)</b>	<b>508</b>	-

the reported root cause of bug #1590179 is related to Keystone [30], this might appear as a ‘dishonest’ test. However, we argue that Keystone is a central component of Openstack - as are the compute and networking components - and as such a natural candidate for testing.

- **C&DN:** Creates and deletes network VM network resources. Loads Openstack’s networking component - Neutron - as well as related components (e.g. Open vSwitch).
- **C&DI:** Creates and deletes VM images. Loads Openstack’s image component, Glance.

**Metric analysis and component rankings.** Table 9 shows that solely based on metric novelty scores, the ‘authenticate user and validate token’ (AU&VT) Rally task - which directly relates to token validation - ranks Memcached 2nd place as a possible root cause (metric novelty rankings correspond to column ‘N’ in Table 9, Memcached is ranked 1st after edge filtering, i.e. the ‘F’ column). Other tasks rank Nova consoleauth (B&D) and Neutron OpenvSwitch agent (C&DN and C&DI) as 1st, which are unrelated to bug #1590179 (at least not directly related to it, according to the description given in [30]).



**Figure 9.** AU&VT Rally task results: (a) cluster novelty score. (b) edge novelty score. (c) no. of components & clusters after edge filtering w/ varying thresholds.

**Table 8.** Openstack versions used in RCA evaluation.

Bug #	Correctness	Openstack ver.	Comments
1533942	Correct (C)	Mitaka (EOL)	Adapted for Ceilometer support (not available by default) <sup>1</sup>
		Kolla 2.0.0.0b3	-
1590179	Faulty (F)	Mitaka (EOL)	(1)
		Kolla 2.0.0.0b2	-
1590179	Correct (C)	Liberty (EOL)	(1)
		Keystone 8.1.0	Adapted to Fernet tokens based on blueprint in [26] <sup>2</sup>
1590179	Faulty (F)	Mitaka (EOL)	(1)
		Keystone 9.0.2	(2)

Still regarding metric novelty, Keystone - one of the Openstack components directly related to bug #1590179 [30] - is lowly ranked according to the AU&VT Rally task. This result is also intuitive: since bug #1590179 relates to performance degradation, metrics should not appear or disappear, rather their values should differ in-between C and F versions, an effect which may be captured by relationships in-between the clusters the metrics belong to.

Since we load Openstack with the ‘boot and delete’ (B&D) tasks in both bugs #1533942 and #1590179, we can compare how the rankings change between bugs, and verify if there is evidence of a dependency between the way in which the system is loaded and the ‘metric novelty’ rankings produced by SIEVE. As shown in Table 10, the average change in rank is  $\sim 5$  positions, providing initial evidence about the lack of such dependency. Further validation of this dependency is left to future work.

**Edge filtering.** The edge novelty statistics results for the AU&VT task - depicted in Figures 9(b) - are of particular interest, since it directly loads Openstack with a task related to bug #1590179.

The edge filtering step identifies 10 new edges no inter-version cluster similarity threshold is applied. These reduce to 2 once the similarity threshold is raised to 0.5. The 2 new edges are between the following components: RabbitMQ > Nova conductor and Memcached > Nova libvirt. The Memcached metrics associated with the edge are not related to cache access metrics, e.g. cache hits or misses. The single edge isolated due to a causality ‘lag change’ pertains to a relationship between RabbitMQ and and Nova libvirt, with no apparent relation to the bug. None of the filtered relationships involves Keystone, and as such metrics such as ‘keystone identity authenticate success rate’, which are known to be related to the bug.

**Final rankings.** In the case of the AU&VT task, Memcached - which effectively is related to the bug - is ranked 1st after the edge filtering steps. However, none of the metrics filtered in the edge filtering steps seems to directly relate to bug #1590179. Also, an intuitive isolation of lag changes in edges involving Keystone clusters did not occur. Further experimentation is required to assess SIEVE’s effectiveness for RCA of performance degradation issues (as opposed to crashes, such as bug #1533942) such bug #1590179.

## 7 Related Work

**Scalable Monitoring.** With the increasing number of metrics exposed by distributed cloud systems, the scalability of the monitoring process becomes crucial. Meng et al. [73] optimize monitoring scalability by choosing appropriate monitoring window lengths and adjusting the monitoring intensity at runtime. Canali et al. [42] achieve scalability by clustering metric data. A fuzzy logic approach is used to speed up clustering, and thus obtain data for decision making within shorter periods. Rodrigues et al. [46] explore the trade-off between timeliness and the scalability in cloud monitoring, and analyze the mutual influence between these two aspects based on the monitoring parameters. Our work is complementary to existing monitoring systems since SIEVE aims to improve the efficiency by monitoring less number of metrics.

**Distributed Debugging.** Systems like Dapper [84] and Pip [80] require the developers to instrument the application to obtain its causal model. X-trace [50] uses a modified network stack to propagate useful information about the application. In contrast, SIEVE does not modify the application code to obtain the call/dependency graph of the application.

Systems such as Fay [49] and DTrace [43] enable developers to dynamically inject debugging requests by developers and require no initial logs of metrics. Pivot Tracing [72] combines dynamic instrumentation with causality tracing. SIEVE can complement these approaches, because it can provide information about interesting components and metrics, so that the developers can focus their efforts to understand them better. Furthermore, SIEVE’s dependency graph is a general tool that can not only be used for debugging, but also for other purposes such as orchestration [89–91].

Data provenance [37, 53, 86] is another technique that can be used to trace the dataflow in the system. SIEVE can also leverage the existing provenance tools to derive the dependence graph.

**Metric reduction.** Reducing the size and dimensionality of the bulk of metric data exposed by complex distributed systems is essential for its understanding. Common techniques include sampling, and data clustering via  $k$ -means and  $k$ -medoids. Kollios et

**Table 9.** Bug #151590179 results: OpenStack components, sorted by the number of novel metrics between correct (C) and faulty (F) versions, for different Rally tasks. Metric novelty rankings correspond to column ‘N’, final rankings (after edge filtering steps) correspond to column ‘F’. Components which are ranked 1st according to the ‘metric (N)ovelty’ ranking are underlined, those ranked 1st after edge filtering (‘(F)inal’ rankings) are emphasized in **bold**.

Component	Rally tasks															
	AU&VT				B&D				C&DN				C&DI			
	(C)hanged (New/Disc.)	(T)otal	Rankings N F		C	T	N	F	C	T	N	F	C	T	N	F
<b>Neutron OvSwitch</b>	<u>42 (15/27)</u>	<u>53</u>	<u>1</u>	-	0 (0/0)	38	14	-	<b><u>40 (14/26)</u></b>	<b><u>52</u></b>	<b><u>1</u></b>	<b><u>1</u></b>	<u>42 (15/27)</u>	<u>53</u>	<u>1</u>	-
<b>Memcached</b>	<b><u>15 (14/1)</u></b>	<b><u>30</u></b>	<b><u>2</u></b>	<b><u>1</u></b>	6 (4/2)	33	5	3	7 (3/3)	33	10	4	5 (4/1)	31	11	3
<b>Nova API</b>	10 (10/0)	25	3	-	3 (3/0)	44	10	6	20 (15/5)	32	3	-	<b><u>18 (13/5)</u></b>	<b><u>30</u></b>	<b><u>3</u></b>	<b><u>1</u></b>
Nova conductor	9 (0/9)	15	4	2	2 (2/0)	30	12	7	27 (0/27)	33	2	-	0 (0/0)	6	18	-
Nova libvirt	8 (0/8)	17	5	3	2 (2/0)	39	13	8	7 (1/6)	17	8	-	12 (6/6)	23	5	2
<b>Neutron server</b>	6 (6/0)	21	5	-	<b><u>12 (5/7)</u></b>	<b><u>52</u></b>	<b><u>2</u></b>	<b><u>1</u></b>	6 (4/2)	34	11	5	18 (13/5)	30	4	-
Neutron L3 agent	4 (0/4)	19	7	-	6 (4/2)	46	4	-	0 (0/0)	15	16	-	2 (2/0)	15	13	-
Nova SSH	3 (3/0)	3	8	-	- (-/-)	-	-	-	3 (3/0)	3	14	-	- (-/-)	-	-	-
Glance registry	3 (1/2)	45	9	-	0 (0/0)	30	17	-	7 (7/0)	7	9	-	8 (8/0)	8	7	-
RabbitMQ	3 (3/0)	3	10	4	4 (2/2)	63	6	4	8 (8/0)	54	6	2	10 (9/1)	57	6	3
Neutron DHCP ag.	2 (2/0)	15	11	-	0 (0/0)	42	15	9	0 (0/0)	15	18	-	0 (0/0)	15	17	-
Nova compute	2 (0/2)	19	12	-	0 (0/0)	45	16	-	6 (2/4)	23	12	-	2 (0/2)	19	14	-
Nova novncproxy	1 (1/0)	4	13	-	3 (0/3)	17	11	-	8 (7/1)	12	7	3	8 (8/0)	13	8	-
Nova scheduler	0 (0/0)	6	14	-	4 (1/3)	31	8	-	14 (14/0)	14	5	-	0 (0/0)	14	15	-
<u>Nova consoleauth</u>	0 (0/0)	6	15	-	<u>18 (4/14)</u>	<u>36</u>	<u>1</u>	-	0 (0/0)	6	17	-	5 (5/0)	20	10	-
Keystone	0 (0/0)	36	16	-	8 (6/2)	42	3	2	2 (1/1)	37	15	7	0 (0/0)	36	16	-
Neutron metadata ag.	0 (0/0)	16	17	-	0 (0/0)	24	18	-	16 (16/0)	16	4	-	20 (20/0)	20	2	-
Glance API	- (-/-)	-	18	-	4 (0/4)	34	7	5	5 (5/0)	15	13	6	6 (6/0)	6	9	-
<b>Totals</b>	<b>103 (50/53)</b>	<b>332</b>	-	-	<b>76 (34/42)</b>	<b>657</b>	-	-	<b>136 (61/75)</b>	<b>378</b>	-	-	<b>128 (81/47)</b>	<b>392</b>	-	-

**Table 10.** Change in ‘metric novelty’ rankings for bugs #1533942 and #1590179, considering the ‘boot and delete’ Rally task.

Component	Rank (bug #)		Ranking change
	1533942	1590179	
Nova API	1	10	-9
Nova libvirt	2	11	-11
Nova scheduler	3	8	-5
Neutron server	4	2	+2
RabbitMQ	5	6	-1
Neutron L3 agent	6	4	+2
Nova novncproxy	7	11	-4
Glance API	8	7	+1
Neutron DHCP ag.	9	15	-6
Nova compute	10	16	-6
Glance registry	11	17	-6
Nova conductor	13	12	+1
Keystone	15	3	+12

al. [67] employ biased sampling to capture the local density of datasets. Sampling based approaches argues for approximate computing [68, 77, 78] to enable a systematic trade-off between the accuracy, and efficiency to collect and compute on the metrics. Zhou et al. [94] simply use random sampling due to its simplicity and low complexity. Ng et al. [74] improved the  $k$ -medoid method and made it more effective and efficient. Ding et al. [47] rely on clustering over sampled data to reduce clustering time.

SIEVE’s approach is unique because of its two-step approach: (1) we first cluster time series to identify the internal dependency between *any* given metrics and then (2) infer the causal relations

among time series. Essentially, SIEVE uses two steps of data reduction for better reduction. Furthermore, SIEVE’s time series processing method extracts other useful information such as the time delay of the causal relationship, which can be leveraged in different use cases (e.g., root cause analysis).

**Orchestration of autoscaling.** Current techniques for autoscaling can be broadly classified into four categories [71]: (i) static and threshold-based rules (offered by most cloud computing providers [3, 4, 17, 27]); (ii) queuing theory [34, 60, 93]; (iii) reinforcement learning [79, 85, 92]; and (iv) time series analysis [44, 65, 82]. Existing systems using these techniques can benefit from the selection of better metrics and/or from the dependencies between components. In this regard, our work is complementary to these techniques: it is intended to provide the developers with knowledge about the application as a whole. In our case study, we showed the benefits of SIEVE for an autoscaling engine using threshold-based rules.

**Root Cause Analysis (RCA).** Large and complex distributed systems are susceptible to anomalies, whose root causes are often hard to diagnose [62]. Jiang et al. [64] compare “healthy” and “faulty” metric correlation maps, searching broken correlations. In contrast, SIEVE leverages Granger causality instead of simple correlation, allowing for richer causality inference (e.g., causality direction, time lag between metrics). MonitorRank [66] uses metric collection for RCA in a service-oriented architecture. It only analyzes pre-established (component, metric) relations according to a previously-generated call graph. SIEVE also uses a call graph, but does not fix metric relations between components, for a richer set of potential root causes. There are other application-specific solutions for RCA (e.g., Hansel [83], Gretel [55]). In contrast, SIEVE uses a general approach for understanding the complexity of microservices-based applications that can support RCA as well as other use cases.

## 8 Experience and Lessons Learned

While developing SIEVE, we set ourselves ambitious design goals (described in §2.2). However, we learned the following lessons while designing and deploying SIEVE for real-world applications.

**Lesson#1.** When we first designed SIEVE, we were envisioning a dependency graph that was clearly showing the relationships between components (e.g., a tree). As a result, not only would the number of metrics that needed to be monitored be reduced, but also the number of components: one would only need to observe the root(s) of the dependency graph, and make the actions of the dependent components according to the established relationships between the root(s) and them. Such a dependency graph would give the orchestration scenario a huge benefit. Unfortunately, our experience has shown us that the relationships between components are usually not linear, making the dependency graph more complex. Also, there was no obvious root. Consequently, we had to adjust our thinking and utilize some application knowledge regarding components and their relations with others. Nevertheless, in our experience, SIEVE provides the developer with a good starting point to improve their workflows.

**Lesson#2.** SIEVE is designed for “blackbox” monitoring of the evaluated application, where SIEVE can collect and analyze generic system metrics that are exposed by the infrastructure (e.g., CPU usage, disk I/O, network bandwidth). However, in our experience, a system for monitoring and analyzing an application should also consider application-specific metrics (e.g., request latency, number of error messages) to build effective management tools. Fortunately, many microservices applications we analyzed already export such metrics. However, given the number of components and exported metrics, this fact can easily create an “information overload” for the application developers. In fact, the main motivation of SIEVE was to deal with this “information overload”. Our experience showed that SIEVE can still monitor the application in the blackbox mode (i.e., no instrumentation to the application), but also overcome the barrage of application-specific metrics.

**Lesson#3.** To adapt to the application workload variations, SIEVE needs to build a robust model for the evaluated application. This requires a workload generator that can stress-test the application thoroughly. To meet this requirement, there are three approaches: (1) In many cases the developers already supply an application-specific workload generator. For instance, we employed the workload generator shipped with the OpenStack distribution. (2) For cases where we did not have an existing workload generator, we implemented a custom workload generator for the evaluated application. For example, we built a workload generator for ShareLatex. Although we were able to faithfully simulate user actions in ShareLatex, such an approach might not be feasible for some applications. Having the ability to utilize existing production traces (e.g., by replaying the trace or by reproducing similar traces) or working in an online fashion to generate the model of the application would certainly help SIEVE. Custom workload generation can then be used to close the gaps in the model for certain workload conditions not covered by the existing traces. (3) We could also explore some principled approaches for automatic workload generation, such as symbolic execution in distributed systems [36].

## 9 Conclusion and Future Work

This paper reports on our experiences with designing and building SIEVE, a platform to automatically derive actionable insights from monitored metrics in distributed systems. SIEVE achieves this goal by automatically reducing the amount of metrics and inferring inter-component dependencies. Our general approach is independent of the application, and can be deployed in an unsupervised mode without prior knowledge of the time series of metrics. We showed that SIEVE’s resulting model is consistent, and can be applied for common use cases such as autoscaling and root-cause debugging.

An interesting research challenge for the future would be to integrate SIEVE into the continuous integration pipeline of an application development. In this scenario, the dependency graph can be updated incrementally [39–41], which would speed up the analytics part. In this way, the developers would be able to get real-time profile updates of their infrastructure. Another challenge is to utilize already existing traffic to generate the dependency graph without requiring the developers to load the system. Using existing traffic would alleviate the burden of developers to supply a workload generator. On the other hand, existing traffic traces might not always capture the stress points of the application. A hybrid approach, in which workload generation is only used for these corner cases, might help to overcome this problem.

**Software availability.** The source code of SIEVE is publicly available: <https://sieve-microservices.github.io/>.

**Acknowledgments.** We would like to thank Amazon AWS for providing the required infrastructure to run the experiments.

## References

- [1] Amazon AWS - Scaling Based on Metrics. [https://docs.aws.amazon.com/autoscaling/latest/userguide/policy\\_creating.html](https://docs.aws.amazon.com/autoscaling/latest/userguide/policy_creating.html). Last accessed: September, 2017.
- [2] Amazon CloudWatch. <https://aws.amazon.com/de/cloudwatch/>. Last accessed: September, 2017.
- [3] Amazon Web Services. <https://aws.amazon.com/documentation/autoscaling/>. Last accessed: September, 2017.
- [4] Google Cloud Platform. <https://cloud.google.com/developers/articles/auto-scaling-on-the-google-cloud-platform>. Last accessed: September, 2017.
- [5] Google Stackdriver. <https://cloud.google.com/stackdriver/>. Last accessed: September, 2017.
- [6] Http Trace of WorldCup98. <http://ita.ee.lbl.gov/html/contrib/WorldCup.html>. Last accessed: September, 2017.
- [7] InfluxDB. <https://influxdata.com/time-series-platform/influxdb/>. Last accessed: September, 2017.
- [8] Introducing Vector. <http://techblog.netflix.com/2015/04/introducing-vector-netflix-on-host.html>. Last accessed: September, 2017.
- [9] Kapacitor. <https://influxdata.com/time-series-platform/kapacitor/>. Last accessed: September, 2017.
- [10] Keystone tokens. <https://docs.openstack.org/keystone/latest/admin/identity-tokens.html>. Last accessed: September, 2017.
- [11] Locust - An Open Source Load Testing Tool. <http://locust.io/>. Last accessed: September, 2017.
- [12] Manualpage of Apache Benchmark. <https://httpd.apache.org/docs/2.4/programs/ab.html>. Last accessed: September, 2017.
- [13] Microsoft Azure Monitor. <https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-overview>. Last accessed: September, 2017.
- [14] Monitoring at Quantcast. <https://www.quantcast.com/wp-content/uploads/2013/10/Wait-How-Many-Metrics-Quantcast-2013.pdf>. Last accessed: September, 2017.
- [15] Nginx. <https://nginx.org/>. Last accessed: September, 2017.
- [16] Observability at Uber Engineering: Past, Present, Future. <https://www.youtube.com/watch?v=2JAnmzVwgP8>. Last accessed: September, 2017.
- [17] OpenStack. <https://wiki.openstack.org/wiki/Heat>. Last accessed: September, 2017.
- [18] Openstack: API References (Response parameters). <https://developer.openstack.org/api-ref/>. Last accessed: September, 2017.
- [19] Openstack: Open source Software for Creating Private and Public Clouds. <https://www.openstack.org/>. Last accessed: September, 2017.



- [20] Openstack: Telemetry. <https://docs.openstack.org/admin-guide/telemetry-measurements.html>. Last accessed: September, 2017.
- [21] Rally. <https://wiki.openstack.org/wiki/Rally>. Last accessed: September, 2017.
- [22] Rancher Container Management. <http://rancher.com/>. Last accessed: September, 2017.
- [23] Sharelatex - A Web-based Collaborative LaTeX Editor. <https://sharelatex.com>. Last accessed: September, 2017.
- [24] Sysdig. <http://www.sysdig.org/>. Last accessed: September, 2017.
- [25] Telegraf: Time-series Data Collection. <https://www.influxdata.com/time-series-platform/telegraf/>. Last accessed: September, 2017.
- [26] Use Fernet tokens with Keystone. <https://blueprints.launchpad.net/kolla/+spec/keystone-fernet-token>. Last accessed: September, 2017.
- [27] Windows Azure. [http://msdn.microsoft.com/en-us/library/hh680945\(v=pandp.50\).aspx](http://msdn.microsoft.com/en-us/library/hh680945(v=pandp.50).aspx). Last accessed: September, 2017.
- [28] Openstack Kolla. <http://docs.openstack.org/developer/kolla/>, 2016. Last accessed: September, 2017.
- [29] Openstack Kolla Launchpad: neutron-openvswitch-agent Bug. <https://bugs.launchpad.net/kolla/+bug/1533942>, 2016. Last accessed: September, 2017.
- [30] Openstack Kolla Launchpad: neutron-openvswitch-agent Bug. <https://bugs.launchpad.net/keystone/+bug/1590179>, 2016. Last accessed: September, 2017.
- [31] *ptrace(2) Linux User's Manual*, 4.07 edition, Aug 2016.
- [32] Scikit Documentation: sklearn.metrics.silhouette\_score. [http://scikit-learn.org/stable/modules/generated/sklearn.metrics.silhouette\\_score.html](http://scikit-learn.org/stable/modules/generated/sklearn.metrics.silhouette_score.html), 2016. Last accessed: September, 2017.
- [33] M. K. Aguilera, J. C. Mogul, J. L. Wiener, P. Reynolds, and A. Muthitacharoen. Performance Debugging for Distributed Systems of Black Boxes. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP)*, 2003.
- [34] A. Ali-Eldin, M. Kihl, J. Tordsson, and E. Elmroth. Efficient Provisioning of Bursty Scientific Workloads on the Cloud Using Adaptive Elasticity Control. In *Proceedings of the 3rd Workshop on Scientific Cloud Computing Date (ScienceCloud)*, 2012.
- [35] Amemiya, Takeshi. *Advanced econometrics*. 1985.
- [36] R. Banabic, G. Candea, and R. Guerraoui. Finding trojan message vulnerabilities in distributed systems. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2014.
- [37] A. Bates, D. J. Tian, K. R. Butler, and T. Moyer. Trustworthy whole-system provenance for the linux kernel. In *24th USENIX Security Symposium (USENIX Security)*, 2015.
- [38] R. Bellman and R. Corporation. *Dynamic Programming*. Rand Corporation research study. Princeton University Press, 1957.
- [39] P. Bhatotia. *Incremental Parallel and Distributed Systems*. PhD thesis, Max Planck Institute for Software Systems (MPI-SWS), 2015.
- [40] P. Bhatotia, P. Fonseca, U. A. Acar, B. Brandenburg, and R. Rodrigues. iThreads: A Threading Library for Parallel Incremental Computation. In *Proceedings of the 20th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2015.
- [41] P. Bhatotia, A. Wieder, R. Rodrigues, U. A. Acar, and R. Pasquini. Incoop: MapReduce for Incremental Computations. In *Proceedings of the ACM Symposium on Cloud Computing (SoCC)*, 2011.
- [42] C. Canali and R. Lancellotti. An Adaptive Technique To Model Virtual Machine Behavior for Scalable Cloud Monitoring. In *Proceedings of the 19th IEEE Symposium on Computers and Communications (ISCC)*, 2014.
- [43] B. Cantrell, M. W. Shapiro, A. H. Leventhal, et al. Dynamic instrumentation of production systems. In *Proceedings of the 2004 USENIX Annual Technical Conference (ATC)*, 2004.
- [44] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, and F. Zhao. Energy-aware Server Provisioning and Load Dispatching for Connection-intensive Internet Services. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2008.
- [45] 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 *Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2014.
- [46] G. da Cunha Rodrigues, R. N. Calheiros, M. B. de Carvalho, C. R. P. dos Santos, L. Z. Granville, L. Tarouco, and R. Buyya. The Interplay Between Timeliness and Scalability In Cloud Monitoring Systems. In *Proceedings of the 20nd IEEE Symposium on Computers and Communications (ISCC)*, 2015.
- [47] R. Ding, Q. Wang, Y. Dang, Q. Fu, H. Zhang, and D. Zhang. Yading: Fast Clustering of Large-scale Time Series Data. In *Proceedings of the 41st International Conference on VERY LARGE DATA BASES (VLDB)*, 2015.
- [48] C. W. S. Emmons, and B. Gregg. A Microscope on Microservices. <http://techblog.netflix.com/2015/02/a-microscope-on-microservices.html>, 2015. Last accessed: September, 2017.
- [49] U. Erlingsson, M. Peinado, S. Peter, and M. Budiu. Fay: Extensible distributed tracing from kernels to clusters. In *Proceedings of the 23th ACM Symposium on Operating Systems Principles (SOSP)*, 2011.
- [50] R. Fonseca, G. Porter, R. H. Katz, S. Shenker, and I. Stoica. X-trace: A pervasive network tracing framework. In *Proceedings of the conference on Networked systems design & implementation (NSDI)*, 2007.
- [51] M. Fowler. Microservices. <http://martinfowler.com/articles/microservices.html>. Last accessed: September, 2017.
- [52] K. P. F.R.S. LIII. On Lines and Planes of Closest Fit to Systems of Points in Space. In *Proceedings of the 17th ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems (SIGMOD)*, 1901.
- [53] A. Gehani and D. Tariq. Spade: Support for provenance auditing in distributed environments. In *Proceedings of the 13th International Middleware Conference (Middleware)*, 2012.
- [54] D. Giles. Testing for Granger Causality. <https://davegiles.blogspot.de/2011/04/testing-for-granger-causality.html>. Last accessed: September, 2017.
- [55] A. Goel, S. Kalra, and M. Dhawan. GRETEL: Lightweight Fault Localization for OpenStack. In *Proceedings of the 12th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*, 2016.
- [56] C. Granger and P. Newbold. Spurious Regressions in Econometrics. *Journal of Econometrics*, 2(2):111–120, 1974.
- [57] C. W. J. Granger. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, 1969.
- [58] W. H. Greene. *Econometric Analysis*. Prentice Hall, 5. edition, 2003.
- [59] E. Haddad. Service-Oriented Architecture: Scaling the uber Engineering Codebase As We Grow. <https://eng.uber.com/soa/>, 2015. Last accessed: September, 2017.
- [60] R. Han, M. M. Ghanem, L. Guo, Y. Guo, and M. Osmond. Enabling Cost-aware and Adaptive Elasticity of Multi-tier Cloud Applications. 2014.
- [61] B. Harrington and R. Rapoport. Introducing Atlas: Netflix's Primary Telemetry Platform. <http://techblog.netflix.com/2014/12/introducing-atlas-netflixs-primary.html>, 2014. Last accessed: September, 2017.
- [62] V. Heorhiadi, S. Rajagopalan, H. Jamjoom, M. K. Reiter, and V. Sekar. Gremlin: Systematic Resilience Testing of Microservices. In *Proceedings of the 2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*, 2016.
- [63] M. A. Jaro. Advances in record-linkage methodology as applied to matching the 1985 census of tampa, florida. 84(406):414–420, 1989.
- [64] M. Jiang, M. A. Munawar, T. Reidemeister, and P. A. S. Ward. Dependency-aware Fault Diagnosis with Metric-correlation Models in Enterprise Software Systems. In *Proceedings of 2010 International Conference on Network and Service Management (NSDI)*, 2010.
- [65] S. Khatua, A. Ghosh, and N. Mukherjee. Optimizing the Utilization of Virtual Resources in Cloud Environment. In *Proceedings of the 2010 IEEE International Conference on Virtual Environments, Human-Computer Interfaces and Measurement Systems (CIVEMSA)*, 2010.
- [66] M. Kim, R. Sumbaly, and S. Shah. Root Cause Detection in a Service-oriented Architecture. In *Proceedings of the ACM SIGMETRICS/international conference on Measurement and modeling of computer systems (SIGMETRICS)*, 2013.
- [67] G. Kollios, D. Gunopulos, N. Koudas, and S. Berchtold. Efficient Biased Sampling for Approximate Clustering and Outlier Detection in Large Data Sets. In *Proceedings of the 2003 IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 2003.
- [68] D. R. Krishnan, D. L. Quoc, P. Bhatotia, C. Fetzer, and R. Rodrigues. IncApprox: A Data Analytics System for Incremental Approximate Computing. In *proceedings of International Conference on World Wide Web (WWW)*, 2016.
- [69] J. M. Liu. Nonlinear Time Series Modeling Using Spline-based Nonparametric Models. In *Proceedings of the 15th American Conference on Applied Mathematics (AMATH)*, 2009.
- [70] R. Lomax and D. Hahs-Vaughn. *Statistical Concepts: A Second Course, Fourth Edition*. Taylor & Francis, 2012.
- [71] T. Lorigo-Botran, J. Miguel-Alonso, and J. A. Lozano. A Review of Auto-scaling Techniques for Elastic Applications in Cloud Environments. In *Proceedings of Grid Computing (CGGrid)*, 2014.
- [72] J. Mace, R. Roelke, and R. Fonseca. Pivot tracing: Dynamic causal monitoring for distributed systems. In *Proceedings of the 25th Symposium on Operating Systems Principles (SOSP)*, 2015.
- [73] S. Meng and L. Liu. Enhanced Monitoring-as-a-service for Effective Cloud Management. In *Proceedings of IEEE Transactions on Computers (TC)*, 2013.
- [74] R. T. Ng and J. Han. Efficient and Effective Clustering Methods for Spatial Data Mining. In *Proceedings of the 19th International Conference on VERY LARGE DATA BASES (VLDB)*, 1994.
- [75] C. H. Papadimitriou, P. Raghavan, H. Tamaki, and S. Vempala. Latent semantic indexing: A probabilistic analysis. *Journal of Computer and System Sciences*, 2000.
- [76] J. Paparrizos and L. Gravano. k-Shape: Efficient and Accurate Clustering of Time Series. In *Proceedings Of the 2016 ACM SIGMOD/PODS Conference (SIGMOD)*, 2016.
- [77] D. L. Quoc, M. Beck, P. Bhatotia, R. Chen, C. Fetzer, and T. Strufe. PrivApprox: Privacy-Preserving Stream Analytics. In *Proceedings of the 2017 USENIX Annual Technical Conference (USENIX ATC)*, 2017.
- [78] D. L. Quoc, R. Chen, P. Bhatotia, C. Fetzer, V. Hilt, and T. Strufe. StreamApprox: Approximate Computing for Stream Analytics. In *Proceedings of the International Middleware Conference (Middleware)*, 2017.
- [79] J. Rao, X. Bu, C.-Z. Xu, L. Wang, and G. Yin. VCONF: A Reinforcement Learning Approach to Virtual Machines Auto-configuration. In *Proceedings of the 6th International Conference on Autonomic Computing (ICAC)*, 2009.
- [80] P. Reynolds, C. E. Killian, J. L. Wiener, J. C. Mogul, M. A. Shah, and A. Vahdat. Pip: Detecting the unexpected in distributed systems. In *Proceedings of the conference on Networked systems design & implementation (NSDI)*, 2006.

- [81] P. J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53 – 65, 1987.
- [82] N. Roy, A. Dubey, and A. Gokhale. Efficient Autoscaling in the Cloud Using Predictive Models for Workload Forecasting. In *Proceedings of the 2011 IEEE 4th International Conference on Cloud Computing (CCIS)*, 2011.
- [83] D. Sharma, R. Poddar, K. Mahajan, M. Dhawan, and V. Mann. Hansel: Diagnosing Faults in OpenStack. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, 2015.
- [84] 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. Technical report, Google, 2010.
- [85] G. Tesauro, N. K. Jong, R. Das, and M. N. Bennani. A Hybrid Reinforcement Learning Approach to Autonomic Resource Allocation. In *Proceedings of the 2006 IEEE International Conference on Autonomic Computing (ICAC)*, 2006.
- [86] J. Thalheim, P. Bhatotia, and C. Fetzer. Inspector: Data Provenance using Intel Processor Trace (PT). In *proceedings of IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2016.
- [87] J. Thalheim, A. Rodrigues, I. E. Akkus, P. Bhatotia, R. Chen, B. Viswanath, L. Jiao, and C. Fetzer. Sieve: Actionable Insights from Monitored Metrics in Distributed Systems. In *Proceedings of Middleware Conference (Middleware)*, 2017.
- [88] N. X. Vinh, J. Epps, and J. Bailey. Information Theoretic Measures for Clusterings Comparison: Is a Correction for Chance Necessary? In *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*, 2009.
- [89] A. Wieder, P. Bhatotia, A. Post, and R. Rodrigues. Brief Announcement: Modelling MapReduce for Optimal Execution in the Cloud. In *Proceedings of the 29th ACM SIGACT-SIGOPS symposium on Principles of Distributed Computing (PODC)*, 2010.
- [90] A. Wieder, P. Bhatotia, A. Post, and R. Rodrigues. Conductor: Orchestrating the Clouds. In *Proceedings of the 4th international workshop on Large Scale Distributed Systems and Middleware (LADIS)*, 2010.
- [91] A. Wieder, P. Bhatotia, A. Post, and R. Rodrigues. Orchestrating the Deployment of Computations in the Cloud with Conductor. In *Proceedings of the 9th USENIX symposium on Networked Systems Design and Implementation (NSDI)*, 2012.
- [92] L. Yazdanov and C. Fetzer. Lightweight Automatic Resource Scaling for Multi-tier Web Applications. In *Proceedings of the 2014 IEEE 7th International Conference on Cloud Computing (CLOUD)*, 2014.
- [93] Q. Zhang, L. Cherkasova, and E. Smirni. A Regression-Based Analytic Model for Dynamic Resource Provisioning of Multi-Tier Applications. In *Proceedings of the Fourth International Conference on Autonomic Computing (ICAC)*, 2007.
- [94] S. Zhou, A. Zhou, J. Cao, J. Wen, Y. Fan, and Y. Hu. Combining Sampling Technique With DBSCAN Algorithm for Clustering Large Spatial Databases. In *Proceedings of the 4th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2000.