Mobile SIP: An Empirical Study on SIP Retransmission Timers in HSPA 3G Networks

Joachim Fabini, Michael Hirschbichler, Jiri Kuthan, Werner Wiedermann

HAL Id: hal-01497039
https://inria.hal.science/hal-01497039

Submitted on 28 Mar 2017

Mobile SIP: An Empirical Study on SIP Retransmission Timers in HSPA 3G Networks

Joachim Fabini[1], Michael Hirschbichler[1], Jiri Kuthan[2], and Werner Wiedermann[3]

[1] Institute of Telecommunications, Vienna University of Technology, Favoritenstr. 9/E389, A 1040 Wien
[2] iptelorg GmbH, Am Borsigturm 11, D 13507 Berlin
[3] Telekom Austria Group, Lassallestrasse 9, A 1020 Wien

Abstract. Mobile packet-switched voice must replace mobile circuit-switched voice latest with the large-scale deployment of LTE. However, short-term full-area LTE coverage is highly unlikely, handover to and interworking with 3G HSPA networks being a must. This raises the question of how the Session Initiation Protocol SIP performs in today's HSPA networks.

Relying on live SIP provider monitoring data and active HSPA measurements, this paper demonstrates that default SIP retransmission timers are inappropriate for real-world deployments of packet switched voice over today's 2G and 3G networks including HSPA. Main reason for this deficiency is the reactive, demand-driven resource allocation strategy of HSPA which results in huge uplink delays. These delays eventually lead to a high number of unnecessary SIP message retransmissions which load networks and servers. SIP clients for mobile devices should therefore implement adaptive, user-configurable SIP retransmission timers and provide appropriate default values for these timers to enable seamless SIP operation in mobile cellular networks.

1 Introduction

From a standardization point of view, the Session Initiation Protocol (SIP) [1] is currently considered to be the most promising candidate for implementing Voice over IP (VoIP) in converged fixed and wireless networks. However, the SIP design and base standard is more than ten years old, being focused on providing best-effort services in deterministically behaving wired networks. During the last years, driven by requirements for deploying SIP in carrier-grade infrastructures, the research community has started to publish and discuss severe SIP deficiencies with respect to unstable behavior in high-load or overload situations.

Latest with the replacement of mobile circuit-switched voice technology with packet-switched-only capable LTE networks, IP-based signaling must operate in the mobile domain. In this paper we demonstrate by means of representative SIP server log files and measurements in state-of-the-art HSPA networks that default SIP timers as recommended by RFC 3261 are inappropriate for SIP operation

in mobile 3G networks. This observation is valid for all SIP messages which are sent after seconds of inactivity, outside of ongoing data sessions, i.e., particularly for registrations and SIP presence related messages. SIP server statistics evaluations show that these SIP message types account for the majority of the overall SIP traffic. Our HSPA measurement results demonstrate that due to the highly reactive nature of HSPA uplink, timers of mobile SIP user agents will deterministically trigger SIP retransmissions, which generate redundant traffic and waste considerable resources in the network and in servers. This is a severe drawback, aggravated by the fact that mobile applications reduce user-configurable options to an absolute minimum and typically do not implement user-configurable timer values.

1.1 Related Work

The SIP protocol and SIP retransmission timers are defined in RFC 3261[1], whereas SIP performance metrics for registration request delay are standardized in RFC 6076[8]. C. Egger et. al. discuss in [5] and [7] the inherent problem of SIP retransmissions and their negative impact on SIP proxy performance and networks, proposing to manipulate SIP retransmission timers to handle SIP overload situations. The authors of [6] conclude that SIP overload control mechanisms are an absolute MUST, implicit overload control being subject to lower overhead but explicit overload control behaving more deterministic in situations of high overload. Vingarzan et. al. discuss in [9] the latency impact of various wireless network technologies, including 2G, 3G, and WiFi onto end-to-end SIP performance for IMS, concluding that GPRS networks are not appropriate for IMS signaling, whereas IEEE 802.11b,g, and CDMA networks can handle IMS traffic adequately in terms of delay and loss. The contribution of different network components (user equipment, NodeB and RNC) on HSUPA end-to-end delay is analyzed in [10], inferring on 30 ms delay for 1 kbyte packets. In earlier work ([12] and [13]) we have dissected difficulties and uncertainty factors to be expected when measuring mobile network delay, focusing on randomness in terms of inter-packet delay and payload size as key factor to reveal load-dependent network behavior.

1.2 Structure of this Paper

The remainder of this paper is structured as follows: section 2 reviews the SIP registration procedure and discusses the methodology used for iptel.org SIP registrar log evaluation and for HSPA measurements. Building on this information, section 3 presents and discusses SIP registrar log and measurement results from live HSPA networks. The paper ends with summary and conclusions in section 4.

2 Protocols and Setup

Adoption of SIP by the 3rd Generation Partnership Project (3GPP) as main signaling protocol for the IP Multimedia Subsystem (IMS) has significantly boosted

the world wide use of SIP. Due to high scalability and freely available standards, several free, open-source SIP client and server implementations are available for download. In the following we restrict discussions to plain SIP registration, arguing that 3GPP SIP signaling messages are typically larger than plain SIP signaling messages because of 3GPP-specific SIP extensions and headers. Therefore, delays are likely to increase for 3GPP IMS signaling even more, plain SIP delays offering a lower boundary on propagation delays in mobile networks. Even if SIP signaling compression is deployed as mandated by 3GPP for the use on wireless cellular networks, compression is unlikely to offer significant delay reductions for sporadically sent messages like SIP registrations.

2.1 SIP Basics: Registration

SIP is a simple, plain-text message based request-reply protocol. Fundamental to SIP is the registration process. During registration a terminal, the so-called SIP user agent, registers the user's generic SIP address – the so called SIP URI – and its current transport address with its assigned SIP server, the SIP registrar. Using the transport address binding stored during the user's registration process, SIP registrars can then forward incoming calls via IP routing to the terminal(s) where the specific user is currently logged in.
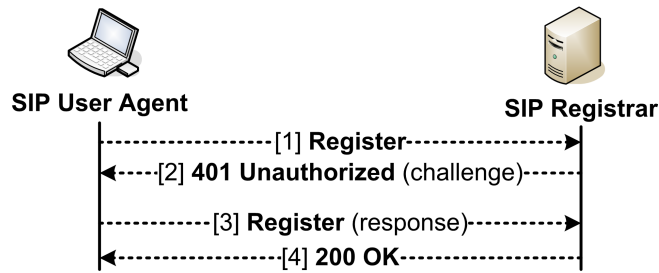

SIP User Agent   **SIP Registrar**

-----------------[1] **Register**---------------→

←-----[2] **401 Unauthorized** (challenge)------

----------------[3] **Register** (response)----------→

←-----------------[4] **200 OK**------------------

Fig. 1. SIP Registration message flow

Fig. 1 illustrates the message sequence required for an authenticated SIP registration. The first registration request (1) is rejected by the registrar using a 401, unauthorized reply (2). This reply message includes a random value, the so-called nonce, which is encrypted by the user agent on receiving the reply using a secret key shared with the registrar. This response is appended as SIP authentication header to a new registration request (3) sent to the registrar. On receiving the authenticated registration request, the registrar repeats the user agent's processing, i.e., encodes the random value with its pre-shared secret. If the result matches the response sent by the user agent in its registration request, the registering user is supposed to be authenticated. The registrar can safely store the binding between the client's SIP URI and its contact (transport) address and confirm successful authentication to the client by sending a 200 OK response (message 4 in fig. 1).

Registrations expire, raising the need for periodical registration message retransmissions to refresh the registration state. The interval between successive re-registrations is always a trade-off between the user agent's reachability for incoming calls, i.e., up-to-date registrar information, and the amount of required signaling load. In a mobile cellular context, mobility and radio provisioning are two factors which can result in IP address changes and therefore demand for relatively small registration intervals. Moreover, mobile firewalls and/or private IP address NAT devices in some cases use these re-registrations as keep-alive messages to maintain their internal bindings. Consequently, re-registration intervals for mobile SIP clients should be in the order of tens of minutes, compared to wired SIP terminals which typically re-register every hour or several hours. However, this high registration frequency causes a high data volume for SIP over mobile networks.

2.2 SIP Registrar Logging

iptelorg GmbH offers free SIP-based IP Telephony services. Users can request their free personal SIP URI @iptel.org and use this URI to register their SIP user agents with the SIP registrar iptel.org. The core functionality behind the iptel.org registrar is implemented by the open source SIP Express Router (SER)[14], which offers complex and detailed logging facilities. The results published in this paper base on 27 hours of anonymized, representative iptel.org SER logs acquired during a typical business day in spring 2012. During these 27 hours, the log has recorded a total of 9,980,000 SIP requests including retransmissions which scales down to an effective traffic of 7,400,000 SIP requests and 7,500,000 SIP replies when eliminating SIP message retransmissions. Prior to evaluation, all SIP messages have been categorized into several orthogonal classes, focusing on the SIP registration process.

First, all SIP requests have been categorized according to their type. A second category differentiates between the user agent classes "fixed" and "mobile" by comparing logged SIP User-Agent headers against a list of well-known user agent strings. Subclass "mobile" aggregates smartphone clients, i.e., software SIP clients for iOS and Android which might face limitations in terms of computing ressources. However, "mobile" does not infer on any specific kind of network connectivity, on the contrary. One of our conclusions in this paper is, that the majority of these mobile clients connect using Wireless LAN networks rather than using cellular technologies. As second option, the subclass "fixed" groups hard-phones, SIP-routers and -proxies, as well as soft-phones for desktop operating systems which typically have abundant ressources available.

Table 1 summarizes the SIP registrar log, classifying SIP requests according to their request type for fixed clients (*Fixed*), mobile clients (*Mobile*), and aggregated (*Total*), considering only first requests (i.e., eliminating all retransmissions). Registrations dominate this statistic, accounting for more than 40% of all requests (3,090,000 registrations in absolute numbers). This confirms our assumption that SIP registration requests generate significant load in mobile

Request	Total	Fixed	Mobile
INVITE	0.19%	0.29%	0.03%
BYE	0.05%	0.08%	0.01%
REGISTER	41.50%	51.45%	25.39%
SUBSCRIBE	2.77%	3.46%	1.65%
NOTIFY	17.81%	28.82%	0.001%
OPTIONS	37.22%	15.21%	72.85%
Others	0.46%	0.70%	0.07%

Table 1. SIP request distribution (no retransmissions)

and fixed networks. However, in a server-based enterprise infrastructure featuring server-based presence services we predict a load shift towards notify and subscribe requests. As third category, because of their distinct handling and response time within the SIP registrar, SIP registration requests must be divided into registrations *without* authentication information (message 1 in fig. 1) and registrations *with* authentication information (message 3 in fig. 1).

More, detailed registration request statistics can be found in the following measurement results section.

2.3 HSPA Measurement Setup

For measuring one-way (OWD) and round-trip delay (RTD) in mobile cellular networks we have used operational High Speed Packet Access (HSPA) networks.
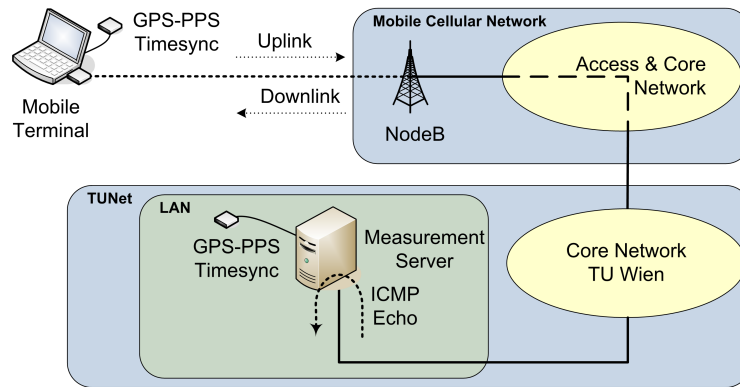

Fig. 2. Measurement Setup for Accurate OWD Measurements

Fig. 2 depicts the setup used for active end-to-end delay measurements in public HSPA networks. Measurement server and mobile terminal are both running Ubuntu Linux 12.04 using a custom-compiled Kernel 3.4.0 at 1 kHz kernel tick with Pulse-per-second (PPS) functionality activated. The measurement server is equipped with an Intel Q6600 CPU at 2.4 GHz and 4 Gbyte of RAM,

connected using its on-board Ethernet interface to the public Internet. The mobile client is a Dell Latitude D630 laptop with an Intel T7100 processor at 1.8 GHz and 2 Gbyte of RAM, connected to the mobile network using a Huawei E870 HSPA ExpressCard 34. For accurate OWD measurements, mobile client and server are both time-synchronized using the network time protocol daemon (NTPD) and local EM 406A based GPS-PPS time sources as recommended by [11]. The maximum measured deviation from global time is below 50 μs, which is sufficiently accurate for our measurements. From a measurement methodology point of view we have followed and extended the recommendations of the IP Performance Metrics Framework[2] to avoid correlation with periodic network behavior. The measurement traffic is random with respect to send time and payload size. ICMP echo request packets of representative random payload size are sent at random time intervals, triggering corresponding ICMP echo replies.

3 Measurement Results

This section presents the iptel.org SIP server log analysis and HSPA delay measurement results for sporadic SIP registration traffic designed according to the iptel.org log evaluation. As main criterion for representativeness we have used the live traces from iptel.org, such that the packet size of our measurement traffic matches the size of real registration messages.

3.1 SIP Registrar Statistics

As mentioned in subsection 2.2, the monitored SIP traffic totals 27 hours of traffic. Out of the 7.4 million unique SIP requests 4.6 million requests are from fixed clients and 2.8 million from mobile clients conforming to the mobile-fixed categorization presented in section 2.2. Table 2 illustrates that mobile clients re-

	Register Requests	Register Requests w Retransmissions	Ratio
Fixed	2,365,404	167,467	7.08%
Mobile	721,303	161,111	22.34%

Table 2. Total register requests vs. registrations which require retransmission

transmit SIP registration requests almost three time as frequent as fixed clients. In average, every fifth register request is retransmitted due to firing of SIP retransmission timer T1 (500 ms). Retransmission messages are evenly distributed over all mobile clients and request sizes, which excludes that this problematic behavior is initiated by one specific user or user-agent implementation. It must be emphasized once again that "mobile" client is NOT a synonym for 3G connectivity. On the contrary, more detailed iptel.org log evaluation indicates that the majority of mobile clients connects using WLAN or faster non-3G technologies.
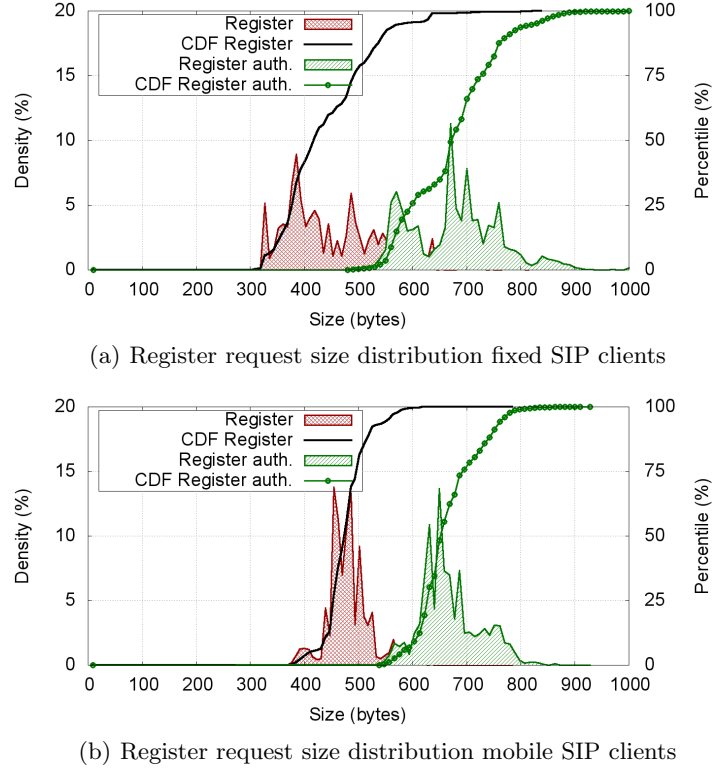
(a) Register request size distribution fixed SIP clients

(b) Register request size distribution mobile SIP clients

Fig. 3. SIP REGISTER size distribution with and without authentication header

As presented in section 2.1 in fig. 1, an authenticated SIP registration requires two SIP register requests. The second request (message 3) differs from the first one (message 1) by the additional SIP authentication header. Fig. 3(a) and fig. 3(b) show the distribution of SIP register request size with and without authentication header for fixed and mobile clients, respectively. Accordingly, the authentication header is increasing SIP request size in average by about 200 bytes.

To eliminate SIP server or core network performance as potential delay sources we have evaluated processing duration of SIP registration requests with and without authentication. Using metrics from [8], nearly 99% of all "401 Unauthorized" responses are sent within 1 ms and 98% of all "200 OK" within 2 ms of registration request arrival. Both values are well below the SIP retransmission timer T1 value of 500 ms, such that SIP server delay can be excluded as reason for timer T1 expiry.

3.2 HSPA Measurement Results

Based on the evaluation of iptel.org's SIP server log we have built a prototypical signaling data stream at IP layer using ICMP messages to measure accurate

round-trip (RTD) and one-way uplink (UL) and downlink (DL) delay measurements for SIP signaling traffic in live HSPA networks. Main motivation of these HSPA delay measurements was to determine the impact and feasibility of VoIP on top of mobile HSPA networks. HSPA was selected as mobile cellular technology because back in Q2 2012, when acquiring iptel.org traces for this publication, no LTE capable terminals have been available on the market. Therefore it is unlikely that mobile clients have accessed iptel.org in spring 2012 using a more advanced technology than HSPA. Today, one year later, Long Term Evolution (LTE) networks are operational in many countries worldwide but typically still limited in their coverage to high-traffic areas.

Using the setup presented in section 2, an extended Internet Control Message Protocol (ICMP) packet generator and an artificial, random signaling data traffic we have measured end-to-end RTD and OWD in public live Austrian HSPA networks. As mentioned earlier, measurement representativeness has been increased by following the recommendations of the IPPM Framework [2] and of its associated metrics for OWD [3] and RTD [4] when designing the measurement traffic. In particular we have created a scenario consisting of 10,000 single ICMP measurement packets. Based on iptel.org log's payload size distribution for mobile clients as depicted in fig. 3(b), the measurement traffic's payload range is uniform distributed between 400 and 800 byte. The measurement traffic's inter-packet interval is uniformly distributed between 1 and 10 seconds, totaling more than 15 measurement hours. The selected delay range has been chosen as a minimum period of mobile network inactivity. We argue this period to be exceeded in most cases for registration requests while the terminal is idle, i.e., has no ongoing packet-switched voice or data sessions.

Figure 4 shows the resulting diagrams and statistics for RTD measurements. The x-y scatter plot in fig 4(a) depicts delay as a function of payload size. Any one of the almost 10,000 dots in the figure correspond to the delay value of one single measurement packet. Only 3 packets out of the 10,000 were lost, meaning that loss can be excluded for the large number of retransmission of mobile clients. The corresponding cumulative distribution function (CDF) in fig 4(c) reveals that – ignoring potential processing overhead and delay in SIP Registrars - slightly more than one third (36.2%) of all measurement samples arrive in less than 500 ms. Accordingly, almost two thirds (63.8%) of the requests will trigger retransmissions. More than 10% of all requests will exceed 1500 ms and trigger at least two retransmissions.

As mentioned earlier in section 2.1 SIP re-registrations must be sent periodically. We argue that there is a high likelihood that intervals of mobile network inactivity preceding a SIP Registration are significantly larger than one second. We have therefore selected from all 10,000 measurement samples the 5441 ones which have inter-packet delays – i.e., network inactivity period before sending the packet – between 5 and 10 seconds. The resulting scatter plot is depicted in fig 4(b) and the distribution function in fig 4(d), respectively.

The result is alarming with respect to user experience, network load and server load. Considering exclusively network delay, less than 4% of all mea-
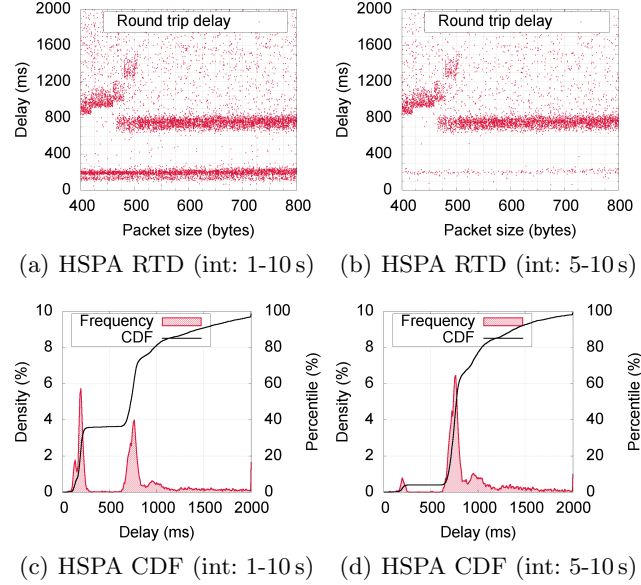
(a) HSPA RTD (int: 1-10 s)　(b) HSPA RTD (int: 5-10 s)

(c) HSPA CDF (int: 1-10 s)　(d) HSPA CDF (int: 5-10 s)

Fig. 4. RTD dependence on inter-packet send interval for HSPA networks.

surement samples arrive within the expected timeframe of 500 ms to avoid SIP retransmissions. More than 96% of all SIP requests will therefore trigger at least one, 8% will exceed the default limit of 1500 ms and trigger at least two retransmissions. These results show that in a passive state with no ongoing call, mobile SIP devices on HSPA networks will effectively – at least – double the effective load on the network and on the registrar. Although resource allocation and -release in HSPA networks depends significantly on operator-specific parameters, measurements in other mobile HSPA networks suggest that the figures we have obtained are relevant with respect to the order of magnitude. It is important to note that these findings are not restricted to Voice over IP but can be generalized for other SIP-based services as Presence or Instant Messaging, too, which all use the default timer T1 value of 500 ms.

For isolating an accurate origin of the large RTD we have dissected measured round-trip-delay samples into UL and DL delay as displayed in fig 5. The figures confirm that the major delay uncertainty factor is HSPA UL. The reason for the high end-to-end delay experienced by a mobile client following a period of network inactivity is intrinsic to HSPA's architecture. Mobile clients must send signaling messages over the wireless link to the network-based scheduler to switch from inactivity or to request allocation of additional capacities in the wireless network. These round-trip messages delay the sending of first messages, whereas subsequent messages can be transmitted much faster due to already allocated capacities.

Comparing one-way UL delay diagrams for 1-10 s inter-packet delay in fig 5, left (5(a) and 5(d)), with the corresponding diagrams with larger inter-packet delay of 5-10 s shown in fig 5, center (5(b) and 5(e)), reveals that the ratio of

requests which unconditionally run into SIP retransmission timeouts increases to more than 90% after more than 5 seconds of inactivity.
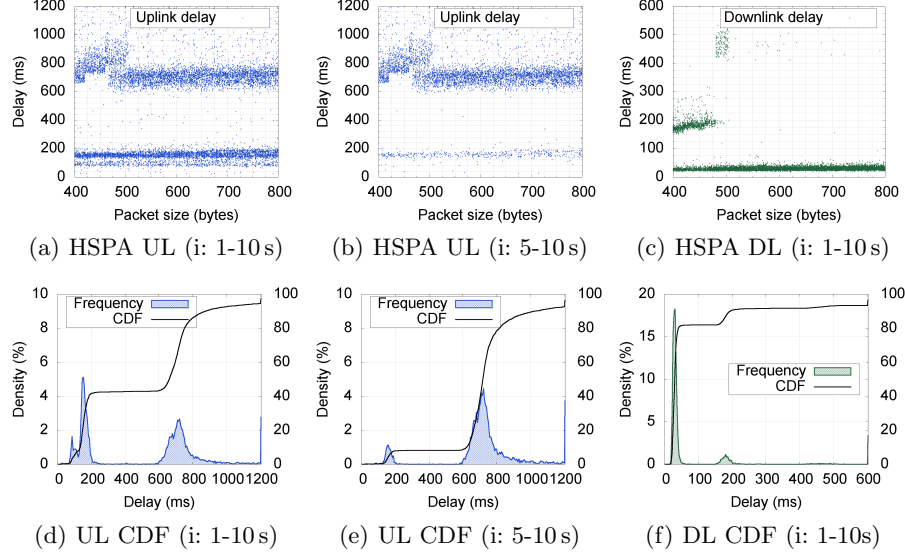

(a) HSPA UL (i: 1-10 s) (b) HSPA UL (i: 5-10 s) (c) HSPA DL (i: 1-10 s)

(d) UL CDF (i: 1-10 s) (e) UL CDF (i: 5-10 s) (f) DL CDF (i: 1-10s)

Fig. 5. UL and DL delay dependence on inter-packet send interval for HSPA networks.

The DL delay depicted in fig. 5(c) and 5(f) recovers much faster from inactivity. On DL data arrival the scheduler is located in the wired network domain and can act instantaneously, which saves time-consuming round-trip signaling messages over the wireless link. More than 80% of DL samples are below 50 ms and more than 90% below 200 ms end-to-end delay

3.3 Discussion

Two main factors may originate SIP retransmissions: mobile client limitations in terms of software and hardware or reactive HSPA UL behavior. Diving deeper into iptel.org's server log details reveals that a) mobile client limitations can be one main reason for increased retransmission ratio on iptel.org and b) as of today few mobile clients use 2G or 3G technologies to connect to iptel.org. These statements can be concluded from table 3 which shows SIP retransmission ratio depending on a presumed mobile link state. Based on the finding of section 3.2 that HSPA clients experience an increased delay after more than 5 seconds of network inactivity, mobile client registrations have been split into two categories. A mobile client's network connection is supposed to be *active* when iptel.org has logged a message from the client's socket within 5 seconds before arrival of the request, or *inactive* otherwise.

Comparing total registration request retransmission ratio for mobile clients with and without authentication in table 3 reveals that more register requests

Connection	Ratio (total)	Ratio (inactive)	Ratio (active)
Fixed (w/o auth.)	7.95%	8.65%	6.07%
Fixed (w auth.)	3.63%	0.00%	4.10%
Mobile (w/o auth.)	18.70%	11.95%	45.11%
Mobile (w auth.)	24.34%	0.00%	24.46%

Table 3. Register request retransmission probability depending on load history

with authentication are retransmitted than *without*. This result conflicts with our HSPA delay measurement results, where 96% of requests sent with inactive mobile network are retransmitted but only 63% of the active ones. Combined with the relatively low retransmission ratio for iptel.org we argue that the majority of clients connect using WiFi or alternative non-cellular technologies. RTD of WiFi networks is typically below 10 ms, such that time-consuming parsing of SIP messages by the mobile client contributes to the retransmission count.

4 Summary and Conclusions

The aggregated finding of SIP server log analysis and HSPA delay measurements is that large-scale deployment and operation of default-configured mobile SIP clients will put today's HSPA networks at risk. As main conclusion of iptel.org SIP server log and HSPA measurement results we recommend all mobile SIP clients to offer user-configurable SIP retransmission timer values. The application-configured default values for HSPA networks must be at least double of the default SIP timer T1 value of 500 ms, i.e. 1 second. For a detailed discussion on retransmission timer effects and fine-tuning we refer the reader to [5]. Alternatively, SIP clients could ideally support automated, mobile-network-technology-dependent values for T1. Using cross-layer information on the network technology which the mobile device currently uses (e.g., WiFi or 3G), mobile SIP clients could configure appropriate SIP timer values to minimize the likelihood of unnecessary retransmissions.

Latest with large-scale deployment of SIP clients in mobile cellular networks, network and SIP server operators will be forced to over-dimension their infrastructure by orders of magnitude and to implement SIP overload control mechanisms. The authors of [7] have demonstrated the catastrophic impact that short-term SIP network overloads can have onto network stability. A SIP retransmission flooding triggered by a short-term overload period eventually can end up in a congestion collapse state for SIP servers and networks. Recovery from SIP congestion collapse requires appropriate algorithms and substantial load reduction. For safeguarding acceptable user experience of SIP in mobile cellular networks it is highly desired to prevent congestion collapse scenarios, adequate SIP timer handling being a first, mandatory step.

Acknowledgment

References

1. J. Rosenberg, H. Schulzrinne, G. Camarillo, et.al., *SIP: Session Initiation Protocol* Network Working Group RFC 3261, June 2002.
2. V. Paxson, G. Almes, J. Mahdavi, M. Mathis, *Framework for IP Performance Metrics* Network Working Group RFC 2330, May 1998.
3. G. Almes, S. Kalidindi, M. Zekauskas, *A One-way Delay Metric for IPPM* Network Working Group RFC 2679, Sept. 1999.
4. G. Almes, S. Kalidindi, M. Zekauskas, *A Round-trip Delay Metric for IPPM* Network Working Group RFC 2681, Sept. 1999.
5. C. Egger, M. Hirschbichler, P. Reichl, *Enhancing SIP Performance by Dynamic Manipulation of Retransmission Timers.* "Leistungs-, Zuverlaessigkeits- und Verlaesslichkeitsbewertung von Kommunikationsnetzen und verteilten Systemen", B. Wolfinger, K. Heidtmann (Hrg.); Universitaet Hamburg, Fachbereich Informatik, Vogt-Koelln-Strasse 30, D-22527 Hamburg, Bericht 298 (2011), pp. 47 - 54.
6. M. Happenhofer, J. Fabini, C. Egger, M. Hirschbichler *An Architectural and Evaluative Review of Implicit and Explicit SIP Overload Handling*; International Journal of Measurement Technologies and Instrumentation Engineering (IJMTIE) (2011), Volume 1, Issue 4; pp. 12 - 27.
7. C. Egger, M. Happenhofer, J. Fabini, P. Reichl *Collapse Detection and Avoidance for SIP Architectures*; Praxis der Informationsverarbeitung und Kommunikation (PIK) (2012), 35, Issue 2, ISSN 0930-5157, pp. 91-100.
8. D. Malas, A. Morton, *Basic Telephony SIP End-to-End Performance Metrics* Network Working Group RFC 6076, January 2011.
9. D. Vingarzan, P. Weik, *End-to-end performance of the IP multimedia subsystem over various wireless networks* , Wireless Communications and Networking Conference, 2006. WCNC 2006. IEEE
10. M. Laner, P. Svoboda, E. . Hasenleithner, M. Rupp, *Dissecting 3G Uplink Delay by Measuring in an Operational HSPA Network*, "Proceedings of the Passive and Active Measurement Conference, PAM'11", Springer, Atlanta, GA (2011), ISSN: 0302-9743; pp. 52 - 61.
11. M. Laner, S. Caban, P. Svoboda, M. Rupp, *Time Synchronization Performance of Desktop Computers*, Precision Clock Synchronization for Measurement Control and Communication (ISPCS), 2011 International IEEE Symposium on
12. J. Fabini, W. Karner, L. Wallentin, T. Baumgartner, *The Illusion of Being Deterministic - Application-Level Considerations on Delay in 3G HSPA Networks*, "Networking 2009", Springer, Lecture Notes in Computer Sciences, Volume 5550/2009 (2009), ISSN: 0302-9743; pp. 301 - 312
13. J. Fabini, L. Wallentin, P. Reichl, *The Importance of Being Really Random: Methodological Aspects of IP-Layer 2G and 3G Network Delay Assessment*, "Communications 2009", ICC '09. IEEE International Conference on
14. *SIP Express Router (SER)*, http://www.iptel.org/ser/