Extended Abstract: Coordinated Communications for Next-Generation Networks

Kiran Makhijani[1], Hamed Yousefi[1], K. K. Ramakrishnan[2], Richard Li[1]
[1]Futurewei Technologies, USA
[2]University of California, Riverside, USA
{kiranm, hamed.yousefi, richard.li}@futurewei.com, kk@cs.ucr.edu

***Abstract*—The current Internet protocol suite, with its best-effort semantics, can result in potentially very different delivery characteristics for packets. Actually, no two paths (or even different packet flows on the same path) can be assumed to have identical properties in terms of bandwidth, delay and jitter. However, multi-site remote collaboration applications are highly inter-dependent and must remain consistent across multiple users. To this end, we introduce a new network capability, called *coordinated communication service*, and propose *coordination points* to support coordinated delivery of multiple flows in the network.**

I. Introduction

While traditional multimedia has performed well over best-effort Internet, the desire to support richer multi-user networked multimedia applications such as AR/VR for interactive holographic communications and for vehicular safety of connected vehicles is growing. The richer multimedia applications seek to fuse the virtual or digital worlds along with the physical or real worlds seamlessly. Potentially allowing multiple users to manipulate the virtual objects in a scene comprising of both the physical and virtual world. All the participating users must observe concurrent changes made in different views (or locations) at the same time.

Thus, a new class of problems in this field is evolving to address the harmonized delivery of co-dependent content among multiple users, which is a vital need in many emerging use-cases such as virtual orchestra, multi-sensory, and multi-party holographic communications. This kind of problem not only requires a greater understanding of dependent aspects of the content by the application but also calls for new networking capabilities to support those dependencies. This is difficult to achieve in today's Internet.

With the above motivation, coordinated communication service was briefly introduced in Network 2030[1] initiative. In this poster, we describe this service in greater detail and propose our solution approach using *coordination points*.

II. Coordinated Communications

A. Research Problem

Consider for example, a virtual orchestra that seeks to provide a large, diverse group of participants the same concert experience as they would have if all of the orchestra members and participants were co-located. This type of application would have stringent low-latency and high-throughput requirements and these characteristics are frequently considered as the needs driving next-generation networks. Little attention has been paid to the coordination aspect of the inter-related multiple streams involved. Since, the network paths for each participating stream may be different, for a perfectly harmonious concert, the application will have to monitor each stream's network conditions continuously and synchronize the streams accordingly.

***Host-mode Coordination Limitation*—**By using this mode, much of the coordination is done by the endpoints from (at) where the flows originate (terminate). Huang *et al.* [1] proposed synchronization of multiple video streams between a single sender and receiver. Their solution does not however address the multi-receiver scenarios. Bi *et al.* [2] proposed adaptation of the delivery rate using DASH. Both are examples of endpoint-side mechanisms. While end-system adaptation can potentially provide coherent delivery of multiple streams at a single receiver, it is inadequate to provide similar coordinated delivery across multiple receivers. We find these approaches limiting particularly in solving multi-receiver path heterogeneity due to the lack of support from the network.

The capability in the network to allow multiple interrelated streams to be transmitted and received in a coordinated fashion will minimize complexity and coordinated rendering requirements across multiple receivers, and enhance user's quality of experience. Delivery of dependent information carried over multiple flows on the same or different paths is challenging as the Internet is a spatiotemporal heterogeneous environment. The path-latency varies and can change independently within each path over time. Thus, the question is "how to guarantee dependent delivery of multiple flows"?

We seek to minimize the effort in the endpoints. Dealing with the coordination of multiple flows being received at a single endpoint can be achieved by receiver-side time-based buffers with sufficient buffering and pacing methods. However, a coordinated transmission of streams to multiple receivers becomes a very complex problem. Endpoints cannot accurately predict network conditions, therefore, in-network coordination is necessary.

[1]https://www.itu.int/en/ITU-T/focusgroups/net2030

978-1-7281-2700-2/19/$31.00 2019 © IEEE

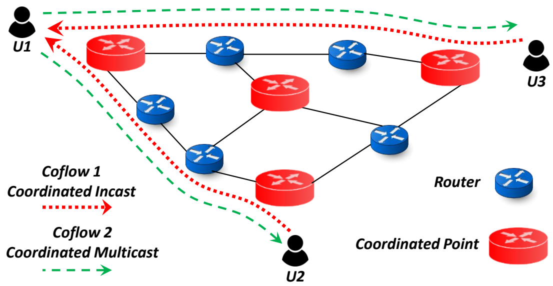

Fig. 1: Co-dependent flows in coordinated service

III. Network-mode Coordination Approach

Our approach is to address the synchronization across dependent flows by creating awareness about flow inter-dependencies in the network.

A simple network topology supporting coordinated service is shown in Fig. 1. It shows 2 coflows: coflow-1 being sent from $U2$ and $U3$, carrying two different pieces of information, which are fused by a coordination point before being delivered to the receiver $U1$; coflow-2 is a flow originating from $U1$, delivered to $U2$ and $U3$ at the same time.

A. Building Blocks

We describe coordinated service by first specifying what supporting building blocks are needed.

1) Co-dependency: The term *coordinated* refers to the need for cooperation among multiple flows with respect to application-defined dependency constraints. This dependency is primarily in terms of *time* to guarantee near-simultaneous delivery of flows. The requirement may also be to ensure *ordering*, to deliver each flow in a specific or relative order. Additionally, dependency may include *QoS fate sharing* so that when the QoS degrades for a particular flow, the other flows may be subjected to the same reduced service level.

2) Coflows: short for co-dependent flows[2], a group of two or more flows that are co-dependent according to the rules mentioned in previous section.

3) Group Communications: The set of flows that form the co-dependent flows may originate from and/or terminate at different endpoints. Thus, the coordinated service is not limited to any specific type of communication pattern between the endpoints. It needs to function on top of one-to-one, one-to-many (multicast), many-to-one (incast), and many-to-many communications.

B. Coordination Points

We propose the use of coordination points that utilize both in-band and out-of-band information as illustrated in Fig 2 to support coordinated delivery in the network. There are three categories of functions performed by coordination points as below:

[2]Not to be confused with the term in cluster applications by Chowdhury *et al.* [3].

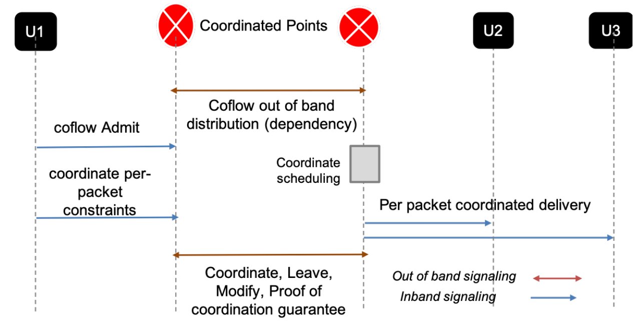

Fig. 2: Distributing coflows in the network

1) Out-of-band Coordination: The role of out-of-band function is to achieve distribution of coflows in the network. The membership of co-dependent flows is conveyed from the end-points potentially when the flows are set up, so that the coarse-grained service (and service level objectives) can be enabled in the network. A distribution graph of coflows and associated dependency constraints may be constructed, and those nodes enhance their scheduling and forwarding by factoring in the timing information in the meta-data of packets.

2) In-band Coordination: Additionally, timing information (timestamp of transmission from sender and time window for delivery to receiver) may be conveyed as meta-data in packets transmitted from the senders. In-band signaling conveys intermediate coordination points about the dependencies and inter-relationship. To formalize these mechanisms to carry them in data path, we utilize Big Packet Protocol (BPP) [4] data plane programmability as below:

$$\texttt{instr}: \textit{coordinate}\ \texttt{when}\ coflowId = x,\ mark = set$$
$$\texttt{param}:\ txTime = T_s,\ delay = \delta_t$$

This instruction is per packet to indicate a member of coflow and dependency in terms of parameters: $marker$, T_s, and δ_t.

3) In-node Coordinated-forwarding: Actual coordination effort is done on the coordination points. The scheduling and forwarding engine should allow packets within sync markers to be sent as per remaining δ_t. It needs to compare the remaining coordination time and accordingly schedule or pace the packet forwarding.

C. Future Work

Currently we are building a simulation platform using BPP to verify that coordinated services help applications perform better. Later a more comprehensive signalling methods will be studied (as in Fig. 2) can be studied.

References

[1] Z. Huang, W. Wu, K. Nahrstedt, A. Arefin, and R. Rivas, “TSync: A New Synchronization Framework for Multi-site 3D Tele-immersion,” in *NOSSDAV’10*, 2010, pp. 39–44.
[2] T. Bi, A. Pichon, L. Zou, S. Chen, G. Ghinea, and G.-M. Muntean, “A DASH-based Mulsemedia Adaptive Delivery Solution,” in *MMVE’18*, 2018, pp. 1–6.
[3] Zhang, Hong and Chen, Li and Yi, Bairen and Chen, Kai and Chowdhury, Mosharaf and Geng, Yanhui, “CODA: Toward Automatically Identifying and Scheduling Coflows in the Dark.” in *SIGCOMM*, 2016.
[4] R. Li, A. Clemm, U. Chunduri, L. Dong, and K. Makhijani, “A New Framework and Protocol for Future Networking Applications,” in *ACM SIGCOMM Workshops (NEAT’18)*, 2018, pp. 637–648.