

# Traffic monitoring and analysis in 3G networks: lessons learned from the METAWIN project

F. Ricciato, P. Svoboda, J. Motz, W. Fleischer, M. Sedlak, M. Karner, R. Pilz, P. Romirer-Maierhofer, E. Hasenleithner, W. Jäger, P. Krüger, F. Vacirca, M. Rupp

A 3G network is a magnificently complex object embedded in a highly heterogeneous and ever-changing usage environment. It combines the functional complexity of the wireless cellular paradigm with the protocol dynamics of TCP/IP networks. Understanding such an environment is more urgent and at the same time more difficult than for legacy 2G networks. Continuous traffic monitoring by means of an advanced system, coupled with routine expert-driven traffic analysis, provides an in-depth understanding of the status and performances of the network as well as of the statistical behaviour of the user population. Such knowledge allows for a better engineering and operation practice of the whole network, and specifically the early detection of hidden risks and emerging troubles. Furthermore, the exploitation of certain TCP/IP dynamic behaviour, particularly the TCP control-loop, coupled with information extracted from the 3GPP layers, provides a cost-effective means to monitor the status of the whole network without requiring access to all network elements. In this article the main lessons are summarized learned from a two-year research activity on traffic monitoring and analysis on top of an operational 3G network.

Keywords: traffic monitoring; traffic analysis; 3G; cellular networks; GPRS; UMTS

## Analyse und Monitoring von Datenverkehr in 3G Netzwerken: ein Erfahrungsbericht aus dem METAWIN-Projekt.

*Ein 3G-Mobilfunknetz stellt ein extrem komplexes Gebilde dar, das in ein stark heterogenes und ständig wechselndes Umfeld eingebettet ist. In ihm werden komplexe Funktionen der Mobilfunkübertragung mit TCP/IP-Protokollen verbunden. Ein tiefes Verständnis dieses Systems ist dringend erforderlich, aber sehr schwierig im Vergleich zu vorhandenen 2G-Netzen. Nur die kontinuierliche Beobachtung dieses modernen Übertragungssystems erlaubt es, sowohl die Leistungsfähigkeit des Netzes zu verstehen als auch das statistische Benutzerverhalten zu interpretieren. Die daraus gewonnenen Erkenntnisse begünstigen sowohl eine verbesserte Netzplanung als auch einen optimalen Betrieb und gestatten es darüber hinaus, verborgene Risiken und vorhandene Schwierigkeiten sehr früh zu erkennen. Insbesondere erlaubt die Beobachtung bestimmter TCP/IP-Dynamik gemeinsam mit der Information, die aus den 3GPP-Schichten gewonnen wird, eine kostengünstige Möglichkeit, den Zustand des Gesamtnetzes zu erfassen ohne jedes Einzelelement des Netzes zu erfassen. In diesem Artikel sind die wesentlichen Erkenntnisse zusammengefasst, die in Folge eines zweijährigen Forschungsprojektes zur Beobachtung und Analyse eines 3G-Mobilfunknetzes gewonnen wurden.*

Schlüsselwörter: Monitoring von Datenverkehr; Analyse von Datenverkehr; 3G; Mobilfunknetze; GPRS; UMTS

Eingegangen am 16. Mai 2006, angenommen nach Revision am 6. Juni 2006  
© Springer-Verlag 2006

## 1. Introduction

Public wide-area wireless networks are now migrating towards third-generation systems (3G), designed to support packet-switched data services. Europe has adopted the Universal Mobile Telecommunication System (UMTS) developed by 3GPP as an evolution of GSM, the worldwide widespread second-generation mobile technology (2G). A 3G network includes two main sections: a Core Network (CN) and one or more Radio Access Network(s) (RAN). The Core Network consists of two distinct domains: Circuit-Switched (CS) for voice telephony and Packet-Switched (PS) for packet data applications. Along with the UMTS RAN (UTRAN) based on W-CDMA, several operators maintain a parallel GPRS RAN evolved from the legacy GSM radio. This structure is sketched in Fig. 1 and explained in Sect. 2. Several UMTS networks became operational since 2003 while first deployments of GPRS date back to 2000. Since then the growing popularity of 3G terminals and services has extended the coverage of Internet access to the geographic area, and 3G networks are becoming key components of the global Internet in Europe. (Keshav, 2005) foresees that cell phones will become the dominant component of future Internet population, while Kleinrock expects this role to be played by "small pervasive devices ubiquitously

embedded in the physical world" (quoted from (Kleinrock, 2005: 112)). Both scenarios underlay that the main access mode in the future Internet will be wide-area wireless. Such access demand will be provided to a large extent by the currently deployed 3G networks and their future evolutions, alongside with other emerging technologies (e.g. WIMAX).

Today the 3G environment is still under evolution, at least along the following dimensions:

- ▶ subscriber population and traffic volumes;
- ▶ terminal capabilities and relative penetration of the various terminal types (handsets, laptops with 3G card, etc.);
- ▶ service portfolio and tariffs offered by the operators.

Ricciato, F., Pilz, R., Romirer-Maierhofer, P., Hasenleithner, E., Vacirca, F., Forschungszentrum Telekommunikation Wien (ftw.), Donaustadtstraße 1, 1220 Wien, Österreich (E-mail: ricciato@ftw.at); Svoboda, P., Rupp, M., Technische Universität Wien, Institut für Nachrichtentechnik und Hochfrequenztechnik, Gußhausstraße 25-29, 1040 Wien, Österreich; Motz, J., Sedlak, M., Jäger, W., Krüger, P., Kapsch CarrierCom, Am Europlatz 5, 1120 Wien, Österreich; Fleischer, W., Karner, M., mobilkom austria AG & Co KG, Obere Donaustadtstraße 29, 1020 Wien, Österreich

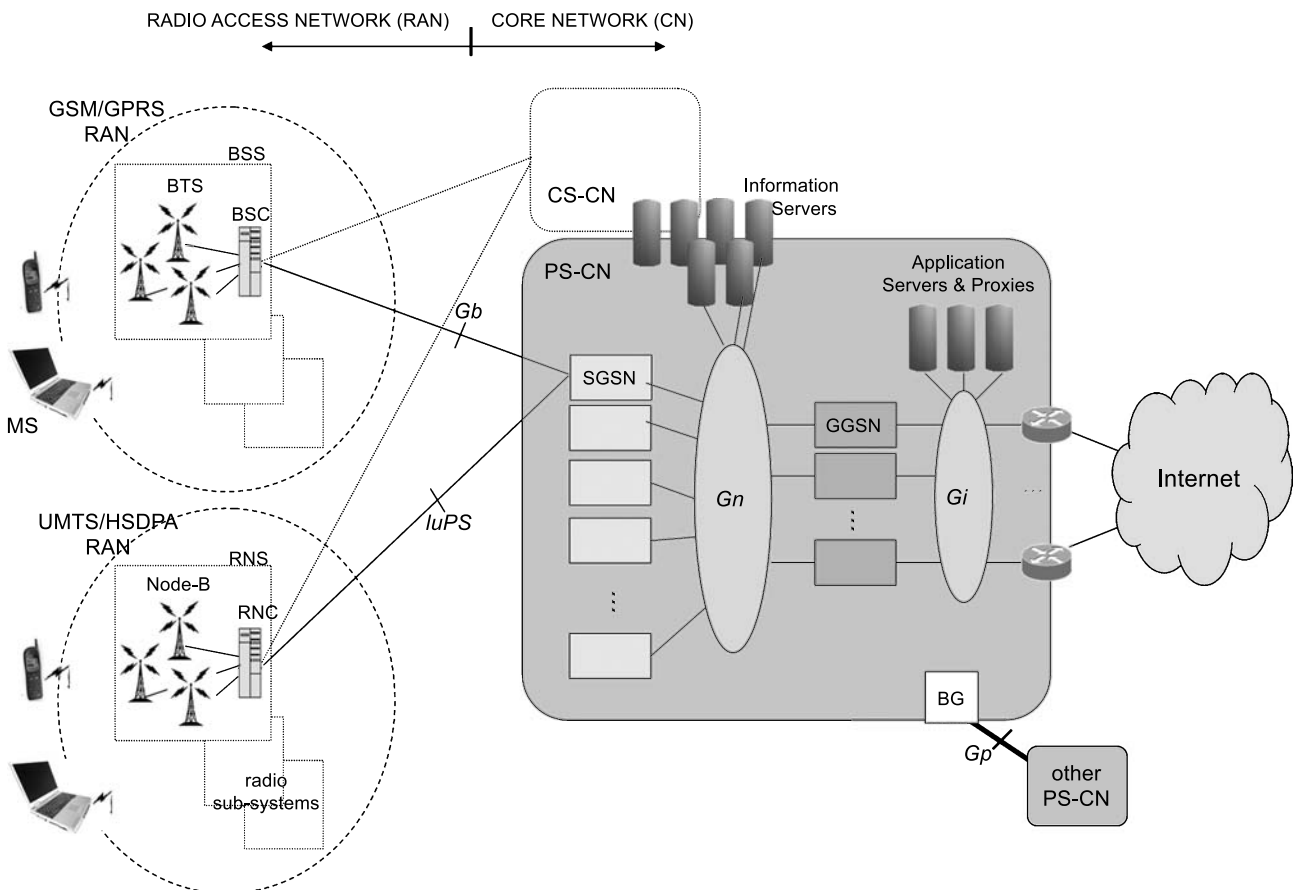


Fig. 1. Structure of a mobile 3G network

Furthermore, technological upgrades are still in the agenda of many operators: EDGE in the GPRS RAN, HSDPA in the UMTS RAN, IMS and new service architectures in the CN (Pospischil, Miladinovich, Kunczler, 2005; Bannister, Mather, Coope, 2004). All these aspects collectively build a potential for changes in the global traffic that can occur at the macroscopic scale (network-wide) and in a relatively short time frame. Hence, the ability to accurately and extensively monitor the network state and the traffic in order to recognize drifts in global performance and usage patterns and/or early detect local troubles is a fundamental pillar of the network operation and optimization process.

The recent years have recorded a surge of studies and research activities based on the passive monitoring of Internet network traffic. On the technological side enabling factors are the availability of high-performance acquisition hardware (e.g. capture cards, wiretaps) and large-capacity storage system (e.g. RAID) at accessible cost. The first large-scale traffic monitoring project on a commercial IP network was run at Sprint (IPMON project, <http://ipmon.sprint.com>). Since then more and more research groups became involved in traffic monitoring activities on top of IP networks. Several methodologies for the analysis of the collected data are being developed, importing concepts and tools from different areas (e.g. signal-processing, data-mining). Some of these techniques can find very concrete applications and provide powerful support for the operation and engineering of real networks, with recognizable benefit in terms of revenue protection and/or cost saving.

Since 2004 we have been involved in a research project aimed at exploiting traffic monitoring and analysis for the engineering of a GPRS/UMTS network. The METAWIN project (Measurements and

Traffic Analysis in Wireless Networks), an applied-research project which runs in close collaboration between industry and academia<sup>1</sup> (METAWIN project, <http://userver.ftw.at/~ricciato/metawin.html>).

The initial idea in METAWIN was to capture packet-level traces from the live network and use them to distil synthetic user models fitted to the observed data. During the project we recognized that the availability of high-quality traces (to be exactly defined later) yields a much higher potential for improving the engineering practice of the real network, far beyond the mere opportunity to fit abstract models to observed patterns. More generally, Traffic Monitoring and Analysis (TMA for short) can play a major role in several technical areas within the running of a real 3G network: operation and maintenance, troubleshooting, planning and optimization, design and engineering, security monitoring and fraud detection, etc. Its reach can be easily extended to less technical areas like for example marketing, service engineering, billing and tariff design.

In this contribution we present a few exemplary applications of TMA to the operation and engineering of 3G network. We present a sample set of findings and results from the exploration of real traces from the operational network. Collectively, all our findings show that advanced traffic monitoring, coupled with in-depth expert analysis, offer interesting and cost-effective opportunities to understand and substantially improve an operational 3G network.

The rest of the paper is organized as follows. In Sect. 2 we briefly introduce the structure of a typical 3G network. In Sect. 3 we relate TMA to the current network monitoring practice. In Sect. 4 we outline the key features of the TMA system developed and deployed

<sup>1</sup> The project partners are: mobilkom austria AG&CoKG, Kapsch CarrierCom, Forschungszentrum Telekommunikation Wien and Technical University of Vienna.

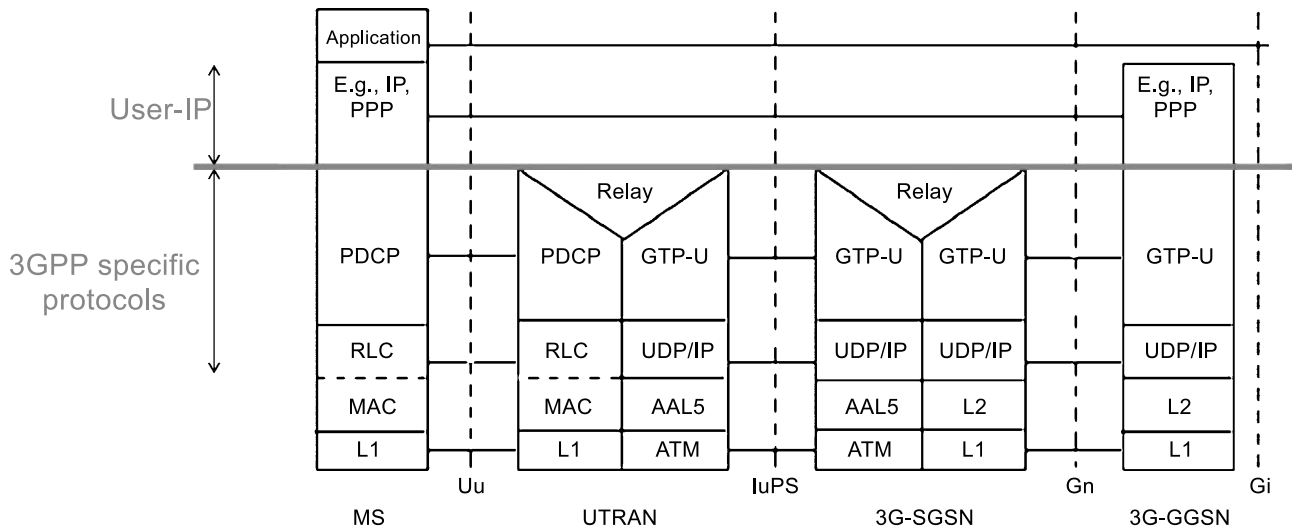


Fig. 2. Partial scheme of the 3G protocol architecture (UMTS data-plane)

during the METAWIN project. In Sect. 5 we provide an overview of selected findings and potential application that emerged during the exploration of the real traffic data in the METAWIN project. Finally in Sect. 6 we conclude and identify the topics of future work.

## 2. A brief introduction to 3G networks

The structure of a 3G network is depicted in Fig. 1. It consists of two Core Networks domains (CS and PS) and one or more Radio Access Network(s) (RAN). Several operators maintain both a GPRS/EDGE and a UMTS/HSDPA RAN. Candidate technologies for additional RANs are WLAN and WIMAX.

The RAN is divided into *subsystems* consisting of one controller (called RNC and BSC respectively in UMTS and GPRS) and a number of transceiver stations attached to it. The latter maintain the air interface in the radio cells, while the former controls the radio connections with the Mobile Stations (MS) and the wired interface to the Core Networks (CS and PS). Radio cells are hierarchically grouped into so called Routing Areas (RA for short). A RA is the PS homologous of the Location Area (LA) in the CS domain. Each RA is included in a single LA. RA changes must be always signaled to the network by means of the so called Routing Area Update procedure. Intra-RA cell changes are signaled only if the MS is in the READY state.

The PS-CN embeds several elements: SGSN, GGSN, and a number of information servers. Some of the latter are shared with the CS-CN of the legacy 2G system, e.g. HLR/AuC. Each SGSN interconnects several BSC/RNC to the PS-CN, and performs functions such as access control, location management, paging, route management (Bannister, Mather, Coope, 2004). The GGSN is the logical gateway between the PS-CN and external packet networks (Internet, private networks) and handles the IP-level connectivity with the MS. The SGSN and GGSN of the same operator communicate through the Gn interface. In addition, the PS-CN of different operators are interconnected through the Gp interface for support of roaming. The Gn protocol stack in Fig. 2 shows that a lower UDP/IP layer is used to carry the user data packets across Gn, with an intermediate encapsulation into GTP, a 3GPP protocol. In fact, the Gn interface is basically a wide-area IP network interconnecting the different SGSN/GGSN sites, and as such it embeds routers, IP subnets, etc. Besides that, the PS-CN is rich in IP-based elements, including servers supporting control and management functions (e.g. DNS, DHCP, RADIUS, see (Bannister, Mather, Coope, 2004)) and application elements (e.g. WAP gateway, proxies, internal servers). The latter are

always located behind the GGSN, on the Gi side (ref. Fig. 1) as they operate directly on the data-plane. Note also that packet filtering and other restriction policies can be located on separate dedicated elements (NAT, IDS, firewalls) at the network boundaries (Gi, Gp) and/or directly configured into the GGSNs.

A mobile 3G network combines two different paradigms: cellular and IP. From the cellular paradigm it inherits the high *functional complexity* required to guarantee quality and efficiency in the management of user mobility and radio resources (e.g. the 3GPP control plane and lower data layers). From the TCP/IP environment it inherits the complexity of the *protocol dynamics* (e.g. TCP) as well as a huge heterogeneity and variability in the usage patterns. This marks an important difference with respect to the previous generation of cellular networks (e.g. GSM). The latter were essentially dedicated to a single service, i.e. voice-based telephony, with a low-dimensionality<sup>2</sup>. Instead in 3G networks, packet-switched and multi-service, traffic sources (users) span a broad range of heterogeneous applications, most of which have high-dimensionality (e.g. WEB). Besides higher *heterogeneity and dimensionality*, 3G users also expose higher *variability*: along each dimension (e.g. transferred volume, number of object in a WEB session) the variability range of the measured values spans several orders of magnitudes.

## 3. Network monitoring: the present and the vision

The classical approach to network monitoring for the purposes of network management relies largely on routine collection of data delivered by the network equipments themselves: built-in counters, logs, SNMP MIBs, etc. This approach has some limitations. First, the quality of the available data is not always adequate: the time granularity and the aggregation level are coarse, the data semantic is limited and in some special cases their reliability can be questioned, e.g. in case of overload or malfunctioning. Second, the process of extracting, gathering and correlating such data involve considerable costs given the broad heterogeneity of equipments types, vendors, software releases, data formats and even data semantic. Moreover,

<sup>2</sup> We use the terms "low-" and "high-dimensionality" to refer to the number of dimensions required to fully describe the traffic generated by an application. Circuit-switched voice telephony has very low dimensionality, as the traffic aggregate can be well defined in terms of *calls* described by few parameters – essentially start-time, duration and bit rate. An example of high-dimensionality service is WEB, where each session is a highly structured entity embedding several HTTP requests, TCP connections, object downloads, etc.

every change in the network (e.g. replacing of equipments or software upgrades) is likely to require changes in the monitoring infrastructure as well. Above all is a general fundamental problem with such approach, namely the lack of decoupling between the monitored system and the monitoring tool, which produces obvious ambiguity problems and makes the latter unreliable in case of equipment malfunctioning.

To complement the routine large-scale data collection from network elements, sporadic fine-grain measurements are performed with small-size network protocol analyzers. These measurement interventions are generally limited in time and space (one or few interfaces) and are often used for troubleshooting actions, after a problem has been detected by external means (e.g. customer complaints).

In summary, the current status of the network monitoring practice is often a combination of large-scale routine data collection from network elements plus sporadic and local fine-grain measurement actions. The future vision of an advanced network monitoring foresees monitoring tools that (1) are capable of performing routine collection of fine-grain measurements in the large-scale, (2) are completely decoupled from the production network and (3) can cleverly process the recently collected data and deliver reports and alarms proactively. This is exactly the vision of a large-scale TMA infrastructure based on passive wiretapping at key network links, with null or very limited interaction with the production equipments. Such a global system would be intrinsically multi-purpose. Observing the traffic allows for the derivation of network-related as well as user-related statistical data. While the former are used by technical departments (operation and maintenance, planning, optimization, etc.) the latter can be useful for marketing purposes.

Large-scale monitoring systems are available on the market and have been used in the last years in the core network of 2G, e.g. GSM. With the deployment of 3G a novel packet-switched (PS) domain has been added to the legacy circuit-switched (CS) one, and some monitoring systems were *extended* to cope with the packet-switched domain. However, the current generation of commercial systems falls short of exploiting the full potential of traffic monitoring in a packet network. To some extent this is perhaps due to the fact that they were conceived and developed as mere extensions of the legacy system for the CS domain. The PS section was implicitly considered just as a set of additional *protocols to be parsed*. Instead, the differences between the two domains are deeper: the PS protocols have completely different *dynamics* on the data plane, some of which can be exploited to the operator advantage (e.g. TCP close-loop dynamics can be exploited in order to infer delays and packet loss, see Sects. 5.1 and 5.2). Furthermore, the PS domain supports a completely different *usage environment* (user populations, terminal types, applications, services, etc.) that is infinitely more heterogeneous and complex than the CS telephony. Accordingly, the choice of meaningful and convenient key performance/quality indicator for the PS can not be reduced to a variation of those successfully adopted in the CS. Instead it requires the application of sound "dual expertise" in both the fields of mobile cellular and packet networking. The latter might not be necessarily present in legacy development groups from the CS world.

#### 4. The METAWIN monitoring system

A pre-requisite to fully exploit the potential of TMA in the context of an operational 3G network is the availability of an advanced monitoring infrastructure. In the METAWIN project we have developed a large-scale monitoring system covering key links on all logical interfaces in the CN. The first prototype was developed in the early stage of the project. Since then it has evolved with additional features, driven by the knowledge about the traffic environment gained

during the exploration of the traces, in a virtuous explore-learn-develop loop involving the whole project staff.

The system is developed entirely on Linux platform. We used Endace DAG acquisition cards (Endace DAG technology, <http://www.endace.com>) and high-end standard PCs equipped with RAID storage. The advanced prototype is now deployed and fully operational in the GPRS/UMTS network of mobilkom austria. It covers selected links on Gi, Gn, luPS, Gb and Gp interfaces (ref. to Figs. 1 and 2).

In the following we outline the most important features of the monitoring system that enable the analysis tasks and applications discussed later in the remaining of the paper.

##### 4.1 Complete capture

The typical rates of 3G traffic on the CN links are such that all frames can be captured with standard hardware equipments, with no need to resort to packet sampling. This is a major simplification compared to TMA in backbone networks with multi-Gbps traffic rates. Also, the removal of application-layer payload (for privacy constraints) reduces the data volume to be stored. The current daily volumes of traffic allow for week-long storage at acceptable costs given the fact that large-storage solutions are relatively inexpensive, nowadays.

##### 4.2 User- and control-plane capture

The system is able to capture and parse the complete frame, including the 3GPP headers below the user IP layer (ref. Fig. 2). Along with user-data packets, it captures and parses signaling frames at each layer. This allows for cross-layer analysis and data/control plane correlation.

##### 4.3 Anonymization

For privacy reasons all subscriber-related fields (e.g. IMSI, MSISDN) are hashed with a non-invertible function. The resulting string univocally distinguishes the Mobile Station (MS) but cannot be referred to the user identity. For simplicity we will maintain the term "IMSI" to refer to the hashed string. This approach preserves packet-to-MS associations, i.e. the possibility to discriminate packets associated to a specific MS, while at the same time protect the identity of the subscriber. Also, to preserve content privacy the user payload at the application layer is stripped away.

##### 4.4 Stateful association tracking

For many applications it is highly desirable to label each packet with certain information associated to the MS it was generated by or directed to. The most important are the MS identifier (hashed IMSI) and the location of the MS at the current time (e.g. cell identifier). For other applications it might be useful also to know the MS type (e.g. handset, laptop card, PDA, etc.) and some additional equipment capabilities. The former can be directly retrieved by the Type Approval Code included in the IMEI, while the latter are usually advertised by the MS during the Attach Request. Similarly, it is often useful to associate individual packets to the PDP-context within which they were generated, and hence to the attributes of the PDP-context (e.g. assigned IP address, APN). The relevance of such associations will become clear in the next section. The problem is that a generic data frame crossing the network does not include all such information in its fields. However, a passive monitoring system can dynamically reconstruct such associations by smartly tracking the message exchange between the MS and the network. More specifically, it is required to inspect every signaling procedure and certain fields of the lower-layer headers of data-plane frames, and to maintain for each entity (MS, PDP-context) a dynamic record of associations. The point to be taken is that in general any attribute that is exchanged between the MS and a generic network element

can be captured and later associated to future packets. The associations between packets, PDP-contexts and (hashed) IMSI can be extracted on any interface between the MS and the GGSN, e.g. Gn. The localization of the terminal can be achieved in GPRS/EDGE EDGE by sniffing the Gb interface (for a detailed description of IMSI-to-cell tracking on Gb see (Borsos et al., 2006, Sect. IV,C). For UMTS/HSDPA sniffing on IuPS would allow for the localization of the MS limited to the Routing Area level, as intra-RA cell changes are not reported to the SGSN, therefore exact cell-level localization for UMTS/HSDPA would require monitoring the Iub interface between the MS and the RNC.

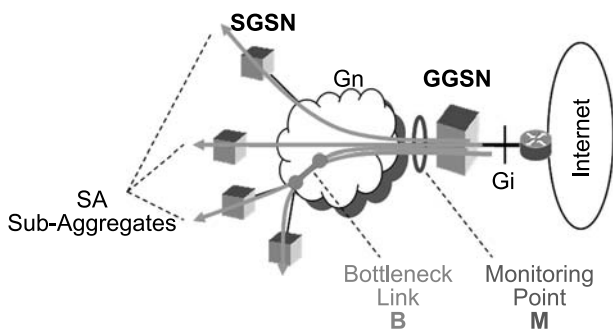
**5. Sample findings and applications**

**5.1 Detecting congestion in the Core Network**

The Gn network (ref. Fig. 2) is the heart of the 3G Core Network. It is an IP-based wide-area network interconnecting the physical sites where the SGSNs and GGSNs are located. Its L1/L2 deployment can include a mix of different technologies: for example the connectivity between GGSN x and SGSN y might pass through a combination of physical circuits, L2 virtual circuits (e.g. ATM, Frame Relay) and LAN segments internal to each site, spanning a chain of links of different nature, dedicated or shared with other paths, and nodal equipments of different technologies.

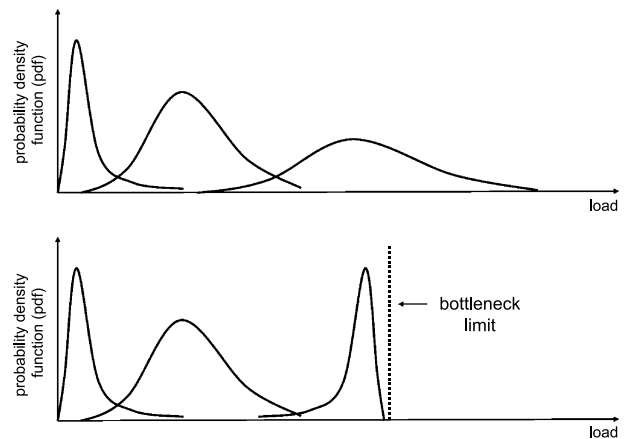
In such a complex and heterogeneous environment it would be highly desirable to verify whether the capacity deployed along the path x-y is adequate to sustain the current traffic volume without accessing all intermediate nodes. This goal can be achieved with TMA by monitoring of few Gn links near the GGSNs and exploiting the properties of the TCP protocol. Most of the current traffic is TCP, and the TCP control-loop introduces a *correlation* between the traffic dynamics and the status of the end-to-end path. The key point is that such correlation induces *observability*: a trouble affecting a TCP flow at some point can be detected by observing the TCP dynamics at a different point along the path. The specific application of this concept to the 3G Core Network enables the detection of capacity restrictions in Gn network by monitoring only the few Gn links near the GGSNs. By using information available at the 3GPP layers – in this case the IP address below the GTP layer (Fig. 2) it is possible to discriminate the traffic components associated to each SGSN-x/GGSN-y pair and analyze each of them separately. A capacity restriction at some point along the path (e.g. link B in Fig. 3) will have two effects on the x-y flow *during the peak hour*:

- ▶ increase the packet loss and/or delay, and hence the TCP retransmission timeouts (RTO) and round-trip-times (RTT) (Ricciato, Vacirca, Karner, 2005);
- ▶ compress the marginal rate distribution of the aggregate traffic against the capacity limit (see Ricciato, Fleischer, 2006).



**Fig. 3. Traffic aggregate to the bottleneck link B captured on a different monitoring point M**

Both effects are observable at a different monitoring point (e.g. point M in Fig. 3) so that capacity alarms can be triggered from the pure analysis of Gn traces with no need to access intermediate nodes. Specifically, an anomalous high level of packet loss and/or RTT in B can be inferred by correlating the TCP DATA and ACK packets observed in M: the difference in the timestamps is an indication of the (semi)-RTT between the monitoring point M and the MS that includes the queuing delay at the bottleneck (if any), while the analysis of the ACK Sequence Numbers and duplicate DATA packets can reveal timeouts and retransmissions (Vacirca, Ziegler, Hasenleithner, 2006). This approach was studied in (Ricciato, Vacirca, Karner, 2005). The second approach to bottleneck detection is to perform a signal analysis of the total traffic rate of the flow x-y as seen in M. The marginal rate distribution measured in small timebins (e.g. 10s) should appear as a bell-shape curve, whose variance and level of symmetry increases with instantaneous load, as depicted in the bottom graph of Fig. 4. Instead when a bottleneck is in place along the path the total aggregate rate at the peak hour will become compressed against the capacity limit, as in the top graph of Fig. 4. Such pattern can be easily detected by the analysis of the variance and skewness at different load levels (see (Ricciato, Fleischer, 2006) for more details).




**Fig. 4. Evolution of the marginal rate distribution of a TCP-controlled aggregate with load: with (bottom) and without (top) congested bottleneck along the path**

It is remarkable that both such approaches *do not assume any external knowledge about the x-y path*, e.g. the provisioned bandwidth, but rely exclusively on the observation of the actual traffic. This increases the robustness of the scheme, as it can reveal capacity restrictions due to configuration errors and equipment malfunctioning, and at the same time reduce the provisioning efforts of the tool itself, as it does not need to be updated upon each re-configuration and/or re-provisioning of the x-y path. Note also that complete capturing of Gn traffic requires only few probes near the GGSNs. Based on such considerations it should be clear that the proposed bottleneck detection schemes are attractive also from the perspective of implementation and maintenance costs.

**5.2 Detecting congestion in the Radio Access Network**

We are currently seeking to extend such approach beyond the Core Network, to detect cell-level congestion in the Radio Access Network. The basic idea is again to merge information extracted at the TCP layer with spatial information derived from the 3GPP layers. With reference to the GPRS/EDGE section, monitoring on the Gb links



**Fig. 5.** Loglog scatterplot of  $N_i$  vs.  $V_i$  for part of the GPRS/EDGE population ( $N_i$  number of daily PDP-contexts;  $V_i$  total volume;  $i$  denotes MS index)

near the SGSNs is sufficient to retrieve the packet-to-cell association for each data packet in both directions. Then the packet-level TCP performance indicators (RTOs and RTTs, see (Ricciato, Vacirca, Kärner, 2005)) can be measured for each cell, and these data used to pinpoint the need for capacity increase in those areas yielding recurrent signs of poor performance. Such novel approach is still in the exploratory phase, but early preliminary results are promising (Ricciato et al., 2006a) about the applicability of such data to complement the current practice of radio network optimization and capacity enhancement (Neubauer, Toeltsch, 2005).

### 5.3 Detecting behaviors on the data plane

One key feature of the METAWIN monitoring system is the ability to correlate information extracted at the TCP/IP layers (e.g. L4 port, IP address, application type) with those extracted at the lower 3GPP layers (e.g. MS, cell, PDP-context attributes), and to extract cross-layer associations (as discussed in Sect. 4). In general terms such feature is certainly “costly” from an implementation point of view, but it pays off as it enables many analysis tasks, some of them relatively simple conceptually, that could not be performed otherwise.

#### 5.3.1 IP spoofing

To start with a simple example, it is possible to detect whether any of the MS is using spoofed IP address by simply matching the source IP address in the packet with that originally assigned during the PDP-context activation. The usage of a spoofed address is a strong indication of malicious activity, performed either deliberately by the user or by malicious agents hijacking his/her terminal (e.g. laptop).

#### 5.3.2 Polling applications

The ability to refer traffic entities (e.g. TCP connections, WEB sessions, PDP-context) to individual users (through the anonymized IMSI) allows for the inspection of individual traffic patterns in the multidimensional space. As a simple example, consider that we are interested in exploring the activity of each MS at the level of PDP-context. With the METAWIN system we can extract the full set of PDP-context along with their attributes, including the unique MS

identifier (hashed IMSI). For each MS  $i$  we can extract the total number of PDP-context generated in a certain period (e.g. one day, denote by  $N_i$ ) and the total transferred volume (cumulated uplink + downlink, denote by  $V_i$ ). We can further discriminate the PDP-contexts activated in the GPRS/EDGE and UMTS/HSPDA areas. In Fig. 5 we report the scatter plot of the  $(N_i, V_i)$  points for a part of the GPRS/EDGE population (logarithmic scale on both axes). Note that these types of visualization are preliminary to any further exploration of the per-MS usage patterns. It can be noticed a linear cluster of MS with very high number of PDP-contexts (ellipsoid in Fig. 5), over 1000 in only 24 h. The arrangement of the data points along a line in the  $(N_i, V_i)$  plane suggests that a constant amount of data is exchanged in each PDP-context. A compact cluster in a multidimensional space is often a strong indication that the involved entities (in our cases MSs) have strong commonalities, i.e. homogeneous behavior. In order to identify the invariant elements within the cluster the next step is to pick one or few sample MS from the cluster and analyze their sub-traces separately. The joint analysis of APN, IP address and L4 ports allowed us to recognize that most of the involved PDP-contexts were opened by the same common application. Further investigations revealed that such application requires periodic polling by the MS to a central server. On a specific class of terminals this caused the periodic activation of short-term PDP-contexts at a very high rate.

There are several lessons to be learned from the above example. The first one is that applications designed and developed for the wired Internet should not be blindly imported into the 3G network. A sound “dual expertise” is required to predict the impact of TCP/IP applications onto the functionally-complex 3GPP layers. Second, we found that the total number of PDP-context generated by such few MSs was extremely high, accounting for a noticeable fraction of the total number of PDP-context (not disclosed), despite the number of involved MS being negligible compared to the total MS population. Note that each PDP-context has a fixed cost in terms of network resources, e.g. signaling load on the GGSN for the activation/deactivation procedures. Furthermore, the billing systems are PDP-context based, i.e. a billing ticket is generated by the GGSN and processed by the billing system for each PDP-context. In this and other similar cases the consumption of some types of network

resources are dominated by very small subpopulations. This is potentially dangerous, at least to the extent that it was not anticipated, as small subpopulations can grow very quickly.

### 5.3.3 Undesired traffic

With similar methods we detected the presence in 3G of so called "undesired traffic". This term and other similar ones (e.g. "unwanted traffic" or "background traffic" as in (Pang et al., 2004)) refer cumulatively to those traffic components originated directly or indirectly by malicious or anyway "non productive" activities (Pang et al., 2004). It includes backscatter traffic associated, to remote DoS attacks, scanning worm probes, spam, exploit attempts, etc. Undesired traffic might have a negative impact onto the underlying network, and in extreme cases drive the network or at least some of its elements to resource exhaustion (Ricciato, 2006).

We found evidence of undesired traffic in the network under study already in December 2004 and discussed its potential impact on a 3G network during periods of large infections (see (Ricciato et al., 2006b) and references therein). The presence of such traffic in the 3G network originates mainly from the fact that an important fraction of 3G terminals are laptops equipped with popular operating systems and exposed to virus and worm infections. During the process of self-replication these agents generate traffic (e.g. scanning SYN's or spam emails) that ultimately consumes network resources. Note that we use the term "resources" in the general sense: there is a broad range of "resource dimensions" in a typical 3G network, examples include memory state in stateful elements (e.g. proxy, GGSN), logical resources in the radio cells (e.g. DCH), signal channel capacity on the radio link, cache-space in the routers, resources of the billing system, etc. Some kinds of traffic sources – particularly those not considered at the time when the network was designed and configured, e.g. scanning worms – might cause large consumption on one or very few resource dimensions. It is likely that each of such resources is monitored independently by the network operator staff, and independently from other traffic measurements. Therefore, the logical correlation between a high level of consumption of a single resource and the presence of a very specific traffic component (if detected at all) is left to the intuition of the network staff. On the other hand large-scale continuous TMA is the only way to ensure that such phenomena are readily recognized and handled by the network operator as they emerge, before they can cause serious damage to the network.

### 5.4 Detecting anomalous behaviors on the signaling plane

Monitoring the traffic on the control-plane allows for the detection of problems and misbehaviors in the 3GPP layers. The most common causes of such events are configuration errors and software bugs in the terminals. Other error sources include misconfigurations and bugs in the network elements (local or foreign), particularly for signaling procedures involving equipments from different vendors. There is also a theoretical risk of attacks on the control plane by malicious users (see (Yang et al., 2006) Sect. III). Such forms of attacks are more difficult to launch than IP-based attacks as they require hacking the control-plane software of the terminals, which is usually encoded in firmware. However, the terminal firmware is much more accessible nowadays than in the past, and several terminal models already allow the installation of modified versions of firmware code available from the Internet.

The detection of anomalies is relatively easier in the control plane than in the data plane. In fact many kinds of problems are explicitly encoded into specific error messages. The monitoring system is able to identify the message-IMSI association also for signaling frames, and to count the occurrence of each error message or erroneous

procedure for each MS. As expected, more errors were found on those interfaces with a higher level of signaling interactions, e.g. Gb. By the analysis of several sample traces we found that most of the error messages are concentrated on a few error types, and that the vast majority of them are generated by a relatively small portion of MSs, typically due to software bug in the terminal. Simple thresholds on the absolute number of signaling messages per-MS are often sufficient to identify buggy terminals, as the frequency of error messages for normal terminal is typically null or extremely sporadic, while buggy terminals tend to iterate erroneous procedures several times. Therefore, the "signaling activity" of buggy terminals is typically several orders of magnitude higher than the rest of the population.

A side effect of such phenomenon is that raw global error ratios (e.g. total number of rejections divided by number of requests) might be misleading if adopted as indicators of the network state, as they can be severely inflated by few buggy terminals. To illustrate we propose a simple numerical example. Consider a GPRS SGSN serving 100,000 users in the peak hour, 50 % of which attach to the network during between 7–9 am, equivalent to approximately 2000 Attach Requests every 5 min. Now consider a single buggy terminal iterating unsuccessful attach procedures. The capacity of the signaling channels on the radio link enforces a minimum period between consecutive Attach Request/Reject cycles of approximately 800 ms. That means in 5 min a *single* buggy terminal can generate 375 Request/Reject pairs, driving the total Rejects-to-Requests ratio over 15 %. Notably such a high reject ratio is not due to a network problem, but just to a single buggy terminal, and should not trigger any network troubleshooting process. Unfortunately global counters are in some cases the only type of indicator provided by the network element. This is a very example of the "semantic limitation" of equipment data discussed above in Sect. 3.

## 6. Conclusions and future work

A 3G network is a magnificently complex object embedded in a magnificently heterogeneous usage environment. It combines the functional complexity of the wireless cellular paradigm with the protocol dynamics of TCP/IP networks. Both the network and the usage environment are continuously evolving.

Understanding such an environment is more urgent and at the same time more difficult than for legacy 2G networks, that were intrinsically simpler and subject to slower changes. Continuous traffic monitoring, coupled with routine expert-driven traffic analysis, should be regarded as a fundamental component in the global process of running a real network. This requires the development and deployment of an advanced TMA system to complement legacy network monitoring systems based on data provided by the network elements. The ideal system will be flexible, extensible and endowed with data-extraction, visualization and summarization, so as to support and amplify the capability of network experts to understand and interpret the current state of the network and its traffic. Such a system will be intrinsically multi-purpose and serve as a centralized source of data for several departments, as envisioned in Fig. 6. We claim that the cost of the TMA process will be paid off by the early recognition of potential troubles and hidden issues: preventing a potential problem or fixing it at an early stage is generally cheaper than restoring the damage after the problem has manifested to the customers.

Undesired traffic and deliberate attacks are the first potential problem that might in principle affect the performance and stability of 3G networks. We are currently working towards scalable methods to early detect infected terminals and recognize mounting infections as early as possible. This is one of the goals of a follow-up research project, DARWIN (Data Analysis and Reporting in Wireless Networks),

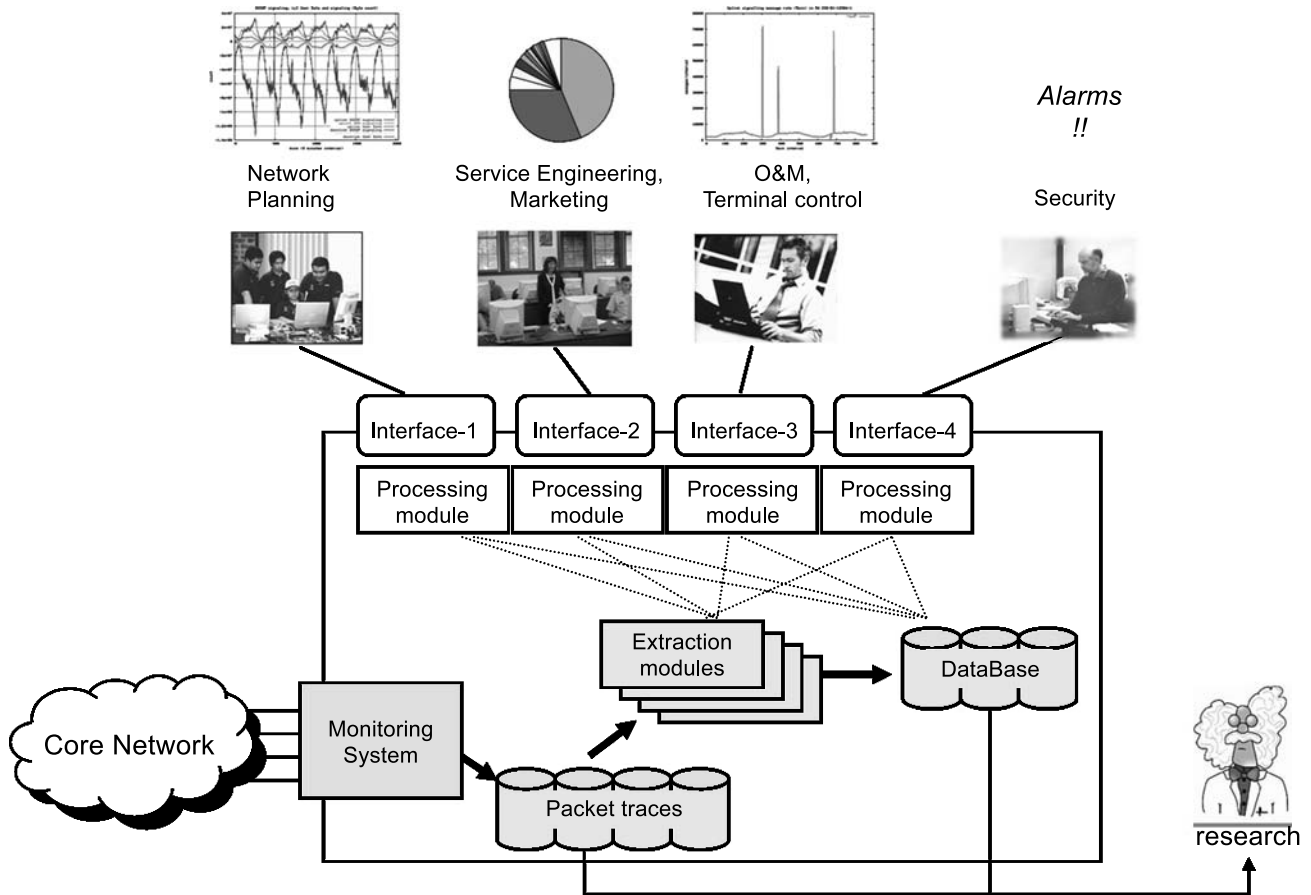


Fig. 6. Scheme of an advanced multi-purpose TMA system

that capitalizes on the expertise and tools developed in the METAWIN project (<http://userver.ftw.at/~ricciato/metawin.html>).

**3GPP acronyms and short description**

3GPP	3rd Generation Partnership Project
APN	Access Point Name
BSC	Base Station Controller: the subsystem controller in the GPRS Radio Access Network
DAG	Data Acquisition and Generation, dedicated capture cards by Endace
DCH	Dedicated Channel
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Server
DoS	Denial Of Service
EDGE	Enhanced Data rates for GSM Evolution
GGSN	Gateway GPRS Support Node
GPRS	General Packet Radio Service
GSM	The second-generation mobile cellular system adopted in Europe
GTP	GPRS Tunneling Protocol
HLR/AuC	Home Location Register/Authentication Center
HSDPA	High Speed Downlink Packet Access
HTTP	Hyper-Text Transfer Protocol
IDS	Intrusion Detection System
IMEI	International Mobile Equipment Number: the unique identifier of the terminal. It includes a different TAC code for each terminal type
IMS	IP Multimedia Subsystem
IMSI	International Mobile Subscriber Number: the unique

(continued)

**3GPP acronyms and short description (continued)**

	identifier of the SIM card (typically associated to a single user)
ISP	Internet Service Provider
MIB	Management Information Base
MS	Mobile Station
MSISDN	The telephone number associated to each SIM (Subscriber Identity Module)
NAT	Network Address Translation
PDP-context	The 3G homologous of a dial-up connection in a modem-based ISPs. IP address are dynamically assigned [released] at the PDP-context activation [deactivation]
RA	Routing Area
RADIUS	Remote Authentication Dial-In User Server/Service
RAID	Redundant Array of Inexpensive Disks
RNC	Radio Network Controller: the subsystem controller in the UMTS Radio Access Network
RTO	Retransmission Timeout
RTT	Round-Trip Time
SGSN	Serving GPRS Support Node
SNMP	Simple Network Management Protocol
TAC	Type Approval Code
TCP	Transmission Control Protocol
TMA	Traffic Monitoring and Analysis
UMTS	Universal Mobile Telecommunications System
W-CDMA	Wideband Code Division Multiple Access
WIMAX	Worldwide Interoperability for Microwave Access, a technology based on the IEEE 802.16 standard (see <a href="http://www.wimaxforum.org">http://www.wimaxforum.org</a> )
WLAN	Wireless Local Area Network



**References**

- Bannister, J., Mather, P., Coope, S. (2004): Convergence technologies for 3G networks. Wiley.
- Benko, P., Malicsko, G., Veres, A. (2004): A large-scale, passive analysis of end-to-end TCP performance over GPRS. Proc. IEEE INFOCOM, Hong Kong, 2004.
- Borsos, T., Szabo, I., Wieland, J., Zarandi, P. (2006): A measurement based solution for service quality assurance in operational GPRS networks. Proc. IEEE INFOCOM, Barcelona, April 2006.
- Keshav, S. (2005): Why cell phones will dominate the future Internet. ACM Computer Comm. Review, vol. 35, n. 2, April 2005.
- Kleinrock, L. (2005): The Internet: history and future. Lectio Magistralis at Politecnico di Torino, October 2005 (<http://www.tlc.polito.it/nordio/seminars>).
- Lee, Y. (2006): Measured TCP performance in CDMA 1x EV-DO network. Proc. of PAM 2006, Adelaide, Australia, March 30–31, 2006.
- Neubauer, T., Toeltsch, M. (2005): UMTS radio network planning – maximizing return on investment, e&i (122), H. 3: 108–113.
- Pang, R., et al. (2004): Characteristics of Internet background radiation. IMC'04, Taormina, Italy, October 2004.
- Pospischil, G., Miladinovich, I., Kunczler, H. (2005): Hot topics for mobile services. e&i (122), H. 3: 102–107.
- Ricciato, F. (2006): Unwanted Traffic in 3G, editorial for ACM Computer Comm. Review, vol. 36, n. 2, April 2006.
- Ricciato, F., Fleischer, W. (2006): Bottleneck detection via aggregate rate analysis: a real case in a 3G network. Short paper at IEEE/IFIP NOMS'06, Vancouver, April 2006.
- Ricciato, F., Pilz, R., Hasenleithner, E. (2006): Measurement-based optimization of a 3G core network: a case study. Accepted to 6th Int. Conf. on Next Generation Teletraffic and Wired/Wireless Advanced Networking (NEW2AN'06), St. Petersburg, Russia, May 29, 2006.
- Ricciato, F., Vacirca, F., Karner, M. (2005): Bottleneck detection in UMTS via TCP passive monitoring: a real case. Proc. of ACM CoNEXT'05, Toulouse, France, October 24–27, 2005.
- Ricciato, F., Vacirca, F., Fleischer, W., Motz, J., Rupp, M. (2006a): Passive tomography of a 3G network: challenges and opportunities. Poster at IEEE INFOCOM, Barcelona, April 2006.
- Ricciato, F., Svoboda, P., Hasenleithner, E., Fleischer, W. (2006b): On the impact of unwanted traffic onto a 3G network. Proc. of 2nd Int. workshop on Security, Privacy and Trust in Pervasive and Ubiquitous Computing (SecPeru06), Lyon, France, June 29, 2006.
- Svoboda, P., Ricciato, F., Hasenleithner, E., Pilz, R. (2006): Composition of GPRS/UMTS traffic: snapshots from a live network. Proc. of 4th Int. Workshop on Internet Performance, Simulation, Monitoring and Measurement, Salzburg (IPS-MOME'06), Austria, February 27–28, 2006.
- Vacirca, F., Ricciato, F., Pilz, R. (2005): Large-scale RTT measurements from an operational UMTS/GPRS network. Proc. of the 1st Int. Conf. on Wireless Internet (IEEE WICON 05), Budapest, Hungary, July 2005.
- Vacirca, F., Ziegler, T., Hasenleithner, E. (2006): An algorithm to detect TCP spurious timeouts and its application to operational UMTS/GPRS networks. To appear in Journal of Computer Networks, Elsevier.
- Yang, H., Ricciato, F., Lu, S., Zhang, L. (2006): Securing a wireless world. Proc. of IEEE – Special Issue on Cryptography and Security Issues, vol. 94, n. 2, February 2006.