



Obscured by the cloud: A resource allocation framework to model cloud outage events



Jonathan Dunne^{a,*}, David Malone^a

Hamilton Institute, Maynooth University, Kildare, Ireland

ARTICLE INFO

Article history:

Received 23 November 2016

Revised 19 April 2017

Accepted 10 June 2017

Available online 12 June 2017

Keywords:

Outage simulation

Resource allocation model

Queueing theory

Cloud computing

ABSTRACT

As Small Medium Enterprises (SMEs) adopt Cloud technologies to provide high value customer offerings, uptime is considered important. Cloud outages represent a challenge to SMEs and micro teams to maintain a services platform. If a Cloud platform suffers from downtime this can have a negative effect on business revenue. Additionally, outages can divert resources from product development/delivery tasks to reactive remediation. These challenges are immediate for SMEs or micro teams with a small levels of resources. In this paper we present a framework that can model the arrival of Cloud outage events. This framework can be used by DevOps teams to manage their scarce pool of resources to resolve outages, thereby minimising impact to service delivery. We analysed over 300 Cloud outage events from an enterprise data set. We modelled the inter-arrival and service times of each outage event and found a Pareto and a lognormal distribution to be a suitable fit. We used this result to produce a special case of the G/G/1 queue system to predict busy times of DevOps personnel. We also investigated dependence between overlapping outage events. Our predictive queuing model compared favourably with observed data, 72% precision was achieved using one million simulations.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

Cloud outage prediction and resolution is an important activity in the management of a Cloud service. Recent media reports have documented cases of Cloud outages from high-profile Cloud service providers ([The 10 biggest cloud outages](#)). During 2016 alone the CRN website has documented the ten highest profile Cloud outages to have occurred so far. Due to the increasingly complex nature of data centre infrastructure, coupled with the rapid continuous delivery of incremental software updates, it seems that Cloud outages are with us for the time being.

For operations teams that maintain a Cloud infrastructure, they rely on state of the art monitoring and alert systems to determine when an outage occurs. Examples of monitoring solutions include: [New relic \(2017\)](#), IBM Predictive Insights ([IBM operations analytics, 2017](#)) and [Ruxit \(2017\)](#). Once a new outage is observed, depending on the outage type (e.g. software component, infrastructure, hardware etc.), additional relevant experts may be called to remediate the issue. The time taken to resolve the issue may depend on a number of factors: ability to find the relevant expert, swift problem diagnosis and velocity of pushing a fix to production systems.

Both SMEs and micro teams within large organisations face a number of challenges when adopting a Cloud platform and a mechanism to deliver products and services. A number of recent studies have outlined that both frequency and duration of outage events are key challenges ([The 10 worst cloud outages, 2015](#)). Almost all European SMEs (93%) employ fewer than ten people ([Muller et al., 2015](#)). Ensuring that adequate skills and resources are available to accommodate incoming outage events is highly desirable.

In this paper we propose a framework that micro teams or SMEs can leverage to best manage their existing resource pool. We chose to base this framework on queueing theory, we believe other similar techniques such as renewal-reward theory might provide similar results.

The core idea of this framework is for operations teams to use the G/G/1 queue to model the inter-arrival and service times of outage events. This article consists of a study of outage event data from a large enterprise dataset. By modelling both inter-arrival and service outage times, a special case of the G/G/1 queue is developed. This G/G/1 queue is then tested against an off-the-shelf queue model (M/M/1) to compare and contrast queue busy time

* Corresponding author.

E-mail address: jonathan.dunne.2015@mumail.ie (J. Dunne).

prediction.¹ Finally, our framework considers dependence between overlapping outage events.

Downtime is bad for business. Whether a company provides a hosting platform, more commonly known as Platform as a Service (PaaS), or for a company that consumes such a platform to deliver their own services, more commonly known as Software as a Service (SaaS). The end result is the same: business disruption, lost revenue and recovery/remediation costs. A recent US study looked at the cost of data centre downtimes and calculated the mean cost to be \$5617 per minute of downtime ([Calculating the cost of data center outages, 2011](#)).

In the current literature, a framework to model Cloud outage events is absent. This study observed that outage events arrive over a period of time and these outages require effort from the team to return a system to a steady state. With these attributes in mind, our literature search focuses on queuing theory and distribution fitting for repairable systems.

Another consideration is the idea of event dependence. Typical off-the-shelf single-server queue models such as M/M/1 and G/G/1 assume that the inter-arrival and service times between events are independent. However if some form of dependence is found between events how useful would a queuing model that assumes independence compare against that of a queuing model with dependence properties.

This study aims to answer a number of questions. First, how are the inter-arrival times of Cloud outage events distributed? Second, how are the service times of Cloud outage events distributed? Third, how can an effective queuing model be built to simulate outage event traffic? Fourth, how are inter-arrival and service times correlated? Fifth, are overlapping outage events related or can we treat each event as independent?

We are further motivated by recent reports and studies into the adoption of Cloud computing. [Carcary et al. \(2014\)](#); [Alshamaila et al. \(2013\)](#) and [Oliveira et al. \(2014\)](#) all conducted studies on Cloud computing adoption by SMEs. The general consensus is that there is no single factor impeding Cloud adoption. The main constraints noted were: Security/compliance/data protection along with geo-restriction and compatibility to legacy systems. It was also noted that manufacturing and services sectors have different concerns in relation to Cloud computing adoption. All of this indicates that some level of customisation is required to meet needs across different industries. Additionally, ([Gholami et al., 2016](#)) provided a detailed review of current Cloud migration processes. One of the main migration concerns mentioned was the unpredictability of a Cloud environment. Factors that led to this unpredictability included: network outages and middleware failures. The study concluded that a fixed migration approach is not possible to cover all migration scenarios due to architecture heterogeneity.

The rest of this article is divided up into the following sections: [Section 2](#) introduces background and related work; [Section 3](#) discusses the dataset collected (and associated study terminology), and outlines the questions that are answered by this study and the limitations of the dataset; [Section 4](#) outlines the experimental approach and associated results; [Section 5](#) discusses the results of our experiments; Finally, in [Section 6](#), we conclude this paper and discuss future work.

2. Background and related work

The following section provides some background information on two common Cloud services: SaaS and PaaS. We then review high-profile Cloud outages that have made media headlines. Finally this section concludes with a in-depth look at relevant stud-

ies in the field of repairable systems modelling, queuing theory and Cloud outage studies.

2.1. Queueing theory

Queueing theory is the study of events that form waiting lines or queues. In queuing theory, a model is constructed so that queue lengths, inter-arrival and service times can be predicted ([Kleinrock, 1975](#); [Gross, 2008](#); [Sundarapandian, 2009](#)). We chose to use a queue with a single server as a model for Cloud outages, we believe that small teams are likely to act like a single server, and if multiple failures occur the team will have to resolve each failure, so work is queued. Naturally, this is an idealised approximation to real operating conditions.

Renewal-reward processes ([Blackwell et al., 1948](#); [Jewell, 1967](#)) were also considered as an alternative model. Queueing theory and renewal-reward processes are closely related techniques. For example, the time between successive outages on a Cloud system could be assigned as the holding times in renewal-reward theory or inter-arrival times in queueing theory. Either are likely to provide useful results, however an aspect of this study is to understand both outage times and backlogged work for a team, so queueing theory provides a suitable tool to aid resource planning of personnel within a micro team or SME. The applications of a renewal reward process might more naturally apply to the age of a software component, the remaining lifetime of a software component, and the replacement reward/penalty of a software component.

2.2. Software as a service

SaaS is defined as a delivery and licensing model in which software is used on a subscription basis (e.g. monthly, quarterly or yearly) and where applications or services are hosted centrally ([SaaS, 2015](#)). The key benefits for software vendors are the ability for software to be available on a continuous basis (on-demand) and for a single deployment pattern to be used. It is this single deployment pattern that can greatly reduce code validation times in pre-release testing, due to the homogeneous architecture. Central hosting also allows for rapid release of new features and updates through automated delivery processes ([From Google to Amazon, 2015](#)).

SaaS is now ubiquitous, while initially adopted by large software vendors (e.g. Amazon, Microsoft, IBM, Google and Salesforce) many micro teams and SMEs are now using the Cloud as their delivery platform and licensing model of choice ([Pole position, 2015](#)).

2.3. Platform as a service

PaaS is defined as a delivery and platform management model. This model allows customers to develop and maintain Cloud-based software and services without the need for building and managing a complex Cloud-based infrastructure.

The main attraction of PaaS is that it allows micro teams and SMEs to rapidly develop and deliver Cloud-based software and services. While focusing on their core products and services micro teams and SMEs are less distracted by having to design, build and service a large complex Cloud-based infrastructure.

However one drawback of PaaS is that a micro team or SME may not have a full view of the wider infrastructure. Therefore if an outage event occurs at an infrastructure level (e.g. network, load balancer) a micro team or SME may be unaware of the problem until the issue is reported by a customer.

Many companies now offer PaaS as their core service. Once seen as the preserve of a large organisation (e.g. Amazon EC2, Google Apps and IBM Bluemix) a number of smaller dedicated

¹ Code is available to simulate the queue and compare G/G/1 to M/M/1.

Table 1
Summary of high-profile Cloud outages in the first half of 2016.

Company	Duration	Date	Outage Details
Office 365	Several days	18th Jan	Users reported issues accessing their Cloud-based mail services. The defect was identified and a software fix was applied. This fix proved unsuccessful. Thereafter, a secondary fix was developed and applied that was successful.
Twitter	8 h	19th Jan	Users experienced general operational problems after an internal software update with faulty code was applied to the production system. It took Twitter 8 h to debug and remediate the defective code.
Salesforce	10 h	3rd March	European Salesforce users had their services disrupted due to a storage problem in their EU Data Centre. After the storage issue was resolved, users reported performance degradation.
Symantec	24 h	11th April	A portal to allow customers to manage their Cloud security services became unavailable. The exact nature of the outage was undisclosed. Symantec were required to restore and configure a database to bring the system back online.
Amazon	10 h	4th June	Local storms in Australia caused Amazon Web Services to lose power. This resulted in a number of EC2 instances to fail, which affected both SaaS and PaaS customers.

companies also offer PaaS (e.g. Dokku, OpenShift and Kubernetes) (PaaS, 2016).

2.4. High-profile Cloud outages

A Cloud outage is characterised the amount of time that a service is unavailable to the customer. While the benefits of Cloud systems are well known, a key disadvantage is that when a Cloud environment becomes unavailable it can take a significant amount of time to diagnose and resolve the problem. During this time the platform can be unavailable for customers.

One of the first Cloud outages to make the headlines was the Amazon outage in April 2011. In summary, the Amazon Cloud experienced an outage that lasted 47 h. The root cause of the issue was a configuration change made as part of a network upgrade. While this issue would be damaging enough for Amazon alone, a number of consumers of Amazon's Cloud platform (e.g. Reddit, Foursquare) were also affected (The 10 worst cloud outages, 2015).

Dropbox experienced two widespread outages during 2013 (Dropbox outage represents first major cloud outage, 2013; Dropbox currently experiencing widespread service outage, 2013). During the first, in January, users were unable to connect to the service. It took Dropbox 15 h to restore a full service. No official explanation as to the nature of the outage was given. The second occurred in May. Again users were unable to connect to the service. This outage lasted a mere 90 min. Again, no official explanation was provided.

Table 1 provides a summary of high-profile Cloud outages observed up to June 2016 (The 10 biggest cloud outages, 2016). While great improvements have been made in relation to redundancy, disaster recovery and ring fencing of critical services, the big players in Cloud computing are not immune to outages.

2.5. Cloud outage and software reliability studies

A number of studies have been conducted in relation to Cloud outages. Additionally, research has been carried out on the observed time to service problems in repairable systems. These studies are discussed below.

Yuan et al. (2014) performed a comprehensive study of distributed system failures. Their study found that almost all failures could be reproduced on reduced node architecture and that performing tests on error handling code could have prevented the majority of failures. They conclude by discussing the efficacy of their own static code checker as a way to check error-handling routines.

Hagen et al. (2012) conducted a study into the root cause of the Amazon Cloud outage on April 21st 2011. Their study concluded that a configuration change was made, to route traffic from one router to another, while a network upgrade was conducted. The backup router did not have sufficient capacity to handle the required load. They developed a verification technique to detect change conflicts and safety constraints within a network infrastructure prior to execution.

Li et al. (2013) conducted a systematic survey of public Cloud outage events. Their findings generated a framework that classified outage root causes. Of the 78 outage events surveyed, they found that the most common causes for outages included system issues (e.g. human error, contention) and power outages.

Sedaghat et al. (2015) modelled correlated failures caused by both network and power failures. As part of the study, the authors developed a reliability model and an approximation technique for assessing a service's reliability in the presence of correlated failures.

Potharaju and Jain (2013) conducted a similar study in relation to network outages, with focus on categorising intra- and inter-data centre network failures. Two key findings are that (i) network redundancy is most effective at inter-datacentre level and interface errors, (ii) hardware failures and unexpected reboots dominate root cause determination.

Bodík et al. (2012) analysed the network communication of a large-scale web application. Then proposed a framework that achieves a high fault tolerance with reduced bandwidth usage in outage conditions.

Snyder et al. (2015) conducted a study on the reliability of Cloud-based systems. The authors developed an algorithm based on a non-sequential Monte Carlo Simulation to evaluate the reliability of large scale Cloud systems. The authors found that by intelligently allocating the correct types of virtual machine instances, overall Cloud reliability can be maintained with a high degree of precision.

Kenny (1993) proposes a model to estimate the arrival of field defects based on the number of software defects found during in-house testing. The model is based on the Weibull distribution, which arises from the assumption that field usage of commercial software increases as a power function of time. If we think of Cloud outages as a form of field defect, there is much to consider in this model. For example, the arrival of Cloud outages in the field could be modelled with a power law distribution (e.g. Pareto distribution) as a starting point.

O'Connor and Kleyner (2011) propose an important thesis regarding reliability engineering. While emphasis is placed on measuring reliability for both mechanical and electrical/electronic systems, the authors do broaden their scope to discuss reliability of computer software. One aspect of interest is their discussion of the lognormal distribution and its application in modelling for system reliability with wear-out characteristics and for modelling the repair times of maintained systems.

Almog (1979) analysed repair data from twenty maintainable electronic systems to validate whether either the lognormal or exponential distribution would be a suitable candidate to model repair times. His results showed that in 67% of datasets the lognormal distribution was a suitable fit, while the exponential was unsuitable in 62% all of datasets.

Adedigba (2005) analysed the service times from a help desk call centre. Her study showed that the exponential distribution did

Table 2
Summary of dataset metrics and factors.

Metric	Value
Number of outage events	331
Data collection duration	18 months (January 2015 to June 2016)
Software components	Business Support System (BSS), collaboration, e-mail, and social
Number of Data Centres	3
Programming Language	Java
Operating System	Linux
Hardware	Intel-based mixture of physical and virtual systems.
Load Balancing	F5

not provide a reasonable fit for call centre service times. However a lognormal distribution was a reasonable fit for overall service times. Her study also showed that a phase-type distribution with three phases provided a reasonable fit for service times for specific jobs within the call centre job queue.

Alsoghayer and Djemame (2014) propose a mathematical model to predict the risk of failures with Grid- and Cloud-based infrastructures. The model uses the observed mean time to failure and the mean time to repair for prediction. The authors found that the best model to predict the time between failures is a Weibull distribution, while the repair (service) times are best modelled by a lognormal distribution.

As can be seen from the literature review, a number of studies have been conducted into Cloud outage failures and the inter-arrival/repair times of computer systems. However, there are no studies that conduct end-to-end research of outage events to build a queue model to predict the likely busy time and resource management of DevOps teams. Our proposed queue model aims to plug the gap that has been identified in the current literature.

3. Dataset and research methodology

Cloud outage studies have been shown to provide an effective way to highlight common failure patterns (The 10 worst cloud outages, 2015). In this and subsequent sections, our study will present a dataset and queuing model. Our aim is to illustrate its efficacy in modelling Cloud outage events.

Our dataset is taken from an enterprise Cloud system that details all outage events over an eighteen month period. In each case, we had access to the full outage report log. For each outage, we observed the arrival time of each event and how long each outage took to repair. With each arrival time known, the inter-arrival time between each outage was computed. We used the repair time duration as the service time.

A number of points related to the dataset are summarised in Table 2.

We note the programming language, OS, hardware and load balancing solution used as part of the Enterprise solution. These factors may have an underlying influence on the root cause of the observed outages, and consequently on the inter-arrival times and repair times. Investigation into the root cause of each outage event is beyond the scope of this article. We discuss this matter in our study limitations subsection below.

Product development follows a Continuous Delivery (CD) model whereby small amounts of functionality are released to the public on a monthly basis. This study focuses on the following aspects of the outage event data: the inter-arrival time between each outage, the time to service each outage event, and whether or not overlapping outage events are related.

The following terminology will now be defined to provide clear context. The majority of these definitions are referenced from a number of sources such as “What is DevOps?” by Loukides (2012) and “The Phoenix Project” by Kim et al. (2014).

Downtime (Outage) The term downtime is used to refer to periods when a system is unavailable. Downtime or outage duration refers to a period of time that a system fails to provide or perform its primary function.

Micro Team In information technology a micro team is typically staffed by three to five persons. Each person within the group has a specific role: developer, test engineer, automation specialist, people manager. With each skill set clearly defined, there is little to no overlap in terms of role definition. Each role within the team can be considered a single point of failure.

Tiger Team A tiger team is a group of experts assigned to investigate and/or solve technical or systemic problems.

DevOps DevOps is a practice that highlights the collaboration between software development and infrastructure personnel. DevOps may also refer to a team which has a core function to build, deploy and maintain a Cloud infrastructure.

Prior to outlining our research questions, it is useful to understand why queuing theory could be used to model Cloud outage events. Outages begin at a specific point in time. The problem is then diagnosed and serviced by tiger and DevOps teams. These characteristics are very similar to the properties of a queue system (i.e. inter-arrival times, service times and queue length).

Both micro teams and SMEs have fewer than ten employees (Executive summary, 2015), yet are adopting the Cloud as a method to deliver software and services. Given the unpredictability of Cloud infrastructure architecture, this study is required to understand whether a micro team/SME has adequate resources to manage future Cloud outage events.

This study aims to answer a number of questions. First, how are the inter-arrival times of Cloud outage events distributed? Second, how are the service times of Cloud outage events distributed? Third, how can an effective queuing model be built to simulate outage event traffic? Fourth, how are inter-arrival and service times correlated? Fifth, are overlapping outage events related, or can we treat each event as independent?

3.1. Inter-arrival time distribution

Probability distributions are used in statistics to infer how likely it is for an event to happen. In the case of Cloud outage inter-arrival times, we can analyse the data and determine which distribution is the best fit. The properties of a distribution can then be used to infer the probability of an event happening. For distribution fitting, we used the R package *fitdistrplus* (Delignette-Muller and Dutang, 2015) to fit various distributions to our dataset. To validate the efficacy of each distribution, the authors used the R package *ADGofTest* (Bellosta), which uses the Anderson-Darling goodness-of-fit test, to determine if the observed data follows a specific distribution (Anderson and Darling, 1952).

3.2. Service time distribution

This study has similar motivations for modeling Cloud outage service times. Being able to determine a probability distribution that best fits this outage event dataset is a useful exercise. By combining both inter-arrival and service time distributions, a queue system can then be built. This queue system can be used to model the arrival and service times of Cloud outage events. The approach to distribution fitting and validation is the same as described in the previous subsection.

3.3. Outage event modelling framework

Queuing models have been used previously across many sciences to simulate the arrival and service times for a collection of

events. Typically observed inter-arrival and service times are used to derive a suitable fitting distribution. Thereafter the distribution parameters (i.e. mean, rate, shape, scale etc.) are used to simulate queue traffic. Simulations allow large experiments that could not be conducted with the real system to be undertaken and for predictions of future behaviour to be made.

For this study we look at how a queuing system can be used to model Cloud outage events. Our queue simulator was developed using the C programming language. An assessment of the usefulness of such simulations is given within the context of resource management within a micro team or SME. Can such simulations provide a reasonable degree of precision to aid resource planning of DevOps/tiger teams with constrained levels of staffing?

3.4. Correlation between inter-arrival and service times

Statistical correlation is used to measure how two variables are related. For this study, we want to check if there is a relationship between the duration of inter-arrival and service times, and if so what is the level of this relationship. As many queueing models, including our G/G/1 model, assume independence of inter-arrival and service times, we look for evidence of dependence.

There are a number of tests that can be conducted to determine correlation. We shall discuss these briefly.

Fisher (1915) and Spearman (1904) ranked coefficient use a single measurement to determine the relationship between two variables. The strength of the relationship is measured between 0 (no correlation) and 1 (high correlation). Additionally, the coefficient can be positive or negative, indicating the type of relationship. Pearson's test is typically used when dealing with variables with a linear relationship while Spearman's test can be used where a relationship is monotonic (whether linear or not).

Linear regression (Galton, 1890) is a method to model the relationship between two variables where one variable is dependent and the other is independent. A hypothesis test is conducted and a p -value is computed. Depending on the size of the p -value, the hypothesis of a relationship/no relationship can be accepted or rejected.

Finally, autocorrelation (Box and Pierce, 1970) is the correlation of a variable with itself (and potentially other variables) at different points over a given time period. The test looks at the time lag between events to infer if a repeating pattern (seasonality) exists. Examining the lags of variables can be useful to determine if there are distinct cyclical patterns between variables, or if these patterns are simply noise.

For our correlation assessment we used the following functions found in the base R package: `cor.test` (Test for association/correlation between paired samples), `lm` (Fitting linear models) and `acf` (Auto and cross-covariance) to test the relationship between inter-arrival and service times.

Correlation tests can also be used to determine dependence between variables, however we shall discuss a specific aspect of dependence in Section 3.5.

3.5. Assessment for no association and linkage between overlapping outage events

As mentioned earlier, the M/M/1 and G/G/1 queuing system assumes that arrivals are independent. This is due to the understanding that both arrival and service times are governed by a Poisson process. We expect that the occurrence of cascading (i.e. dependent) outage events will play a role in real outage events. Therefore, for the final part of our statistical analysis tests whether overlapping outage events are independent or not.

The following method will be used. First, defect outage reports will be analysed to determine if an arrival overlaps with the ser-

vice time of a prior outage event. Next, the outage reports will be examined to determine if the two overlapping outages are related by component area and root cause. The outage counts will then be arranged in a 2 by 2 contingency table format. Fisher's exact test for independence (Fisher, 1922; 1925) is then conducted. For the actual test the authors used the R library `fisher.test` (R package `fisher.test`).

3.6. Study limitations and threats to validity

The dataset has a number of practical limitations, which are now discussed. The event data collected for this study is comprised of outage reports from an enterprise system deployed over three data centres. For the purposes of event modelling, the authors have assumed a simple queue: in other words, a queuing system with one "server". Given the lack of studies in the area of modelling Cloud outage events within a queuing framework, the authors wanted to validate such a framework in the context of a simple queue initially.

Bearing in mind that an outage of a specific type (e.g. hardware, network, high availability) will typically affect a single data centre at a time, there is a class of outage (e.g. software, configuration) that may effect multiple systems concurrently and in parallel. Our queue model could be extended to a queue with multiple servers using the following approach. The inter-arrival and service times for each data centre could be modelled separately. Using this result we could obtain a per data centre queue busy time. For outages that affect multiple data centres, we can model these events in isolation to determine the probability of such an event occurring. Next, we can analyse within the context of an overlapping event within the entire multi-server data centre infrastructure.

For the M/M/1 simulation, in the absence of a suitably good fitting mean parameter we used the means from our inter-arrival and service time distributions.

The outage events that form part of this study are from an enterprise Cloud system. We recognise that the outage events observed, can have a variety of root causes (e.g. configuration-manual, concurrency-contention, hardware, high availability and network). A number of concurrency-contention outage events may be tied to the code base which is specific to this Cloud system. However, concurrency and contention issues can occur in distributed computing systems irrespective of the underlying software code.

The hardware used in the Cloud environment is Intel based (i.e. x86_64), therefore, any hardware failures will be applicable to this hardware form factor. That said, this is a common hardware type. Additionally, the programming language used to code the software services is Java. Any outage events that are related to a software defect should be compared to similar Cloud outage data where Java is used.

The outage events are applicable to the software domain of BSS, collaboration, e-mail and social applications. As a consequence, the analysis may not be relevant outside of this realm.

4. Results

The results of the study are now discussed. This section shall follow the same format as the methodology section.

4.1. Inter-arrival time distribution

Table 3 shows a summary of the seven distributions fitted against the observed inter-arrival time data. Each distribution is listed along with its corresponding Anderson Darling test statistic and p -value. Fig. 1 shows four goodness-of-fit plots for a fitted Pareto distribution: Density, Cumulative Distribution Function

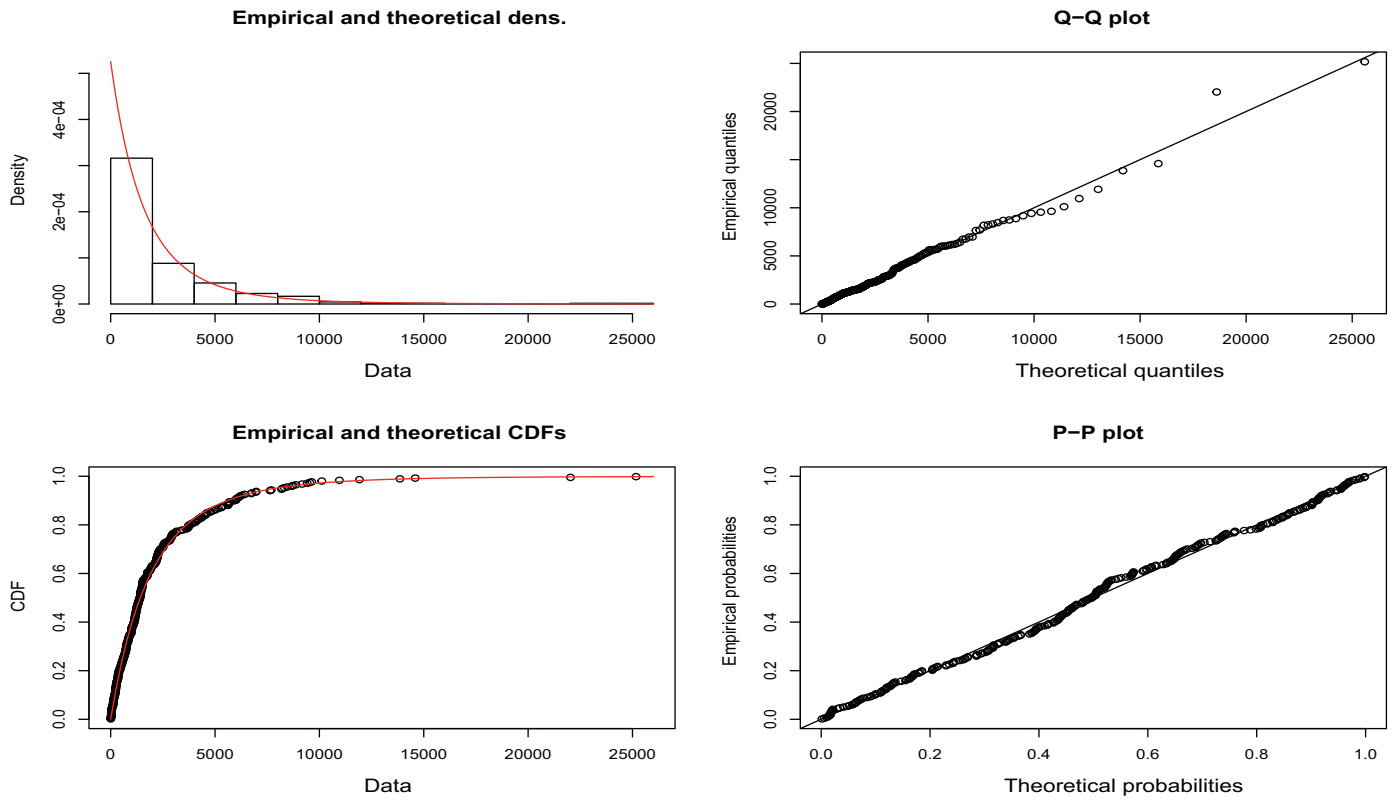


Fig. 1. Density, CDF, P-P and Q-Q plots for a fitted Pareto Distribution against inter-arrival time data.

Table 3
Inter-arrival time distributions : goodness-of-fit summary.

Distribution	AD Test Statistic	<i>p</i> -value
Pareto	0.53	0.72
loglogistic	1.93	0.10
lognormal	3.79	0.01
gamma	632.89	1.83e-06
exponential	Infinity	1.83e-06
logistic	Infinity	1.83e-06
weibull	Infinity	1.83e-06

Table 4
Service time distributions : goodness-of-fit summary.

Distribution	AD Test Statistic	<i>p</i> -value
lognormal	0.34	0.90
loglogistic	0.74	0.53
Pareto	1.60	0.15
weibull	6.82	4.00e-04
gamma	272.44	1.83e-06
exponential	Infinity	1.83e-06
logistic	Infinity	1.83e-06

(CDF), Probability (P-P) (Gibbons and Chakraborti, 2011) and Quantile (Q-Q) (Wilk and Gnanadesikan, 1968).

To answer the question of how the inter-arrival times of Cloud outage events are distributed, seven continuous distributions were fitted against the dataset. To test the goodness of fit, an Anderson-Darling (AD) goodness-of-fit test was conducted against each probability distribution.

Our hypothesis question asks: does the dataset follow a specified distribution? We then use a corresponding *p*-value to test if the dataset comes from a chosen distribution. If the *p*-value is less than 0.05, we can reject the hypothesis as being unlikely (i.e. our dataset is unlikely to follow a specific distribution). If the *p*-value is greater than 0.05, we have not found a strong statistical reason to reject the hypothesis, and can continue to consider the hypothesis plausible (i.e. our data is likely to follow a specific distribution).

With the exception of Pareto and loglogistic distribution, all others were a poor fit, indicated by the low *p*-value and the very large AD test statistic. Pareto was found to be the best fit, with an AD test statistic of 0.53 and a *p*-value of 0.72. It is worth noting that as the AD test statistic becomes large, the corresponding *p*-value remains fixed, which explains why the four worst fitting distributions have identical *p*-values.

Fig. 1 graphically illustrates how well the Pareto distribution fits our dataset. The Q-Q plot shows that the majority of data fits the distribution model line, with the exception of a number of large quantiles residing outside the model line. Additionally, the P-P and CDF plots indicate that our dataset is a good fit for Pareto with the majority of points positioned along the model line/curve. By and large, all points reside on the model line with the exception of the probability values between 0.37 and 0.57. However this observation does not undermine the assumption that the Pareto distribution is a reasonable fit for our dataset.

4.2. Service time distribution

Table 4 shows a summary of the seven distributions fitted against the observed service time data. Each distribution is listed along with its corresponding Anderson Darling test statistic and *p*-value. Fig. 2 shows four Goodness-of-fit plots for a fitted lognormal distribution: Density, (CDF), (P-P) and (Q-Q).

For the second research question of how the service times of Cloud outage events distributed, again seven continuous distributions were fitted against the dataset. Using the same method for inter-arrival times, an AD goodness-of-fit test statistic and a *p*-value were computed for each distribution. We also pose the same

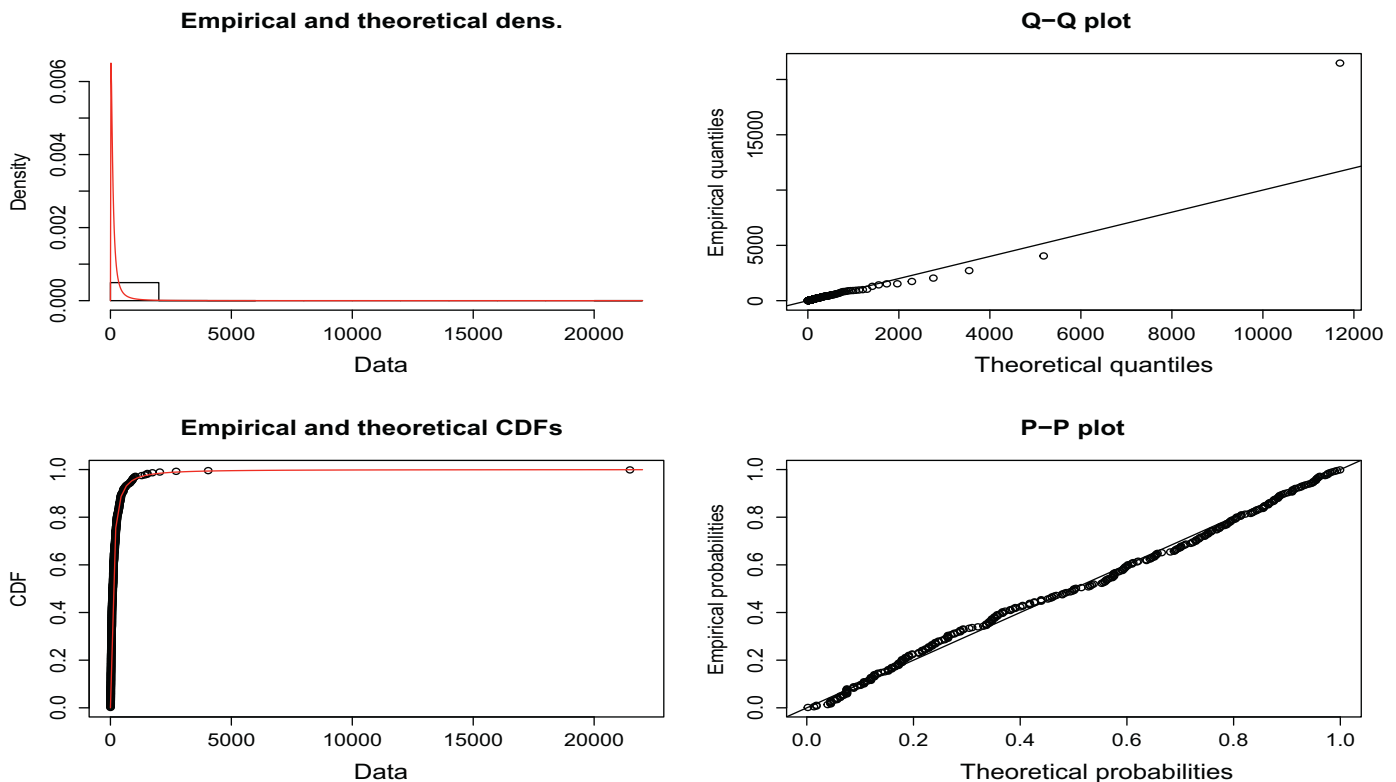


Fig. 2. Density, CDF, P-P and Q-Q plots for a fitted lognormal Distribution against service time data.

hypothesis question: Does the dataset follow a specified distribution?

Both loglogistic and Pareto scored well, however lognormal was found to be the best fitting, with an AD test statistic of 0.34 and a p -value of 0.90. All other distributions had a p -value of ≤ 0.15 . Once again, we can see that as the AD test statistic becomes large, the corresponding p -values become fixed around a value of $1.83e-06$.

The plots contained in Fig. 1 show how the lognormal distribution is a good fit to our dataset. For the Q-Q, plot the majority of values fit the distribution line. That said, there are a very small number of values that stray from the line, with one obvious extreme value. By and large the fit is very good. Additionally, for the P-P plot, the values from our dataset either reside on or very close to the line which illustrates the quality of fit.

4.3. Outage event modelling framework

Now that we have shown the results of distribution fitting, we shall use these distributions to test the special case of our G/G/1 queue model for 1M events. Our parameters were based on fitted sample distributions. For the Pareto distribution, our rate and shape parameters were computed to be 4.94 and 9404.06 respectively. Our lognormal service distribution had computed location and scale parameters of 4.58 and 1.30 respectively. We also conducted the same number of simulations against an M/M/1 queue. For the mean inter-arrival and service times we used the computed means from our Pareto and lognormal distributions.

Table 5 shows a summary of the queue model experiments conducted as well as details of the observed outage data over an eighteen month period. The model type defines whether observed data or a simulation was used. The type of simulation is also included. The % Busy and % Free columns relate to the number of overlapping events in the queue. Specifically, we counted the number of times an outage event (either observed or simulated) entered the

Table 5

Summary of results from queue modelling experiments and observed overlapping outage events.

Model Type	% Busy	% Free
Observed Data	7.9	92.1
Simulation (G/G/1)	5.7	94.3
Simulation (M/M/1)	3.0	97.0

queue system while an existing outage was currently being serviced. This value is presented as an overall percentage.

As we can see from Table 5, for the observed data the queue was free approximately 92% of the time (i.e. either 0 or 1 outage was being served), and the queue was busy approximately 8% of the time (i.e. while an outage was being served another outage event arrived). Comparing the results of both simulations: the G/G/1 simulation compared favourably with the observed results with approximately 94% and 6% free and busy time. However the M/M/1 simulation compared less well, with 97% and 3% free and busy time. Clearly, the G/G/1 model gives a better prediction than the M/M/1 model. However, the model is still a little optimistic in terms of its forecasting of busy and free times.

4.4. Correlation between inter-arrival and service times

Fig. 3 shows the results of the autocorrelation test between the sequences of inter-arrival and service times. Starting with the inter-arrival times, we can see that the lags at positions 0, 25 and 240 respectively cross the confidence interval. With only three values passing the confidence line (and the lag at position 0 being expected), clearly there is little evidence of a seasonality in the values of inter-arrival times. For service times, we observed a number of lags outside the confidence interval at positions 0–10, 100–120 and 160. While the correlation at lag 0 is expected, there is weak evidence of seasonality for lower and middle values of ser-

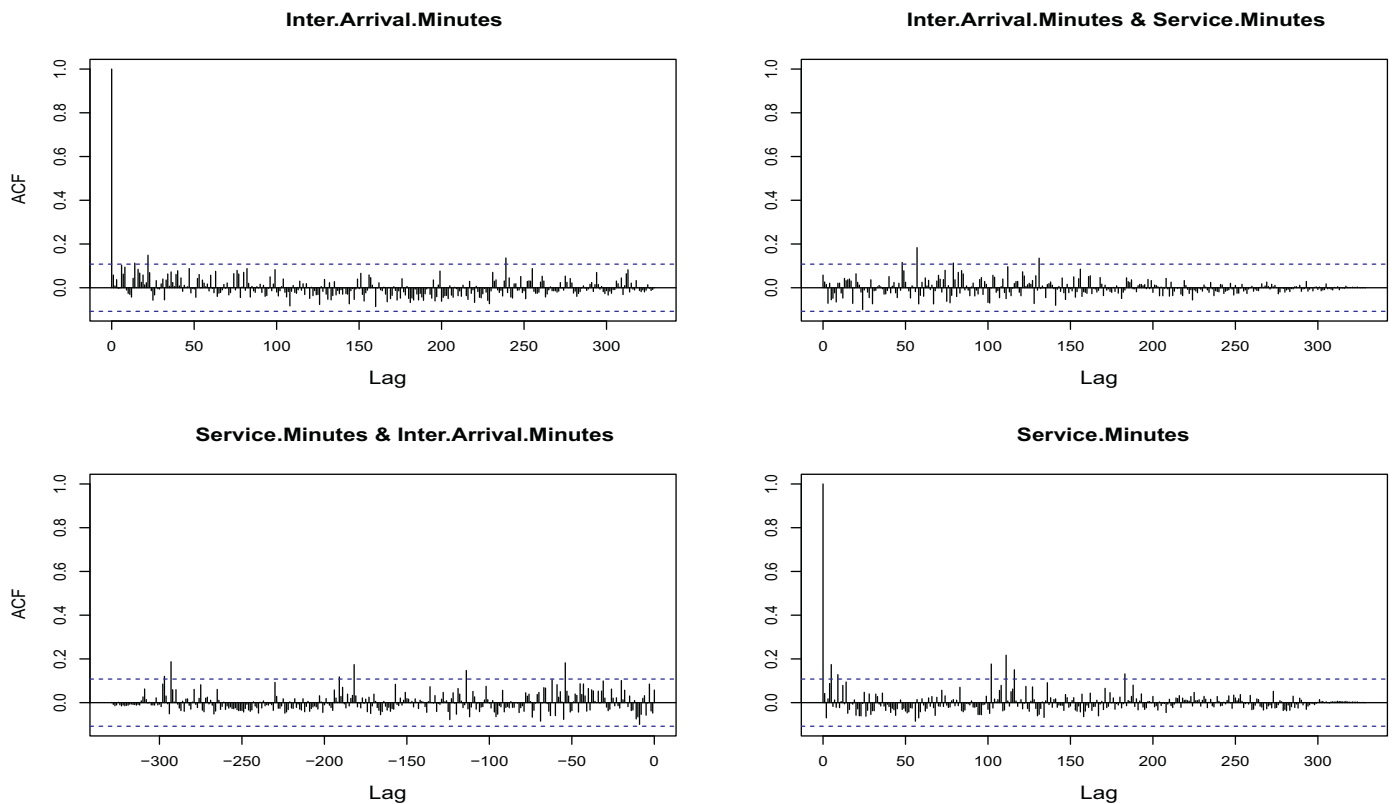


Fig. 3. Autocorrelation plots for inter-arrival and service times.

vice times. Finally, looking at the graphs of both inter-arrival and service times, we can see two lags at approximately positions 53 and 140 passing the confidence line. This suggests there is weak evidence of association between the two variables.

In addition to looking at the autocorrelation plots, we also performed tests on corresponding pairs of inter-arrival times and service times to determine their relationship. Both Pearson and Spearman tests of correlation were executed. R^2 values were computed as 0.06 (Pearson) and 0.06 (Spearman). These results indicate there is a minute positive correlation between inter-arrival and service times.

Finally, we ran a linear regression test using inter-arrival times as the dependent variable and service times as the independent variable. Our hypothesis states: there is no association between inter-arrival and service times. A p -value of 0.297 was computed. This does not give us a strong reason to reject our hypothesis.

In summary, it appears that inter-arrival and service times are at most weakly correlated, and so match our queueing assumptions relatively well.

4.5. Assessment for no association and linkage between overlapping outage events

Analysis of the inter-arrival times between each of the 331 outages was conducted to determine how many outage events overlapped. In other words, if an outage was currently being serviced by a DevOps resource, did a subsequent outage occur, and if so were these overlapping events linked. Our analysis found 26 overlapping outage events. We inspected each outage report to determine if there was a link between these outages and outages already in the queue. As part of this study we looked at the component affected and the root cause to determine whether a link between events was present.

We found evidence of a link (i.e. a common failure pattern) between 7 overlapping outages. It is worth noting that in 4 cases a temporal network outage was the root cause. In 5 cases, the e-mail component was the component affected. While no formal regression analysis was conducted, we can conjecture that there is a correlation between network failures and the e-mail component. Table 6 contains additional analysis of this work.

Table 7 shows a 2-by-2 contingency table that contains counts of overlapping, non-overlapping, linked and non-linked outage events. Fishers exact test was carried out on the table data. Our null hypothesis states that there is no association between overlapping and linked outages. A p -value of <0.001 was calculated. Given the low p -value, we can reject the null hypothesis. In other words, based on our observations, there is evidence to suggest that overlapping outages are linked to a common failure event.

5. Discussion

Section 4 presented the results of distribution fitting, queue modelling, tests for correlation between inter-arrival and service times and tests for no association between overlapping outage events. The following section provides deeper analysis and discussion of these results. In each section, references will be made to each research question asked in Section 3.

Prior to a detailed discussion of our results, we summarise the results along with each corresponding research question. Table 8 provides this summary.

5.1. Inter-arrival time distribution

The results section has shown that the Pareto distribution is a good fit to model the inter-arrival times of Cloud outage events, which answers our first research question.

Table 6
Summary details of overlapping outages with analysis of component area, root cause and linkage assessment.

Outage #	Component	Root cause	Outage details
1	E-mail	Network	Cascade network failures were observed in the e-mail component. A second network failure was observed due to latency caused by the first network failure. Assessment: outages linked.
2	E-mail	Network/Configuration	A network bottleneck was observed. A configuration change was made to alleviate the bottleneck. This change caused additional bottlenecks. Assessment: Outages linked.
3	E-mail	Concurrency	A failover operation failed to work correctly, which caused an outage. While the system was in a failed state, crash log information was not output correctly. Assessment: Outages linked.
4	Social	High availability	A number of nodes in the social component failed due to a server crash. While these nodes were down, extra load was added to the available nodes in the cluster, which caused a subsequent outage. Assessment: Outages linked.
5	E-mail	Network	A temporal network outage occurred in the e-mail system. Most of the nodes failed gracefully and returned to normal operations, however a number nodes did not fail gracefully, which caused a secondary outage. Assessment: Outages linked.
6	E-mail	Configuration	A service on the e-mail system failed due to contention. A config change was made to remediate the initial contention. The config change caused additional contention further along the service stack. Assessment: Outages linked.
7	Collaboration	Network	A temporal network outage occurred in the collaboration component, which caused all nodes to fail gracefully. Almost all nodes returned to normal when the network was restored. A number of nodes, however, were in a hung state from the initial outage, which caused a secondary outage. Assessment: Outages linked.

Table 7
Test for no association between overlapping and linked outages using Fisher's exact test.

Outage type	Linked	Non-linked
Non-overlapping	0	305
Overlapping	7	19

The decision to use Pareto as an inter-arrival time distribution is an interesting choice. The Pareto distribution is a power law distribution and has applications in many fields of science. However the field where a Pareto distribution is typically used is in the area of finance. Specifically for modelling income and wealth (Arnold, 2015).

The characteristics of our data that make Pareto so attractive is the number of values within a specific range. Inter-arrival times range from 3 to 1,057,122 min. Using 2500 min as an arbitrary point of delineation, 71% of inter-arrival times were below 2500 min, while 29% were above 2500 min. While this split does not conform to the textbook “80-20” rule (Chen et al., 1993), it does illustrate that our dataset contains a significantly higher proportion of shorter inter-arrival times than longer ones. Given this specific trait, it is not unsurprising that the Pareto distribution is such a good fit.

This study has answered our first research question: what distribution can be used to model inter-arrival times of Cloud outage events? DevOps teams can use the shape and scale parameters of their inter-arrival distribution to compute a mean and standard deviation, which can provide an expected time between outage

events. Additionally, this result can be used to compute the proportion of inter-arrival times above or below a specific duration. These results can be used to aid resource planning. For example, if the expected inter-arrival time is known, or if a high proportion of outages is known to occur within a specific duration, duty rosters can be generated to ensure adequate staffing is available when an outage occurs. Finally, this result can be used as a component in a wider queue model framework to infer team busy time.

5.2. Service time distribution

We have learned from our results that the lognormal distribution is an excellent fit to model the service times of Cloud outage events recorded in our dataset.

The lognormal distribution is important in the description of natural events. Many naturally occurring processes are modelled by the culmination of incremental changes. Such processes include general system usage, vehicle mileage per year, count of switch operations and wear-out characteristics of machines and systems. This distribution is also versatile in that, depending on the location and scale parameters, a number of different distribution shapes can be accommodated.

This result adds to the wealth of existing studies that support the notion that service times for repairable systems can be modelled using a lognormal distribution. We noted previously the work done by O'Connor and Kleynner (2011). However, a number of additional recent studies have observed similar results in their studies of repairable systems, such as Apostolakis et al. (1980), Ananda (2003) and Ananda and Gamage (2004).

Table 8
Summary of research question, results and techniques used.

Research question	Results	Techniques
1. How are the inter-arrival times of Cloud outage events distributed?	Pareto distribution is the best fit. AD test statistic = 0.53 p -value = 0.72	Anderson-Darling goodness-of-fit test. Anderson and Darling (1952)
2. How are the service times of Cloud outage events distributed?	lognormal distribution is the best fit. AD test statistic = 0.34 p -value = 0.90	Anderson-Darling goodness-of-fit test. Anderson and Darling (1952)
3. How can an effective queuing model be built to simulate outage event traffic?	Our simulation achieved 72% precision in predicting queue busy/free times compared to observed data.	Queuing theory. Kleinrock (1975); Gross (2008)
4. How are inter-arrival and service times correlated?	Weak evidence of association between inter-arrival and service times. p -value = 0.297	Ranked coefficients. Fisher (1915); Spearman (1904) Linear regression. Galton (1890) Autocorrelation. Box and Pierce (1970)
5. Are overlapping outage events related, or can we treat each event as independent?	Strong evidence of association between overlapping outage events. p -value < 0.001	Fishers exact test for independence. Fisher (1922); 1925)

This study has answered our second research question: what distribution can be used to model service times of Cloud outage events? DevOps teams can employ the location and scale parameters of their service distribution to compute a mean and standard deviation, which can provide an expected duration service time. With the service times known, teams can create schedule plans to determine expected engagement times. Finally this result can be used in conjunction with the result from the previous section to model the idle and busy times of a team as part of a queue modelling exercise.

5.3. Outage event modelling framework

We asked the question: how can an effective queuing model be built to simulate outage event traffic? The result from our experimental model shows that by combining well fitted distributions for both inter-arrival and service times a model that provides a good level of precision can be built.

Table 5 provides a summary of the % busy and free time for our observed data and for two sets of simulations: G/G/1 and M/M/1 queues. It is unsurprising that M/M/1 lacks precision. There are two factors to consider here. First, neither the inter-arrival nor service distributions could be adequately modelled using an exponential distribution. We recall from Tables 2 and 3 the AD test statistic for the exponential distribution was infinite. For the purposes of the simulation we used the computed means from the Pareto and lognormal distributions. Second, the M/M/1 queue assumes that the arrival times between events are independent. We have shown that a small proportion (2%) of overlapping outage events are linked. These two factors make the M/M/1 queue unsuitable for queue simulation based on the observed data.

Conversely, the G/G/1 queue provided a greater degree of precision than the M/M/1 queue. This is due to the fact that the two distributions selected were a good fit against the observed data compared to the exponential distribution. There is still a minor lack of fidelity between our G/G/1 simulation and our observed data. We can surmise that while the goodness-of-fit for the service time distribution was excellent (p -value = 0.90), the goodness-of-fit for the inter-arrival time distribution was only very good (p -value = 0.72). Moreover, there is the question of independence between arrivals. We must conclude that a small number of dependent outages coupled with the less-than-exact fit of the inter-arrival distribution may skew the precision of our simulation. We discuss improvements to this model in future work.

Now let us look at the practical application of such a simulation model. The core idea of this paper is to produce a model that is effective in simulating the arrival of Cloud events. We have previously mentioned the challenges that both micro teams and SMEs have when working in the area of Cloud computing. One of the key challenges is the deployment of resources, and how one can position these resources where they are most needed. Let us consider the following scenario as an example of our simulation framework.

Table 9 shows the output from a G/G/1 simulation using the parameters from our dataset. There are two columns: Time (Measured in minutes) and Queue length. Let assume that we have uptime from 12:00 1st of January. Looking at the output below, we can see that we will need one resource to service the first thirteen outage events. These thirteen outages will arrive and be serviced in sixteen days. Looking at the fourteenth outage event we can see this event will arrive at approximately 11:18 on the 16th of January. DevOps Management have a good indication that two DevOps resources will be required at this time: one to service the thirteenth outage and a second resource to service the overlapping fourteenth outage. DevOps management can also infer that both resources will be required for only for a short duration. In this

Table 9

Sample output from an G/G/1 simulation.

Duration (Min)	Queue length	Date & time
2	1	2017-01-01 00:02
142	0	2017-01-01 02:22
3744	1	2017-01-03 14:24
3761	0	2017-01-03 14:41
5577	1	2017-01-04 20:57
5644	0	2017-01-04 22:04
11,043	1	2017-01-08 16:03
11,048	0	2017-01-08 16:08
14,989	1	2017-01-11 09:49
15,186	0	2017-01-11 13:06
19,566	1	2017-01-14 14:06
19,605	0	2017-01-14 14:45
22,249	1	2017-01-16 10:49
22,278	2	2017-01-16 11:18
22,286	1	2017-01-16 11:26

case, eight minutes approximately. Thereafter, one resource will be needed to debug and remediate subsequent outage events.

Another application of the queue simulation model is to assess staffing requirements over a calendar year. By knowing the duration of a year in minutes (525600), we can easily check to see how many events will occur during a calendar year. In a simulation conducted for the purposes of this example, we found the queue length was greater than 1 on 28 occasions. 27 times the queue length = 2 and once the queue length = 3. A final application is to look at the queue busy time in a given calendar year. If we add the times the queue is busy (i.e. time difference between the queue length being 1 or more and 0) for outages that occur over the period of a year, we can see that the queue will be busy for approximately 144 days. Clearly, these types of what if scenarios are very useful for resource planners.

5.4. Correlation between inter-arrival and service times

Using a myriad of tests we have answered our fourth research question: how are inter-arrival and service times correlated? Our results show that there is little evidence to suggest a correlation between inter-arrival and service times.

Fig. 3 showed graphically how both variables were correlated not only with themselves but each other. Given the low number of lags crossing the confidence interval, coupled with the sparse positioning of these lags, there is little evidence to suggest any meaningful correlation. Likewise, we saw similar results from both the Pearson, Spearman, and linear regression tests.

DevOps teams can use this result in a number of ways. An ideal goal for a Cloud-based business is to have as near to 100% uptime as possible, while ensuring that when an outage does occur, the time to service such an outage is as short as possible. In other words, having very long inter-arrival times between outage events and very short service times is highly desirable. The goal for each requires a separate solution: in the case of long inter-arrival times, ensuring that when a system does fail, it fails gracefully without any loss of service; in the case of service times, having an advanced suite of system monitoring solutions coupled with a simple system of rollback to prior code versions and/or configuration changes is key.

Given the lack of correlation between inter-arrival and service times, DevOps teams can be confident that process changes to reduce service times will not lead to a reduction in inter-arrival times. Moreover, with increased reliability brings longer inter-arrival times, this, in essence, will not lead to longer service times.

5.5. Assessment for no association and linkage between overlapping outage events

Our results section has highlighted that there is an association between outage events that overlap and outages that are linked (i.e. cascade failures).

Table 6 provides a good insight into the nature of linked outage events. We saw that in five of the seven linked outages e-mail was a common component. Likewise, we observed that network and configuration issues were the root causes in four of the seven outages. This may not be a coincidence. In one other case we saw that in a disaster recovery scenario, server node failover did not work as expected, which caused a cascade failure due to high concurrency.

Table 7 highlights that overlapping outages are uncommon, with approximately 8% of all outages recorded over an eighteen month period overlapping. Additionally, linked outages are rarer still, with only approximately 2% recorded over the same duration. However, as demonstrated by the results of Fisher's test, there is overwhelming evidence to suggest that both events are associated. Therefore overlapping outage events are related. We can calculate that when an overlapping event occurs there is approximately a 25% probability that these events are linked. Removal of these types of failures is key to the success of a business, and will lead to increased customer satisfaction by increased up time.

DevOps teams can learn from these results. Linked failures cause additional workloads for small teams. From a remediation perspective, DevOps teams can work with their software development counterparts to ensure their infrastructure and software are more resilient to temporal network outages. By conducting a series of negative tests teams can determine how gracefully their systems fail under scenarios like temporal network outages. Additionally, setting invalid parameters within a large distributed system can have knock on effects. It is worth pointing out that introducing a system of managed configuration changes (similar to developer code reviews prior to check in) can help alleviate the problems encountered with invalid configuration changes.

6. Conclusion

One purpose of this research was to examine which probability distributions could be used to best model inter-arrival and service times of outages. By using the best-fitting distributions as part of a special case of the G/G/1 queue modelling system, this study demonstrated how this model can be used to determine the busy time of a Cloud outage queue system. Additionally, this study examined the correlation between inter-arrival and service times. Furthermore, we observed whether overlapping outage events are linked.

It was found that inter-arrival and service times of Cloud outage events could be reasonably modelled with a Pareto and lognormal distribution respectively. Additionally, by using these distributions, a queue model framework could be built to infer the percentage busy time of this queue with a good degree of accuracy. Furthermore, we found no evidence of correlation between inter-arrival and service times. Finally, our research showed that there is evidence to suggest that overlapping outage events are linked.

The findings of this study support previous work specifically in the field of repair times of maintainable Cloud-based software systems. This work provides more comprehensive analysis of the inter-arrival times of Cloud outage events and how a useful special case of the G/G/1 can be developed to determine queue busy time.

However, the main application of this research is to DevOps and project planners within an SME or micro team. Both can leverage this framework to build an accurate resource planning model that can identify both skill and personnel gaps. Identification and reme-

diation of these gaps will greatly benefit teams in the challenging area of Cloud outage resolution.

By using the simple queue as a starting point, future work is planned to validate the framework in the context of a complex queuing system (e.g. a queue with multiple "servers"). (John, 1963) discusses dependencies between inter-arrival and service times within a queue system, as the assumption of independence between the two times are not always valid.

Also of interest is the wear-out characteristics of specific software components. Conducting a study of Cloud outages as part of a wider renewal process study, we can understand the mean time to failure of a given component. With the age (uptime) of a component known, how likely is such a component to fail? Finally, how often does a component need to be replaced, and what role does hygienic recycling play in system stability?

Future work will also include research into the relationship between inter-arrival and service time durations, and how these durations relate to service level agreement impact.

References

- Adedigba, A., 2005. Statistical distributions for service times. Citeseer Ph.D. thesis.
- Almog, R., 1979. A study of the application of the lognormal distribution to corrective maintenance repair time. Monterey, California. Naval Postgraduate School.
- Alshamaila, Y., Papagiannidis, S., Li, F., 2013. Cloud computing adoption by SMEs in the north east of England: a multi-perspective framework. *J. Enterp. Inf. Manag.* 26 (3), 250–275.
- Alsoghayer, R., Djemame, K., 2014. Resource failures risk assessment modelling in distributed environments. *J. Syst. Softw.* 88, 42–53.
- Ananda, M.M., 2003. Confidence intervals for steady state availability of a system with exponential operating time and lognormal repair time. *Appl. Math. Comput.* 137 (2), 499–509.
- Ananda, M.M., Gamage, J., 2004. On steady state availability of a system with lognormal repair time. *Appl. Math. Comput.* 150 (2), 409–416.
- Anderson, T.W., Darling, D.A., 1952. Asymptotic theory of certain "goodness of fit" criteria based on stochastic processes. *Ann. Math. Stat.* 193–212.
- Apostolakis, G., Garribba, S., Volta, G., 1980. Synthesis and analysis methods for safety and reliability studies, 106. Springer.
- Arnold, B.C., 2015. Pareto distribution. Wiley Online Library.
- Bellosta, C. J. G., R package adgofest.
- Blackwell, D., et al., 1948. A renewal theorem. *Duke Math. J.* 15, 145–150.
- Bodík, P., Menache, I., Chowdhury, M., Mani, P., Maltz, D.A., Stoica, I., 2012. Surviving failures in bandwidth-constrained datacenters. In: Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication. ACM, pp. 431–442.
- Box, G.E., Pierce, D.A., 1970. Distribution of residual autocorrelations in autoregressive-integrated moving average time series models. *J. Am. Stat. Assoc.* 65 (332), 1509–1526.
- Carcary, M., Doherty, E., Conway, G., 2014. The adoption of cloud computing by Irish SMEs—an exploratory study. *Electron. J. Inf. Syst. Eval.* 17, 1–12.
- Chen, Y.-S., Pete Chong, P., Tong, Y., 1993. Theoretical foundation of the 80/20 rule. *Scientometrics* 28 (2), 183–204.
- Delignette-Muller, M.L., Dutang, C., 2015. Fitdistrplus: an R package for fitting distributions. *J. Stat. Softw.* 64 (4), 1–34.
- Fisher, R.A., 1915. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika* 10 (4), 507–521.
- Fisher, R.A., 1922. On the interpretation of χ^2 2 from contingency tables, and the calculation of P. *J. R. Stat. Soc.* 85 (1), 87–94.
- Fisher, R.A., 1925. Statistical Methods for Research Workers. Genesis Publishing Pvt Ltd.
- Galton, F., 1890. Kinship and correlation. *North Am. Rev.* 150 (401), 419–431.
- Gholami, M.F., Daneshgar, F., Low, G., Beydoun, G., 2016. Cloud migration process—a survey, evaluation framework, and open challenges. *J. Syst. Softw.* 120, 31–69.
- Gibbons, J.D., Chakraborti, S., 2011. Nonparametric statistical inference. Springer.
- Gross, D., 2008. Fundamentals of Queuing Theory. John Wiley & Sons.
- Hagen, S., Seibold, M., Kemper, A., 2012. Efficient verification of it change operations or: how we could have prevented Amazon's cloud outage. In: Network Operations and Management Symposium (NOMS), 2012 IEEE. IEEE, pp. 368–376.
- Jewell, W.S., 1967. A simple proof of: $L = \lambda w$. *Oper. Res.* 15 (6), 1109–1116.
- John, F.I., 1963. Single server queues with dependent service and inter-arrival times. *J. Soc. Ind. Appl. Math.* 11 (3), 526–534.
- Kenny, G.Q., 1993. Estimating defects in commercial software during operational use. *IEEE Trans. Reliab.* 42 (1), 107–115.
- Kim, G., Behr, K., Spafford, G., 2014. The phoenix project: a novel about IT DevOps, and Helping Your Business Win.
- Kleinrock, L., 1975. Queuing Systems. Wiley.
- Li, Z., Liang, M., O'Brien, L., Zhang, H., 2013. The cloud's cloudy moment: a systematic survey of public cloud service outage. [arXiv:1312.6485](https://arxiv.org/abs/1312.6485) arXiv preprint.
- Loukides, M., 2012. What is DevOps? O'Reilly Media, Inc.

- Muller, P., Caliendo, C., Peycheva, V., Gagliardi, D., Marzocchi, C., Ramlogan, R., Cox, D., 2015. SME performance review European SME's.
- O'Connor, P., Kleyner, A., 2011. *Practical Reliability Engineering*. John Wiley & Sons.
- Oliveira, T., Thomas, M., Espadanal, M., 2014. Assessing the determinants of cloud computing adoption: an analysis of the manufacturing and services sectors. *Inf. Manage.* 51 (5), 497–510.
- Potharaju, R., Jain, N., 2013. When the network crumbles: an empirical study of cloud network failures and their impact on services. In: *Proceedings of the 4th annual Symposium on Cloud Computing*. ACM, p. 15.
- Sedaghat, M., Wadbro, E., Wilkes, J., De Luna, S., Seleznev, O., Elmroth, E., 2015. Die-hard: reliable scheduling to survive correlated failures in cloud data centers.
- Snyder, B., Ringenberg, J., Green, R., Devabhaktuni, V., Alam, M., 2015. Evaluation and design of highly reliable and highly utilized cloud computing systems. *J. Cloud Comput.* 4 (1), 1.
- Spearman, C., 1904. The proof and measurement of association between two things. *Am. J. Psychol.* 15 (1), 72–101.
- Sundarapandian, V., 2009. *Probability, statistics and queuing theory*. PHI Learning.
- Wilk, M.B., Gnanadesikan, R., 1968. Probability plotting methods for the analysis for the analysis of data. *Biometrika* 55 (1), 1–17.
- Yuan, D., Luo, Y., Zhuang, X., Rodrigues, G.R., Zhao, X., Zhang, Y., Jain, P.U., Stumm, M., 2014. Simple testing can prevent most critical failures: an analysis of production failures in distributed data-intensive systems. In: *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, pp. 249–265.
- Auto and cross-covariance and correlation function estimation. <http://bit.ly/2dKfLZl>.
- Fitting linear models. <http://bit.ly/2dvqYEt>.
- R package fisher.test. <http://bit.ly/1NU3c5y>.
- Test for association/correlation between paired samples. <http://bit.ly/2djPSA7>.
- Calculating the cost of data center outages. <http://bit.ly/2blnDgh> 2011.
- Dropbox currently experiencing widespread service outage. 2013.
- Dropbox outage represents first major cloud outage of 2013. 2013.
- The 10 worst cloud outages. 2015.
- Executive summary - final report - annual report on european smes - 2014 / 2015 - smes start hiring again. 2015.
- From Google to Amazon - the rise of the cloud catalog. 2015.
- Pole position : Ranking the top 5 IaaS, PaaS and private cloud providers. 2015.
- Why multi-tenancy is key to successful and sustainable software-as-a-service (SaaS). 2015.
- The 10 biggest cloud outages of 2016. 2016.
- Best platform as a service (PaaS). 2016.
- IBM operations analytics - predictive insights. 2017.
- New relic - application and performance monitoring. 2017.
- cloud native monitoring. 2017.