

Optimal Resource Provisioning for Data-intensive Microservices

Roland Mark Erdei, Laszlo Toka

MTA-BME Network Softwarization Research Group, Faculty of Electrical Engineering and Informatics

Budapest University of Technology and Economics, Hungary

erdeiroland@edu.bme.hu, toka.laszlo@vik.bme.hu

Abstract—With the continuous progress of cloud computing, many microservices and complex multi-component applications arise for which resource planning is a great challenge. For example, when it comes to data-intensive cloud-native applications, the tenant might be eager to provision cloud resources in an economical manner while ensuring that the application performance meets the requirements in terms of data throughput. However, due to the complexity of the interplay between the building blocks, adequately setting resource limits of the components separately for various data rates is nearly impossible. In this paper, we propose a comprehensive approach that consists of measuring the resource footprint and data throughput performance of such a microservices-based application, analyzing the measurement results by data mining techniques, and finally formulating an optimization problem that aims to minimize the allocated resources given the performance constraints. We illustrate the benefits of the proposed approach on Cortex, an extension to Prometheus for storing monitored metrics data. The data-intensive nature of this illustrative example stems from real-time monitoring of metrics exposed by a multitude of applications running in a data center and the continuous analysis performed on the collected data that can be fetched from Cortex. We present Cortex’s performance vs resource footprint trade-off, and then we build regression models to predict the microservices’ resource consumption and draw a mathematical programming formulation to optimize the most important configuration parameters. Our most important finding is the linear relationship between resource consumption and application performance, which allows for applying linear regression and linear programming models. After the optimization, we compare our results to Cortex’s recommendation, leading to a CPU reservation reduced by 50-80%.

Index Terms—cloud-native, microservices, performance, resource footprint, optimization

I. INTRODUCTION

Cloud computing is widely used in the digital industry as a technology that enables cheap and easy deployment not only for online web services but also for big data processing, machine learning, and storing data for the long term. The cloud is backed by physical data centers worldwide that host virtual machines offered to customers and users. The cloud concept empowers users to replace their hardware with cloud server instances and creates a new economic model where the customer pays only for the usage and not for the entire hardware itself. However, decision-making in cloud environments can be complex due to the diversity in pricing models and service offerings. There are no rules of thumb as each customer could

have a specific set of constraints and requirements of their cloud application when it comes to selecting the perfect cloud environment [1].

Microservices have recently gained striking popularity due to being highly maintainable and easy to develop. Microservices-based applications allow the deployment of each microservice (or component) to be in physically separated virtual machines if needed. Microservices, being separately manageable, have a massive advantage in scaling, which is one of the most critical features of the cloud context. Instead of launching multiple instances of a whole application, there is a possibility to only scale in or scale out a specific microservice component [2]. Hence creating microservices with the help of container technologies can result in robust and easy deployment with a small footprint. However, the goal is not only to create a small footprint application, but also to match the Quality of Service (QoS) of the original monolith performance-wise. When it comes to data-intensive applications, data throughput is the most crucial aspect. The components of the application communicate with each other via APIs. Wrongly configuring a microservice component where the data throughput is consequently low can result in a bottleneck in terms of the whole application’s performance. Scaling the bottleneck components until the QoS is met will mitigate the issue. The problem we tackle is how to correctly configure the resource provisioning of each and every microservice component so that the overall application performs as required.

To create a resource-efficient microservice, the most critical settings to tune are the CPU and memory limits that control the resource usage of a microservice instance. These limits can prevent components from using more resources than needed, and with the help of these limits, the app provider or operator makes sure in a private cloud that concurrent applications use the cluster as efficiently as possible, whereas, in a public cloud, limits help to pre-estimate costs. Setting the limits for several components at once raises a complex optimization problem when it comes to microservices. Application components have their tasks while working closely with other components, and if one of the components lags, the whole microservices-based application will suffer. On the other hand, resource provisioning must consider the shared resources: over-provisioning one component can cause starvation at other components. Furthermore, each component may have several configuration

parameters which need to be tested in terms of resource usage to find the optimal limit values making the specific component work as efficiently as possible.

In this paper, we tackle this exact challenge: we propose an optimization framework that estimates resource usage vs performance of microservices-based, data-intensive applications with the help of machine learning, then minimizes the allocated resources given the QoS constraints. The rest of the paper is organized as follows. Section II presents the related work, and in Section III we propose the methodology to solve the problem at hand. In Section IV we introduce an illustrative example for microservice-based applications, Cortex, then we train regression models on the performance and resource consumption measurements data and create a linear programming formula for resource optimization, completed with numerical analysis. Section V concludes our findings.

II. RELATED WORK

Performance modeling of microservices-based applications allows us to determine the capacity distribution among each microservice. This enables planning for applications and the detection of the bottleneck in microservices. [3] proposes to apply statistical models, e.g., Theil-Sen estimator or Support Vector Regression, for this purpose. After analyzing the data acquired by their approach on the example applications, it was identified that the microservices follow a typical performance versus workload relationship pattern, which suggests the performance degrading with the increase of workload up until a certain point when all the virtual resources are used. In all tests, the CPU utilization increased linearly with the number of requests sent to the microservices. The approach was tested on several test applications, including a compute-intensive application, a database accessing an application, and a web accessing application. Similarly to [3] we also want to create a model to detect and avoid bottlenecks; however, we want to provision the resources individually for each microservice using as minimal resources as possible.

Zhang et al. [4] presented Sinan, a machine learning-driven resource manager for microservice-based applications. Sinan presents the challenges of managing complex microservices and leverages a set of scalable and validated models to reduce resource usage while meeting the end-to-end QoS. Sinan trains two models with the traces: a CNN (Convolutional Neural Network) model for short-term performance prediction and a Boosted Trees model to evaluate long-term performance evolution. Combining the two models allows Sinan to be effective in both near-future and distant future resource management. Similarly to [4], our goal is to meet the QoS requirements with an analytical model which can predict the expected average resource usage.

The detection of a bottleneck component is important if we want to increase the performance of a complex application deployed in the cloud. [5] presents an analytical model that can detect bottlenecks and predict the performance of a multi-tier application. The suggested approach consists of two resource provisioning steps: a predictive one for long-term scales and

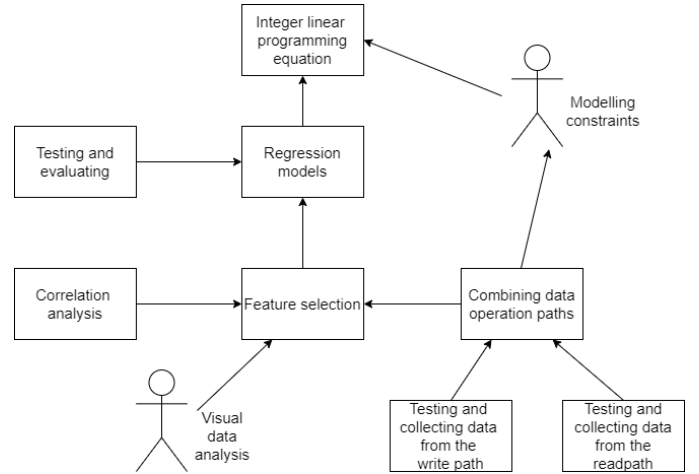


Fig. 1. Methodology diagram for the thought process.

a reactive one for short timescales. Long-term prediction is useful when the load can be predicted, e.g., a typical daily pattern, but only reactive provisioning can handle unexpected high workloads. Using both and combining them can create an effective provisioning scheme. With this approach in a scenario where the workload of a three-tier application has been doubled, the technique showcased in [5] was able to double the application capacity within five minutes while maintaining the QoS targets. Unlike [5], our focus is to predict and create resource limits for the average resource usage of microservices of any cloud-native application.

Our model shows similarities to CostHat [6], an approach to model the deployment costs, including compute and IO costs. CostHat is a graph-based model of the deployment costs and can be used for applications implemented on top of AWS Lambda. Just like our results, [6] demonstrates that between the used features, there are linear dependencies. Similarly to [6] we focus on reducing costs if the model is used in a public cloud environment. However, our model calculates CPU usage as output, which is extremely useful for private clouds.

III. METHODOLOGY AND MODEL

Our goal is to minimize the CPU limit provisioned for each component in a microservice-based application for a static load. To this end, we formulate mathematical programming problems with user-defined constraints. Then the result of the optimization yields the expected CPU usage based on which the resource limits of each microservice can be appropriately set. In order to formulate the optimization model, first, we need to find the most important features of the operation of the microservices, as well as their coefficients. Our approach is to build a regression model from a measurement dataset that incorporates the inference between configuration settings, application performance and resource footprint. Fig. 1 shows the steps of our proposed approach.

The measurement dataset contains the monitored microservices' resource usage. Assuming a data-intensive application at hand, there are typically two paths that must be monitored:

the writing path (data ingestion) and the reading path (data retrieval). The cloud-native application is typically constituted by various microservices which are closely coupled together. These microservices (or components) can be customized with configuration parameters and settings. While measuring the effects of these settings and scaling in and out the components, a data collector pipeline stores the microservices’ resource usage and performance indicators. The important features that are later used in the data mining step can be, for example, the replica number of each component, the number of incoming service requests, or the number of served queries per second. After measuring several scenarios with various settings and analyzing the key differences between each setup, we propose to train a regression model.

We create a separate regression model for every main component of the application. Feature selection is important for keeping the model simple. Most of the potential features show linear dependencies with the resource usage, as pointed out in [3] and [6], therefore in this paper, we make a case for linear relationships so that we can later formulate the optimization with linear programming. By keeping only the highly correlating features, the linear programming solvers can find the solution quickly even if some variables are integers, turning the problem into non-deterministic polynomial-time (NP) complexity [7]. After training the regression model on the measurement dataset, given the coefficients of the regression, we can continue with linear programming. Note that in case the features exhibit a non-linear effect on the resource usage or on the application performance, then polynomial regression might be used, but then the respective optimization problem will not be a linear programming instance.

The decision variables are the important features from the regression model. Using these variables, one must create the algebraic expression which describes the target function; in our case, we strive to minimize the CPU resource usage. As the regression model is trained for precisely that dependent variable, the target function is produced by the regression formula: features as variables, their coefficients and the intercept. Based on the measurement data, one can create constraints on the features to be taken into account in the optimization problem. Typical constraints include but are not limited to high availability requirements on microservice instances, maximal memory consumption, application performance in terms of data throughput or rate of requests served. In general, it is beneficial to minimize the search space as much as possible, resulting in faster optimization. Non-negativity constraints can be applied for most of the features.

With the optimization model formulated, any solver can optimize the problem with the given user inputs and suggest the best possible setup for all the features. In the next section, we illustrate how this approach can save resources for the operation of a cloud-native microservices-based application.

IV. ILLUSTRATIVE EXAMPLE OF CORTEX

Cortex [8] is a horizontally scalable, multi-tenant, long term storage for Prometheus [10] written in Go. It is a Cloud

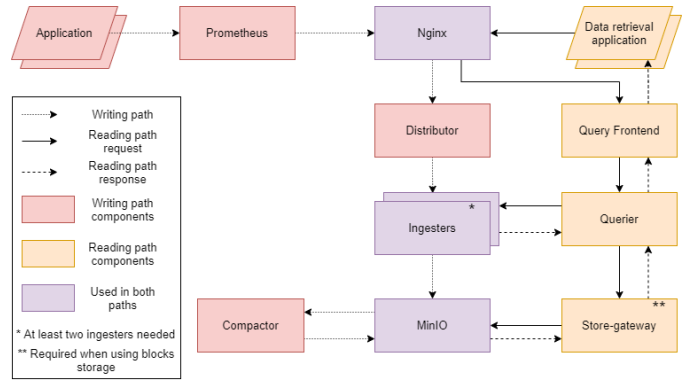


Fig. 2. Cortex’s architecture including the external applications [8]

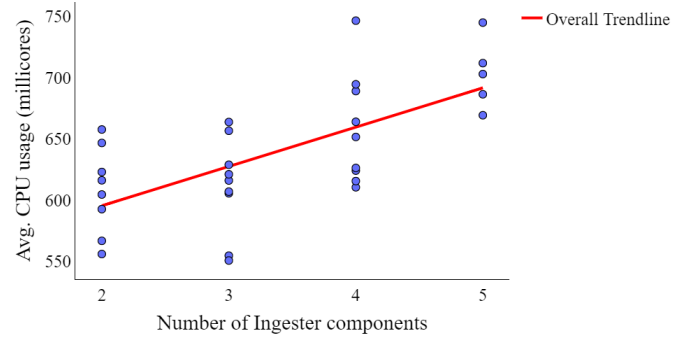


Fig. 3. Write path CPU usage vs. number of Ingesters used with the various numbers of time series stored by Cortex.

Native Computing Foundation (CNCF) project [11]. Cortex can receive metrics from multiple Prometheus servers and serve aggregated queries across all data in a single place which is useful when the cloud application operator wants to get an overview of all the metrics and data collected by these Prometheus servers. Our long-term storage choice was MinIO [12]. Fig. 2 illustrates our setup with Cortex components.

A. Cortex measurements

We deployed Cortex in a Kubernetes [13] cluster and monitored the CPU usage of each component. We ran data writing measurement experiments for various configuration scenarios for a couple of hours each in order to measure the whole pipeline’s CPU usage. Some components are only active at the beginning or at the very end of the pipeline. Also, we tested query performance multiple times and created a data retrieval application to perform different queries to avoid serving data from the cache. Based on our experiments, we selected the most important parameters of the Cortex microservices in order to find the features which accurately infer total resource usage. Our target function is the overall CPU usage, including every component of Cortex and Prometheus.

During our data ingestion experiments, we wrote several blocks into MinIO to make sure we tested the whole data write pipeline numerous times for all tested settings. We generated




Fig. 4. Histogram of measurements for the average CPU used. The number of time series ingested by Cortex is illustrated by coloring.

60,000-300,000 time series to be ingested and deployed 1-5 replicas of each microservice on the write path of Cortex. We found the number of time series continuously written into Cortex and the number of Nginx [14], Distributor and Ingestor components to be the most important features affecting CPU consumption. These components contribute majorly to the write path performance of Cortex; scaling out other components does not change the CPU usage or the ingestion capacity. Fig. 3 shows the relationship between the CPU usage and the number of Ingesters used, while the coloring indicates the number of time series ingested by Cortex. The histogram of measurements, distinguished again by colors for various numbers of time series ingested by Cortex, against the average CPU usage, can be seen in Fig. 4.

For testing the query performance, we fetched 60,000-300,000 time series with various time aggregation levels (between 10 seconds and 1 hour) from Cortex. With every test, we queried the last 4 hours of data from Cortex while continuously writing data to ensure we always have data both in the long term storage and in the Ingesters' memory. The requests were sent to Nginx ingress; it forwarded those to the Query frontend, which could optionally transform a query into smaller queryset batches and send it to the Querier that performed the data retrieval and aggregation. From the read (query) path, we found the number of time series queried in Cortex and the aggregation (or step as in Prometheus' nomenclature) to be the important features. Fig. 5 presents the measured CPU usage of the microservices and the level of data aggregation in seconds.

The correlation matrix shown in Fig. 6 shows the most dominant features of both paths. The write path model includes the number of ingested time series and the number of required components of Nginx, Distributors and Ingesters. The read path model uses the number of time series and the aggregation used while querying the data stored in MinIO. The conclusion from the measurements is that the resource footprint of the data ingestion pipeline of Cortex is mainly affected by the number of time series stored and later queried. Using several Ingestor instances increases the capacity of the write path but also induces higher resource consumption: this component stores

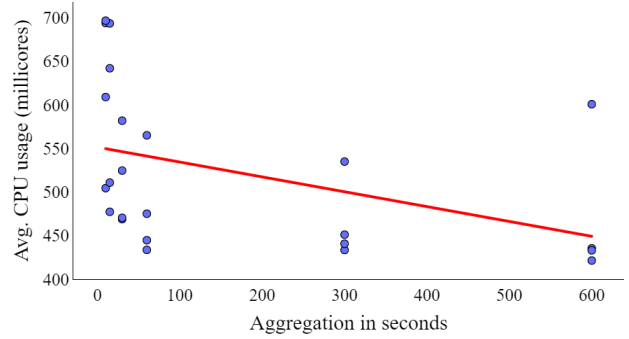


Fig. 5. Read path CPU usage vs. level of data aggregation in seconds.

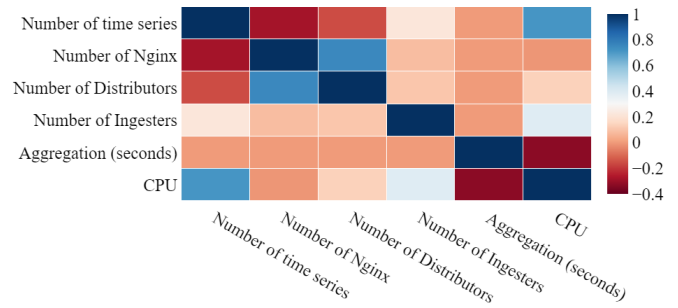


Fig. 6. The two paths' correlation matrix.

TABLE I
REGRESSION COEFFICIENTS.

Features	Coefficients
Number of time series	9.28
Number of Distributors	32.54
Number of Ingesters	24.21
Number of Nginx	1.91
Query aggregation	-0.182
Intercept	129.3

the metrics and time series before those get written into the long term storage, e.g., MinIO. On the read path, higher data aggregation leads to less CPU usage.

B. Linear regression models

After the measurements had been performed, we created different regression models for the CPU usage of each component and trained them. We used linear regression models because the chosen features showed linear dependencies with the target-dependent variable, i.e., CPU consumption. The models are trained on data from the two paths: write path and read (query) path for each microservice, as a component might be used continuously by both paths, so creating an aggregate model which takes features and information about the different paths at the same time is important.

Table I shows the coefficients of the features used in the model. The performance of the regression model is best




Fig. 7. The two paths combined CPU usage by the number of time series used by Cortex.

reflected by R^2 - the coefficient of determination - which was around 0.75. The Mean Absolute Error (MAE) is around 58.2 [millicores]. For bringing the coefficients to the same order of magnitude, we divided the number of time series by 10,000. With this data transforming step, the coefficient (9.28) is for every 10,000 time series used. Fig. 7 shows all the data points used for this model: the overall CPU consumption against the stored and queried number of time series.

C. Integer linear programming for optimization

With a complex regression model including both the read and write paths in our hand, we created an integer linear programming (ILP) optimization for the application with specific constraints. Finally, we configured the microservices based on the results of the optimization.

We minimize CPU usage for the user inputs, i.e., the number of time series used in the system, the available memory for the microservices, and the data aggregation given in seconds for the query part. The output is the average CPU utilization expected by components, the number of Ingesters, Nginx, and Distributors suggested to use, and whether the given memory is enough for the microservices.

We use the following notation:

- T = number of time series used [10,000]
- I = number of Ingestor instances
- D = number of Distributor instances
- N = number of Nginx instances
- Q = level of data aggregation [sec]
- M = memory resource available [MB]

The ILP minimization target is:

$$9.28T + 32.54D + 24.21I + 1.91N - 0.182Q, \quad (1)$$

subject to

$$I > 1 \quad (2)$$

$$D > 1 \quad (3)$$

$$N > 1 \quad (4)$$

$$T > 0 \quad (5)$$

TABLE II
THE COMPONENTS' CPU CONSUMPTION SHARES.

Component	% of total CPU usage
Ingestor	32 %
Distributor	26 %
Querier	14 %
Query frontend	13 %
Store gateway	9 %
Nginx	6%

$$Q > 0 \quad (6)$$

$$1500 + 128N + 512D + 1500I + 1000T/10 < M \quad (7)$$

$$I - T/10 > 0. \quad (8)$$

The objective function (1) stands to minimize the CPU resource given by the user's input and the regression coefficients. We use (2), (3), (4) constraints for high availability of the components: multiple replicas are helpful for load balancing on the write path but are also important for providing reliability, e.g., a Distributor failure may result in a temporary halt on the write path. It can potentially lead to data loss if not handled correctly. As for (5) and (6), we assume that T, Q are provided by the user, currently set to be non-negative. We also assume that the memory limit M is also provided by the user. From the measurements, we managed to formalize the main Cortex components' memory consumption, and we formalized it in Constraint (7). For the components which are not incorporated in the model, the total memory need is approximately 1,500 MB, not significantly varied by the load on the application. As for the three components used in the model, an Ingestor needs approximately 1,500 MB memory when used by both paths parallel, an Nginx only needs 128 MB memory, while the Distributor needs approximately 512 MB. Furthermore, from the tests, we can confirm that not only the CPU usage is being influenced by the number of time series but also the memory usage: for every 100,000 metrics, there is a need for another 1,000 MB by the microservices in total. Constraint (8) stands to ensure high availability: from the measurements, we found that launching an additional Ingestor replica for every 100,000 time series provides a safe operation with the load balance applied.

When the optimization yields the minimal CPU usage results with the optimal I, D, N replica numbers, one must translate the results to resource limits for every microservice instance. In the Cortex case, Table II indicates the fraction of the total CPU usage among the respective microservices.

D. Illustrative user input examples and optimization results

In order to demonstrate the applicability of the suggested optimization, we created a few user input examples and calculated the respective results. We used PuLP [15] for the ILP optimization. We used the COIN-OR Branch and Cut Solver (CBC), which is an open-source mixed-integer linear programming solver written in C++. Other available solvers include GNU Linear Programming Kit (GLPK), LP Solve, Coin-or linear programming (Clp). Given the size of the

TABLE III
ILLUSTRATIVE EXAMPLES SUMMARY TABLE WITH CORTEX'S RECOMMENDATION.

Time series used [1,000]	Memory available (GB)	Aggregation (sec)	Ingestor replica number	Distributor replica number	Nginx replica number	Predicted CPU usage (millicores)	Cortex planned CPU (millicores) [9]	Relative saving
300	12	10	3	2	2	1,300	6,000	78%
400	16	60	4	2	2	1,405	8,000	82%
120	9	15	2	2	2	1,105	2,400	54%
250	8	120	Infeasible				5,000	-
250	12	120	3	2	2	1,218	5,000	76%

problem space and the carefully crafted constraints, the ILP examples were solved in seconds.

In the first example, the inputs are 300,000 time series with 12 GB memory, and the aggregation should be 10 seconds. In the second example, the inputs are 400,000 time series with 16 GB memory available, and the aggregation should be 60 seconds. In the third example, the inputs are 120,000 time series with 9 GB memory available, and the aggregation should be 15 seconds. In the fourth example, the inputs are 250,000 time series with 8 GB memory available, and the aggregation should be 120 seconds. Here the model shows that using 8 GB memory will not be enough, so we should solve the problem by reducing the number of time series ingested or by using more memory, just as it is shown in the fifth example. We summarized these examples in Table III. These tests show that the Distributor and the Nginx replica should be 2 for such inputs. We ourselves set the minimum value for them to be 2 because of high availability, but using more replicas for these components will not give significant benefits performance-wise, so the optimization effort results in the lowest number of instances possible in the presented cases. The Ingestor number depends on the number of ingested time series and on the available memory.

Although Cortex's documentation recommends capacity planning [9], those resource figures are not optimal when we strive to reduce and keep the resource usage as low as possible. Along with the illustrative examples, in Table III we also show the recommended CPU provisions based on Cortex's documentation and the relative saving which we managed to achieve with the regression and optimization models. In the examples, this saving ranges from 54-82%. If we translate these numbers to cost, i.e., we deploy these microservices in a public cloud and pay for the reserved resources, then we can save significant amounts by applying the proposed method.

V. CONCLUSION

In this paper we argued that resource provisioning is essential and, in fact, possible to perform scientifically when using cloud-native microservices. We proposed an approach that we demonstrated on a data ingestion and storage cloud-native application, Cortex. After measuring the performance vs resource footprint trade-offs with several configuration settings and creating a regression model for resource usage, we showed the most critical features and their quantitative effects on CPU utilization, which proved to be linear. Then we introduced an

integer linear programming formulation that can be solved to minimize the expected CPU usage for given user inputs. With some illustrative examples, we showed that a prudent resource reservation might halve the costs paid for cloud resources while meeting the QoS constraints set for the application.

ACKNOWLEDGMENT

This work was supported by the Ministry of Innovation and Technology of Hungary from the National Research, Development and Innovation Fund through projects *i*) no. 135074 under the FK_20 funding scheme and *ii*) 2019-2.1.13-TÉT_IN-2020-00021 under the 2019-2.1.13-TÉT-IN funding scheme. L. Toka was supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences. Supported by the ÚNKP-21-5 New National Excellence Program of the Ministry for Innovation and Technology from the source of the National Research, Development and Innovation Fund.

REFERENCES

- [1] Faiza Samreen, Yehia Elkhatib, Matthew Rowe, Gordon S. Blair. Daleel: Simplifying cloud instance selection using machine learning NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium
- [2] Chris Richardson, What are microservices?, <https://microservices.io> Accessed 28 Oct 2021.
- [3] Anshul Jindal, Vladimir Podolskiy, Michael Gerndt, Performance Modeling for Cloud Microservice Applications. ICPE '19: Proceedings of the 2019 ACM/SPEC International Conference on Performance Engineering
- [4] Yanqi Zhang, Weizhe Hua, Zhuangzhuang Zhou, Edward Suh, Christina Delimitrou, Sinan: Data-Driven Resource Management for Interactive Multi-tier Microservices, Workshop on ML for Computer Architecture and Systems (MLArchSys), 2020
- [5] B. Urgaonkar, P. Shenoy, A. Chandra, P. Goyal, Dynamic Provisioning of Multi-tier Internet Applications, Second International Conference on Autonomic Computing (ICAC'05), 2005
- [6] P. Leitner, J. Cito and E. Stöckli, Modelling and Managing Deployment Costs of Microservice-Based Cloud Applications, 2016 IEEE/ACM 9th International Conference on Utility and Cloud Computing (UCC), 2016
- [7] Madhu Sudan, The P vs. NP problem, <https://madhu.seas.harvard.edu/papers/2010/pnp.pdf>, 2010
- [8] Cortex, <https://cortexmetrics.io>, Accessed 28 Oct 2021.
- [9] Capacity Planning <https://cortexmetrics.io/docs/guides/capacity-planning> Accessed 30 Oct 2021.
- [10] Prometheus, <https://prometheus.io>, Accessed 28 Oct 2021.
- [11] Cloud Native Computing Foundation, <https://www.cncf.io>, Accessed 28 Oct 2021.
- [12] MinIO, <https://min.io>, Accessed 28 Oct 2021.
- [13] Kubernetes, <https://kubernetes.io>, Accessed 28 Oct 2021.
- [14] Nginx, <https://www.nginx.com> Accessed 28 Oct 2021.
- [15] PuLP, <https://pypi.org/project/PuLP>, Accessed 28 Oct 2021.