

Real-time Enhancement of VoIP Media Quality using SAVI SDI

Lilin Zhang, Jieyu Lin, Hadi Bannazadeh and Alberto Leon-Garcia

ECE Department, University of Toronto



Abstract

Traditional IMS/VoIP services lack real-time network provisioning and/or reconfiguration to improve the VoIP quality. In this work, we use SAVI SDI to prototype a novel method of enhancing VoIP media quality by using its converged and centralized manner of resource management. We will demonstrate that with a reconfiguration operation from the SDI platform, the media quality of a bandwidth thirsty HD video call can be elevated to a higher level in real time to achieve visibly significant quality improvement compared to a call made over best effort connection.

Introduction

VoIP service is not tolerant of network packet loss, and also requires large end-to-end bandwidth for HD media flows between end users. The problem gets worse quickly with a number of VoIP sessions, or file transfer, or Video streaming simultaneously on-going in the access network. As the access network resources are soon used up, it leads to longer queues/packet drops along the congested area, which significantly degrades the media quality experience of the end users affected.

Traditional IP networks do not support a fast re-configuration to resolve such congestion-lead VoIP media degradation in a real-time manner. OpenFlow technology provides a solution:

1. Forwarding logic \rightarrow Controller;
2. Forwarding processing \rightarrow OpenFlow-enabled switches;
3. Controller \longleftrightarrow OF-enabled switches (in OpenFlow protocol).

SAVI testbed has provided the application platform for the VoIP service. It has:

- created the compute and storage pool by leveraging opensource virtualization technologies;
- also expanded with the *virtual network functions* to configure/reconfigure the OpenFlow-enabled networks.

In this work, we demonstrate the capability of SAVI SDI platform that with one network path reconfiguration from the platform, the media flow of an on-going video call session can be re-routed through a less congested path, and the end users are able to see visibly significant quality improvement in media fluidity of the call session.

Determinants of Media Quality

The media quality of VoIP call sessions is determined by:

1. Video codecs \rightarrow efficiency, bitrate \rightarrow lossiness, image capture size \rightarrow pixels, and frame rate \rightarrow fluidity.
2. Network condition - In particular, VoIP is not tolerant of large packet loss. See more in Table 1
3. User client - as the real-time video/audio are captured, compressed, transcoded and transmitted on the fly, the more capable the user client, the better the media quality.

Fig. 1 shows the side-by-side snapshots of the video call encoded in low bitrate and the one encoded in high bitrate.

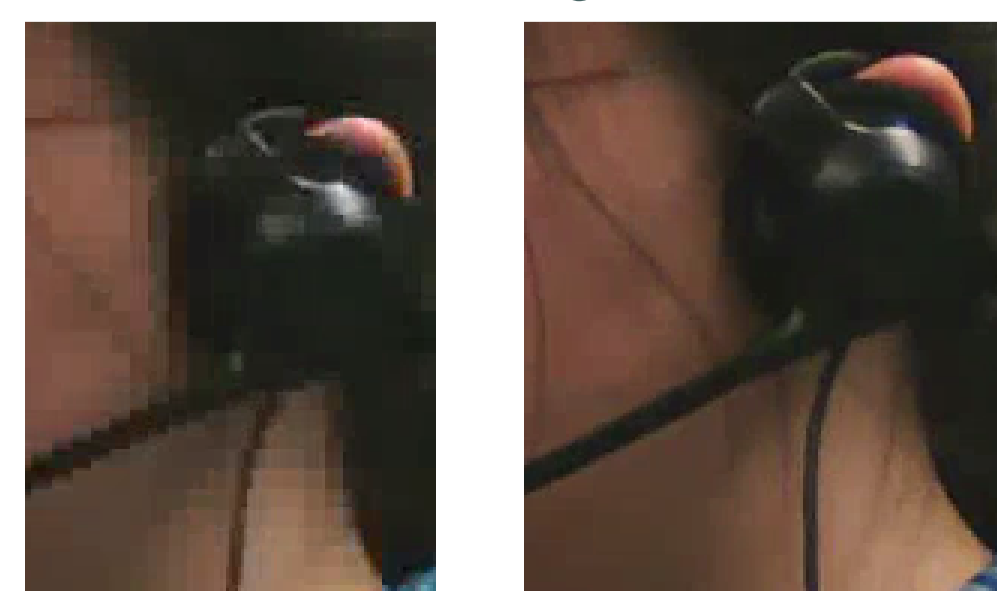


Figure 1: Bitrate = 297Kbps (left) Vs. Bitrate = 2Mbps (right), both captured by Logitech c930e 1080p webcam ("p" means progressive scanning, "1080p" means that the webcam is capable of scanning up to 1600 x 1080 pixels every $\frac{1}{60}$ second)

Table 1 highlights the network requirements for VoIP services. Generally for audio codec G.711, 1% packet loss can "significantly degrade" a VoIP call. Other codecs can tolerate even less packet loss. For instance, the default G.729 codec requires packet loss far less than 1% to avoid audible errors. Video codec has similar stringent requirements in low packet loss.

QoS	Requirement
Latency	≤ 150 ms one-way [1]
Packet loss	$\leq 1\%$
Bandwidth	\geq datarate

Table 1: VoIP QoS Requirements

Table 2 lists the datarates of video calls of various resolution sizes, that as the image capture size increases, the average datarate also increases. As pointed out in Table 1, the more pixels were captured by the webcam, the more bandwidth will be required for the end-to-end media flow.

Size	Avg. datarate (1.0e + 07)	σ (1.0e + 06)
176 x 144	0.0340 bps	0.0090
352 x 288	0.1086 bps	0.0219
640 x 480	0.2917 bps	0.0442
1280 x 720	0.4824 bps	0.1608
176 x 144	0.5955 bps	0.1055
352 x 288	1.6173 bps	0.6679
640 x 480	2.7497 bps	0.9270
1280 x 720	2.5734 bps	2.0420

Table 2: Various datarate for different capture sizes encoded in 297 Kbps (top) and encoded in 2 Mbps (bottom)

IMS Network Embedding

The IMS network consists of two parts:

- **Signalling plane** - The flows of SIP [2] messages sent between UE \longleftrightarrow IMS domain;
- **Media plane** - The flows of actual audio/video RTP packets sent between UE \longleftrightarrow UE.

In Fig. 2, the logical IMS network (top) is mapped to the physical switches and routers in the access network (bottom).

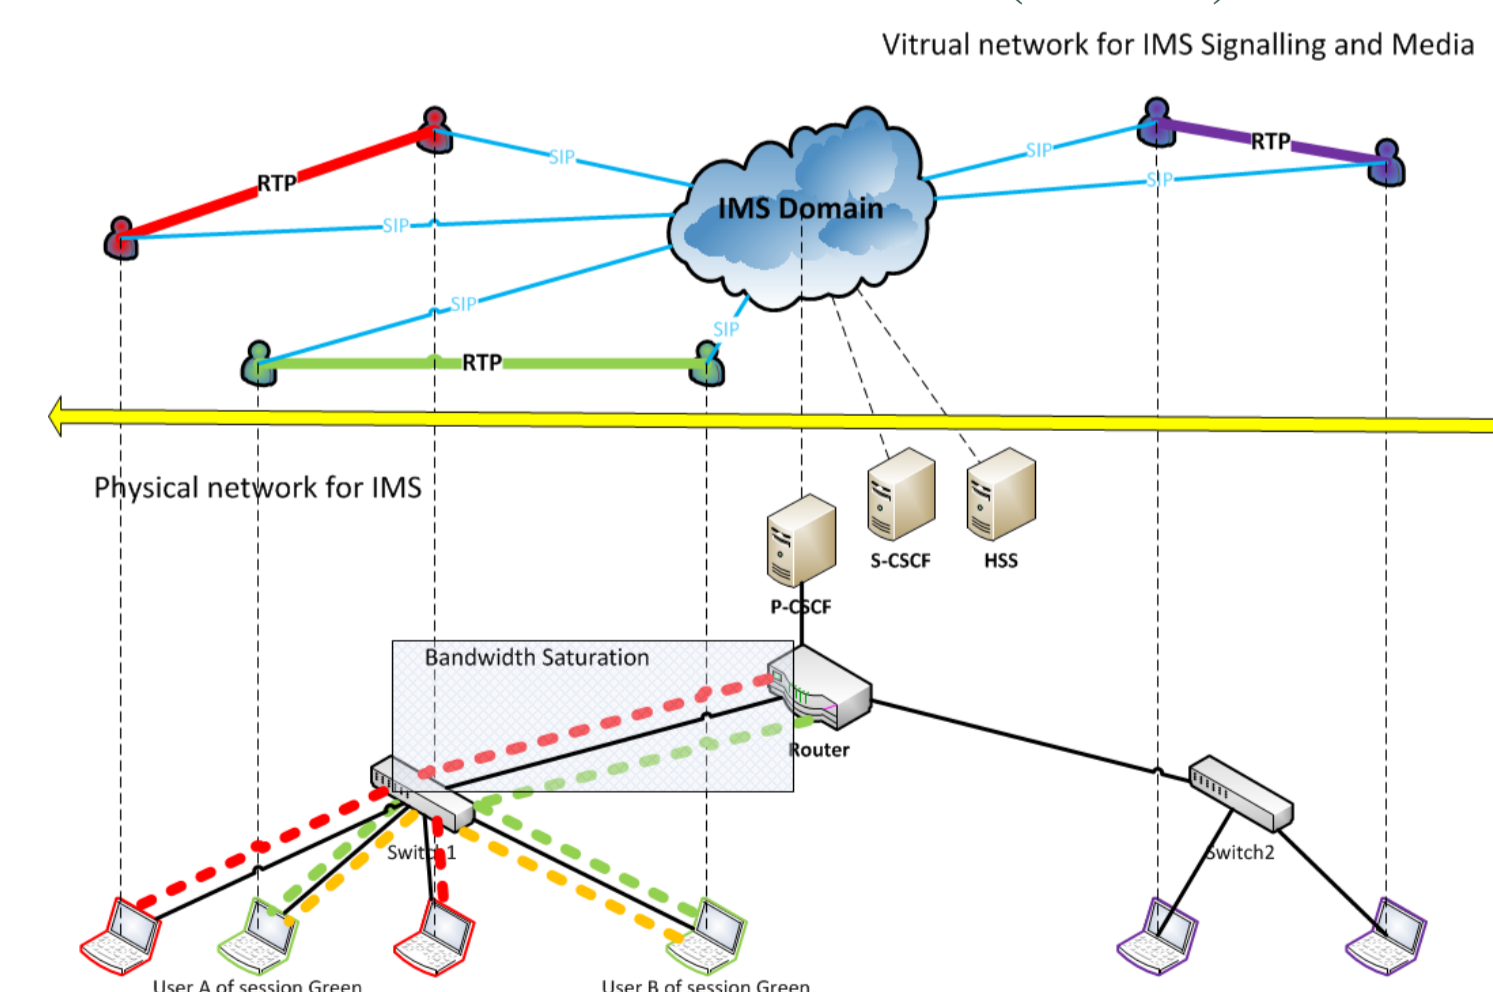


Figure 2: Virtual Network Mapping for IMS service

The **default forwarding path** of the GREEN session (the green dash line) leads to potential *network congestion* as it uses the shared link with the RED session, and consequently downgrades the media quality for both RED session and GREEN session. Its **alternative forwarding path** (orange dash line) could enhance the media quality for both call sessions.

VoIP Service on SAVI Testbed

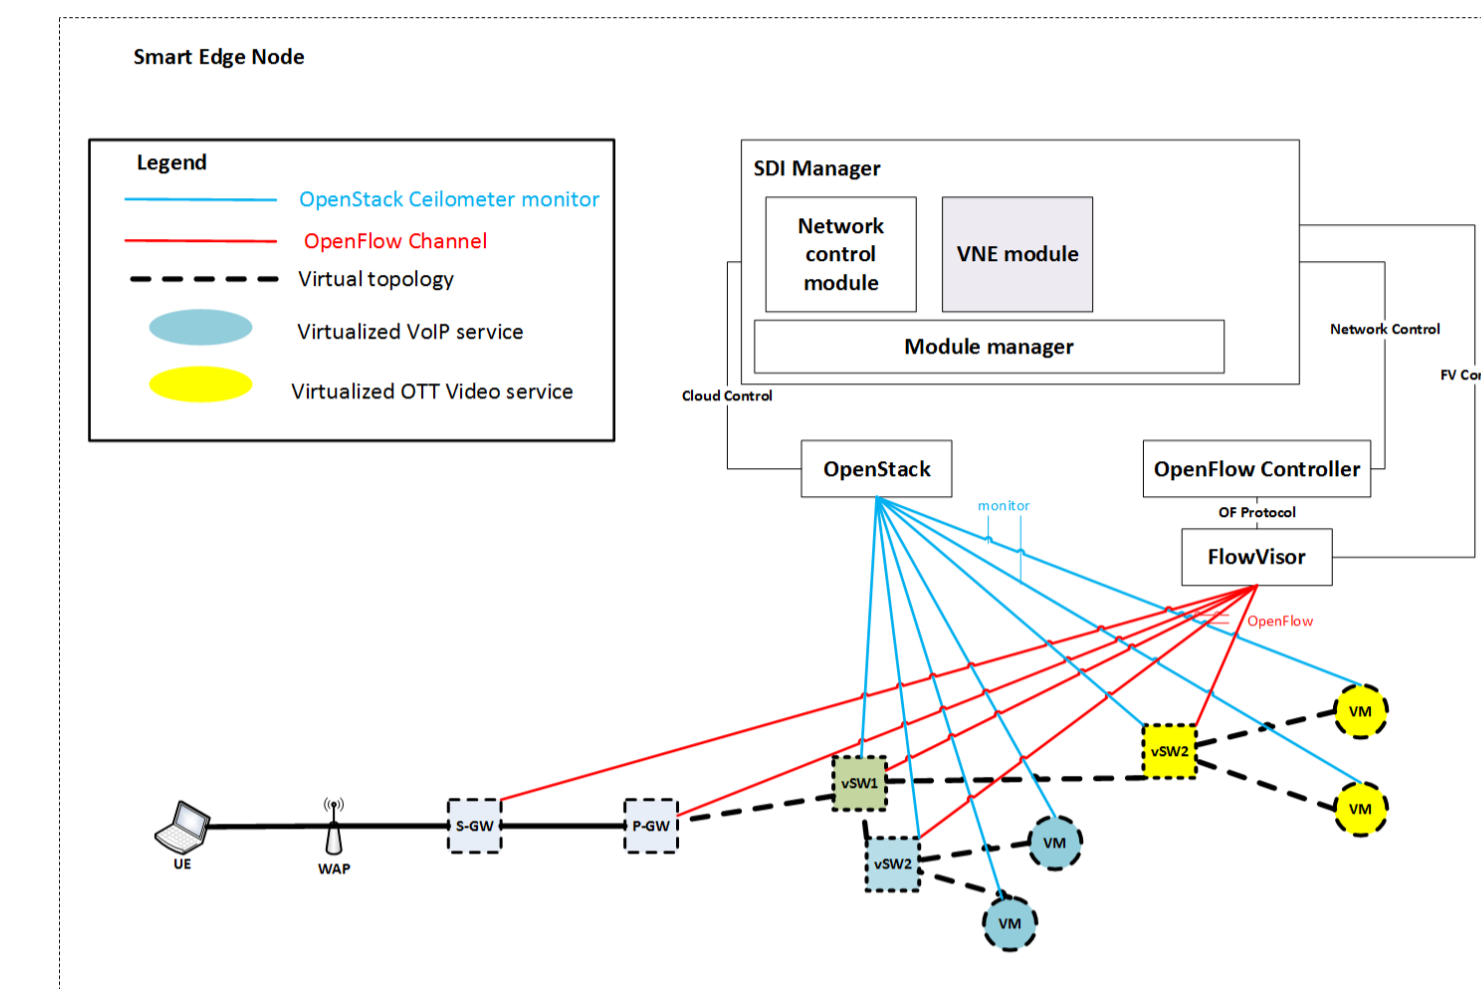


Figure 3: Virtualizing Multimedia Services in SAVI testbed

SAVI testbed leverages OpenStack to provide compute and storage pool and measurement for applications:

- It monitors the usage statistics of each physical and virtual servers and switches (blue lines in Fig. 3).

This feature could be made use of to locate the "hotspot" in the network.

SAVI testbed has converged the configuration of network resources on to the platform as well:

- The centralized controller communicates in OpenFlow protocol (red lines in Fig. 3) to each and every switch, including the soft switches and physical Pronto OpenFlow switches.

Enhancement Achieved using SAVI SDI

The topology of our demo is shown in Fig. 2. We use a pair of traffic nodes (generating up to 928Mbps), to simulate bandwidth-hungry applications, such as simultaneous video call sessions, or file transfers, or video streaming.

We use two computers (belonging to different subnets) to simulate the GREEN session, making a point-to-point video call. The visible difference brought by the SDI enhancement is the **fluidity** of the video call session - the media freezes (using the default path) and media restores (using the alternative path) in presence of the traffic congestion brought out by the RED session.

The media flow is uniquely identified:

- < 17 (protocol type for UDP), ip.src, ip.dst, eth.src, eth.dst

To enhance the media quality - bypassing the congested area, the centralized Controller issues the OpenFlow actions to Switch 1 (*manually specified*):

- the flow identified as above are forwarded between two ports of Switch 1, without going through the Router.

We also sniff RTP traces between the two computers in four scenarios to measure the quantifiable enhancement of the end-to-end RTP packet loss rate:

- NC(green) - the video call media flows through the default path, without congestion in the access network;
- NC(orange) - the video call media flows through the alternative path, without congestion in the access network;
- C(green) - the video call media flows through the default path, with congestion in the access network;

- C(orange) - the video call media flows through the alternative path, with congestion in the access network.

Table 3 summarizes one-way loss statistics of the end-to-end path: the SAVI SDI enhancement is able to decrease the packet loss rate from 1.3% to 0.0%, consistent with the media fluidity enhancement observed at user side in the screencast.

Path	# of packets	# of lost	Loss rate
NC(green)	408,724	0	0.0%
NC(orange)	396,600	2	0.0%
C(green)	373,055	4959	1.3%
C(orange)	411,762	1	0.0%

Table 3: Summary of one-way end-to-end loss statistics for each path for video sessions of 5 minutes long

Demo ScreenCast

We pre-recorded a screencast of a VoIP video call session 1'10" long. The guy keeps walking back and forth at constant speed to provide as a reference of the video fluidity. In the first 20", the video call is made without traffic congestion in the access network; the following 30" is the same session with congestion happening (image frozen); the last 20" is the same VoIP session with SDI enhancement operation done, while the congestion is still on going.

Conclusions

In this work, we demonstrated the capability of SAVI SDI platform through a multimedia application - VoIP service. As SAVI testbed has converged with the virtual network functions to configure/reconfigure OpenFlow-enabled network elements, it provides a real-time quality enhancement solution to VoIP sessions adversely affected by congestions in the access network.

With one path reconfiguration from the platform, the end users can experience visibly improvement in media quality of their video call session. We have also seen the quantifiable improvement in end-to-end RTP packet loss rate in the point-to-point HD video call experiment.

References

- [1] Shehla Abbas, Mohamed Mosbah, Akka Zemmari, and Universit Bordeaux. ITU-T Recommendation G.114, One Way Transmission Time. In *International Conference on Dynamics in Logistics 2007 (LDIC 2007)*, Lect. Notes in Comp. Sciences. Springer-Verlag, 1996.
- [2] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler. SIP: Session Initiation Protocol, 2002.

Acknowledgements

The above works are or have been funded in part or completely by the Smart Applications on Virtual Infrastructure (SAVI) project funded under the National Sciences and Engineering Research Council of Canada (NSERC) Strategic Networks grant number NETGP394424-10.