



Characterization of IoT Workloads

Uma Tadakamalla^(✉) and Daniel A. Menascé^(✉)

Department of Computer Science, George Mason University, Fairfax, VA, USA
{utadakam,menasce}@gmu.edu

Abstract. Workload characterization is a fundamental step in carrying out performance and Quality of Service engineering studies. The workload of a system is defined as the set of all inputs received by the system from its environment during one or more time windows. The characterization of the workload entails determining the nature of its basic components as well as a quantitative and probabilistic description of the workload components in terms of both the arrival process, event counts, and service demands. Several workload characterization studies were presented for a variety of domains, except for IoT workloads. This is precisely the main contribution of this paper, which also presents a capacity planning study based on one of the workload characterizations presented here.

Keywords: Workload characterization · Internet of Things · Capacity planning · G/G/n queue · Quality of Service in edge computing

1 Introduction

Siegel et al. [35] argue that scalability is needed to support the continued expansion of the Internet of Things. Therefore, performance engineering studies are very important for understanding tradeoffs between security, availability, and response time of various types of IoT applications.

Workload characterization is a fundamental and necessary step in carrying out any performance engineering study [26]. The workload of a system is defined as the set of all inputs received by the system from its environment during one or more time windows. The characterization of the workload entails determining the nature of its basic components (e.g., transactions, I/O requests, IoT device requests) as well as a quantitative and probabilistic description of the workload components in terms of both the arrival process, event counts, and service demands (e.g., arrival rate of requests and interarrival time distributions, distribution of the number of IoT device signals received, distribution of the file sizes returned by an HTTP request) [26].

General methods for workload characterization have been discussed in [11, 12, 26]. Specific applications of these techniques to a variety of domains were developed by many researchers (see examples in Sect. 5). However, there is a need for workload characterization studies for IoT applications.

The recent development of Internet of Things (IoT) and edge/fog computing demands models for this new environment. Our prior work includes the development of an analytic model, called FogQN, based on queuing networks [37] and an autonomic controller that uses FogQN to dynamically determine the optimal breakdown of processing between fog and cloud servers [38].

Any modeling effort of fog and cloud computing calls for workload characterization studies of IoT workloads. The understanding of the characteristics of IoT workloads can be used to perform capacity planning studies. These are the main contributions of this paper. More specifically, we (1) describe the methodology we used to analyze IoT traces; (2) describe and analyze three publicly available IoT datasets: NY city taxi trips, GPS trajectories of taxis in Beijing, Chicago taxi trips; and (3) present a capacity planning study based on the workload characterization of the NY city taxi trips. Our workload characterization includes counts of events, i.e., IoT device signals, at various time scales (e.g., hour of the day, day of the week) and a characterization of the interarrival time of signals received from IoT devices.

The rest of this paper is organized as follows. Section 2 describes the general data collection and analysis methodology used in this paper. Section 3 has one subsection for each of the datasets we analyzed. Each subsection describes the dataset and presents the results of the workload characterization for that dataset. Section 4 provides an example of how a queuing model can be used to answer what-if questions using the workload of NY city taxi trips. Section 5 discusses related work. Finally, Sect. 6 presents concluding remarks and future work.

2 General Data Collection and Analysis Methodology

The data collection and analysis methodology presented here can be applied to a variety of IoT workloads. This paper analyzed several publicly available IoT datasets. Some existing datasets are from applications in which data is sent by a set of sensors at regular intervals (e.g., every 5 min) in a synchronous way. We did not consider these datasets because they are not very interesting from the point of view of workload analysis. The applications we considered in our study have IoT devices that are independent of each other and send signals at irregular intervals (e.g., signals sent by a taxi cab whenever a passenger is dropped off).

Our analysis methodology consisted of the following steps:

1. Data is aggregated from all the files that make up the dataset.
2. The aggregated data is cleansed by removing any invalid and duplicate data, and any outliers.
3. The cleaned up data is sorted based on the timestamp of the records.
4. The sorted data is filtered based on characteristics such as days, hours, month, latitude/longitude of the IoT device.
5. The filtered data is characterized by computing event counts by hour of the day on a daily and monthly basis, and by day of the week.
6. The distribution of the interarrival time of signals generated by IoT devices is characterized. We used Quantile-Quantile (Q-Q) plots and Cumulative Distribution Functions (CDF) to that effect [21].

A Q-Q plot is a graphical tool that helps determine if the data points in a given data set come from the same distribution as a given theoretical distribution. A Q-Q plot is a scatter plot that plots two sets of quantiles (from the dataset and from the theoretical distribution) against each other. If both quantiles come from the same distribution, the points in the Q-Q plot form a roughly straight line. We experimented with several candidate theoretical distributions for each dataset and did a linear regression on the points. The distribution that had a coefficient of determination R^2 closest to 1 was chosen as the best fit theoretical distribution for the dataset. The candidate distributions can only be those that can take non-negative values because an interarrival time cannot be negative. For that reason we selected the lognormal, Weibull, and Gamma distributions. Note that the Weibull distribution has the exponential distribution as a special case, depending on the value of its parameters.

Table 1 presents the expressions for the probability density function (pdf) and the expressions used to compute the parameters of the three considered distributions as a function of \bar{X} , S and $C = S/\bar{X}$, the mean, standard deviation and coefficient of variation of the interarrival times, respectively, computed from the datasets.

Table 1. Features of the lognormal, Weibull, and Gamma distributions.

Distribution	Pdf	Parameters
Lognormal	$\frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{[\ln x - \mu]^2}{2\sigma^2}}$	$\mu = \ln(\bar{X}) - \ln(\sqrt{(1+C^2)}), \sigma = \sqrt{\ln(1+C^2)} \mu \in (-\infty, +\infty), \sigma \geq 0$
Weibull	$\frac{k}{\lambda} \left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k}$	$k \approx C^{-1.086}, \lambda = \bar{X}/\Gamma(1+1/k) k, \lambda > 0$
Gamma	$\frac{x^{k-1} e^{-x/\theta}}{\theta^k \Gamma(k)}$	$k = 1/C^2, \theta = S^2/\bar{X} k, \theta > 0$

The theoretical distribution quantile data is generated using the inverseCumulativeProbability method in the Java Apache Commons Math3 distribution package [2] with parameters computed using the equations in Table 1.

3 IoT Datasets

We describe and analyze in this section, three IoT datasets: NY city taxi trips, GPS trajectories of taxis in Beijing, and Chicago taxi trips.

3.1 New York City Taxi Trip Data

The New York City taxi trip data is provided by Illinois Data Bank, which is operated by the University of Illinois at Urbana Champaign. This dataset [15] contains records of four years (2010–2014) of taxi operations in New York City including 697,622,444 trips. The data is stored in the CSV format, organized by

year and month. Each month’s data is stored in a separate file. Each row in the file represents a single taxi trip. Each trip records the pickup and drop-off dates, times, and coordinates, and the metered distance reported by the taximeter. For this analysis, we only considered the drop-off date and time, drop-off latitude and longitude fields. We assumed that a fog node is at Grand Central Terminal, whose latitude and longitude coordinates are (40.7527, -73.9772), and it serves all the IoTs devices (taxis) that are within a one-mile radius. This means that signals received from the taxis at drop off locations that are within a 1-mile radius are served by the Grand Central Terminal fog node. Therefore, we selected all the records that are within 1 mile radius from the fog node for this analysis. We cleaned up the data by removing duplicate and invalid entries and used the cleaned up data to generate interarrival times. We then removed the outliers (interarrival times greater than 2000 s) from the interarrival times dataset.

Figure 1(b) shows the variation of the number of taxi signals by hour of the day for Sunday, February 7, 2010 and Monday, February 8, 2010. It is apparent that taxi cabs are utilized more on Mondays (weekday) than on Sundays (week-end), with the exception of 12:00 am through 5:00 am. This may be because more people in New York use cabs on weekdays to move around. The number of taxi signals on the early hours of Sunday exceeds the taxi cab requests during the same time on Monday because people are more likely go out on Saturday nights, and they utilize taxi cabs to get back home during the wee hours on Sunday. However, at the same time on Monday, most people are at home resting for the next work day. Also, the number of taxi signals is higher during the morning (5:00 am to 9:00 am) and evening rush hours (4:00 pm to 6:00 pm) during a Monday because between these peaks most people are more likely to be working in their offices.

Next, we analyzed the number of taxi signals for the entire month of February, 2010 grouped by hour of the day as shown in Fig. 1(a). The figure shows that the number of taxi signals is lower during non-working hours compared to those of working hours. Also, there is a clear rise in the number of signals during morning and evening rush hours from 5:00–9:00 am and 4:00–7:00 pm, respectively.

Next, we studied the variation of the number of taxi signals by days of the week and aggregated the data for each day of the week of February, 2010 as shown in Fig. 2. The figure shows that the lowest signal counts are recorded on Sundays.

We now turn our attention to the characterization of interarrival times of taxi signals using Q-Q plots and CDFs as explained in Sect. 2. To determine the best fit distribution, the quantiles of interarrival times of taxi signals were plotted against those of various theoretical distributions (i.e., lognormal, Weibull and Gamma). Table 2 shows the parameters used for each distribution and the corre-




Fig. 1. (a) Left: NY Grand Central Terminal taxi signal counts aggregated by hour of the day for the entire month of February, 2010, (b) Right: NY Grand Central Terminal taxi signal counts by hour of the day for Sunday, February 7, 2010 (weekend) and Monday, February 8, 2010 (weekday)

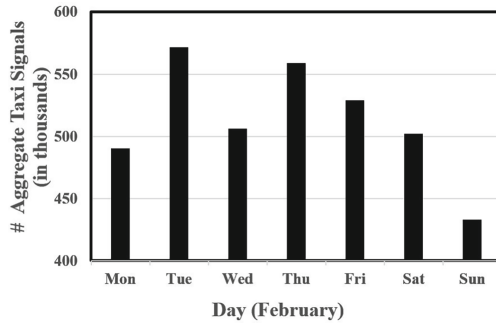


Fig. 2. NY Grand Central Terminal taxi signal counts aggregated by days of the week for February, 2010

Table 2. Fitting February 8, 2010 NY City taxi signal interarrival time data.

Distribution	Parameters	R^2
Lognormal	$\mu = -1.630, \sigma = 1.494$	0.941
Weibull	$k = 0.316, \lambda = 0.081$	0.902
Gamma	$k = 0.120, \theta = 4.976$	0.895

sponding R^2 value. The lognormal distribution has the best fit for the data with an R^2 value equal to 0.941. The corresponding Q-Q plot is shown in Fig. 3(a). The CDF plots of taxi signal interarrival times and the lognormal theoretical distribution are shown in Fig. 3(b). They both match very closely. Based on the R^2 value from the Q-Q plot and CDF plots, we can conclude that the data best fits the log-normal distribution.




Fig. 3. (a) Q-Q plot (left) and (b) CDF plots (right) using NY Grand Central Terminal February 8, 2010 taxi signal interarrival times data and theoretical lognormal distribution data with $\mu = -1.630$ and $\sigma = 1.494$.

3.2 Microsoft T-Drive Trajectory Dataset

The Microsoft T-Drive Trajectory dataset [41] is provided by Microsoft for research purposes. This dataset contains the GPS trajectories of 10,357 taxis (one file per taxi) during the period of February 2–8, 2008 within Beijing. We ignored the data for February 2 and February 8 because they are incomplete. Each file of this dataset contains the trajectory of one taxi. The total number of points in this dataset is about 15 million and the total distance of the trajectories reaches about 9 million kilometers. We assumed that the fog node is located at Tiananmen Square, whose latitude and longitude are (39.9055, 116.3976), and that this node will serve the IoT devices (i.e., taxis) within a one-mile distance. We then selected all the records that are within a 1-mile radius from that node and used that data to generate the interarrival times of the signals. We then removed the outliers from the interarrival times data.

Figure 4(b) shows the variation of the number of taxi signals by hour of the day for Sunday, February 3, 2008, and Monday, February 4, 2008. It is apparent that taxi cabs are utilized less over the night hours than during day time. Also, there are more taxis utilized during evening hours on weekends than weekdays.

Next, we analyzed the number of taxi signals from February 3–7, 2008 grouped by hour of the day as shown in Fig. 4(a). The figure shows that the number of taxi signals is lower during night hours than during day time. A similar trend was seen in Fig. 5. This figure shows the variation of the number of taxi signals by days of the week from February 3–7, 2008. The highest number of taxi signals on weekdays can be seen on Mondays and it decreases through the week. The second highest number is observed on Sundays maybe because Tiananmen Square is a popular place for visitors and there are more visitors on weekends than on weekdays.




Fig. 4. (a) Left: Beijing Tiananmen Square taxi signal counts aggregated by hour of the day for February 3–7, 2008, (b) Right: Beijing Tiananmen Square taxi signal counts by hour of the day for Sunday, February 3, 2008 (weekend day) and Monday, February 4, 2008 (weekday).




Fig. 5. Beijing’s Tiananmen Square taxi signal counts aggregated by days of the week.

Next, we characterized the interarrival times of taxi signals using Q-Q plots and CDFs as explained in Sect. 2. To determine the best fit distribution, the quantiles of interarrival times of taxi signals were plotted against those of various theoretical distributions (i.e., lognormal, Weibull and Gamma). Table 3 shows the parameters used for each distribution and the corresponding R^2 value.


The lognormal distribution has the best fit for the data with an R^2 value equal to 0.986. The corresponding Q-Q plot is shown in Fig. 6(a). The CDF plot of taxi signal interarrival times and lognormal theoretical distribution is shown in Fig. 6(b). They both match very closely. Based on the R^2 value from the Q-Q plot and CDF plots, we can conclude that the data best fits a lognormal distribution.

3.3 Chicago Taxi Trips Dataset

The Chicago taxi trips dataset provided by the City of Chicago’s open data portal [1] contains information on taxi trips in Chicago reported to the City of

Table 3. Fitting February 5, 2008 Tiananmen Square taxi signal interarrival time data.

Distribution	Parameters	R^2
Lognormal	$\mu = -0.130, \sigma = 1.111$	0.986
Weibull	$k = 0.616, \lambda = 1.119$	0.974
Gamma	$k = 0.410, \theta = 3.970$	0.946

**Fig. 6.** Q-Q plot (left) and CDF plots (right) using Beijing Tiananmen Square February 5, 2008 taxi signal interarrival times data and theoretical lognormal distribution data with $\mu = -0.130$ and $\sigma = 1.111$.

Chicago. We exported February 2015 data in a CSV format using their API. Each record in the file represents a single taxi trip and includes pickup and drop-off dates, times, and coordinates, and trip duration (in sec). The pickup and drop-off times are rounded to the nearest 15 min and the trip duration is rounded to the nearest minute, meaning that the trip durations are in multiples of 60s. For this analysis, we only considered the trip end time (trip start time + trip duration), drop off latitude and longitude fields. We assumed that the fog node is at Millennium Park, whose latitude and longitude are (41.8826, -87.6226), and it serves all the IoT devices (taxis) that are within one-mile radius. Therefore, we selected all taxi trip records whose drop off location is within one-mile radius from the fog node for this analysis. We cleaned up the data by removing records with missing data and used the clean data for taxi trip count analysis. To compute the interarrival times, we grouped the taxi signals reported each minute and computed the interarrival times by distributing them uniformly within that minute.

Figure 7(b) shows the variation of the number of taxi signals by hour of the day for Sunday, February 22, 2015 and Monday, February 23, 2015. It is apparent that taxi cabs are utilized more on Mondays (weekday) than on Sundays (weekend), with the exception of 12:00 am through 6:00 am. This may be because more people in Chicago use taxis on weekdays to move around than on

weekends. The number of taxi signals on the early hours of Sunday exceeds the taxi signals during the same time on Monday because more people are likely to go out on Saturday nights than on Sunday nights, and they utilize taxis to get back home in the early hours of the next day. Also, the number of taxi signals is higher during the morning (6:00 am to 9:00 am) and evening rush hours (3:00 pm to 6:00 pm) during a Monday (weekday) because people are more likely to use taxis to go to work and go back home during these times.

Next, we analyzed the number of taxi signals for the entire month of February, 2015 grouped by hour of the day as shown in Fig. 7(a). The figure shows that the number of taxi signals is lower during non-working hours compared to those of working hours. Also, there is a clear rise in the number of signals during morning and evening rush hours from 5:00 am to 9:00 am and 3:00 pm to 6:00 pm, respectively.




Fig. 7. Chicago Millennium Park taxi signal counts. (a) Left: aggregated by hour of the day for the entire month of February 2015, (b) Right: by hour of the day for Sunday, February 22, 2015 (weekend) and Monday, February 23, 2015 (weekday).

Next, we studied the variation of the number of taxi signals by days of the month and aggregated the data for each day of the month of February as shown in Fig. 8(a). The figure shows that the signal counts are higher on weekdays than on weekends and the lowest signal counts are seen on Sundays every week.

Next, we studied the variation of the number of taxi signals by day of the week and aggregated the data for each day of the week of February 2015 as shown in Fig. 8(b). The figure shows that the weekday counts are higher than the weekend counts and increase from Monday to Friday. Also, lowest signal counts are recorded on Sundays.

We then characterized the interarrival times of taxi signals using Q-Q plots and CDFs as explained in Sect. 2. To determine the best fit distribution, the quantiles of interarrival times of taxi signals were plotted against those of various theoretical distributions (i.e., lognormal, Weibull and Gamma). Table 4 shows the parameters used for each distribution and the corresponding R^2 value.

The R^2 for lognormal and Weibull distributions are very close. However, the lognormal distribution has the best fit for the data with an R^2 value equal to

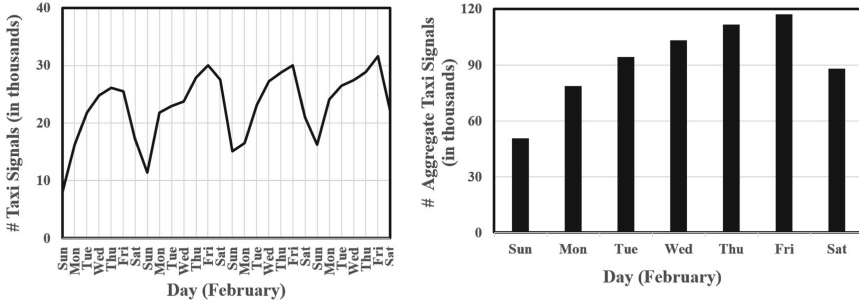


Fig. 8. Chicago Millennium Park taxi signal counts. (a) Left: for each day in February 2015 (b) Right: aggregated by days of the week in February 2015.

Table 4. Fitting February 23, 2015 Chicago taxi signal interarrival time data

Distribution	Parameters	R^2
Lognormal	$\mu = 0.241, \sigma = 1.439$	0.9621
Weibull	$k = 0.35, \lambda = 0.71$	0.9618
Gamma	$k = 0.144, \theta = 24.809$	0.7977




Fig. 9. Q-Q plot (left) and CDF plots (right) using the Chicago Millennium Park February 23, 2015 taxi signal interarrival times and theoretical lognormal distribution data with $\mu = 0.241$ and $\sigma = 1.439$.

0.9621. The corresponding Q-Q plot is shown in Fig. 9(a) and the plots for the CDF of interarrival times and the lognormal theoretical distribution are shown in Fig. 9(b). They both match very closely. Based on the R^2 value from the Q-Q plot and CDF plots, we can conclude that the data best fits a lognormal distribution even though a Weibull distribution would be a good fit also.

4 Workload Characterization Use in Capacity Planning

As indicated above, workload characterization is an essential step for capacity planning purposes. Consider the following what-if question: How many fog servers are required to support a given load with an average response time below a certain value? We show here how we can answer this type of question using the NY City taxi workload. Let n be the number of fog servers that handle signals received from taxis within a one-mile radius of a given location. All arriving signals join a single queue and are dispatched to the first available fog server when they reach the head of the line.

The average response time of a taxi signal was computed using the approximate G/G/ n queuing equation given below [26]

$$T \approx E[S] + \frac{C(\rho, n)}{c(1 - \rho)/E[S]} \times \frac{C_a^2 + C_s^2}{2} \quad (1)$$

where $E[S]$ is the average processing time of a taxi signal, $\rho = \lambda E[S]/n$ is the utilization of the set of n fog servers that receive a collective average arrival rate of λ taxi signals/sec, C_a is the coefficient of variation (i.e., the ratio of the standard deviation by the mean) of the interarrival time, C_s is the coefficient of variation of the service time, and $C(\rho, n)$ is the Erlang formula given by

$$C(\rho, n) = \frac{(n\rho)^n/n!}{(1 - \rho) \sum_{j=0}^{n-1} (n\rho)^j/j! + (n\rho)^n/n!}. \quad (2)$$

Because the utilization ρ must be less than 1, we have that $\lambda < n/E[S]$, i.e., the average arrival rate cannot exceed $n/E[S]$. Our data showed that the maximum rate of signals received from taxis within a one-mile radius from Grand Central Terminal during the date of February 8, 2010 was approximately 4 signals/sec. We used the G/G/ n equations above to compute the variation of the average signal response time as a function of the average arrival rate λ for five values of n (see Fig. 10). We used the following numerical values for Fig. 10: $E[S] = 0.2$ s, $C_a = 2.88$, $C_s = 0.94$ (from 2/8/2010 data). As expected, the figure shows that the maximum arrival rate of signals that can be handled increases in proportion to the number of fog servers. For example, when $n = 1$, the maximum arrival rate the system can handle has to be less than 5 signals/sec whereas for $n = 5$, the maximum arrival rate the system can handle has to be less than 25 signals/sec. Additionally, the average response time decreases as n increases for a given arrival rate. For example, at an arrival rate = 4.5 signals/sec the average response time with one server is 9.13s whereas with 5 servers the average response time is 0.2s. If we want the average response time not to exceed 1s for an average arrival rate of taxi signals of 10 signals/sec we need at least 3 fog servers.

5 Related Work

Workload characterization studies have been conducted for various types of applications and systems. Some examples include: e-commerce [25], auction




Fig. 10. Average response time vs. arrival rate for $n = 1, 2, 3, 4, 5$

sites [5], WWW [24], networking [28,30], live streaming media [39], spam traffic [19], storage systems [36], data centers [32], cloud computing [23], grid computing [14], memory systems [8], and database systems [16]. [27] quantifies a Poisson process approximation for IoT aggregate arrival processes. The studies above have shown that different domains have their own specific workload characteristics. Our paper fills a much needed gap in terms of understanding and characterizing IoT workloads.

The vision and challenges of edge computing were discussed in [9,34]. There are some very good IoT and fog/edge computing surveys: a survey of mobile edge computing was presented in [3]; a survey of architecture, enabling technologies, security and privacy, and IoT applications was presented in [22]; and Ngu et al. presented a survey on IoT middleware [29]. Cruz et al. presented a reference model for IoT middleware [13]. [33] presents an IoT architecture based on transparent computing to build scalable IoT platforms. Transparent computing enables users to select services on-demand, without being concerned with the installation and management of services.

Similarly to [38], the work in [40] aims at reducing the response time of IoT applications by offloading the load of fog-capable devices to the cloud. Another work along the same vein is [10]. Fan and Ansari [17] presented an application aware scheme to allocate IoT-based workloads to edge servers in order to minimize the response time of IoT applications. The work in [4] proposes a method for reducing latency and device energy consumption using the fog, which is based on computational offloading and network utility optimization. The work in [18] presents a vision of human-centered edge-device based computing, known as Edge-centric Computing and the research challenges associated with its implementation. The work in [7] proposed a new technique called Home Edge Computing, a three-tier edge computing architecture that provides data storage and processing near the users (home server) to achieve ultra-low latency.

The work in [20] analyzed a motion dataset to characterize the kinetic energy that can be harvested by an IoT node and developed energy allocation algorithms

for such nodes. The work by Pereira et al. [31] discusses an experimental evaluation of latency in IoT service composition with mobile gateways and assesses the capabilities and limitations of a standard machine-to-machine middleware. IoT devices with security flaws are attractive targets for attacks. [6] discusses HoneyScope, a network centric approach to protect vulnerable IoT devices by creating virtualized views of the network and nodes.

None of the studies cited above present a comprehensive workload characterization of actual IoT applications.

6 Concluding Remarks and Future Work

Understanding and quantitatively characterizing the workload generated by IoT devices is key to being able to analyze the performance of edge/fog computing environments. Our study analyzed three datasets that contain information generated by taxis in three big cities. Our workload characterization, which can be applied to other IoT workloads, included counts of events, i.e., IoT device signals, at various time scales (e.g., hour of the day, day of the week) and a characterization of the interarrival time of signals received from IoT devices.

Our results indicated that the interarrival time of IoT signals for all three datasets can be very well approximated by a lognormal distribution. We also observed that the count of events for the three taxi-related datasets can be well explained by expected daily routines of habitants of large cities. We also showed that workload characterization results can be used for capacity planning studies of edge computing environments.

In the future, we intend to apply our characterization methodology to IoT datasets that deal with other types of IoT devices. We are also investigating the sensitivity of our results with respect to the location of the fog node, and how it may affect the probability distribution and parameters of the request interarrival times.

References

1. Chicago data portal. <https://data.cityofchicago.org/>
2. Package org.apache.commons.math3.distribution. <http://commons.apache.org/proper/commons-math/javadocs/api-3.5/org/apache/commons/math3/distribution/package-summary.html>
3. Abbas, N., Zhang, Y., Taherkordi, A., Skeie, T.: Mobile edge computing: a survey. *IEEE Internet Things J.* **5**(1), 450–465 (2018)
4. Ahn, S., Gorlatova, M., Chiang, M.: Leveraging fog and cloud computing for efficient computational offloading. In: 2017 Undergraduate Research Technology Conference (URTC), IEEE MIT, pp. 1–4. IEEE (2017)
5. Akula, V., Menasce, D.: Two-level workload characterization of online auctions. *Electron. Commer. Res. Appl.* **6**, 192–208 (2007)
6. Al-Shaer, E., Wei, J., Hamlen, K.W., Wang, C.: HONEYScope: IoT device protection with deceptive network views. *Autonomous Cyber Deception*, pp. 167–181. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-02110-8_9

7. Babou, C.S.M., Fall, D., Kashihara, S., Niang, I., Kadobayashi, Y.: Home edge computing (HEC): design of a new edge computing technology for achieving ultra-low latency. In: Liu, S., Tekinerdogan, B., Aoyama, M., Zhang, L.-J. (eds.) *EDGE 2018*. LNCS, vol. 10973, pp. 3–17. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-94340-4_1
8. Barroso, L.A., Gharachorloo, K., Bugnion, E.: Memory system characterization of commercial workloads. In: *Proceedings of 25th Annual International Symposium Computer Architecture, ISCA 1998*, pp. 3–14. IEEE Computer Society, Washington, DC (1998)
9. Bonomi, F., Milito, R., Zhu, J., Addepalli, S.: Fog computing and its role in the Internet of Things. In: *Proceedings of MCC Workshop on Mobile Cloud Computing, MCC 2012*, pp. 13–16, New York, NY, USA. ACM (2012)
10. Brogi, A., Forti, S.: QoS-aware deployment of IoT applications through the fog. *IEEE Internet Things J.* **4**(5), 1185–1192 (2017)
11. Calzarossa, M., Massari, L., Tessera, D.: Workload characterization issues and methodologies. In: Haring, G., Lindemann, C., Reiser, M. (eds.) *Performance Evaluation: Origins and Directions*. LNCS, vol. 1769, pp. 459–482. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-46506-5_20
12. Calzarossa, M., Serazzi, G.: Workload characterization. *Proc. IEEE* **81**, 1136–1150 (1993)
13. da Cruz, M.A.A., Rodrigues, J.J.P.C., Al-Muhtadi, J., Korotaev, V.V., de Albuquerque, V.H.C.: A reference model for Internet of Things middleware. *IEEE Internet Things J.* **5**(2), 871–883 (2018)
14. Di, S., Kondo, D., Cirne, W.: Characterization and comparison of cloud versus grid workloads. In: *2012 IEEE International Conference Cluster Computing*, pp. 230–238, September 2012
15. Donovan, D., Work, D.B.: *New york city taxi trip data (2010–2013)* (2016)
16. Elnaffar, S., Martin, P., Horman, R.: Automatically classifying database workloads. In: *Proceedings of 11th International Conference Information and Knowledge Management, CIKM 2002*, pp. 622–624, New York, NY, USA. ACM (2002)
17. Fan, Q., Ansari, N.: Application aware workload allocation for edge computing-based IoT. *IEEE Internet Things J.* **5**(3), 2146–2153 (2018)
18. Garcia Lopez, P., et al.: Edge-centric computing: vision and challenges. *SIGCOMM Comput. Commun. Rev.* **45**(5), 37–42 (2015)
19. Gomes, L.H., Cazita, C., Almeida, J.M., Almeida, V., Meira, Jr., W.: Characterizing a spam traffic. In: *Proceedings of 4th ACM SIGCOMM Conference Internet Measurement, IMC 2004*, pp. 356–369, New York, NY, USA. ACM (2004)
20. Gorlatova, M., Sarik, J., Grebla, G., Cong, M., Kymissis, I., Zussman, G.: Movers and shakers: kinetic energy harvesting for the Internet of Things. In: *The 2014 ACM International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS 2014*, pp. 407–419, New York, NY, USA. ACM (2014)
21. Jain, R.: *The Art of Computer Systems Performance Analysis*. Wiley, Hoboken (1991)
22. Lin, J., Yu, W., Zhang, N., Yang, X., Zhang, H., Zhao, W.: A survey on Internet of Things: architecture, enabling technologies, security and privacy, and applications. *IEEE Internet Things J.* **4**(5), 1125–1142 (2017)
23. Magalhaes, D., Calheiros, R.N., Buyya, R., Gomes, D.G.: Workload modeling for resource usage analysis and simulation in cloud computing. *Comput. Electr. Eng.* **47**, 69–81 (2015)

24. Menascé, D., Abrahao, B., Barbará, D., Almeida, V., Ribeiro, F.: Fractal characterization of web workloads. In: Eleventh International World Wide Web Conference, Honolulu, HI, pp. 7–11 (2002)
25. Menasce, D., Almeida, V., Fonseca, R., Mendes, M.: A methodology for workload characterization of e-commerce sites. In: Proceedings of 1st ACM Conference on Electronic Commerce, EC 1999, pp. 119–128, New York, NY, USA. ACM (1999)
26. Menasce, D.A., Almeida, V.A.F., Dowdy, L.W.: Performance by Design: Computer Capacity Planning by Example. Prentice Hall, Upper Saddle River (2004)
27. Metzger, F., Hofeld, T., Bauer, A., Kounev, S., Heegaard, P.E.: Modeling of aggregated IoT traffic and its application to an IoT cloud. *Proc. IEEE* **107**(4), 679–694 (2019)
28. Nedyalkov, I., Stefanov, A., Georgiev, G.: Characterization of the traffic in IP-based communication networks. In: 2018 International Conference on High Technology for Sustainable Development (HiTech), pp. 1–4. IEEE (2018)
29. Ngu, A.H., Gutierrez, M., Metsis, V., Nepal, S., Sheng, Q.Z.: IoT middleware: a survey on issues and enabling technologies. *IEEE Internet Things J.* **4**(1), 1–20 (2017)
30. Paxson, V., Floyd, S.: Wide area traffic: the failure of poisson modeling. *IEEE/ACM Trans. Netw.* **3**(3), 226–244 (1995)
31. Pereira, C., Pinto, A., Ferreira, D., Aguiar, A.: Experimental characterization of mobile IoT application latency. *IEEE Internet Things J.* **4**(4), 1082–1094 (2017)
32. Postema, B.F., Geuze, N.J., Haverkort, B.R.: Fitting realistic data centre workloads: a data science approach. In: Proceedings of the Ninth International Conference on Future Energy Systems, e-Energy 2018, pp. 486–491, New York, NY, USA. ACM (2018)
33. Ren, J., Guo, H., Xu, C., Zhang, Y.: Serving at the edge: a scalable IoT architecture based on transparent computing. *IEEE Netw.* **31**(5), 96–105 (2017)
34. Shi, W., Cao, J., Zhang, Q., Li, Y., Xu, L.: Edge computing: vision and challenges. *IEEE Internet Things J.* **3**(5), 637–646 (2016)
35. Siegel, J.E., Kumar, S., Sarma, S.E.: The future Internet of Things: secure, efficient, and model-based. *IEEE Internet Things J.* **5**(4), 2386–2398 (2018)
36. Smirni, E., Reed, D.: Lessons from characterizing the input/output behavior of parallel scientific applications. *Perform. Eval.* **33**(1), 27–44 (1998)
37. Tadakamalla, U., Menasce, D.A.: FogQN: an analytic model for fog/cloud computing. In: Proceedings of 1st Workshop on Managed Fog-to-Cloud (mF2C), joint with 11th IEEE/ACM International Conference on Utility and Cloud Computing. IEEE/ACM (2018). <https://www.cs.gmu.edu/~menasce/papers/mF2C2018TM.pdf>
38. Tadakamalla, U., Menasce, D.A.: Autonomic resource management using analytic models for fog/cloud computing. In: Proceedings of IEEE International Conference on Fog Computing. IEEE (2019)
39. Veloso, E., Almeida, V., Meira, W., Bestavros, A., Jin, S.: A hierarchical characterization of a live streaming media workload. In: Proceedings of 2nd ACM SIGCOMM Workshop on Internet Measurement, IMW 2002, pp. 117–130, New York, NY, USA. ACM (2002)
40. Yousefpour, A., Ishigaki, G., Gour, R., Jue, J.P.: On reducing IoT service delay via fog offloading. *IEEE Internet Things J.* **5**(2), 998–1010 (2018)
41. Zheng, Y.: T-drive trajectory data sample, August 2011. <https://www.microsoft.com/en-us/research/publication/t-drive-trajectory-data-sample/>