Improving Mobile VoIP Quality Through Bandwidth Optimisation

Rafael Dantas
Faculty of Science and Engineering
University of Limerick
Limerick, Ireland
Email: rafael.dantas@ul.ie

Dr. Chris Exton
Faculty of Science and Engineering
University of Limerick
Limerick, Ireland
Email: chris.exton@ul.ie

Dr. Andrew Le Gear
One Horizon Group
Nexus Innovation Centre
University of Limerick
Limerick, Ireland
Email: andrew.legear@horizon-globex.ie

***Abstract*—Maintaining consistent VoIP quality is a challenging task, especially where it is carried using a mobile internet connection. With greater than 50% of the world's mobile user population using older GSM incarnations, this is very much a present technical research challenge. This paper presents an approach to improving mobile VoIP telephony by drastically lowering the bandwidth consumption while maintaining acceptable call quality when compared to competitor solutions. Although initial test had good results, we've found that a detailed experiment might be necessary to benefits of the implementation.**

***Index Terms*—VoIP, Mobile, Bandwidth, Consumption, Network**

I. Introduction

As the number of smart mobile phones grows so too does the stress on existing network infrastructure. While old hardware can be upgraded and new equipment added, it can quickly become an expensive task to replace every device in order to provide proper 3g or faster access to the internet. With the majority of the world's mobile user population continuing to use 2g GSM networks[1], network traffic optimisations can still provide benefits for mobile VoIP applications where modern GSM infrastructures are not available.

The Horizon Global Exchange is a complete solution for ultra-low bandwidth mobile *Voice over Internet Protocol (VoIP)* communications developed by the One Horizon Group[2]. This solution includes a mobile application for Android and IOS systems, node servers that handle the signalling for both ends of the call and VoIP servers for handling the call itself.

This research will deliver the following contributions:

- A base comparison of the Horizon Global Exchange against competing solutions which will provide enough data to test the following assertions:
  - Assertion 1: The Horizon Global Exchange uses less bandwidth compared to other mobile VoIP solutions.
  - Assertion 2: The Horizon Global Exchange can deliver a stable and reliable connection between the call participants even under poor network conditions, where other solutions can not.
  - Assertion 3: These features can be achieved without great sacrifices to the perceived quality of the call.
- Further optimisation of both codec and protocol, to further improve bandwidth consumption and the overall stability of the connection.

II. Related Work

Performance and quality comparisons between applications are commonplace in mobile VoIP literature.

Often the objective of these tests is to evaluate the performance of existing technologies to select the best solution for the product under development. For example, [3] performed such an analysis for selecting the best codec for the Inmarsat mini-M system. Similarly, [4] also performed comparisons for the S3C241O micro-controller.

The objective may also be to test the performance of a solution against the market standards and competitor solutions. To decrease the degradation of quality on VoIP systems caused by packet loss, [5] proposed a new AMR codec and tested it against the usual codecs used with the H.323 protocol. A performance evaluation of Skype was made by [6] comparing it to industry standards. To improve bandwidth utilisation on mobile networks, [7] tested his TCP implementation with four alternatives and found significant improvements.

Lastly, quality tests may also be used to propose new ideas. [8] used *Mean Opinion Scoring (MOS)* [9] to test if subjects could distinguish between many narrowband and wideband codecs, and found a correlation between higher scores and higher bit rates.

III. The Horizon Global Exchange

The Horizon Global Exchange is a complete mobile VoIP network solution including PSTN breakout capabilities. It can achieve up to 10 times the efficiency of the competing applications by employing the following optimisations.

Figure 1 shows the basic workings of a normal VoIP call using Horizon Global Exchange.

First the application contacts a central node, which will try to find the other user inside the network. The central node will then assign a VoIP server for handling the call, sending this information to both participants. Both ends will try to connect with the VoIP server. This is done to avoid any NAT or proxy problems that would arise otherwise. Once connected, the VoIP

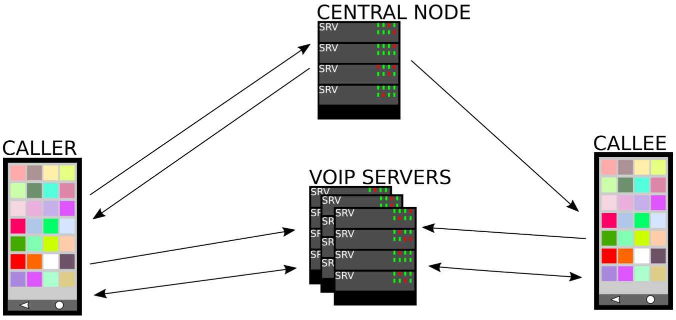

Fig. 1. Typical workflow of a Horizon Global Exchange call.

server becomes responsible for relaying the call between the participants.

If one or both the participants are not able to connect to the VoIP server, it will transcode the disconnected end to a SIP server, which will call this user using the normal telephony infrastructure.

A. Header elimination

Most protocols and codecs carry some metadata alongside the payload. This header information is used by the receiver to understand and access the information contained inside, especially in cases where the configuration parameters can be changed for each individual packet.

As an example, SPEEX parameters can be quite complex; as detailed by [10].

If both sides of the communication agree over a configuration parameters of the payload the header section that carries this information becomes redundant. It is then possible to reduce header size, and ultimately remove it altogether, by standardising every parameter of the transmission.

The Horizon Global Exchange uses only one codec with a very specific set of parameters, making payload headers completely obsolete and allowing every packet to carry only the actual audio signal.

B. Header-to-payload ratio reduction

Even if the payload does not carry a single bit of header information, every single network protocol will envelop all data produced by the layer above and append its own headers to it. Since those protocols are usually outside the scope of the application, the changes proposed in III-A cannot be applied here.

In this situation, Nagle's algorithm becomes an option to further reduce header-to-payload ratios.

C. Silence detection

In a normal conversation, usually only one of the participants is speaking while the other is only listening to what is being said. If an application is not aware of this special case, it will always record and send audio information both ways.

Using silence detection algorithms, it is possible to avoid sending useless data through the wire.

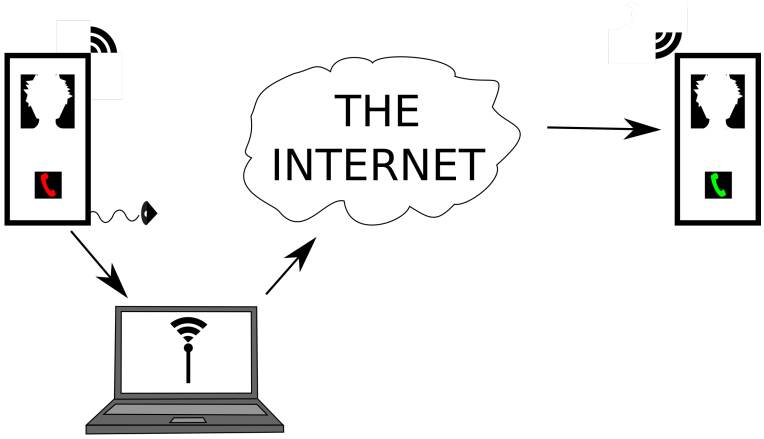

Fig. 2. Representation of Experiment 1.

The Horizon Global Exchange has a special silence packet, which is sent by the application whenever it detects its user has gone silent. This packet also doubles as a keep-alive message, being sent periodically by the silent side in order to maintain the connection open.

IV. Evaluation

In this section, all applications selected will be tested against the assertions mentioned on section I. The following subsections will detail the experimental design, tools used and the results obtained by the tests realised.

A. Experimental Design

The first experiment, represented in figure 2, will provide data for a quantitative analysis over network data collected.

One cell phone will be connected to a laptop acting as an eases point. This laptop will be able to collect any packet passing trough the network while transparently forwarding them to their proper destinations. It will also be able to replicate real-world networking issues, by artificially increasing the number of packets lost between the mobile device and the network layer.

One phone will be used to call the second. A minute-long audio file was played into the first cell phone microphone, while the second phone remained silent for the entire call.

If the application uses silence detection to reduce data consumption, it would reflect on the total ammonia of data exchanged through the network, as the silent side of the call would not send as much data as the busy side.

With all incoming and outgoing packets being recorded on the devices, it is possible to analyse the first and second assertions discussed in section I.

The second experiment is a normal call but with actual users holding a minute-long conversation using each application selected for this test. At the end of each call, both users are going to evaluate the perceived quality using a MOS scale between 1 (poor) and 9 (excellent). The objective of this experiment is to test the third assertion of section I.

Application	Wire throughput (kb/s)	Payload throughput (kb/s)	Packet rate (p/s)	Frame size (bytes)	Payload size (bytes)	Payload-to-header ratio
Horizon (UDP)	12.76	9.06	13.53	121	86	2.45
WhatsApp (UDP)	35.75	27.15	32.26	142	108	3.15
Viber (TCP)	56.52	21.77	67.30	107	41	0.63
Facebook (TCP)	67.10	47.31	38.30	224	158	2.39
WeChat (TCP)	86.68	54.14	63.09	176	110	1.66
Hangouts (TCP)	90.65	51.00	76.68	151	85	1.29
QQ (TCP)	106.88	52.91	101.88	134	66	0.98
Skype (UDP)	113.37	86.69	100.44	144	110	3.25

TABLE I
COMPARISON BETWEEN VARIOUS APPLICATIONS PERFORMANCES

Laptop	Lenovo Thinkpad T420
Operating system	Debian GNU/Linux (stretch)
Kernel version	4.6.0-1-amd64
Cell phone 1	Motorolla Moto G3 (Android 6.0)
Cell phone 2	Sony-Ericsson Xperia (Android 6.0)

TABLE II
EQUIPMENT USED ON THE EXPERIMENTS

B. *Tool Support*

The audio sample was reproduced using mplayer[11], a command-line interface player which was called by the script before changing the network properties.

Tcpdump[12] was used for gathering network packets and storing them on disk.

For data analysis, tshark[13] and wireshark[14] were selected for filtering the results, in order to remove unrelated packets made by other applications.

Two bash scripts were written:

- A coordinator script used to play the audio sample and change the network properties, while making sure all data was being recorded into separated files.
- A filter script which called tshark using the correct parameters for each application, in order to improve the output of Tcpdump for analysis.

A python script was also written for producing graphical representations of the data.

C. *Quantitative Analysis: Bandwidth Usage*

Table I has a comparison between all applications analysed as part of our evaluation. Wire throughput is an average data rate going through the network, including the headers for the physical, datalink, network and transport layers. Payload throughput is also an average data rate, including only the data carried by the transport layer, without its headers. Packet rate is the average number of packets sent per second. Frame size is an average of each individual packet size, including the entire network overhead while payload size removes the same overhead sizes. Finally, payload-to-header ratios are a simple ratio between the payload size and the average network overhead size (frame size - payload size).

The results are striking - using less than half bandwidth than the second closest solution, the Horizon Global Exchange had a significant network performance lead in this experiment.

V. FUTURE WORK

The next step in this research is to test the assertions 2 and 3 from section I, which are going to require different experiments than the one described in the previous section. After that, we intend to analyse ways to improve the current protocol used by the Horizon Global Exchange. The following changes are going to be tested next:

- A dictionary that will store common patterns produced by the codec to further compress the information being transmitted,
- A permanent personal audio profile to improve codec compression when talking with the same person more than once,

VI. ACKNOWLEDGEMENTS

This work was supported with the financial support of the Science Foundation Ireland grant 13/RC/2094 and co-funded under the European Regional Development Fund through the Southern & Eastern Regional Operational Programme to Lero - the Irish Software Research Centre (www.lero.ie). (SOW2016-034)

This work was possible thanks to Brian Collins, CEO of the One Horizon Group, Inc., for permitting technical analysis of the product.

REFERENCES

[1] Cisco, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2015–2020 White Paper," Tech. Rep., Feb. 2016. [Online]. Available: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html.

[2] B. Collins and C. Dziedzic, "Method and devices for routing in a satellite-based communication system," Patent US 9137729, Sep. 2015. [Online]. Available: https://www.lens.org/lens/patent/US_9137729_B2.

[3] S. Dimolitsas, "Evaluation of voice codec performance for the Inmarsat mini-M system," in *10th International Conference on Digital Satellite Communications (ICDSC-10)*, vol. 1995, IEE, 1995, pp. 101–105, ISBN: 0 85296 635 0. DOI: 10.1049/cp:19950015. [Online]. Available: http://digital-library.theiet.org/content/conferences/10.1049/cp%7B%5C_%7D19950015.

[4] J. Zhou, T. Wu, and J. Leng, "Research on voice codec algorithms of SIP phone based on embedded system," *Proceedings - 2010 IEEE International Conference on Wireless Communications, Networking and Information Security, WCNIS 2010*, pp. 183–187, 2010. DOI: 10.1109/WCINS.2010.5541916.

[5] J. W. Seo, S. J. Woo, and K. S. Bae, “Study on the application of an AMR speech codec to VoIP,” *2001 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings Cat No01CH37221*, vol. 3, pp. 1373–1376, 2001, ISSN: 15206149. DOI: 10.1109/ICASSP.2001.941184. [Online]. Available: http://ieeexplore.ieee.org/document/941184/%20http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=941184.

[6] K. Kim and Y.-J. Choi, “Performance comparison of various VoIP codecs in wireless environments,” *Proceedings of the 5th International Conference on Ubiquitous Information Management and Communication*, 89:1–89:10, 2011. DOI: http://doi.acm.org/10.1145/1968613.1968718. [Online]. Available: http://doi.acm.org/10.1145/1968613.1968718.

[7] K. Liu and J. Lee, “On Improving TCP Performance Over Mobile Data Networks,” *IEEE Transactions on Mobile Computing*, vol. 1233, no. c, pp. 1–1, 2015, ISSN: 1536-1233. DOI: 10.1109/TMC.2015.2500227. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7328303.

[8] A. Ramo, “Voice quality evaluation of various codecs,” *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 4662–4665, 2010, ISSN: 1520-6149. DOI: 10.1109/ICASSP.2010.5495201. [Online]. Available: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=5495201.

[9] R. C. Streijl, S. Winkler, and D. S. Hands, “Mean opinion score (MOS) revisited: methods and applications, limitations and alternatives,” *Multimedia Systems*, vol. 22, no. 2, pp. 213–227, Mar. 2016, ISSN: 0942-4962. DOI: 10.1007/s00530-014-0446-1. [Online]. Available: http://dx.doi.org/10.1007/s00530-014-0446-1%20http://link.springer.com/10.1007/s00530-014-0446-1.

[10] E. Touloupis, A. Meliones, and S. Apostolacos, “Implementation and evaluation of a voice codec for ZigBee,” *Proceedings - IEEE Symposium on Computers and Communications*, pp. 341–347, 2011, ISSN: 15301346. DOI: 10.1109/ISCC.2011.5983861.

[11] MPlayer Team, *Mplayer (v1.3.0/debian)*, Feb. 2016. [Online]. Available: http://mplayerhq.hu/.

[12] Tcpdump Team, Apr. 2015. [Online]. Available: http://www.tcpdump.org.

[13] Combs, G. and Wireshark Team, Oct. 2016. [Online]. Available: http://www.wireshark.org.

[14] ——, Oct. 2016. [Online]. Available: http://www.wireshark.org.