PreQuEst: a Scalable and Proactive Quality Enrichment for Presence Services

D. Costantini, S. Niccolini
NEC Europe Ltd. – Heidelberg, Germany
Email: {costantini; niccolini}@nw.neclab.eu

P. Bellavista
Universita' degli studi di Bologna, Italy
Email: paolo.bellavista@unibo.it

***Abstract*—The market growth of VoIP infrastructures and services is pushing towards additional facilities to enrich service offer in an open way, thus representing also a differentiating aspect and a competitive advantage for service providers. We claim that the proactive provisioning of estimations about runtime VoIP quality is a crucial facility still missing in VoIP services: for instance, such a facility could be integrated within a Presence Service to inform users of expected quality of VoIP communications towards their contacts before calling them. The paper proposes a novel and scalable solution for proactive quality monitoring, specifically designed for VoIP traffic and based on proper runtime decisions about network partitioning. On top of this monitoring layer, we have designed and implemented an effective and flexible application-level facility for end-to-end QoS estimation, exploited by an enriched Presence Service in the examined case study. The paper describes the practical experience and the lessons learned in implementing an industrial prototype of the proposal, which is integrated with standard and widespread mechanisms (e.g., SIP Voice Quality Report Event, SIP Answer Mode and SDP Media Loopback extensions) in order to facilitate market diffusion and acceptance. First preliminary experimental results show that the proposed prototype achieves accurate quality evaluations and is largely more scalable than traditional end-to-end approaches.**

I. Introduction

The recent and expected future growth of VoIP is due to several factors, among which its proneness to easily provide additional services otherwise unavailable with Public Switched Telephony Network (PSTN), by synergically exploiting the same and ubiquitously available IP-based infrastructure. Differently from PSTN, the quality of a VoIP call is not guaranteed and easily predictable because it depends on several parameters, e.g., available bandwidth in the network segments involved at runtime, packet loss, and End-to-End (E2E) delay. These parameters are typically destination-dependent and unknown in advance to the caller. According to the value of such parameters, calls may have very high quality or be almost a waste of time for involved users. The opportunity of having an accurate estimation of call quality before call starting could represent a crucial VoIP service facility: for instance, it could make users choose accordingly the most appropriate tool/way to communicate with their buddies (video, chat, email, fax, ...). Therefore, such a facility for proactive VoIP quality estimation would constitute an important asset in VoIP infrastructures and an important competitive advantage for VoIP providers.

The primary innovative idea of this paper is to show how it is possible to effectively and scalably exploit quality monitoring information, usually employed for operator-side back-office evaluations, such as network dimensioning and traffic engineering, to seamlessly provide final users with value-added applications with VoIP quality estimations. Of course, in addition to the original provisioning of VoIP quality estimations to end users, the same facility may represent a valuable investment for operators in order to proactively estimate several network parameters according to a similar approach, thus enabling operation management services, e.g., troubleshooting and bottleneck identification.

Proactive VoIP quality estimation is a technical challenge mainly because of scalability issues. We claim that proper network partitioning at runtime could be the key to achieve reasonably accurate quality estimations while imposing a limited overhead, by dynamically exploiting locality considerations and by sharing quality measurements performed on previous VoIP calls. In short, the idea is to suitably split the E2E paths between all potential pairs of users (the user to whom we are offering quality estimations and the list of her potential contacts) into different network segments. We claim the suitability of evaluating the expected quality of each segment based both on the available monitoring data about recent calls in the same network partition (locality) and on scarcely intrusive injection of small synthetic traffic in the same locality when needed.

The paper describes the design and implementation of "Presence Enrichment through VoIP Quality Estimation (PreQuEst)", a novel facility for effective and proactive quality estimations, integrated in off-the-shelf presence services as a case study. PreQuEst relies on an innovative VoIP-oriented monitoring layer that exploits an original network-partitioning approach instead of the traditional E2E one. The advantages of network partitioning resides in a huge scalability gain, primarily deriving from the possibility to partially reuse network parameter information collected from passive monitoring of other calls over partly-joint paths, thus avoiding unnecessary active measurements. This novel approach to quality monitoring is required because general network monitoring systems usually do not provide a VoIP-oriented set of measurements or do not focus at all on users' access network values. In addition, given their E2E approach, they lack the proactivity and scalability necessary in the targeted VoIP scenarios. Furthermore, and very important to the purpose of rapid spreading over the real market, PreQuEst takes advantage of standard mechanisms and tools, thus seamlessly integrates with already installed equipment, devices, and applications. In particular, given the wide installed base of VoIP-related equipment, PreQuEst integrates with SIP and its correlated facilities, e.g., for testing purposes and event reporting.

The remainder of the paper is organized as follows. Section II overviews most relevant related work, by pointing out the motivations for our original approach. The presentation of the PreQuEst architecture follows in Section III, while Section IV details the implementation of the developed industrial prototype, by also reporting some first preliminary results and experimental evidences about how PreQuEst achieves accurate VoIP quality estimations by outperforming traditional E2E approaches in terms of scalability. Conclusive remarks, lessons learned from the practical experience of prototype implementation, and directions of future work end the paper.

II. Related Work

To the best of our knowledge, the problem of proactively providing quality estimations to VoIP users, by adopting suitable application-layer metrics, has not been sufficiently investigated yet.

The few VoIP-specific monitoring solutions in the literature cannot provide final users with proactive quality estimations for E2E communications. For instance, the Brix Network's service TestYourVoIP [6] allows users to test their VoIP traffic (recently also IPVideo traffic) towards few central servers (see [7]). However, that does not enable practical and direct estimations of user-to-user quality. Another relevant but partial work in VoIP call quality estimation is the set of monitoring mechanisms currently included in Skype [5]. But the Skype solution evaluates quality only for ongoing calls, without any possibility to offer estimations for potential user-to-user communications, e.g., for all the currently online users in the personal contact list. That has a technical reason: the number of monitoring measurements in Skype quadratically grows with the number of examined communication endpoints due to the fact that Skype adopts a full-mesh network approach for reactive quality evaluation. Other related but partial solutions can be found in literature, anyway unsuitable for accurate per-user E2E estimations. For instance, [9] proposes a method for accurate quality assessment that is however quite intrusive, even if it does not actually inject any synthetic traffic: it replaces packet payloads of received flows and applies objective measurement methods; unfortunately, it cannot measure an important parameter such as delay, cannot proactively provide quality values, and needs specific software on users' terminals. [10] instead proposes combined active/passive probes but cannot measure impairments in users' access networks. The same problem occurs in [11], which presents a very similar solution embedded into Cisco routers.

About network partitioning and monitoring metrics compositions, these research topics have been widely addressed at the IP layer in the recent years. For instance, related research issues have been the subject of inter-domain monitoring projects [19] and are currently addressed by the IP Performance Metrics Working Group within its work on multi-metric composition [20]. Furthermore, the IETF recently started working on the definition of monitoring architectures for RTP [4]. To the best of our knowledge, our proposal is original in proposing the adoption of a service-level approach that permits the runtime definition and deployment of application-specific metrics and in specifically targeting the vertical domain of VoIP quality evaluation.

The only research activity in the literature that is close to our proposal is the recent [8] paper, which partially addresses scalability issues by adopting a "per network-segment" approach. [8] proposes to divide the path between end users into two segments, each one with specific quality metrics. The first segment is the local access network of one user, the second segment consists of all the remaining part of the E2E path. However, [8] does not address the general issue of providing a reusable and extensible facility for service-level quality awareness and, most important, only supports network partitioning into two, statically pre-defined, segments. As better detailed in the following, we claim the need of a more granular segmentation approach in order to achieve a sufficiently accurate view of the network status for VoIP quality estimations, with positive effects also on solution scalability.

III. The PreQuEst Architecture

The PreQuEst project intends to demonstrate the feasibility of proactive solutions based on multi-segment network partitioning for the estimation of potential VoIP call quality. In particular, it aims at effectively and scalably achieving accurate estimations for all users' contacts (e.g., in a user's personalized buddy list) by relying on a VoIP-oriented monitoring system and on application-specific monitoring metrics. To practically exemplify how our approach can relevantly enhance users' experience, Fig. 1 depicts the case study of a quality-enriched presence service. The left column in the contact list contains the usual presence status information (online, offline, busy, . . .), while the right one contains the representation of the quality our solution proactively estimates for each contact.

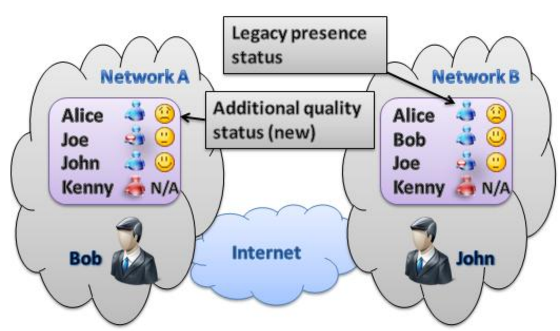

Fig. 1: Enriched Presence Service overview.

PreQuEst has been designed and implemented according to a layered architecture (Fig. 2). In short, starting from the

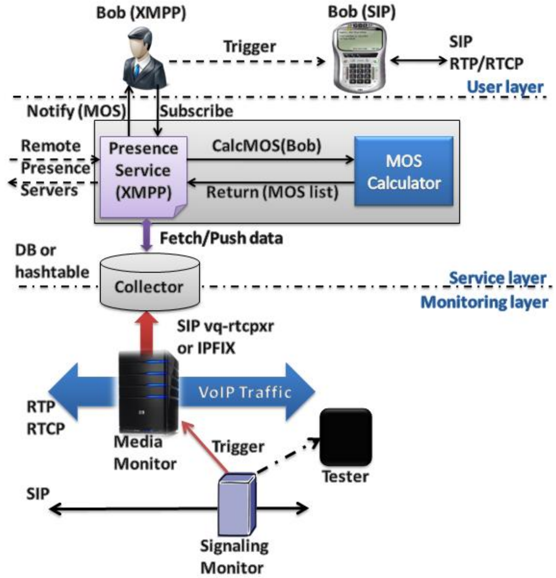

Fig. 2: Prototype architecture.

bottom layer, PreQuEst Monitor is responsible of detecting and monitoring VoIP media traffic. Monitoring data will be stored into the Collector, which is in charge of gathering them and cleanly separating the service layer from the implementation details of the underlying monitoring. Quality information can be fetched and properly used either by an application-specific quality facility (usual way, as in our quality-enriched presence service) or directly by final users/applications (cross-layer access). Several motivations justify this cross-layer possibility: for instance, we claim the usefulness of such a visibility to facilitate fully aware applications to automatically select a network access when multiple ones are available (multi-homing). Let us rapidly note that PreQuEst could be easily re-used also for non-VoIP-related applications, for instance in an IPTV scenario where each user could maintain a list of preferred video sources and a PreQuEst-based IPTV-specific quality facility could offer proactive channel evaluation. To this purpose, there is only the need to instantiate a different quality facility on top of our Collector, by extending/refining the VoIP quality facility presented in this paper.

Different applications, in fact, need different metrics, and our layered approach allows us to change/extend the underlying monitoring system to provide any kind of new metrics, and at the same time to change/extend the quality-related blocks at service layer to satisfy the needs of the applications running on top. The separation of layers also facilitates the runtime deployment of new components when and where needed.

A. *Scalable Monitoring through Network Partitioning*

VoIP traffic is basically divided in signaling and media. We assumed SIP [12] as signaling protocol and RTP/RTCP [13] as media protocol. In general, QoS management in VoIP covers different aspects, based on both kinds of traffic. Call quality primarily depends on media traffic and can be calculated from a minimum set of network-based parameters, namely delay, packet loss, and jitter. To gather the aforementioned parameters, one or multiple monitors should be activated along the media path to analyze the traversing traffic, and the placement of such monitors should provide the needed metrics in both a carrier and an enterprise scenario.

The entities involved in a VoIP scenario are users' terminals, proxy servers, and, optionally, media gateways. In this scenario, we could place monitors on either the terminals or the network infrastructure. The former solution would recall the Skype E2E topology. In this case, the only way to assess the quality between two users is to directly calculate it. Manifestly this solution cannot scale well. Placing the monitoring probes on the network infrastructure, instead, is more scalable but existing solutions (see section II) do not fit our targeted scenario. In fact, as already stated, the approaches in the literature do not offer enough granularity to fully re-use the segmented monitoring information; on the contrary, our idea is to partition the interested paths into 3 segments (see Fig. 3), thus taking full advantage of user locality and often re-using monitoring data to evaluate E2E quality without any injected traffic. In fact, in PreQuEst each user constantly has an associated local quality (sometimes thanks to active injection); by dynamically combining this users' quality with the one evaluated between interested localities, the overall E2E quality is assessed.

Suitable monitoring hosts could be Session Border Controllers (SBC) or Edge Nodes, which are normally placed at the edge of the operators' network. Their location allows the partitioning of the media path in the 3 aforementioned segments (see Fig. 3), thus enabling the monitoring of users' local network impairments (1 and 3) and inter-operator ones (2). These values are crucial for the service layer to have a fine-grained monitoring view, so to combine them in the appropriate way depending on application-level purposes.

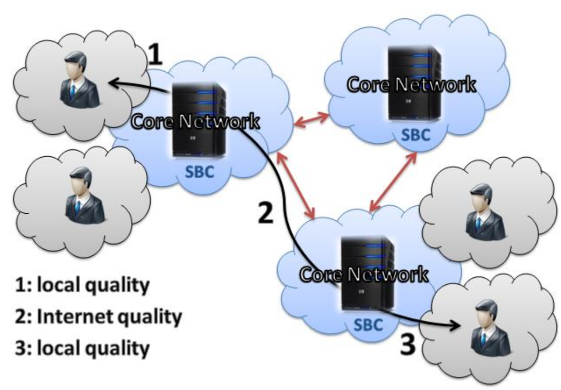

Fig. 3: Network partitioning.

3-segment partitioning permits to save a considerable amount of monitoring traffic if compared with an E2E approach. In fact, both solutions induce an overhead depending on the total number of users (N): in our case that overhead grows linearly with the number of users, while in the other case the growth is also quadratic with the average size of users' presence lists (x).

Such formulas are:

$$\frac{x \times (x-1)}{2} \times \frac{N}{x} \qquad \text{(E2E approach)}$$

$$N + \frac{N_{SBC} \times (N_{SBC}-1)}{2} \qquad \text{(partitioned approach)}$$

where N_{SBC} represents the number of operator networks (largely smaller than the number of users or the average presence list size). This simple scalability analysis will be corroborated experimentally later on in a real deployment scenario.

The raw monitoring data of each network segment is stored into PreQuEst Collectors (one for each domain - e.g., t-online.de). To collect monitoring data, Collectors could either exploit the IPFIX protocol or the SIP Event Package for Voice Quality Report [1]. Given the big overhead associated with the latter, our current prototype exploits it only when the protocol is already deployed and used by other applications. Otherwise, anytime there is the need to aggregate monitoring data for different users, our Collector exploits the non-session-oriented IPFIX solution.

However, the above partitioning-based passive monitoring is sometimes insufficient because PreQuEst should perform its evaluations also in time intervals when there are no active calls in the interested localities. Therefore, PreQuEst also actively injects synthetic network traffic to estimate quality parameters over non-utilized network segments, where and when needed. In order to achieve accurate quality estimations, a central issue is the proper choice of the most suitable kind of traffic to inject. It is well known that different protocols are treated in different ways, e.g., ICMP Ping could be filtered by firewalls and TCP is subject to Congestion Control. Given its VoIP-oriented nature, PreQuEst injects traffic by performing VoIP calls through a testing entity placed on SBCs. This way, the monitored traffic represents with a good approximation the kind of real traffic expected from the users.

PreQuEst active injection raises two primary issues: minimization of testing traffic and transparency for final users. To these aims, we take advantage of the SDP Media Loopback [2] and the SIP Answering Mode extensions [3]. Such extensions allow to perform the needed test calls looping back the synthetic traffic without requiring any user intervention.

B. Proactive VoIP Quality Estimation Service

The Collector monitoring data are used by our PreQuEst application-specific quality facility to calculate E2E relevant values, namely Mean Opinion Score (MOS [14]) for our VoIP scenario. The actual VoIP quality estimation is delegated to a separated component, possibly application-specific, called MOS Calculator. By delving into finer details and describing the process step by step: (1) MOS Calculator fetches user location and quality from the local Collector; (2) based on user's presence list, it asks other presence servers for the same information about the user's buddies; (3) it can retrieve the inter-operator metrics based on users' locations; and finally (4) it processes these data by using the E-Model [15] method to assess the E2E quality.

The computed MOS is delivered to final users via SIP messages, with a notification-based approach relying on the Publish/Subscribe paradigm. Event messages carry an "Event: presence" header and contain the MOS score into PIDF [17], which we have extended in our prototype to include this new value (e.g., <mos>4.1</mos>). The quality assessment method is based on the ITU E-Model, which returns a score *R* in a range between 0 and 100, where 100 represents perfect quality [16]. Such score is then converted into the commonly used scale of MOS values through the conversion formula recommended by ITU in [15].

Finally, let us note that PreQuEst has been designed to allow flexible deployment of new quality-related services. For example, it is possible to deploy an IPTV favorite channels monitoring system by adding a few components to the current PreQuEst prototype. E.g., by instructing the monitoring layer to monitor the needed information (i.e. the time when each user starts and ends watching a channel) and by replacing the quality assessment function that proactively estimates the channel quality using video quality measurement techniques. To perform the latter, system/service administrators can simply replace the default quality estimation agent by invoking PreQuEst with an explicit URI to indicate the different agent to be exploited (non-default option); PreQuEst adopts XML encoding to ask and return quality results to facilitate agent replacement and system extensibility.

IV. PreQuEst Implementation Insights and Preliminary Experimental Results

We have implemented and experimentally validated the PreQuEst prototype by carefully taking into consideration interoperability with existing standards and widespread solutions. In particular, the developed prototype has been deployed over standard off-the-shelf Linux machines acting as SBCs, with no specific hardware requirements. Such hosts were emulating different network impairments for the traffic using the Linux iproute2/netem functionalities. We used SIP as the signaling protocol, while the codecs for RTP media traffic were G.711 and G.729. Finally, we used Snom 360 IP Phones, which allowed us to visualize the presence service also through their embedded mini-browser. All the PreQuEst components (monitors, Collector, and VoIP-specific quality facility) have been implemented in the C language; the Sofia-SIP Library [21] was used to manage the SIP stack.

In the current implementation, the PreQuEst monitor receives a copy of SIP messages from the OpenSER SIP Server [22]. Using these data, it can detect the set of users currently connected in its area and notify them to the Tester component, which will take care of proactively verifying users' presence and of periodically scheduling test calls. Given that off-the-shelf equipment does not support yet the SDP Media

Loopback and SIP Answer Mode extensions, in the validation tests reported in the following, active testing calls were replaced by a simple set of 100, periodically repeated, ping requests. The monitoring of ongoing calls is triggered by SIP INVITE/200 OK (start) and BYE (end). To this purpose, the pcap library is used both to retrieve the negotiated RTP/RTCP ports and to instruct the monitor to collect the data to send to the Collector. The Collector accesses monitored values by using the SIP vq-rtcpxr package and the Publish method. It receives per-segment monitored data via the RTCP-XR extensions (RTCP-XR [23] natively supports only E2E reports; we originally extended it with segmented quality values).

The operations needed to compute MOS quality estimations are detailed in [24] and [25]. [24] was chosen by following the ITU default values and by considering random packet loss [26] and no playout discarding.

Finally, from the point of view of the integration of the presence service (e.g., XMPP or SIMPLE) with our VoIP quality facility, two different interfaces showing the MOS score were implemented. The first one offers a Web-based interface with a presence list including quality values; the other shows the same presence list formatted according to the text-based display of the Snom360 IP Phones [18]. The presence list content is sent to each interested phone through a SIP Notify message containing an XML body that is reproduced on the Snom360 mini-browser (see Fig. 4).

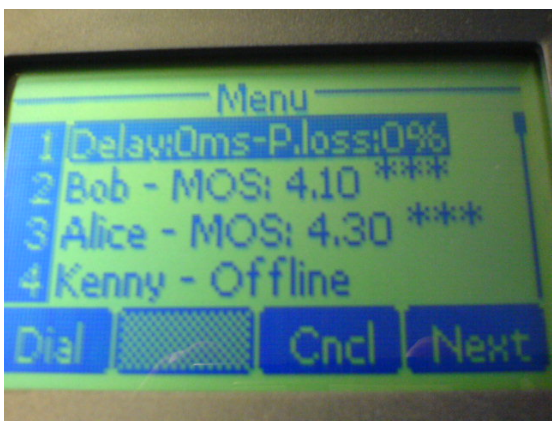

Fig. 4: Snom360 IP Phone display.

The implemented prototype has been tested over our lab testbed, consisting of two separated LANs (domains). A host, placed at the edge of each domain, is used as SIP proxy, media gateway, presence server, and monitor. Another host for each domain emulates network impairments via iproute2/netem, and finally a switch gathers the Snom phones connected to the same domain (see Fig. 5), which can be moved to the remote LAN (roaming).

The extensive validation of the PreQuEst prototype over wide-scale real deployment environments is currently undergoing. Anyway, the preliminary experimental results we have already collected are encouraging and corroborate the claim of effectiveness and scalability of our network partitioning

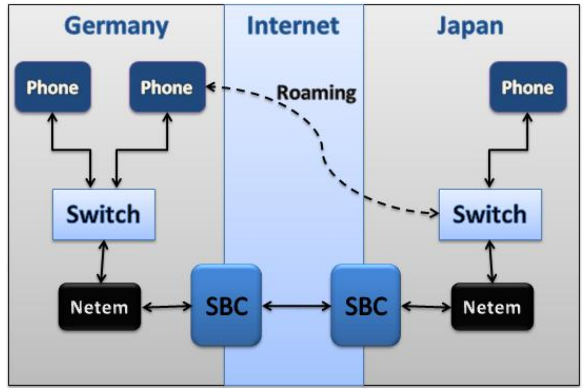

Fig. 5: Testbed overview.

approach. First of all, about the accuracy of estimations, we have verified that the quality values visualized on the presence lists always coincided with the actual voice quality experienced during the calls, independently of the impairments emulated by netem during tests. The only few discrepancies corresponded to short time intervals of high dynamicity, i.e., when some network parameters had abrupt variations due to emulated traffic jam or temporary bursts of VoIP call requests. In addition, the periodical injection of few seconds of RTP traffic, even with a loose interval of 30 minutes, has demonstrated to be sufficient to accurately estimate call quality in usual working conditions (emulated traffic and network parameters corresponding to statistically common conditions for an operator network [17]).

In addition, we have experimentally verified that the PreQuEst active injection of synthetic traffic does not significantly affect network performance. That is mainly due to the fact that traffic injection occurs only in localized network segments and when these segments are particularly under-loaded. The number of test calls, in fact, decreases during peak hours, when ongoing calls provide enough data through passive monitoring, and grows during low hours (e.g., during the night). Our testbed has shown that in the worst case of no ongoing calls over a network segment, PreQuEst tends to produce a maximum number of test calls per period that is proportional to the number of users. Although it is certainly a non-negligible network traffic, the calls have to be performed only for few seconds and the load is distributed between all the involved operators.

To give a quantitative estimation of this traffic, let us consider a carrier example using formulas from section III-A. The case of Germany as a federation of operators adopting PreQuEst was considered (approximately 50 millions active phone lines, each one potentially with a SIP account). Such lines are managed, directly or indirectly, by 6 big operators. With an assumed average presence list size of 50 buddies, and a test interval of 30 minutes, the needed test calls per minute for the E2E approach would be $\sim 41M$ against the $\sim 1.6M$ for the partitioned approach, one order of magnitude bigger. Furthermore, these numbers represent the worst case, when no ongoing call is providing passive monitored data, and passive monitoring relevantly decreases network overhead due to our original partitioning approach.

V. Conclusions, Lessons Learned, and Future Work

Given the expected relevance of VoIP-related communications in the years to come, the possibility to proactively estimate E2E VoIP quality between any potential pair of users is a crucial facility for operators and final users. The PreQuEst project has demonstrated the feasibility of achieving accurate and effective quality estimations based on careful network partitioning and synergic exploitation of passive monitoring and active injection of synthetic traffic. In addition, our application-level approach has allowed to design a general quality estimation framework, which can be easily exploited to enhance different kinds of off-the-shelf applications, from Presence Service to IPTV, with quality estimations that possibly depend on application-specific metrics.

During our design/implementation activities we identified some lacks in the existing related standards: RTCP XR and the associated SIP vq-rtcpxr package do not consider partial quality values but only E2E; the PIDF document, containing user personal information, does not include currently quality-related fields; SIP Answer Mode and SDP Media Loopback extensions are still immature from the point of view of security support. In addition, we experienced several inadequacies of libpcap for our real-time monitoring purposes, especially because of unreliability (e.g., we experienced a number of ignored packets shortly after the beginning of captures, affecting packet loss calculations), thus we started developing a different monitoring library from scratch.

The results achieved by our first PreQuEst prototype are encouraging our further research activities. In particular, we are currently working on continuing to simplify the integration of PreQuEst with off-the-shelf Presence Service clients by proposing an extension to the standard PIDF+XML data format [17]. In addition, we are extensively validating the accuracy and scalability of our prototype by comparing its experimental performance with alternative solutions in the literature, over real and wide-scale deployment environments. Finally, we are extending our monitoring layer with specific mechanisms for IPTV traffic and for locality-based data aggregation.

References

[1] Pendleton, A., Johnston, A., Sinnreich, H., Clark, A.: Session Initiation Protocol Package for Voice Quality Reporting Event. IETF Internet Draft draft-ietf-sipping-rtcp-summary (work in progress). (2007)

[2] Hedayat, K., Jones, P., Roychowdhury, A., SivaChelvan, C., Stratton, N.: An Extension to the Session Description Protocol (SDP) for Media Loopback. IETF Internet Draft draft-ietf-mmusic-media-loopback (work in progress). (2008)

[3] Willis, D., Allen, A.: Requesting Answering Modes for the Session Initiation Protocol (SIP). IETF Internet Draft draft-ietf-sip-answermode (work in progress). (2008)

[4] Hunt, G., Arden, P.: Monitoring Architectures for RTP. IETF Internet Draft draft-hunt-avt-monarch (work in progress). (2008)

[5] Skype, http://www.skype.com

[6] Brix Networks, http://www.testyourvoip.com

[7] Sylor, M., Venna, N., Ripps, H.: Voice Quality on the Internet in 2005 as Measured by www.TestYourVoIP.com. In: 17th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management. Dublin (2006)

[8] Yamada, H., Fukumoto, N., Isomura, M., Uemura, S., Hayashi, M.: A QoE based service control scheme for RACF in IP-based FMC networks. In: The 9th IEEE International Conference on E-Commerce Technology and The 4th IEEE International Conference on Enterprise Computing, E-Commerce and E-Services (CEC-EEE 2007), pp. 611–618. IEEE Computer Society, Tokyo (2007)

[9] Conway, A.E.: A passive method for monitoring voice-over-IP call quality with ITU-T objective speech quality measurement methods. In: IEEE International Conference on Communications, 2002. ICC 2002, pp. 2583–2586. New York (2002)

[10] Agrawal, S., Ramamirtham, J., Rastogi, R.: Design of active and passive probes for VoIP service quality monitoring. In: 12th International Telecommunications Network Strategy and Planning Symposium, 2006. NETWORKS 2006. New Delhi (2006)

[11] Cisco Systems: Measuring Delay, Jitter, and Packet Loss with Cisco IOS SAA and RTTMON. http://www.cisco.com/en/US/tech/tk869/tk769/technologies_white_paper09186a00801b1a1e.shtml

[12] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., Schooler, E.: SIP: Session Initiation Protocol. IETF RFC3261. (2002)

[13] Schulzrinne, H., Casner, S., Frederick, R., Jacobson, V.: RTP: A Transport Protocol for Real-Time Applications. IETF RFC3261. (2003)

[14] ITU: Recommendation P.800. Methods for subjective determination of transmission quality. (1996)

[15] ITU: Recommendation G.107. The E-model, a computational model for use in transmission planning. (2005)

[16] ITU: Recommendation G.109. Definition of categories of speech transmission quality. (1999)

[17] Sugano, H., Fujimoto, S., Klyne, G., Bateman, A., Carr, W., Peterson, J.: Presence Information Data Format (PIDF). IETF RFC3863. (2004)

[18] Snom Technology, http://www.snom.com

[19] INTERMON IST project, http://www.ist-intermon.org

[20] IETF IPPM WG, http://www.ietf.org/html.charters/ippm-charter.html

[21] Nokia Research Center, http://opensource.nokia.com/projects/sofia-sip/

[22] OpenSER SIP Server, http://www.openser.org

[23] Friedman, T., Caceres, R., Clark, A.: RTP Control Protocol Extended Reports (RTCP XR). IETF RFC3611. (2003)

[24] Perlicki, K.: Simple analysis of the impact of packet loss and delay on voice transmission quality. In: Journal of Telecommunications and Information Technology. Nr. 2, pp. 53–56. (2002)

[25] Balan, H. V., Eggert, L., Niccolini, S., Brunner, M.: An Experimental Evaluation of Voice Quality over the Datagram Congestion Control Protocol. In: INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE. (2007)

[26] ITU: Recommendation G.113. Transmission impairments due to speech processing. (2001)

[27] Dischinger, M., Gummadi, K. P., Haeberlen, A., Saroiu, S.: Characterizing Residential Broadband Networks. In: Proceedings of the ACM IMC07, 2007. San Diego (2007)