

Device-Specific Traffic Characterization for Root Cause Analysis in Cellular Networks

Peter Romirer-Maierhofer^(✉), Mirko Schiavone, and Alessandro D'Alconzo

Forschungszentrum Telekommunikation Wien (FTW), Vienna, Austria
{romirer,schiavone,dalconzo}@ftw.at

Abstract. Nowadays mobile devices are highly heterogeneous both in terms of terminal types (e.g., smartphones versus data modems) and usage scenarios (e.g., mobile browsing versus machine-to-machine applications). Additionally, the complexity of mobile terminals is continuously growing due to increases in computational power and advances in mobile operating systems. In this scenario novel traffic patterns may arise in mobile networks, and it is highly desirable for operators to understand their impact on the network performance. We address this problem by characterizing the traffic of different device types and Operating systems, analyzing real traces from a large scale mobile operator. We find the presence of highly time synchronized spikes in both data and signaling plane traffic generated by different types of devices. Additionally, by investigating a real case, we show that a device-specific view on traffic can efficiently support the root cause analysis of some type of network anomalies. Our analysis confirms that large traffic peaks, potentially leading to large-scale anomalies, can be induced by the misbehavior of a specific device type. Accordingly, we advocate the need for novel analysis methodologies for automatic detection and possibly mitigation of such device-triggered network anomalies.

1 Introduction

In the last decade operators of mobile networks have witnessed the spread of heterogeneous mobile devices. Their heterogeneity stems from several respects: Mobile devices include different terminal types, operating systems and support a large variety of different applications. The level of device complexity is increased by the permanent evolution of their computational power, the introduction of novel mobile services and upgrades of the respective mobile APPs and Operating Systems (OSs). In such a heterogeneous and complex network scenario it is very appealing for mobile network operators *i*) to verify that device-specific traffic patterns are conform with the assumptions taken during dimensioning of network capacity, and *ii*) to identify device-specific traffic patterns that may induce undesirable events (such as temporary overload due to large device populations issuing synchronized downloads). This aspect is even more critical due to today's dynamics of mobile OSs making the assessment of device-specific traffic patterns a moving target.

In this study we present a device-specific view on traffic behavior in an operational cellular network by first categorizing device types and mobile OSs present in the network, and second by characterizing traffic patterns at different protocol layers for the resulting device- and OS- categories. Additionally, we show how device-specific traffic characterization can be exploited for the identification of the root causes of detected network anomalies, a task which is typically resource-intensive in operational networks [1]. Monitoring device-specific traffic behavior enables the discrimination of anomalies caused by large populations of specific devices from other, network-triggered anomalies, such as e.g., malfunctioning of network equipment typically affecting all device types at the same time.

The remainder of this paper is organized as follows. We present our monitoring setup and approach to device categorization in §3. In §4.1 we characterize device-specific traffic behavior at the data plane. Additionally, we report on observed similarities and differences compared to the previous work in [2]. We investigate the device characteristics in Third Generation Partnership Project (3GPP)-specific signaling at sub-minute granularity in §4.2. To the best of our knowledge this is the first work investigating device-specific 3GPP signaling at sub-minute granularity. In §4.3 we report about a device-specific network anomaly detected from the analysis of Domain Name System (DNS) traffic for different device categories and mobile OSs. Finally, we discuss lessons learned and derive guidelines for future work in §5.

Our findings collectively suggest that the extraction and analysis of device-specific traffic patterns becomes more and more important for dimensioning and troubleshooting mobile networks. Accordingly, novel methodologies to detect and mitigate undesirable device-specific traffic behavior are required in the near future.

2 Related Works

In order to enable a consistent taxonomy for describing heterogeneous mobile devices Gansemer *et al.* presented a scalable database approach for classifying devices, according to their features [3], into three categories: smartphones, mobile phones and Personal Digital Assistants (PDAs). While the focus of that work lies on enabling the selection of the optimal devices for development of specific mobile applications, we use our device categorization for studying their network-layer behavior. The authors in [4] presented the classification of mobile devices by means of fingerprinting of their specific traffic behavior at the network and transport layer. To this purpose the authors deploy supervised machine learning techniques for processing passively monitored traffic features such as Time-To-Live or TCP congestion-window size. However, the experimental measurements presented in [4] cover only a limited number of devices, while we characterize the traffic behavior of the entire population of mobile devices within an operational mobile network. Kumar *et al.* report about large spatio-temporal differences between smart-phones and laptops in [5]. Their study relies on MAC address-based device classification and has been conducted by means of large-scale trace analysis within a campus Wireless Local Area Network (WLAN). Our work includes also additional device categories and is based on characterizing the

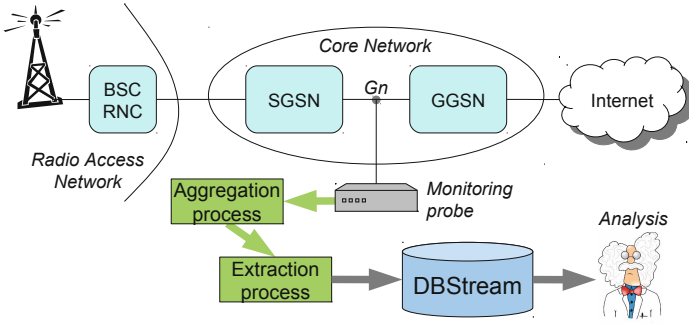


Fig. 1. Monitoring infrastructure and measurement setup

traffic behavior of a device population that is several orders of magnitude larger than the one studied in [5].

The characterization of Machine-to-Machine (M2M) device-specific traffic behavior in operational cellular networks has been presented earlier in [2, 6, 7]. The work presented in [6] includes the characterization of traffic volume behavior and radio performance of M2M devices. The authors in [2] study volume time series, session durations, mobility, applications and network performance of M2M devices. While our approach to device categorization is similar to these earlier studies, we do not limit ourselves to the characterization of M2M devices, but rather we encompass additional device categories such as tablets and USB modems. Additionally, we complement the studies presented in [2, 6, 7] by *i*) characterizing the traffic behavior of two different types of mobile operating systems, *ii*) studying the 3GPP-signaling protocol behavior and *iii*) investigating traffic behavior at shorter time-scales enabling us to characterize device synchronization at sub-minute granularity. Finally, we compare the traffic volume characteristics in the network under study with the characteristics presented in [2] and discuss observed similarities and differences.

3 Methodology

Traces were collected at the core network of a large scale European mobile operator. The results reported in the paper are based on several consecutive weeks of observation in the fourth quarter of 2013. The measurement setup is depicted in Fig. 1. Packet-level traces are captured at the Gn links between the GGSN and SGSN — for details about 3G network structure please refer to [8]. We analyze anonymized data captured by the METAWIN monitoring system developed in a previous research project [9]. This monitoring system relies on Endace DAG cards [10], and packets are recorded with Global Positioning System (GPS)-synchronized timestamps offering an accuracy of ± 100 ns.

Our measurement setup enables the investigation not only of data traffic, but also of 3GPP-specific signaling traffic, e.g., mobile data sessions set-up and tear-down, that is PDP contexts in the terminology of 3GPP (see [11] for more details). Signaling information is exploited for associating an anonymized Mobile

Device Identifier (MDID) and the observed device type to each monitored packet. This latter is achieved by extracting the Type Allocation Code (TAC) digits from the International Mobile Equipment Identity (IMEI), while the serial number digits of the IMEI are anonymized by means of an irreversible hash function. Additionally, anonymized traffic data is transformed into high-level records of e.g. aggregate IP flows or signaling events (ref. “Aggregation Process” in Fig. 1). Finally, the recorded measurements are forwarded within an extraction process (ref. Fig. 1) to DBStream [12], a PostgreSQL-based parallel data processing system. In this paper, we use DBStream as our main analysis system. Note that all results presented in this work report normalized values to protect confidentiality of business-relevant information. Therefore, we cannot report the exact number of devices in our dataset, nor the number of devices per category.

3.1 Device Categorization

Studying device-specific traffic behavior relies on a proper categorization of the device types observed in the monitored traffic. To this purpose we follow the approach presented in [2, 6]. To determine the hardware model we started inspecting the device TAC codes and matching them with the publicly available GSM Association database [13]. For categorizing M2M devices we adopt the base template scheme defined in [14]. Then, similarly to [2] we used public information (e.g., production brochures, specification sheets) to manually supplement and verify this template. Furthermore, we distinguish the categories *tablet*, *feature phone* and *smartphone*. In particular, we label as feature phone all devices that are only capable of transferring data via 2G, i.e., General Packet Radio System (GPRS) or Enhanced General Packet Radio System (E-GPRS), while we consider smartphones all devices that are at least Third Generation (3G)-capable, i.e., support data transfer via Universal Mobile Telecommunications System (UMTS) or High Speed Packet Access (HSPA). We also distinguish between *USB modems* and *modems*. The first category includes devices connected via USB to PCs or Laptops, the latter refers to PCs or Laptops built-in modems. Finally, devices specifically designed for offering 3G-connectivity via local WLAN connections are labeled as *router*.

By relying on this approach we managed to label up to 96% of the devices observed within our study. As there is no standardized definition of M2M devices some might have multiple uses (e.g., PC built-in card and M2M device modem), thus the accuracy of our categorization is affected. In these cases, we adopted a conservative approach and we labeled as M2M, devices which were exclusively advertised as such by vendors.

3.2 Operating System Classification

In addition to previous works, our study includes the characterization of traffic behavior for two different types of mobile OS in §4.3. The different OS types have been derived from publicly available TAC information. For example, the TAC codes assigned to Apple devices can be associated to the iOS operating system, whereas TAC codes assigned to Nexus devices can be labeled as Android OS).

While we are aware that such a manual labeling provides only moderate classification coverage, as shown in §4.3, it still enables us to study important OS specific traffic characteristics. For obtaining larger classification coverage, information from higher protocol layers would be required — refer to the earlier study presented in [4] for an illustrative example.

4 Results

We start our study of device-specific traffic characteristics by analyzing aggregate traffic volume behavior for different types of devices.

4.1 Device-specific Characteristics at the Data Plane

As reported in [2] the introduction of new device types, as e.g., M2M devices, induce novel traffic usage patterns that may question current assumptions for optimization of network capacity. For instance, Shafiq *et al.* report that careful allocation of network resources is required due to uplink-heavy M2M devices contradicting optimization approaches relying on downlink asymmetry of network traffic [15].




Fig. 2. CDF of ratio $\log(\text{Uplink}/\text{Downlink})$, 7 day aggregate

For investigation of this aspect in the network under study, we plot the ratio of uplink traffic volume to downlink traffic volume for different device categories over a period of 7 days in Fig. 2. We plot the ratio for the device category M2M and all other device categories separately. The latter categories are named according to their ranked share of the overall traffic volume during the observation period¹. For enabling a comparison of our results with the findings presented in [2, Fig. 2c] we also plot the ratios after taking their logarithm, referred to as Z . Negative values of Z indicate larger volume in downlink than in uplink, while positive values of Z refer to larger uplink volume than downlink

¹ In order to obfuscate business-sensitive information the specific category labels have been substituted by their rank.




Fig. 3. Downlink and uplink traffic volume time series, time bins of 5 minutes

volume. In Fig. 2 we observe a clear separation of the distribution of M2M devices and all other devices. This is consistent with the finding presented in [2, Fig. 2c]. However, while Shafiq *et al.* report that $\approx 40\%$ of smartphones have $Z < -0.4$ and $\approx 20\%$ of smartphones have more uplink than downlink traffic (i.e. $Z > 0$), in our study we cannot identify any device category exhibiting the reported qualitative shape (compare Fig. 2 and [2, Fig. 2c]). Another dissimilarity to the work in [2] is observed in the volume share of M2M devices. While we observe $Z \leq 0$ for almost 40% of M2M devices, Shafiq *et al.* report less than 10% of M2M devices exhibiting larger downlink volume than uplink volume.

For further study of the device-specific data volume behavior, in Fig. 3 we plot data volume time series for different device categories, in time bins of 5 minutes, over a period of one week in the fourth quarter of 2013. We observe diurnal patterns for all time series of Fig. 3. In case of smartphones (ref. Fig. 3a) the daily peak volume is reached in the evening hours of each day. This represents a strong difference compared with the observation provided in [2, Fig. 3a] where the authors report a decrease of overall traffic volume for the evening hours. This might be explained by the different pricing models used in the two networks. In the case of [2] users may be encouraged by higher tariff models to switch from cellular connectivity to WLAN connections available in private homes.

The local volume peaks observed in Fig. 3a at noon during lunch time are a further deviation from the results presented in [2]. This specific characteristic is not visible for any other device type (ref. Fig. 3c and Fig. 3d). Our comparison with the work in [2] suggests that even conceptually simple metrics such as e.g., volume time series may exhibit significant dissimilarities for different networks.

An aspect which should be specifically regarded when establishing generic traffic models based on observations derived from real networks.

From the time series of M2M devices (ref. Fig. 3b) we may establish following findings which are consistent with the results presented in [2, Fig. 3b]: *i*) there is a clear difference of the time series between working days and week days, *ii*) M2M devices are the only device types that show higher aggregate uplink volume than downlink volume, and *iii*) we observe distinct spikes in the traffic volume indicating the presence of M2M devices sending data traffic in a synchronized manner. The aspect of synchronized traffic patterns will be further addressed in the next section.

Summarizing our investigation of traffic volume behavior, we may confirm earlier works reporting the presence of uplink-heavy, synchronized M2M terminals. However, we report distinct differences in the aggregate time series, specifically for the smartphone device category.

4.2 Device-Specific Characteristics at the Signaling Plane

As reported in [16], 3G wireless networks are currently designed for traditional human-type communication rather than for synchronized, machine-type communication patterns as they may be triggered by M2M applications or mobile APPs relying on periodic transfers of data. Accordingly, it is highly desirable for mobile network operators to assess the effect of synchronized communication patterns onto available network capacity. This is specifically important since large increases of M2M-devices are expected until 2020. Such a perspective makes an evolution from high data rate networks to M2M-optimized low cost networks very appealing for network operators [16].

In this section we characterize the signaling behavior of different device classes. In this section we characterize the signaling behavior of different device classes. In particular we focus on the PDP-context establishment procedure. Whenever a mobile client requires the establishment of a mobile data session, it is required to send a PDP-context create request message (ref. to [8] for more details).

Fig. 4 shows the Cumulative Distribution Function (CDF) of the number of Packet Data Protocol (PDP)-context create requests for the device class M2M and the two most popular device categories, computed over an aggregation period of 7 consecutive days². We observe similar characteristics for M2M devices and the top category. While the second category converges quickly to 1, suggesting that the devices of this category are operated by means of lower amounts of PDP context create procedures. Moreover, we observe that the CDF of M2M devices converges slower to 1 than in case of other device types. This suggests that a small fraction of M2M devices relies on frequent session establishments, an aspect which we address further below.

² In order to obfuscate business-sensitive information the specific category labels have been again substituted by their rank according to the overall volume during the observation period.




Fig. 4. CDF of number of PDP-context create, different categories, 7 days (rescaled)




Fig. 5. Time series of PDP-context creates, 1 day, time bins of 10 sec

Fig. 5 depicts the overall time series of PDP-context create request messages over a period of one day. We observe a time-of-day variation in the PDP-context create request messages indicating more context establishments during day-time and less connection establishments during night hours. However, in addition to this daily cycle, we report distinct short-term spikes arising every full-hour and also smaller spikes occurring at a periodicity of 15 minutes.

For further investigation of these spikes in the connection establishment process, we plot the time series of PDP-context creates for different device classes over a period of two days in Fig. 6. Comparing Fig. 5 and Fig. 6 we discover that the reported peaks are triggered by different device classes. In particular M2M devices are responsible for the higher peaks (e.g. refer to the largest spike in the morning hours depicted in Fig. 6b). Additionally, we observe that M2M devices exhibit their largest spikes during night hours indicating that a certain population of M2M devices relies on periodic session establishments during off-peak hours. It is interesting to note that also smartphones exhibit some sort of synchronized behavior recognizable by the first three prominent spikes persistently re-occurring

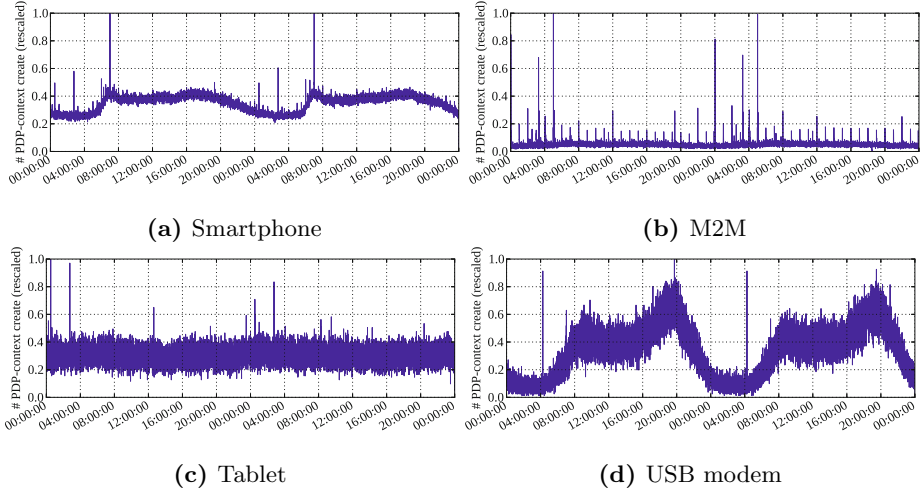


Fig. 6. Signaling: PDP-context create event time series, time bins of 10 seconds

in the morning hours (ref. Fig. 6a). By isolating these spikes for the different OS types derived from the publicly available TAC identifier, we find that those are caused by different types of mobile OSs. This suggests that specific types of OS may in fact exhibit different behavior in the context establishment process. For instance, the first two prominent spikes can be mapped to iOS-related TAC identifiers while the third peak can be associated to Android OS-specific TAC identifiers. The aspect of OS-specific signaling behavior will be further discussed in §4.3. Additionally, Fig. 6d shows distinct spikes every day at 04:00 for the device category of USB modems. This behavior may be triggered by an M2M-like application relying on periodic data transfers once a day. For further investigation of the synchronization level exhibited by the different device classes, we introduce the following indicator functions and sets.

Be e and E the generic device and the set it belongs, respectively. We denote by $d = 1, 2, \dots, 7$ the days of the week, by $m = 1, 2, \dots, 1440$ the minutes of the day, and by $w = 1, 2, \dots, 52$ the weeks of the year. For each device, let define the function:

$$\Theta_e(m, d, w) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } e \text{ creates a PDP-context within} \\ & \text{minute } m \text{ of day } d \text{ on week } w \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

For counting the number of days in the week the device e is creating a PDP-context, at minute m , we introduce the function:

$$\Sigma_e(m, w) = \sum_{d=1}^7 \Theta_e(m, d, w) \quad (2)$$

For filtering on the devices active in more than τ days per week, we define the following indicator function:

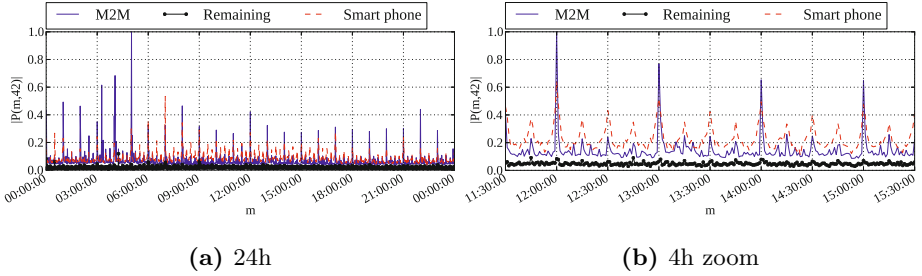


Fig. 7. Cardinality of $P_{m,42}$, different zoom levels

$$\Lambda_e(m, w) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } \Sigma_e(m, w) > \tau \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

The set of devices active at least τ days at the minute m of the week w , is defined as follows:

$$C_{m,w} \stackrel{\text{def}}{=} \{e \in E : \Lambda_e(m, w) = 1\} \quad (4)$$

In order to filter those devices that exhibit stable activity patterns across two consecutive weeks, we consider the sets $P_{m,w}$ obtained intersecting the $C_{m,w}$ with the conform set (i.e., the same minute) of the previous week:

$$P_{m,w} = C_{m,w} \cap C_{m,w-1} \quad (5)$$

Fig. 7a depicts graphical representation of the cardinality of the sets $P_{m,42}$ for the different minutes at week 42. From this plot we notice that M2M and smartphones exhibit similar behavior since both the classes tend to be synchronized with the full hour, and partially synchronized also at a periodicity of 30 and 15 minutes. We found that more than 20% of M2M devices exhibit this behavior while the share of smartphones is around 1%. In Fig 7b we report a zoom on 4 hours of Fig. 7a. We observe that the variation around the spikes is larger for smartphones compared to the case of M2M devices. This suggests that smartphones are synchronized at a coarser level, while M2M devices are tightly synchronized in time. This tight time-synchronization is highly undesirable for mobile operators, since it requires allocation of large network resources when the network is dimensioned to keep up with capacity demands during traffic peaks. Consequently, the expected future increase of M2M devices [16] suggests that a continuous monitoring of M2M-triggered communication patterns is needed, such that adequate countermeasures can be timely enforced (e.g., de-synchronization of M2M devices).

4.3 Investigation of a Device-Specific Anomaly

In this section we report about a large-scale network anomaly detected while studying the characteristics of the Domain Name System (DNS) traffic generated by each device category and for different OS.




Fig. 8. Time series of DNS requests per device category, 2 days, time bins of 1 minute

Since the vast majority of Internet applications rely on the proper functionality of the DNS, ensuring the availability and performance of DNS servers is an essential task for operators of mobile networks. Due to the emerging of synchronized traffic behavior, as e.g., shown in §4.2, it is important to know whether a large device population queries the operator’s DNS servers in a highly synchronized manner, potentially impairing DNS servers performance. To this purpose, in Fig. 8 we plot the time series of DNS queries for different device categories, over a period of two days, aggregated in time bins of 1 minute. In Fig. 8b we observe that M2M devices trigger several time-synchronized peaks in the DNS query count across both days, showing the largest peaks around midnight. Further manual investigation of the spikes depicted Fig. 8b showed that these spikes are deterministic and re-occur every day at the same time of day. Since network capacity is typically dimensioned in a peak-oriented manner, the presence of short-term peaks results in higher capacity demands and wastage of resources in non-peak intervals. Both aspects are highly undesirable at the network dimensioning stage. Our observation suggests the need for deploying randomization strategies for de-synchronization of network traffic in order to optimize network dimensioning. The expected large increase of (synchronized) M2M devices by 2020 [16] urges for the adoption of such mitigation strategies in the near future. However, in Fig. 8 we observe that not only M2M devices, but also smartphones and tablets exhibit synchronized peaks in the DNS query counts. For instance, we report smaller spikes throughout the whole day and larger, persistent spikes in the night hours every day (ref. Fig. 8a and Fig. 8c).

Interestingly, Fig. 8 shows a sudden and large increase of DNS queries in the morning hours of the second day. This clear anomaly only affects smartphones and tablets, but no other device types (ref. e.g., USB modem in Fig. 8). For further investigation of the anomaly, we plot the time series of DNS query




Fig. 9. Time series of DNS requests per operating system, 2 days, time bins of 1 minute

counts separately for TAC identifiers related to Android OS and iOS (recall §3) in Fig. 9. From this figure we observe that only iOS-based devices are affected by the anomaly. This confirms that the misbehavior of a large population of a certain device type may in fact trigger macroscopic, network-wide anomalies. As mobile devices nowadays rely on frequent updates of mobile apps and operating systems, it cannot be excluded that novel device-specific (mis-)behavior are induced by such software updates over time. Accordingly, the assessment whether software-updates result in new device-specific traffic characteristics or even might induce novel network anomalies, requires continuous monitoring of device-specific behavior across different types of devices and OSs. Hence, we consider the the analysis of device-specific traffic behavior a moving target.

We further investigate the iOS-based anomaly depicted in Fig. 9 as follows. Our monitoring system allows for the extraction of IP-level flow counters aggregated in time bins of 1 minute. That is, in each time bin traffic volumes (i.e., IP packets transferred in uplink and in downlink) are aggregated per flow, where a flow is identified by the anonymized MDID, destination IP address, source port and destination port. These flow aggregates enable the analysis of three different flow types: flows where traffic has been transferred exclusively in uplink, i.e., data has been sent from the mobile devices towards the Internet. Such flows are referred to as “uplink-only” flows. Conversely, “downlink-only” refers to flows where traffic has been transferred only from servers in the Internet towards mobile devices, and “two-way” flows refer to the case where traffic volume has been transferred in both directions. As mentioned in [17] one-way traffic is caused by different sources, such as scanning, peer-to-peer applications, back-scatter and unreachable services. As noticed in [17] the latter type is particularly helpful for large scale monitoring of network and service outages.

In Fig. 10 we plot the (rescaled) number of different flow types versus time during the interval of the iOS-based anomaly. Fig. 10a depicts the flows involving servers located in the anomalous IP range, while Fig. 10b shows flows involving the residual server IP addresses. In order to enable direct comparison both figures have been rescaled by the same undisclosed factor. In Fig 10b we observe that all flow types are rather stable over time. The number of two-way flows shows two slight bumps during the anomaly, which are likely caused by the anomalous increase of DNS queries that has been presented in Fig. 9b. In contrast, we observe two distinct bumps for all flows of the anomalous IP range in Fig. 10a.




Fig. 10. Time series of flow types during observed anomaly, 1 min. time bins (rescaled)

The majority of these flows is unanswered by the destined servers within the same time bin, resulting in two bumps of uplink-only flows (indicated by the solid line). This finding and the fact that we observe a simultaneous increase of uplink DNS queries, suggest that the anomaly is induced by a large set of iOS-based devices issuing connection requests in uplink (e.g., by sending TCP SYN packets). While we find that still some of the flows are answered by the servers (see “two-way” flows in Fig. 10a), we also report two bumps for the class of downlink-only flows. Such downlink-only flows may be present if connection requests in uplink are answered late (e.g. by TCP reset packets) such that the downlink packet is counted in one of the subsequent time bins after the corresponding uplink flow. The pattern shown in Fig. 10a suggests that the servers located in the anomalous IP range suffered from temporary faults leading to a large number of unanswered and rejected client connection requests. In fact the mobile devices reacted upon these unsuccessful connection requests by triggering even more uplink-only flows (indicated by the persistently increasing height of the bumps in Fig. 10a). Such a behavior is highly unwanted, since the scheduling mechanisms in the Radio Access Network (RAN) typically rely on assigning higher bandwidth to devices exceeding certain traffic rate thresholds [18]. As a result of this scheduling strategy, anomalies like the reported one may lead to large bandwidth demands for sending unsuccessful connection requests and, hence, to a significant wastage of network resources in the RAN. Since RAN resources are allocated in a shared manner, such incidents may even impair the network performance of those devices not directly involved into an anomaly.

Our findings suggest that new methodologies are required for detecting and mitigating device-specific anomalies. Current protection mechanisms, such as e.g., IP firewalls or intrusion prevention systems, are mainly designed for mitigating anomalies and attacks originating from *external* sources located in the Internet. Accordingly, such protection infrastructure is placed towards the edge of the core network between the Gateway GPRS Support Node (GGSN) and the Internet (ref. Fig 1). However, methodologies for efficiently mitigating device-triggered anomalies should be enforced in the RAN section close to the mobile devices, in order to avoid the propagation of anomalies and their negative effects towards the core network. Furthermore, our study illustrates how root causes of anomalies can be efficiently carried out by relying on device- and OS-specific traffic characteristics. Further work along this direction has been documented

in [19]. There, we provide guidelines for designing an automatic diagnosis system for network anomalies, which relies on entropy metrics calculated from device-specific traffic features.

5 Conclusions and Future Work

In this paper we present a device-specific view on the traffic within an operational cellular network, taking into account both, data plane and signaling plane traffic. By discussing previous works on traffic characterization in mobile networks we show that comparing data sets collected from different networks is not trivial, since even simple metrics such as, e.g., traffic volume time series, may exhibit significantly different patterns across different networks. This is specifically critical in case generic traffic models are derived from network-specific traffic behavior.

Moreover, we report the presence of time-synchronized peaks at different protocol layers (e.g., DNS and PDP-context create) and at different time of the day. In particular we find that such spikes are not only triggered by synchronized M2M devices, but are also observed in other device categories (e.g., smartphones, tables and USB modems). This finding suggests that de-synchronization of mobile devices for mitigating the negative effects of synchronized traffic peaks should not only be carried out for M2M devices, but also for all the other device classes.

Finally, we also report about a device-specific network anomaly detected from the analysis of DNS traffic. Analyzing this anomaly we show that device-specific traffic characterization supports investigation of device-induced network anomalies, and the identification of their root causes.

Acknowledgments. This work has been performed in the framework of the EU-IP project mPlane, funded by the European Commission under the grant 318627.

References

1. D'Alconzo, A., Coluccia, A., Romirer-Maierhofer, P.: Distribution-Based Anomaly Detection in 3G Mobile Networks: From Theory to Practice. *International Journal of Network Management*, September 2010
2. Shafiq, M.Z., et al.: Large-Scale Measurement and Characterization of Cellular Machine-to-Machine Traffic. *IEEE/ACM Transactions on Networking* (2013)
3. Gansemer, S., Groner, U., Maus, M.: Database classification of mobile devices. In: *Intelligent Data Acquisition and Advanced Computing Systems Technology and Applications (IDAACS)* (2007)
4. Granell, E., et al.: Smart devices fingerprint detection. In: *IEEE Globecom Workshops* (2012)
5. Kumar, U., Kim, J., Helmy, A.: Changing patterns of mobile network (WLAN) usage: smart-phones vs. laptops. In: *Wireless Communications and Mobile Computing Conference, IWCMC 2013* (2013)
6. Marjamaa, J.: A measurement-based analysis of machine-to-machine communications over a cellular network. Master's thesis, Aalto University, Helsinki, June 2012

7. Baer, A., Svoboda, P., Casas, P.: MTRAC - discovering M2M devices in cellular networks from coarse-grained measurements. In: International Conference on Communications, ICC (2015)
8. Bannister, J., Mather, P., Coope, S.: Convergence technologies for 3G networks: IP, UMTS. John Wiley and Sons, EGPRS and ATM (2004)
9. Ricciato, F., et al.: Traffic monitoring and analysis in 3G networks: lessons learned from the METAWIN project. *Elektrotechnik und Informationstechnik* (2006)
10. Endace measurement systems. <http://www.endace.com>
11. ETSI. 3GPP TS 129.060, version 7.9.0 (2008)
12. Baer, A., et al.: Large-scale network traffic monitoring with DBStream, a system for rolling big data analysis. In: International Conference on Big Data (2014)
13. GSMA IMEI Database. <http://imei.db.gsm.org/imei/>
14. AT&T, Florham Park, NJ, USA. AT&T specialty vertical devices. http://www.rfwel.com/support/hw-support/ATT_SpecialtyVerticalDevices.eps
15. Law, L.K., Krishnamurthy, S.V., Faloutsos, M.: Capacity of hybrid cellular-ad hoc data networks. In: The 27th Conference on Computer Communications on IEEE INFOCOM 2008 (2008)
16. Laner, M., et al.: Traffic models for machine type communications. In: International Symposium on Wireless Communication Systems, ISWCS 2013 (2013)
17. Glatz, E., Dimitropoulos, X.: Classifying internet one-way traffic. In: Proceedings of the 2012 ACM Conference on Internet Measurement Conference (2012)
18. Laner, M., et al.: A comparison between one-way delays in operating HSPA and LTE networks. In: Symposium on Modeling and Optimization in Wireless Networks, WiOpt (2012)
19. Schiavone, M., et al.: Diagnosing device-specific anomalies in cellular networks. In: ACM CoNEXT 2014 Workshop, Sydney, Australia (2014)