Signaling Traffic in Internet-of-Things Mobile Networks

Stefan Geissler*, Florian Wamser*, Wolfgang Bauer‡, Michael Krolikowski‡, Steffen Gebert‡, Tobias Hoßfeld*
*Chair of Communication Networks, University of Würzburg, Germany
Email: {stefan.geissler|florian.wamser|hossfeld}@informatik.uni-wuerzburg.de
‡EMnify GmbH, Germany
Email: {wolfgang.bauer|michael.krolikowski|steffen.gebert}@emnify.com

***Abstract*—The Internet of Things (IoT) is expected to play an increasingly important role in next-generation mobile networks. Detailed knowledge of this type of traffic and the associated signaling in those networks is of interest to operators and standardization in order to run these networks successfully. The aim of this study is to characterize IoT signaling traffic for mobile communication networks (i) from a network operator's point of view, (ii) at device level, and (iii) for the establishment of data connections. General statistics on signaling and device behavior, such as the observed IoT traffic volume and message composition, are given before device classes with statistically different signaling behavior are identified. Finally, by characterizing the aggregated signaling traffic, we show that the Markov assumption, widespread in standardization and literature, regarding the aggregated arrival process for data connections does not apply in reality.**

***Index Terms*—IoT, mobile networks, signaling, dataset, modeling, measurement, GTP, SIGTRAN**

I. Introduction

Assessing the traffic carried in a communication network is an important task in order to operate a network successfully. In particular with new types of traffic that have fundamentally different characteristics, such as the traffic occurring in the context of the Internet of Things (IoT) [1–5], those assessments are essential to plan and operate networks accordingly.

With IoT traffic, a number of new characteristics arise in the networks. By 2022, it is expected to have ~18 billion IoT devices, 1.5 billion of which use cellular connectivity [6, 7]. These devices span heterogeneous areas such as industrial, healthcare, residential, automotive, sports, and entertainment. Between 2016 and 2022 there is an average growth of 21% per year, driven by new use cases [7]. This growth in numbers as well as the heterogeneity raises the question of the scalability of the underlying infrastructure. Further, the characteristics of traffic are fundamentally changing with IoT, especially given the growth in machine-to-machine communications [8]. For 2020 about 41% or 12.86 billion IoT devices are installed as smart home devices [9]. The traffic generated by smart home IoT devices differs from the traffic generated by conventional devices [10]. In general, a mixture of machine-driven and event-driven traffic patterns is expected for upcoming IoT traffic [2].

From a technical perspective, this change in traffic is met by a different handling of IoT devices in the networks compared to previous mobile use cases. Newly emerging IoT mobile virtual network operators (MVNOs) run specialized service platforms providing worldwide device connectivity by leveraging already existing infrastructure of mobile network operators (MNOs). By providing global coverage to their customers through roaming agreements, MVNOs can provide their service without setting up a dedicated physical cellular network infrastructure. This technical setup results in a mixture of traffic with unknown characteristics in the mobile core, since traffic comes from different locations and devices worldwide, different applications. In addition, different customers are being aggregated through the usage of few carrier networks [11]. This leads to unknown signaling traffic in the control plane, which results in significant uncertainties for the operation of networks [12, 13]. Overall, technical questions such as the volume of expected signaling traffic, the identification of certain classes of devices, or the identification of possible approaches to system optimization become important.

To provide a first step towards these issues, this paper aims at characterizing IoT signaling traffic for mobile networks (i) from a network operator's point of view, (ii) at device level, and (iii) for data connection establishment.

A large scale dataset is obtained by monitoring signaling transport (SIGTRAN) as well as 3GPP GPRS Tunneling Protocol (GTP) signaling traffic of over 270k IoT devices in cooperation with an MVNO that provides global IoT connectivity through over 500 roaming partners in 192 countries. In this work, we dissect the signaling behavior of devices using 2G/3G network connectivity and provide a broad overview regarding the signaling volume for both protocol stacks as well as a detailed evaluation regarding the occurrence of specific signaling patterns in the case of IoT traffic. Furthermore, we identify features characterizing the signaling behavior of single IoT devices and perform a device classification based on the identified signaling characteristics. We evaluate the signaling behavior exhibited by devices of different device classes and show that the devices of different classes exhibit statistically significant differences regarding their signaling traffic. Finally, we dissect the arrival process of new data connections as a proxy for system load and present approaches to model the aggregated arrival process of the observed devices. We show that the Markov assumption, widespread in standardization and literature, regarding the aggregated arrival process for data

connections does not apply in reality. The methodology used to obtain the results as well as the structure of the document is shown in Figure 1.

The work is based on two major contributions. First, a large-scale dataset is presented containing signaling traffic of a global deployment of IoT devices from various use cases. Second, the dissection of signaling behavior and a subsequent proposal of classification metrics is given to distinguish IoT devices via their signaling traffic with an analysis of the corresponding arrival processes.

The remainder of this paper is structured as follows. Section II provides an overview of research work within the area of mobile IoT traffic. Section III introduces background information on MVNOs and their architecture including relevant signaling procedures. A general description of the dataset and preprocessing steps for preparing the raw data is provided in Section IV. Section V dissects the dataset and presents detailed general statistics. In Section VI, the classification mechanism and results are presented. Section VII deals with the analysis of the arrival process of new PDP context establishments. Finally, Section VIII concludes this work.

II. Related Work

Previous related research in the field of IoT traffic falls into two categories. Investigation of traffic characteristics and datasets related to IoT deployments on the one hand and research regarding the signaling efforts of both machine to machine (M2M)-centric mobile networks as well as common mobile networks on the other hand. Studies investigating the traffic characteristics of IoT devices at a large scale, i.e. evaluating global scenarios including roaming devices, are still lacking from the general literature. Further, while all of the following studies investigate traffic patterns of IoT devices, they all focus on data plane traffic.

Regarding research on IoT traffic characteristics, the authors of [2] have compiled a taxonomy of available traffic patterns observed in IoT networks and investigated the applicability of the Poisson approximation in the context of IoT traffic. Further research in the area of IoT traffic modelling compares human generated traffic to M2M traffic [10, 14], proposed source traffic [5] and traffic models [15] as well as classification mechanisms for smart devices using WiFi [4] or general communication networks [1]. Finally, there exist studies analysing M2M communication [16]. However, all of the above focus on data plane traffic. Hence, when it comes to analyzing signaling traffic of mobile IoT devices beyond a single MNO, to the best of our knowledge, no other studies are available at the time of writing.

III. Background

This section introduces the general architecture of an MVNO and provides a brief background on the signaling traffic that occurs when data roaming is carried out as in networks with MVNOs.

A. MVNO Platform Description

Figure 2 shows an abstract scenario of a device roaming in 2G/3G. Here, the roamer connects to the VPLMN[1] SGSN[2] and VLR[3] components for data connectivity with GTP[4] and SIGTRAN[5] connectivity with MAP[6], respectively. The visiting network components subsequently communicate with the corresponding services, HLR[7] for SIGTRAN, GGSN[8] for GTP, in the home network. The colored marks in Figure 2 indicate the points of measurement at which signaling traffic has been monitored for the dataset used in this work since these are the only points for incoming signaling traffic for an MVNO. Note that the shown architecture is strongly simplified and only shows messaging flows relevant for the collected dataset, namely signaling between SGSN and GGSN, HLR as well as between VLR and HLR.

B. Roaming Signaling Procedures

During mobile roaming, devices need to perform two main procedures to establish either network connectivity or data connectivity. With only network connectivity, IoT devices can only send or receive calls or text messages. Based on preliminary investigations, both modes have proven to be representative of IoT MVNOs, as some devices save energy and only perform network attachment, while others establish a full connection to the network with data connectivity.

For network attachment, devices perform an *attach procedure* to register with the visiting network VLR. The roamer sends an attach request to the VLR that includes, among other information, its identity (i.e. IMSI). Subsequently, the VLR initiates the *authentication procedure* with the corresponding home network HLR. After authentication, the VLR uses the *update location procedure* to update the location information stored at the HLR and thereby completes the procedure to establish network connectivity without data connectivity.

Analogously, signaling during data roaming consists of similar procedures between the visited SGSN (VSGSN) and HLR. Instead of an *update location procedure*, the VSGSN initiates an *update GPRS location procedure*. Finally, the roamer can activate a new PDP context by requesting a home network APN (Access Point Name). The roamer's request triggers the VSGSN to resolve the APN and a following *PDP Context Request* towards the home network GGSN (HGGSN). After successful completion, the roamer is able to use data connectivity.

IV. Dataset

A. Data Description and Processing

The raw data collected in this work consists of all SIGTRAN and GTP messages between HLR and VLR as well

[1] Visiting Public Land Mobile Network
[2] Serving GPRS Support Node
[3] Visitor Location Register
[4] 3GPP GPRS Tunneling Protocol
[5] SIGnaling TRANsport protocols (extension of the SS7 protocol family)
[6] ETSI/3GPP Mobile Application Part
[7] Home Location Register
[8] Gateway GPRS Support Node

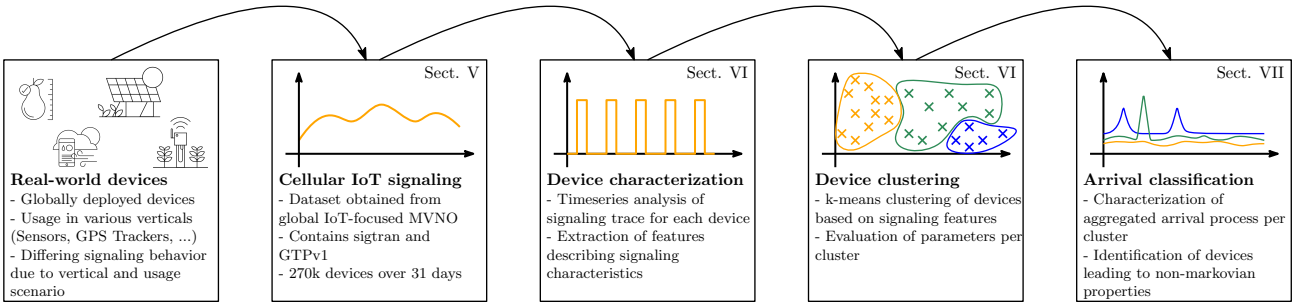

Fig. 1: Methodology used to obtain the results from the creation of a dataset to the characterization of important key figures for IoT signaling traffic in mobile networks.

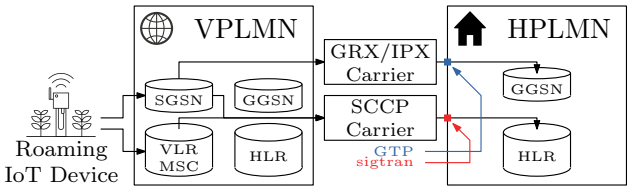

Fig. 2: Network architecture overview (roaming), left: visited public land mobile network, right: subscriber's home network with aggregated traffic ingress and measuring points.

as between HLR and VSGSN in a roaming scenario. Consequently, all messages belong to the ETSI/3GPP Mobile Application Part (MAP) protocol as part of the SS7 signaling system (ITU-T Q.700-series) or the 3GPP GPRS Tunneling Protocol (GTP) respectively. This includes signaling for *authentication*, *network attachment*, *data connectivity* as well as *mobility*. Table I provides a list of messages captured and parsed in the scope of this work. The table also shows the corresponding dialog classification based on message sequences.

The messages contained in the raw dataset are assembled into *Dialogs* using Apache Spark. To this end, a state-machine keeps track of the current signaling state of each device throughout the dataset and matches messages belonging to the same dialog. Here, a dialog is defined as a single signaling interaction between serving and home network and each dialog consists of all messages related to the initial request. A dialog is considered finished when the corresponding response has been captured. Hence, in accordance with the specifications of the SS7 signaling system as well as the GTP protocol, the dialog assembly process generates the dialog types presented in Table I.

These dialog types have been selected for further study as they, as well as their corresponding errors, contribute 95% of the total observed traffic volume. All remaining dialogs have been combined into groups labeled *OTHER* in the case of SIGTRAN and *PDP_OTHER* in the case of GTP. These dialogs are mostly related to SMS transmission and interrogation requests, which have not been evaluated in this work.

B. Dataset Overview

The dataset has been collected between 01.01.2020 and 31.01.2020. Figure 3 shows a timeseries over the whole month with the number of million dialogs per hour depicted along the y-axis. It can be seen that the timeseries exhibits gradual growth in signaling traffic over the monitored period. The blue line shows a linear regression with a constant of 0.71 million dialogs and a coefficient of 0.00042 million dialogs per hour. Furthermore, the data exhibits a cyclic pattern with 31 peaks, exactly the number of observed days. On January 8th an operator outage lead to a significant signaling incident, inducing roughly fourfold signaling traffic for about 20 minutes before returning to baseline. The specific reason for the incident is unfortunately unclear. In total, signaling traffic from 346 different mobile networks in 192 countries has been observed during the measurement period.

Table II presents an overview of key characteristics regarding the dataset. Between the two monitored signaling types GTP and SIGTRAN, roughly 274k devices have been observed, 84.1% of which have established at least one GTP tunnel in the observed time frame. In total, devices have generated around 1.4 billion signaling messages, 72.3% of which relate to successful signaling procedures. 11.9% are related to technical errors and 14% have been actively rejected by the system. After assembling the raw data into dialogs, about 650 million dialogs could be identified, 69.2% of which have been successful, 12.5% have failed and 15.2% have been rejected by the system. A more detailed dissection of errors and rejected messages is provided in Section V-B.

V. Global IoT Statistics for the Dataset

This section summarizes general IoT statistics for the dataset before presenting a detailed decomposition of the dataset. Finally, we present temporal correlations within the dataset and show common signaling patterns that have been observed.

A. General Statistics

On average, 172k unique IoT devices have been observed per day. A device is counted as active if either a successful updateLocation (UL) dialog after authentication or PDP context creation (PDP_CREATE) has occurred. Broken down to

TABLE I: Dialog types generated by assembling corresponding messages. op code in brackets. $n \in [1, \infty[$.

Dialog	Abbreviation	Contained Messages	Short Description
sendAuthenticationInfo	SAI	sendAuthenticationInfo Request(56)→ sendAuthenticationInfo Response(56)	VLR, VSGSN request authentication vectors
updateLocation	UL	updateLocation Request(2)→ $n\times$(insertSubscriberData(7) Request→ insertSubscriberData Response(7))→ updateLocation Response(2)	VLR updates location information at HLR
updateGprsLocation	UL_GPRS	updateGprsLocation Request(23)→ $n\times$(insertSubscriberData Request(7)→ insertSubscriberData Response(7))→ updateGprsLocation Response(23)	VSGSN updates location information at HLR
cancelLocation	CL	cancelLocation Request(3)→ cancelLocation Response(3)	HLR deletes location information at VLR, VSGSN
create pdp context	PDP_CREATE	createPDPContext Request(16)→ createPDPContext Response(17)	VSGSN establishes new GTP tunnel
update pdp context	PDP_UPDATE	updatePDPContext Request(18)→ updatePDPContext Response(19)	VSGSN updates existing GTP tunnel
delete pdp context	PDP_DELETE	deletePDPContext Request(20)→ deletePDPContext Response(21)	VSGSN closes existing GTP tunnel

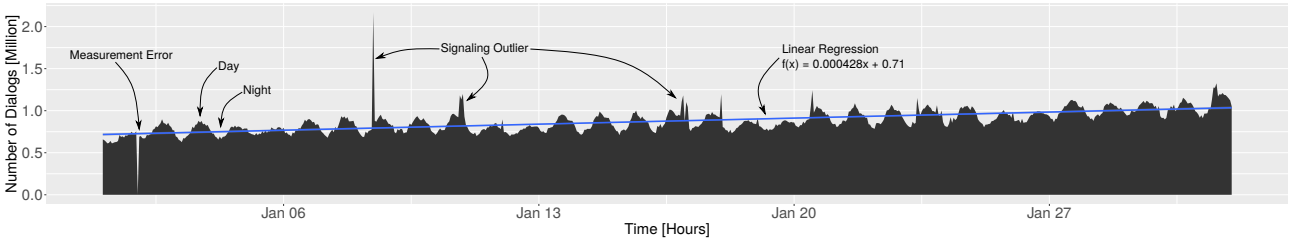

Fig. 3: Number of signaling dialogs over time in January 2020.

TABLE II: Dataset overview and key data points.

Protocol	Devices	Operators	Countries	Type	Messages Abs.	Messages Frac.		Dialogs Abs.	Dialogs Frac.	
SIGTRAN	274 184 (100%)	346	192	Error	14 400 550	0.010	0.587	3 849 272	0.006	0.551
				Reject	198 425 402	0.140		99 212 701	0.152	
				Success	595 878 305	0.419		235 523 346	0.361	
				Unknown	26 165 873	0.018		20 829 090	0.032	
GTPv1	230 602 (84.1%)	191	152	Error	154 987 535	0.109	0.413	77 593 654	0.119	0.449
				Success	431 698 276	0.304		215 848 695	0.331	
				Unknown	186	0.000		9	0.000	

hours, roughly 55k devices are active on average. In total, all devices generate an average traffic volume of 20 million signaling dialogs per day, 875k per hour, or 244 dialogs per second within our dataset.

With regard to modeling purposes, the average number of signaling dialogs per day and device amounts to 103 with a standard deviation (sd) of 1 485 and a coefficient of variation (c) of 14.4. The high coefficient of variation indicates high heterogeneity among devices.

Out of the 230k devices using data connectivity, on average, a device establishes 21.9 connections (sd: 84.7, c: 3.86) per day. The average duration is 2 890 seconds (sd: 20 205, c: 6.99). The high variation here is a further indicator for highly heterogeneous behavior of the devices active in the dataset.

B. Errors and Rejected Dialogs

Of the 1.4 billion signaling messages, 27.7% relate to unsuccessful signaling procedures, with errors and rejected messages contributing 12.5% and 15.2%, respectively. An error is defined as an invalid sequence of signaling messages, such as incomplete or out of order interactions. As it is nearly impossible to identify the reason for incomplete dialogs, the observed errors are not evaluated in more detail at this point.

Some dialogs have been actively rejected by the home network. These contain mostly requests from devices equipped with SIM cards that have not been activated, are no longer ac-

TABLE III: Dialog composition of signaling trace.

Protocol	Dialog Type	Abs.	Frac.
SIGTRAN	SAI	172 618 050	0.26
	SAI_REJECT	87 962 499	0.13
	UL	31 013 855	0.05
	CL	21 343 648	0.03
	OTHER	20 829 090	0.03
	UL_REJECT	11 049 617	0.02
	UL_GPRS	10 547 793	0.02
	UL_GPRS_ERROR	1 726 272	0.00
	UL_ERROR	1 200 272	0.00
	SAI_ERROR	860 920	0.00
	UL_GPRS_REJECT	200 585	0.00
	CL_ERROR	61 808	0.00
GTPv1	PDP_CREATE	93 915 264	0.14
	PDP_DELETE	93 236 365	0.14
	PDP_CREATE_ERROR	77 556 648	0.12
	PDP_UPDATE	28 697 066	0.04
	PDP_DELETE_ERROR	36 009	0.00
	PDP_UPDATE_ERROR	997	0.00
	PDP_OTHER	9	0.00

TABLE IV: Dialog sequences with at least 1% contribution (73% of total dialogs).

Dialog Sequence	Abs.	Frac.
SAI_REJECT	46 265 888	0.157
SAI	32 634 374	0.110
PDP_DELETE	19 907 081	0.068
SAI_REJECT→SAI_REJECT	17 088 982	0.058
PDP_CREATE	14 012 643	0.048
PDP_DELETE→PDP_CREATE	13 215 440	0.045
PDP_CREATE→PDP_DELETE	8 477 924	0.029
SAI→SAI	8 376 366	0.028
UL	8 059 256	0.027
PDP_UPDATE	6 094 100	0.021
SAI→PDP_CREATE	5 682 473	0.019
PDP_CREATE_ERROR	5 046 938	0.017
UL_REJECT	4 624 730	0.016
SAI→UL	4 315 840	0.015
UL_GPRS	4 051 486	0.014
SAI→PDP_CREATE→PDP_DELETE	3 710 986	0.013
SAI→PDP_DELETE→PDP_CREATE	3 705 860	0.013
SAI→SAI→PDP_CREATE	3 170 721	0.011
PDP_CREATE_ERROR→PDP_CREATE_ERROR	3 123 481	0.011
UL_REJECT→SAI_REJECT	2 937 871	0.010

tive or are not allowed to establish data or phone connectivity. In this context, two reasons for rejected dialogs are prevalent.

Inactive SIMs. Of 650 million dialogs, roughly 13.4% are rejected due to devices with inactive SIM cards. These dialogs are generated by 8.3% of the devices (23 002 devices) that have at least one of their dialogs rejected with *UnknownSubscriber*. Hence, this relatively small number of devices is responsible for 13.4% of total signaling dialogs.

Invalid Roaming Attempts. Accordingly, 11 million dialogs (1.7%), generated by 15.2% of the devices (41 874 devices), are rejected due to invalid roaming partner selection. This occurs if a device selects a visited network that, e.g. due to policy reasons or customer configuration, cannot be used as a roaming partner. These dialogs fail with the message `RoamingNotAllowed`.

C. Sequence of Messages

Table III decomposes the dataset according to the observed dialog types as well as how much of the total signaling volume each type contributes. The last column shows the arrival rate for each dialog type. Note that arrivals during the incident shown in Figure 3 have been removed here.

The table shows the significant portion of messages attributed to inactive devices (SAI_REJECT) as well as invalid roaming attempts (UL_REJECT). Furthermore, it can be seen that a large fraction of errors is related to GTP context creation. These errors occur mostly due to the APN rejecting the device (10.7%) or `APN Congestion` (1%).

Temporal Correlation of Signaling Dialogs. In order to improve the understanding of device behavior, Table IV shows all dialog sequences observed in the dataset that contribute at least 1% to the total number of dialog sequences. A dialog sequence is thereby defined as a sequence of dialogs without a significant pause inbetween dialogs. More specifically, the timeseries of each device has been divided into bins of one minute with a device being active if at least one dialog occurs within each bin. The resolution of one minute has been selected as it coincides with the timeout for activity in the monitored mobile core. Based on this activity diagram, a sequence is defined as all dialogs occurring in bins with activity without there being a bin without activity inbetween.

Table IV shows that a significant portion of the resulting sequences consist of three or less dialogs with the majority only featuring a single dialog. We can also observe that a significant fraction of sequences occurs due to devices closing an already established PDP context, directly followed by the creation of a new tunnel, meaning devices don't establish a tunnel, send data and close the tunnel. Instead, devices close and reestablish PDP contexts, send their data and leave the tunnel open until the next iteration.

VI. Device Classification

As already observed earlier, the signaling behavior differs significantly between devices. Hence, in the following section, we establish a set of device features extracted from purely evaluating signaling traffic and show that devices can be clustered using the k-means algorithm. Table V shows the selected features. Thereby, the error as well as reject rate denote the fraction of dialogs resulting in errors or being rejected by the system, respectively. Further, the grade of periodicity of a device is defined as the sum of the autocorrelation values of the three most significant lags observed while calculating the autocorrelation for lags between 1 and 1500 in minutes. Thus, this evaluation is able to identify periods of up to 24 hours for devices with up to three significant periods.[9] Although further features have been evaluated, our investigations have shown that this minimal feature set is enough to classify devices observed in the trace. Note that Section VII introduces an

[9]Devices with single period p also exhibit high autocorrelation values at $2p$ and $3p$.

TABLE V: Features used for device classification.

Featurename	Description
Error Rate	Percentage of dialogs resulting in errors
Reject Rate	Percentage of dialogs resulting in reject
Periodicity	Sum of autocorrelation of three most significant lags

additional feature to further refine the clustering performed here.

Evaluating the within-cluster sum of squares using the elbow-method, $k = 5$ has been decided to be a suitable number of clusters to perform the k-means algorithm on. The resulting clusters correlate with the expected outcome according to expert knowledge contributed by the MNVO. In order to visualize the results of the clustering, Figure 4 shows the biplot resulting from the primary component analysis and plotting the two most significant primary components against each other. Each point represents one device and the colors represent the assigned cluster. Additionally, the arrows and labels indicate the influence of the original features on the shown primary components. This visualization allows a visual identification of the relation between features and the resulting clusters. We identify the five clusters as *Non-Periodic* (Green), *Semi-Periodic* (Pink), *Periodic* (Blue), *High Error Rate* (Red) and *High Reject Rate* (Yellow).

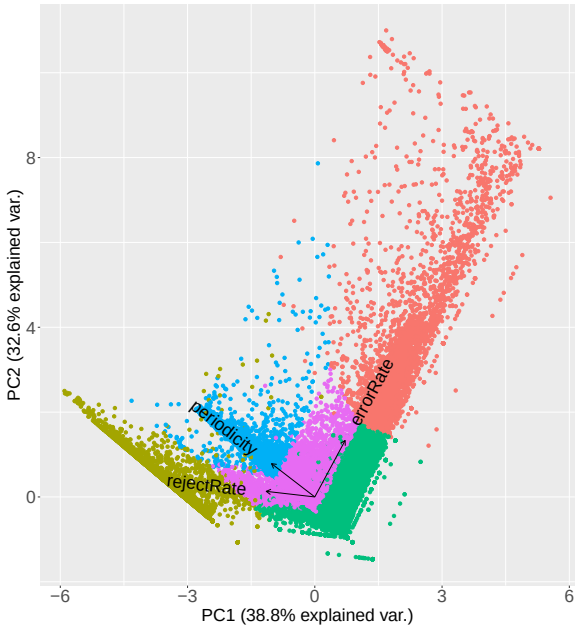

Fig. 4: Biplot of two most significant primary components and influencing original features.

Figure 5 shows the distribution of devices among those five identified clusters. It can be seen that a significant portion of devices exhibits non-periodic or semi-periodic behavior, with the other classes containing only 16% of devices altogether.

A. Signaling Characteristics by Device Class

In order to validate the classification with respect to differing device behavior within the identified classes we now compare signaling characteristics between classes. Figure 6 shows the empirical cumulative distribution function (ECDF) of inter sequence inter arrival times, meaning the distribution of the time between activity phases per device over all devices within one device class.

It can be seen that the different device classes exhibit differing behavior regarding their sequence inter arrival times. Periodic devices show clear peaks at inter arrival times of 30, 60, and 120 minutes. The two sided Kolmogorov-Smirnov (KS) test confirms statistically significant differences regarding the sequence interarrival times.

Analogously, Figure 7 shows the ECDF of the GTP context duration for the same set of device classes. Note that the classes *High Error Rate* and *High Reject Rate* have been omitted here, since their signaling behavior is strongly influenced by their high number of erroneous and rejected dialogs and a comparison to regular devices regarding their GTP context durations would be invalid. Additionally, the y-axis has been limited to $[0.7, 1]$ to make the effect more visible. The KS test once again shows statistically significant differences for each pair of distributions.

While less prominent, periodic devices again show peaks at context durations 30, 60 and 120 minutes, indicating a correlation between the two values. As already shown in Table IV, a significant fraction of sequences intend to either create or delete PDP contexts, further validating this observation.

Finally, when it comes to the distribution of dialog types used by devices of each class, with the exception of *high error rate* and *high reject rate*, no significant difference between classes could be observed. In the case of the error and reject classes, the fraction of rejected and erroneous dialogs is, by definition, higher compared to other classes.

VII. Aggregated PDP Context Arrival Process

Finally, in order to better understand the underlying arrival process of the devices resulting from the performed classification, we examine the aggregated arrival process of newly established data connections. To this end, we dissect the arrival process of PDP_CREATE dialogs. This subset of messages has been selected as it acts as a proxy for the general system load for MVNOs since a successful PDP_CREATE requires successful SAI, UL, and UL_GPRS dialogs beforehand.

As already stated before, large parts of literature and standardization [2, 15, 17, 18] assume Markov properties when dealing with the aggregated arrivals in large scale IoT environments. The Markov assumption is suitable as the superimposed traffic of an infinite number of sources exhibits memorylessness, according to Palm-Khintchine [19]. However, in the following, we show that the assumption does not hold true in reality due to the presence of time synchronous devices, but can be restored through additional device classification and filtering. Note that the synchronous behavior of devices is expected to stem from firmware implementation specifics

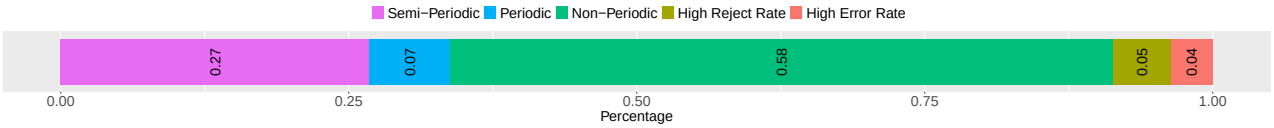

Fig. 5: Devices classified by signaling behavior via k-means.

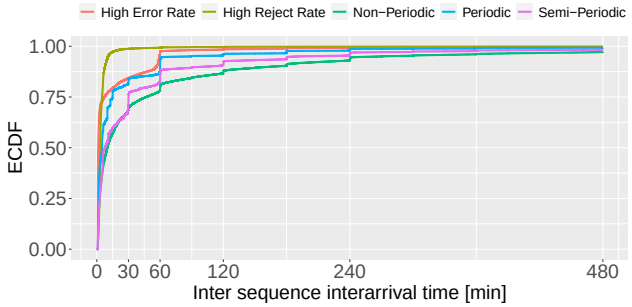

Fig. 6: ECDF of inter sequence interarrival times by device class.

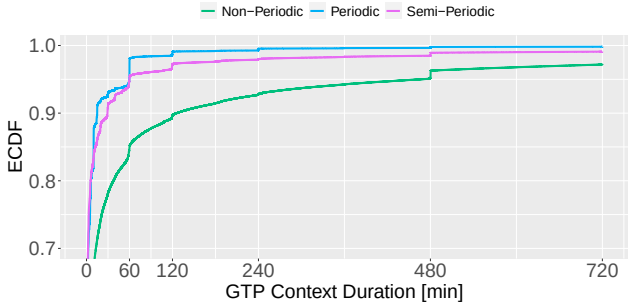

Fig. 7: ECDF of GTP context duration by device class.

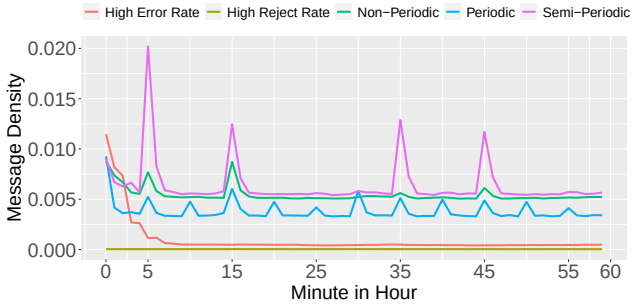

Fig. 8: Message probability density function of PDP_CREATE dialogs for the probability that messages occur at a specific minute within any hour during the entire trace.

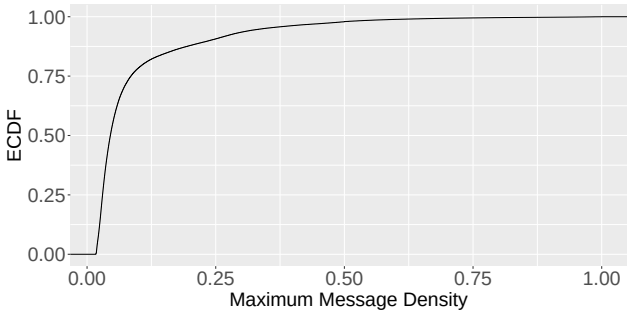

Fig. 9: ECDF of maximum message density over all devices with at least 30 activity phases.

rather than actual synchronization between independent devices. Devices seem to be programmed to transmit data at fixed times, instead of fixed intervals, resulting in quasi synchronous behavior in the aggregated traffic.

To this end, Figure 8 shows the message density over one hour for all device classes for the whole trace. Hence, the y-axis shows the probability of an arbitrarily selected dialog happening within the corresponding minute along the x-axis. The plot is shown to visualize the synchronized behavior of devices within the trace. Although devices are expected to not synchronize their signaling behavior with other devices, clear synchronization patterns can be observed within the trace. Specifically, increased density between 0 and 5 minutes as well as peaks at minutes 0, 5, 15, 35 and 45 can be observed. This behavior violates the assumption of memorylessness, as in a process exhibiting the Markov property, the message density would need to be constant.

However, based on the message density function of single devices, we are able to identify time synchronous devices and hence divide the arrival process in synchronized and non-synchronized devices. To this end, we examine the maximum message density for messages of a specific device, as is shown in Figure 8 for all devices. This classification is based on the assumption that the message density function of a non-periodic, non-synchronized device would follow a uniform distribution. Figure 9 shows the ECDF of the maximum message density over all devices with at least 30 activity phases. This limitation is introduced as the density value is not significant for devices with less than one activity per day.

Based on the distinct knee observed in the figure, devices with a maximum message density of at least 0.075 are classified as synchronized. This results in 23% of devices being classified as synchronized. Note that this classification has been performed in addition to the clustering performed earlier, so each device can be classified by both their signaling behavior as well as their synchronicity. Figure 10 shows the same plot as Figure 8 with devices split into synchronized and non-synchronized classes.

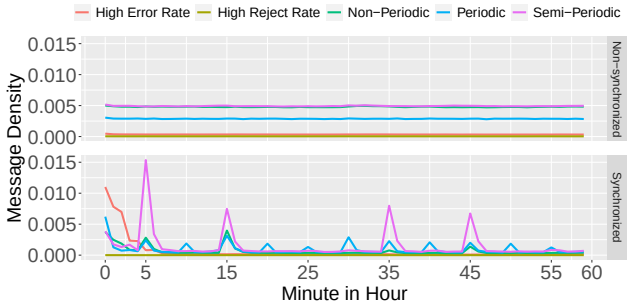

Fig. 10: Message probability density of PDP_CREATE dialogs over minutes within anyab hour of the trace per synchronization class.

It can be seen that non-synchronized devices now exhibit a uniform message density distribution while synchronized devices feature the peaks observed before as well as low baseline density.

Base on these observations we can now assume the arrival process of non-synchronous devices to exhibit memorylessness and consequently negative exponentially distributed interarrival times. To this end, Figure 11 shows the Q-Q-plot of the empirical interarrival times of all non-synchronized devices, irrespective of device class, and the corresponding negative exponential fit. The red marks show the 10% to 90% quantiles, the gray marks show the 1% to 99% quantiles. It can be seen that the interarrival times of the aggregated process of all non-synchronized devices can be closely approximated using an exponential distribution as it is the case for Markov processes. Furthermore, the interarrival times exhibit no significant autocorrelation with the largest observed value being 0.017 for lags between 1 and 1000.

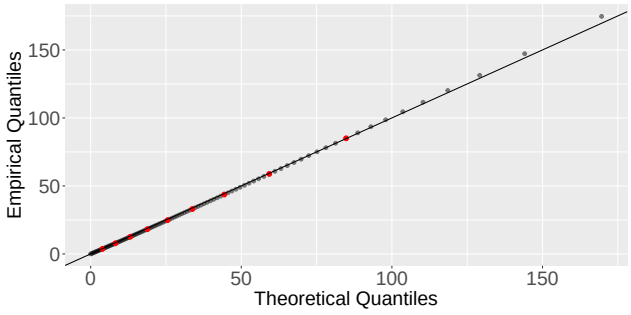

Fig. 11: Q-Q-plot of interarrival times of non-synchronized devices.

VIII. Conclusion

In this work, we presented the results of our analysis of a 31 day IoT signaling trace containing more than active 270k IoT devices of an MVNO operating worldwide. We found that about 84% of the observed 2G/3G devices use data connectivity with the remaining 16% only using network connectivity for phone and text messaging as they do with circuit-switched services. We have shown that devices exhibit significant differences regarding their signaling behavior and extracted features that allow the modelling of different device classes based on the rate of erroneous and rejected dialogs as well as the periodicity of devices. We have shown that devices identified by this classification mechanism exhibit statistically significant differences when it comes to their signaling behaviour. Finally, by evaluating the aggregated arrival process of new data connections, we have shown that the often assumed memorylessness does, surprisingly, not hold true in reality due to the presence of time synchronous devices, but can be restored through additional classification and filtering steps. Next steps include the behavioral classification of different IoT devices to better understand their behavior and be able to move towards a detailed source traffic model of single devices as well as performance models of large scale systems.

Acknowledgment

This work is funded by the Bavarian Ministry of Economics, Regional Development and Energy within the project *5CALE* as part of the R&D program for information and communication technology in the field of 5G mobile communications. The authors alone are responsible for the content.

References

[1] H. Tahaei, F. Afifi, A. Asemi, F. Zaki, and N. B. Anuar, "The rise of traffic classification in iot networks: A survey," *Journal of Network and Computer Applications*, vol. 154, p. 102 538, 2020.

[2] F. Metzger, T. Hoßfeld, A. Bauer, S. Kounev, and P. E. Heegaard, "Modeling of aggregated iot traffic and its application to an iot cloud," *Proceedings of the IEEE*, vol. 107, no. 4, pp. 679–694, 2019.

[3] A. Sivanathan, D. Sherratt, H. H. Gharakheili, A. Radford, C. Wijenayake, A. Vishwanath, and V. Sivaraman, "Characterizing and classifying iot traffic in smart cities and campuses," in *2017 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, IEEE, 2017, pp. 559–564.

[4] A. Sivanathan, H. H. Gharakheili, F. Loi, A. Radford, C. Wijenayake, A. Vishwanath, and V. Sivaraman, "Classifying iot devices in smart environments using network traffic characteristics," *IEEE Transactions on Mobile Computing*, vol. 18, no. 8, pp. 1745–1759, 2018.

[5] M. Laner, P. Svoboda, N. Nikaein, and M. Rupp, "Traffic models for machine type communications," in *ISWCS 2013; The Tenth International Symposium on Wireless Communication Systems*, VDE, 2013, pp. 1–5.

[6] Global Mobile Data Traffic Forecast, "Cisco visual networking index: Global mobile data traffic forecast update, 2017–2022," *Update*, vol. 2017, p. 2022, 2019.

[7] P. Cerwall, P. Jonsson, R. Möller, S. Bävertoft, S. Carson, I. Godor, *et al.*, "Ericsson mobility report," *On the Pulse of the Networked Society. Hg. v. Ericsson*, 2017.

[8] T. Taleb and A. Kunz, "Machine type communications in 3gpp networks: Potential, challenges, and solutions," *IEEE Communications Magazine*, vol. 50, no. 3, pp. 178–184, 2012.

[9] Statista, "The Internet of Things (IoT)* units installed base by category from 2014 to 2020 (in billions)," Tech. Rep., 2018. [Online]. Available: https://www.statista.com/statistics/370350/internet-of-things-installed-base-by-category/.

[10] I. Cvitić, P. Zorić, T. Kuljanić, and M. Musa, "Analysis of network traffic features generated by iot devices," in *XXXVII Simpozijum o novim tehnologijama u poštanskom i telekomunikacionom saobraćaju*, Dec. 2019.

[11] A. Lutu, B. Jun, A. Finamore, F. Bustamante, and D. Perino, "Where things roam: Uncovering cellular iot/m2m connectivity," *arXiv preprint arXiv:2007.13708*, 2020.

[12] C. Schwartz, F. Lehrieder, F. Wamser, T. Hoßfeld, and P. Tran-Gia, "Smart-phone energy consumption vs. 3g signaling load: The influence of application traffic patterns," in *2013 24th Tyrrhenian International Workshop on Digital Communications-Green ICT (TIWDC)*, IEEE, 2013, pp. 1–6.

[13] G. Gorbil, O. H. Abdelrahman, M. Pavloski, and E. Gelenbe, "Modeling and analysis of rrc-based signalling storms in 3g networks," *IEEE Transactions on Emerging Topics in Computing*, vol. 4, no. 1, pp. 113–127, 2015.

[14] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, and J. Wang, "A first look at cellular machine-to-machine traffic: Large scale measurement and characterization," *ACM SIGMETRICS performance evaluation review*, vol. 40, no. 1, pp. 65–76, 2012.

[15] F. Wamser, P. Tran-Gia, S. Geißler, and T. Hoßfeld, "Modeling of traffic flows in internet of things using renewal approximation," in *Advances in Optimization and Decision Science for Society, Services and Enterprises*, Springer, 2019, pp. 483–492.

[16] F. Malandra, S. Rochefort, P. Potvin, and B. Sansò, "A case study for m2m traffic characterization in a smart city environment," in *Proceedings of the 1st International Conference on Internet of Things and Machine Learning*, 2017, pp. 1–9.

[17] R. R. Tyagi, F. Aurzada, K.-D. Lee, and M. Reisslein, "Connection establishment in lte-a networks: Justification of poisson process modeling," *IEEE Systems Journal*, vol. 11, no. 4, pp. 2383–2394, 2015.

[18] 3GPP, "Geran improvements for machine-type communications (mtc)," *Technical report (TR)*, 2014.

[19] C. Palm, "Intensitätsschwankungen im Fernsprechverker," *Ericsson Technics*, 1943.