Demo: VR Video Conferencing over Named Data Networks

Liyang Zhang
Northeastern University
Boston, MA
liyangzh@ece.neu.edu

Syed Obaid Amin, Cedric Westphal
Huawei Research Center
Santa Clara, CA
{obaid.amin,cedric.westphal}@huawei.com

ABSTRACT

This demo shows an implementation of 360/virtual reality video conferencing system implemented over NDN, including producing content, formatting into NDN format, transmitting over NDN network, managing the flow of interest/content requests, and displaying in a web browser so as to show 360 degree rotation and zoom in/out features.

CCS CONCEPTS

• **Computing methodologies** → **Virtual reality**; • **Networks** → *Naming and addressing*; *Network experimentation*;

KEYWORDS

Virtual Reality, Named Data Networking

ACM Reference Format:

Liyang Zhang and Syed Obaid Amin, Cedric Westphal. 2017. Demo: VR Video Conferencing over Named Data Networks. In *Proceedings of ICN '17*. ACM, New York, NY, USA, 3 pages.
https://doi.org/10.1145/3125719.3132092

1 INTRODUCTION

AR/VR will grow to become a $100 billion market in 2025 according to a Goldman Sachs forecast. It is predicted that the deployment of 5G with faster networks with very short round trip times will enable the deployment and wide adoption of AR/VR applications.

However, on its own, 5G may not be sufficient to support AR/VR applications properly. The peak 5G bandwidth of 1Gbps and sub-millisecond RTTs could support some AR/VR applications. Yet, this level of performance may not be achieved by most users most of the times. We believe architectural support will be required to provide most consumers of AR/VR with the proper QoE.

ICN provides some benefits which will be useful for AR/VR. In particular, for a VR application, ICN offers native multicast support which allows multiple users to participate in a virtual or remote environment, multi-point to multi-point communication semantics, and potentially the sharing of the common tiles which compose a Field of View. As for AR, the ability of ICN to cache some content near the users could store enhancements that augment the reality locally.

To demonstrate the benefits of ICN for AR/VR, we demonstrate here the implementation of such a VR application using NDN. Namely, we implemented an immersive 360 video stream over NDN. One potential application, and the goal of this demo, is to enable VR video conferencing, where the remote participant feels like he/she is in the room with the other participants. VR video conferencing is a challenging application with strict delay and scalability requirements, which makes it an interesting first step.

ICN works on a pull-based model i.e. a consumer needs to send out a request, called Interest, to fetch the desired content, called Data. Using pull-based model for implementing a real-time audio/video conferencing system, which has stringent latency requirements[1], requires re-engineering the existing IP-based solutions. Existing work in this regard, like NDN-RTC [3] work in a peer-to-peer model, hence may face scaling issues. Our solution follows the design of our previous work [1] that is based on specialized service edge routers, that can scale well above 40 participants.

Here, the goal is to demonstrate the feasibility of ICN-based architecture for VR traffic, not the orchestration of services. We therefore simplify the network architecture, at the cost of more configuration at the end-hosts. Our solution utilizes two channels – one for control signals and one for data communication. The control channel provides mechanisms to synchronize the consumer and producer state; whereas the data exchange leverages ICN features, resulting in bandwidth efficiency. The proposed conferencing system provides higher reliability (i.e., faster recovery from the transient network conditions) and better performance with respect to media name sync among participants.

ICN '17, September 26–28, 2017, Berlin, Germany

ACM ISBN 978-1-4503-5122-5/17/09…$15.00
https://doi.org/10.1145/3125719.3132092

[1] Less than 150ms and 350ms, for audio and video traffic respectively [2]

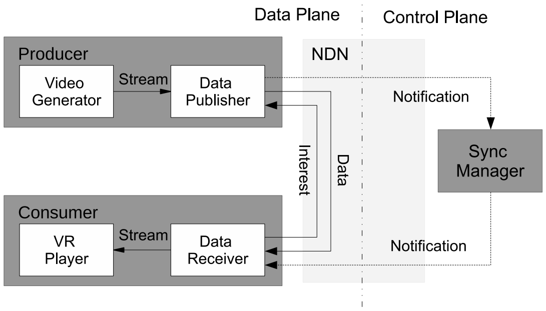

Figure 1: Framework of the NDN VR Conferencing System

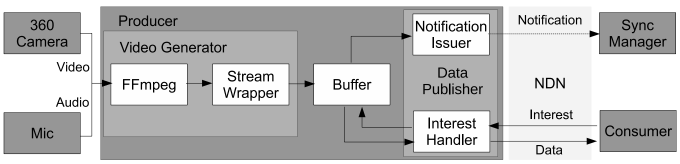

Figure 2: Block Diagram of Producer

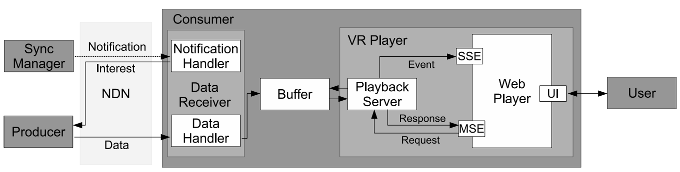

Figure 3: Block diagram of the consumer

2 SYSTEM OVERVIEW

Our objective is to demonstrate a real-time NDN-based VR videoconferencing. There are several goals: (1) Support real-time video streaming between producer and consumer; (2) Fully compatibility with NDN features, including application-suitable naming, native multicast support, etc; (3) Support seamless VR playback on a wide range of platforms; (4) Support multi-party conferencing; (5) Support rate adaption.

To this end, we implemented the framework shown in Figure 1. The system comprises of three components, the producer, the consumer, and the sync manager, which are in charge of video generation, video playback, and signaling, respectively.

At the producer side (see Figure 2), we use FFmpeg to capture and encode video on-the-fly and feed the stream to the buffer, from which the producer segments, packetizes, and names the data on a per frame base. Notifications about new chunks are sent out. It also responds to consumer-sent interests for video chunks by sending out the corresponding data.

The consumer (see Figure 3) receives notifications, and decides which chunks to fetch based on a prefetching protocol[2]. Interests for these chunks are sent in subsequence. Meanwhile, the received data chunks are sorted in correct order according to their names, and fed to the buffer. The VR player reads from the buffer, and renders the video in spherical (or equirectangular) mode. It also interacts with the user's behavior, such as selecting the Field of View (FoV) and zooming in and out.

The sync manager works as the signaling server. Its main job is to receive notifications about new video chunks from the producers and to broadcast them to the subscribing consumers. It is also in charge of user joining/leaving. When such an event happens, the involved user must send a notification to the sync manage, which then notifies other users so they can properly handle functionalities such as display window management.

There are several features of the designed system. First, it works in a centralized manner, in the sense that the sync manager coordinates all the participants. Compared with peer-to-peer solutions such as NDN-RTC, this abstracts a distinct control plane, over which the management of a conference can be easily implemented. Second, on the data plane, we adopt a prefetching protocol that allows a consumer to "reserve" data chunks that are to be generated. In this way, unnecessary latency is avoided. Third, we adopt a lightweight VR player that is universally available, and easily substituted by other players.

3 DEMO SET-UP

Our demo will display the whole system and allow demo participants to interact with it. It will consist of 2 Ricoh Theta S cameras, each connected to a laptop, and each representing a meeting space; the laptops will be connected through ethernet cable to an NDN network (emulated on another laptop). The server and customer sides will run on both ends of the demo, as each end transmit and receives a video stream to the other party. The Sync Manager runs on the laptop that also emulates the NDN network.

The demo will show a 360 video view of the other camera on both laptops. Partipants will be able to rotate the view, zoom in, pan out and interact with other participants on the other side of the demo. A short video of the demo is available at http://users.soe.ucsc.edu/~cedric/AVConfDemo.mp4

[2]Prefetching means here that the Interests for the chunks are issued for chunks not-yet-created, to be delivered at a future time.

REFERENCES

[1] Asit Chakraborti, Syed Obaid Amin, Aytac Azgin, Ravishankar Ravindran, and Guo-Qiang Wang. 2017. SRMCA: A Scalable and Resilient Real-time Multi-party Communication Architecture. (2017). arXiv:arXiv:1703.03070

[2] Yan Chen, Toni Farley, and Nong Ye. 2004. QoS Requirements of Network Applications on the Internet. *Inf. Knowl. Syst. Manag.* 4, 1 (Jan. 2004), 55–76. http://dl.acm.org/citation.cfm?id=1234242.1234243

[3] Peter Gusev and Jeff Burke. 2015. NDN-RTC: Real-Time Videoconferencing over Named Data Networking. In *Proceedings of the 2Nd ACM Conference on Information-Centric Networking (ACM-ICN '15)*. ACM, New York, NY, USA, 117–126. https://doi.org/10.1145/2810156.2810176