Bottleneck Footprints in TCP over Mobile Internet Accesses

Philipp Svoboda, Fabio Ricciato, and Markus Rupp

***Abstract*— In this paper we propose a novel method to detect a bottleneck in the core section of a mobile network using TCP related counters. The dataset used in this work consists of the full TCP statistics for two different UMTS SGSNs over four different one-day periods: one day in March 2006 and three consecutive days in September 2006. We derived the number of packets and the number of retransmissions per individual user during the peak hours. These two quantities were put into scatterplots to derive "footprints" of global TCP behavior in the SGSN areas. One of the traces contained a capacity bottleneck on an SGSN link. This trace was taken as the reference footprint for bottleneck presence. Based on such reference we developed a method to infer the presence of future bottlenecks based on footprint similarity. We tested two different distance metrics. It turns out that the simple correlation performs similar to a much more sophisticated function based on a ratio of symmetriezed Kullback-Leibler distances.**

***Index Terms*— 3G, bottleneck, TCP, footprint.**

I. Introduction

In this paper we propose a novel method to detect capacity bottlenecks in the core section of a mobile cellular network using TCP measurements. The TCP-protocol guarantees a reliable exchange of data between sender and receiver by means of retransmissions of lost packets. Therefore the frequency of retransmission can be taken as an indicator for the level of packet loss along the path of the TCP flows. In wired networks the loss probability due to a link error is very small, therefore a high number of observed retransmissions would likely indicate network congestion. In wireless cellular networks the bit errors on the radio link are recovered by FEC/ARQ methods. However, in practice there is a hard limit on the maximum number of packet retransmissions (e.g. three), leaving room for a certain level of residual packet loss. This effect, together with the handover procedures due to terminal mobility, cause a certain "physiological" level of packet loss (see [1]). Therefore in such a context the bottleneck detection based on loss statistics becomes more difficult. The basic goal of our work is to find an efficient and simple method to decide whether there is congestion in the network based purely on passive measurements. This approach is basically an example of *passive tomography* technique [2].

Mobile data traffic can be referred to individual terminals using information available at the lower 3GPP layers [3], therefore it is possible to extract TCP statistics on a per user basis from complete packet-level traces. We denote by n_i and N_i respectively the number of retransmissions and total packets observed for the generic ith user in a given time period (e.g. 30 min). Since 3G mobile networks have a hierarchical tree-like deployment (ref. Fig. 1) we can consider the traffic aggregate flowing through a certain network section, e.g. a SGSN. In case that a capacity bottleneck is in place along its path, the whole aggregate will suffer a relatively high level of packet loss, hence will yield a higher frequency of retransmissions than in normal operational conditions. A simple approach would be to use the global indicator $(\sum_i n_i / \sum_i N_i)$ as a basis to detect a bottleneck. The problem with such an approach is that the distributions of n_i and N_i span several orders of magnitude and are heavy-tailed, so that the global averages are prone to bias by a few large contributors, e.g. terminals that are affected by high loss due to local conditions. In a previous work [3] we tried to work around the problem by means of an ad-hoc filtering procedure. Here we propose an alternative approach based on bi-dimensional patterns. The advantage of this novel method is a better scalability for finer space granularity. The method itself can not exactly locate the bottleneck element. However, it can decide that a certain traffic aggregate is affected by a bottleneck along the common path. This should serve as an initial trigger for the network operator to start further extensive investigation along the traffic path.

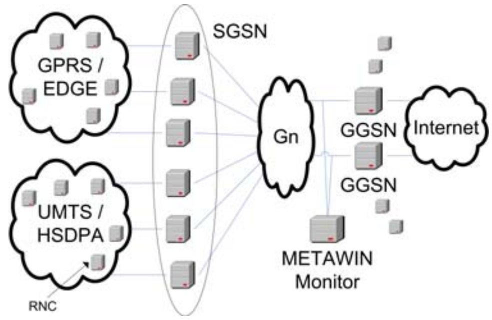

Fig. 1. Measurement setup.

II. Measurement Setup and Data Extraction

The reference network scenario is depicted in Fig. 1. As most access networks, the 3G mobile network has a hierarchical tree-like deployment. The mobile stations and base stations are geographically distributed. Going up in the hierarchy (first BSC/RNC, then SGSN, ultimately GGSN) the level of concentration increases, involving a progressively smaller number of equipments and physical sites. In a typical network there are relatively few SGSNs and even fewer GGSNs. Therefore it is possible to capture the whole data traffic from home subscribers on a small number of Gn/Gi links. For further details on the structure of a 3G mobile network referee to [4]. To meet privacy requirements traces are anonymized by hashing all fields related to user identity at the lower 3G layers (e.g. IMSI, MSISDN), while the user payload above the

Manuscript received April 11, 2007. The associate editor coordinating the review of this letter and approving it for publication was Christos Douligeris.

P. Svoboda and M. Rupp are with the INTHFT Department, Vienna University of Technology, Austria (email: {psvoboda, mrupp}@nt.tuwien.ac.at).

F. Ricciato is with ftw. Forschungszentrum Telekommunikation, (email: ricciato@ftw.at).

Digital Object Identifier 10.1109/LCOMM.2007.070559.

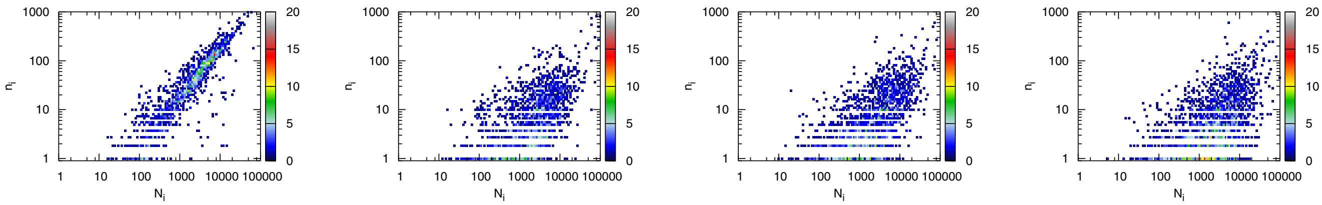

(a) $\mathtt{R_A}$ (Mar., with bottleneck) (b) $\mathtt{R_B}$ (Mar., without bottleneck) (c) $\mathtt{S1_A}$ (Sep., Day1, SGSN-A) (d) GPRS

Fig. 2. Scatterplot of N_i over n_i in the peak hours (log-binning, lin. scale).

TCP/IP layer is removed. The TCP statistics were extracted using a modified version of `tcptrace` [5].

The core dataset used in this work consists of the full TCP statistics for two different UMTS SGSNs (“A” and “B”) over four different one-day periods: one day in March 2006 and three consecutive days in September 2006. The measurements were taken from the live network of a major mobile provider in Austria, EU. Hereafter the March dataset will be indicated by `R` and the September ones by `S1`, `S2` and `S3`. The SGSN will be indicated in the index: therefore the symbol $\mathtt{S2_A}$ refers to the data collected on the second day of September on the first SGSN. A capacity bottleneck was in place on a link to SGSN-A in March 2006, i.e., in the $\mathtt{R_A}$ dataset.

III. Matching TCP Footprints

The TCP statistics $\langle N_i, n_i \rangle$ were extracted for 30 minutes time bins. As the impact of a bottleneck is more evident during the peak hour we focus our analysis only on the period from 7 to 9 pm, for a total of four bins. We used scatterplots to visualize the process $\langle N_i, n_i \rangle$, i.e., to create “TCP footprints”. As both variables span several orders of magnitude we introduced a logarithmic binning with 150 bins on each axis. In a first step we filtered the datasets using a median filter to remove outliers. In a second run we applied a mean filter with a window size of five to smooth the footprint. In a final step we normalized the value of the bins by the number of events in the trace. The datasets are now equal to an empirical binned bidimensional probability density function. This step compensated the varying amount of user traffic in the different measurement periods.

Figure 2 visualizes the scatterplots for the three datasets $\mathtt{R_A}$, $\mathtt{R_B}$, $\mathtt{S1_A}$ and a GPRS sample. The latter is related to the footprint of a GPRS SGSN, and is depicted to show the similarity of the footprints for UMTS and GPRS/EDGE. This is an interesting observation, since the two radio technologies have different capacity and usage patterns. The number of events found within one bin is given by the color of the pixel. The color bar on the right is in linear scale.

A visual comparison reveals that footprint $\mathtt{R_A}$ differs from the other traces shown in Figs. 2(b) and 2(c). Additional we show a footprint from a different radio access network, namely GPRS. Even this footprint looks more similar to Figs. 2(b) and 2(c) than Fig. 2(a). In fact $\mathtt{R_A}$ is the sample associated to the bottleneck. In Fig. 2(a) we see a strong positive correlation between N_i and n_i, especially for larger values of N_i. This was expected: In fact, a capacity bottleneck can be modeled as an element introducing *random* packet loss with a certain probability p on all flows. Hence the absolute number of retransmissions for each flow will be roughly proportional to the flow size, i.e. $n_i \propto p \cdot N_i$. In contrast to this all other footprints in Fig. 2 (without bottleneck) yield a much weaker correlation. Note that all other footprints not depicted here look similar to $\mathtt{S1_A}$, indicating a certain stability of the “normal” footprint.

TABLE I
Correlation coefficient

SGSN-A	$\mathtt{R_A}$	$\mathtt{S1_A}$	$\mathtt{S2_A}$	$\mathtt{S3_A}$	GPRS
$\mathtt{R_A}$	1.00	0.29	0.32	0.33	0.55
$\mathtt{R_B}$	0.41	0.91	0.93	0.94	0.88
SGSN-B	$\mathtt{R_B}$	$\mathtt{S1_B}$	$\mathtt{S2_B}$	$\mathtt{S3_B}$	
$\mathtt{R_A}$	0.41	0.26	0.34	0.33	
$\mathtt{R_B}$	1.00	0.88	0.92	0.93	

TABLE II
Modified KL distance

SGSN-A	$\mathtt{S1_A}$	$\mathtt{S2_A}$	$\mathtt{S3_A}$
$\Delta_{mKL}(x)$	0.21	0.15	0.13
SGSN-B	$\mathtt{S1_B}$	$\mathtt{S2_B}$	$\mathtt{S3_B}$
$\Delta_{mKL}(x)$	0.23	0.14	0.13

Our goal is to define a metric indicating if a new footprint is more similar to $\mathtt{R_A}$ or $\mathtt{R_B}$. We first evaluated the 2-D correlation coefficient on our datasets, called Δ_C. This metric is often applied in image processing to identify the similarity between two images. It outputs a dimensionless value. Table I gives the results for all permutations between the datasets.

It is constructed from two parts. The upper part shows the correlation between $\mathtt{R_A}$ in the first row respectively $\mathtt{R_B}$ in the second row and all datasets collected from SGSN-A, while the lower part holds the results for datasets from SGSN-B. The correlation in the group of non congested datasets is above 0.8. Comparing the datasets with the congested dataset $\mathtt{R_A}$ yields values below 0.4. The variance of the correlation with $\mathtt{R_B}$ is relatively small and the values strongly discriminate against the results for $\mathtt{R_A}$. Therefore the correlation serves well to detect the congestion in this case. The sensitivity of this methods will be discussed later in this section. In a next step we evaluated the Kullback-Leibler (KL) distance between the datasets. The KL distance is a measure between a given probability distribution P and an arbitrary distribution Q. Often P represents some reference data obtained by measurements and Q is generated by a model approximating P. Eq. (1) shows the KL distance metric.

$$\Delta_{KL}(P\|Q) = \sum_{i,j} P(i,j) \cdot \ln \frac{P(i,j)}{Q(i,j)} \tag{1}$$

Note that the KL distance is not symmetric. We used a

TABLE III
METRIC BENCHMARK WITH A GENERATED DATASET

Metric	$\mathtt{mix}_{10}$	$\mathtt{mix}_{20}$	$\mathtt{mix}_{30}$	$\mathtt{mix}_{40}$	$\mathtt{mix}_{50}$	$\mathtt{mix}_{60}$	$\mathtt{mix}_{70}$	$\mathtt{mix}_{80}$	$\mathtt{mix}_{90}$
$\Delta_C(\mathtt{mix}_x, \mathtt{R}_\mathtt{A})$	0.34	0.41	0.47	0.56	0.65	0.73	0.82	0.90	0.91
$\Psi_{KL}(\mathtt{mix}_x)$	36,9	8,89	3.50	1.66	0.88	0.47	0.27	0.17	0.14

symmetrized version of the metric proposed in [6]. It is shown in Eq. (2).

$$\Delta_{sKL}(P,Q) = \frac{1}{2} \cdot (\Delta_{KL}(P\|Q) + \Delta_{KL}(Q\|P)) \quad (2)$$

The output of Eq. (2) is $\Delta_{sKL}(P,Q) \geq 0$, with equality if P equals Q. As the equation uses a logarithmic fraction, bins of value zero in either P or Q will lead to an undefined result. Therefore we added a constant offset to the footprints and reduced the sum over i,j to bins were both P and Q were non zero. In a next step we defined a ratio Ψ_{KL} as in Eq. (3). Here we calculate the ratio of $\Delta_{sKL}(\mathtt{S1}_\mathtt{A}, \mathtt{R}_\mathtt{B})$, the distance between the new data and the non bottleneck trace, and $\Delta_{sKL}(\mathtt{S1}_\mathtt{A}, \mathtt{R}_\mathtt{A})$, the distance between the new data and the bottleneck trace. In other words we normalized the results gained in Eq. (2).

$$\Psi_{KL}(\mathtt{S1}_\mathtt{A}) = \frac{\Delta_{sKL}(\mathtt{S1}_\mathtt{A}, \mathtt{R}_\mathtt{B})}{\Delta_{sKL}(\mathtt{S1}_\mathtt{A}, \mathtt{R}_\mathtt{A})} \quad (3)$$

Table II holds the results of the symmetrized and normalized KL distance. Again the upper row is for SGSN-A and the lower for SGSN-B. Note that the value of (3) approaches infinity if the input sample is similar to $\mathtt{R}_\mathtt{A}$. Therefore an increasing value of Ψ_{KL} indicates a bottleneck. As all the values in Table II are without a bottleneck, values below 0.2 can be taken as a physiological limit for non congested traces. The values for the two SGSNs are very close indicating a high similarity between their footprints. Hence a performance estimation can only be archived using a normal dataset mixed with congested traffic.

In the following we investigate the case that only a portion of the SGSN area is affected by a capacity restriction (partial bottleneck). This could be the case if for example only one RNC link is congested, instead of the SGSN link. In this case only a fraction of the total samples will distribute as in the congested footprint, while the rest will stay in normal patterns. To analyse this effect we take a linear function to combine $\mathtt{S2}_\mathtt{B}$ and $\mathtt{R}_\mathtt{A}$ by randomly choosing datapoints from one of the datasets. The new trace is called $\mathtt{mix}_{\{0...100\}}$. The number in the index indicates the percentage of $\mathtt{S2}_\mathtt{B}$ in $\mathtt{mix}$. Table III shows the correlation coefficient and the modified KL distance for $\mathtt{mix}$ and $\mathtt{R}_\mathtt{A}$ with the share of $\mathtt{S2}_\mathtt{B}$ changed in steps of 10%. The first row of the table holds the results of the correlation. Dataset $\mathtt{mix}_{100}$, identical to $\mathtt{S2}_\mathtt{B}$, has a correlation of 0.92 (see Table I). We set the detection limit to 0.88, which is the lowest value of Δ_C for all datasets correlated with $\mathtt{R}_\mathtt{A}$. This value is met by $\mathtt{mix}_{75}$. With a margin of 10%, added for more reliable detection, approximately 30% of saturated traffic triggers the alarm.

The results for the Ψ_{KL} function are presented in the second line of Table III. The minimal distance is equal to the distance between $\mathtt{S2}_\mathtt{B}$ and $\mathtt{R}_\mathtt{A}$ that is 0.14. We set the detection limit to the largest result from Table II, that is 0.23, and added again the margin for more reliable detection. This led to a trigger point of approximately 25% congested traffic.

Comparing the two metrics we learn that the performance for small shares, e.g., less than 20%, of congested traffic is similar. While the Ψ_{KL} function is much more sensitive in case that the bottleneck traffic increases even more.

IV. CONCLUSIONS

We presented two different metrics for bottleneck detection inside a mobile core network based on simple TCP statistics, namely the number of TCP retransmissions and the total number of packets for each user. We used the correlation coefficient and a modified version of the KL distance. Both metrics allow for a clear detection of bottleneck footprints. The correlation coefficient proves to be as sensitive as the modified metric based on the KL distance for a small amount of bottleneck traffic. Both metrics need 30% of congested datatraffic to trigger an alarm. We conclude that the correlation coefficient is sufficient to detect this kind of capacity bottlenecks and trigger re-provisioning. Further investigation is required to see if this approach also applies to other bottlenecks. The general lessons to be learned from such preliminary work is that when using TCP statistics to infer performance problems resorting to multi-dimensional pattern matching is a more effective approach than using averages and global ratios.

ACKNOWLEDGEMENTS

This work is part of the DARWIN project supported by the K-Plus program. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the funding partners.

REFERENCES

[1] F. Vacirca, T. Ziegler, and E. Hasenleithner, "An algorithm to detect TCP spurious timeouts and its application to operational UMTS/GPRS networks," *Computer Networks*, vol. 50, no. 16, pp. 2981–3001, 2006.
[2] F. Ricciato *et al.*, "Passive tomography of a 3G network: challenges and opportunities," in *INFOCOM 2006*, pages 116–122, 2006.
[3] F. Ricciato, F. Vacirca, and P. Svoboda, "Diagnosis of capacity bottlenecks via passive monitoring in 3G networks: an empirical analysis," *Computer Networks*, vol. 57, pp. 1205–1231, March 2007.
[4] H. Holma and A. Toskala, *WCDMA for UMTS, Radio Access For Third Generation Mobile Communications, Third Edition*. Wiley, 2004.
[5] Modified tcptrace. `http://userver.ftw.at/~vacirca`, FTW.
[6] D. Johnson and S. Sinanovic, "Symmetrizing the Kullback-Leibler distance," *Technical Report*, 2001.