Research Project Proposals

CS 397/497 Selected Topics in Computer Networks, Spring 2024

Sen Lin

sen.lin@u.northwestern.edu

Project 1: Multi-Flow

From metaverse streaming to a more general network pattern

P1: Co-Flow VS **Multi-Flow** VS Multi-Path

	Co-Flow	Multi/Parallel-Flow	Multi-Path
Topology	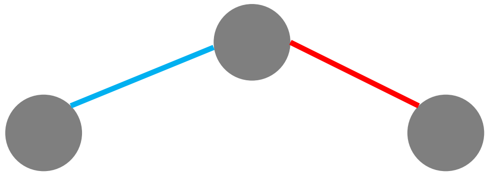	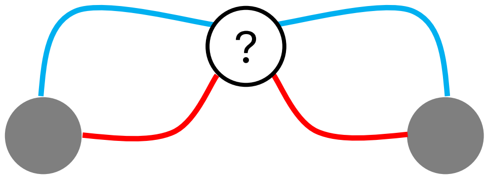	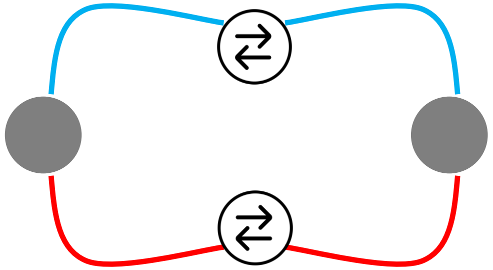
# Connections	N	1	1 (aggregation, N subflows)
Assumption	• a group of network flows that share a common performance goal	• Application layer multiplexing (correlation, dependency, priority, reliability)	• Huaizhong Han (2005): source routing (overlay router OR multi-homing) • Olivier Bonaventure, IETF (after ~2010): multiple interfaces and links (e.g., Cellular & WiFi)
Notes		• No need multi-path assumption • Lack of specialized congestion control	• IETF MPTCP is implemented in major OSs • Specialized congestion control algorithms (e.g., Linked Increase Algorithm, RFC 6356) • No consideration of application layer multiplexing (a drop-in replacement of TCP)

P1: Metaverse Streaming Case

To enable immersive digital twins and sensitive interactive experience → Network Characteristics

Bidirectional Intensive

Video, Audio, Haptics...

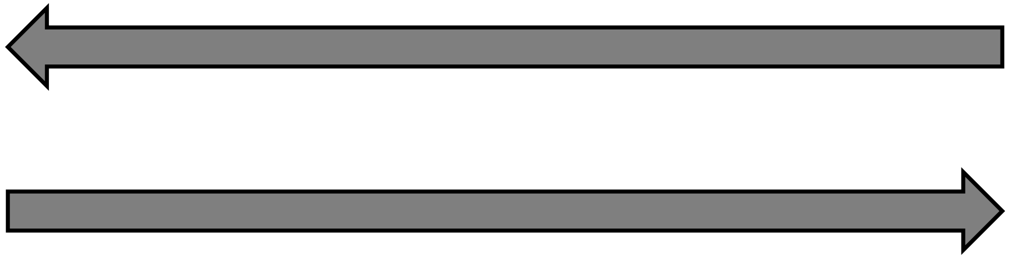

Diverse Traffic

Audio, 3D point clouds, Actions, Orientations, Positions...

Dependency

Image credits: Apple Vision (MacRumors); Disney HoloTile Floor (IGN)

P1: Metaverse Streaming Case

To align application flows with transport flows

We find the key to the imminent problem from the traffic diversity in metaverse applications, where different flows exhibits **distinct priorities**, **reliability requirement**, and **correlation** and **dependency**. We propose ruling heterogeneous traffic in one connection and optimize the bi-direction transmission in a comprehensive way by **aligning the application-layer awareness with transport-layer controlling**.

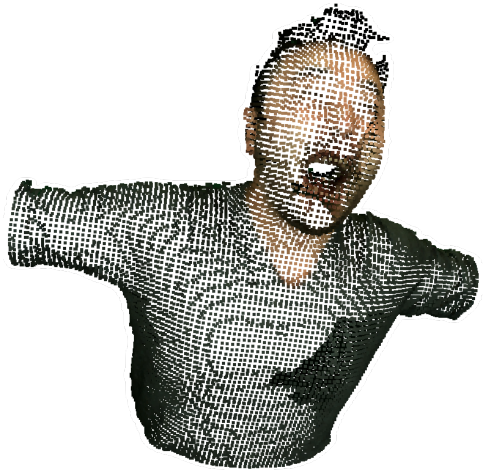

3D PC credit: Microsoft Research

P1: More Questions

- Flow scheduling algorithms for multi-flow connections
- Smarter packet assembling in multi-flow connections
  - Frame placement in each packet
- Multi-flow connections in publish-subscribe network patterns
- Multi-flow connections with in-network computation

Project 2: Unbinding IP

To maximize connection flexibility

P2: VPN with Multiple Proxy Nodes Available

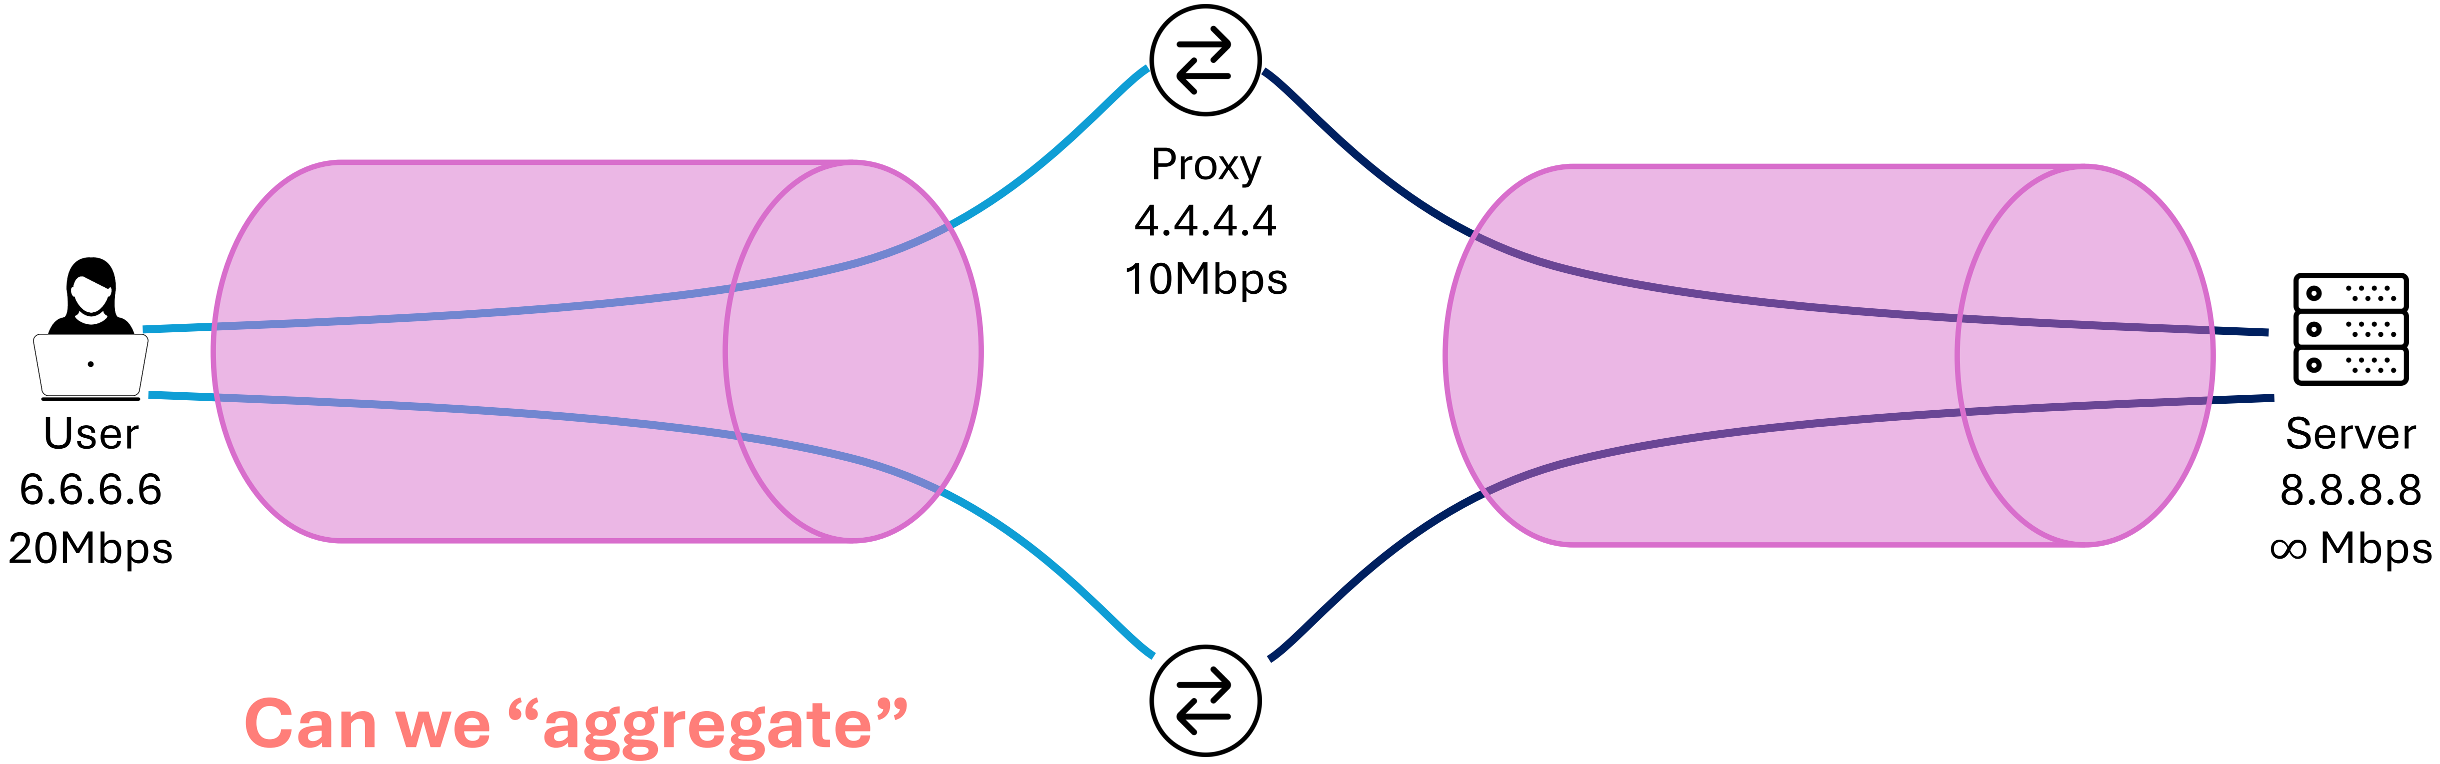

Can we "aggregate" multiple links and their bandwidths?

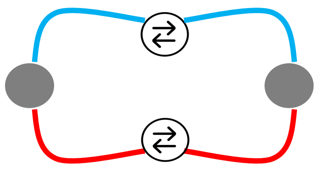

Multipath failed due to its assumption of end-to-end

P2: Background - Unbind IP (Server-Side)

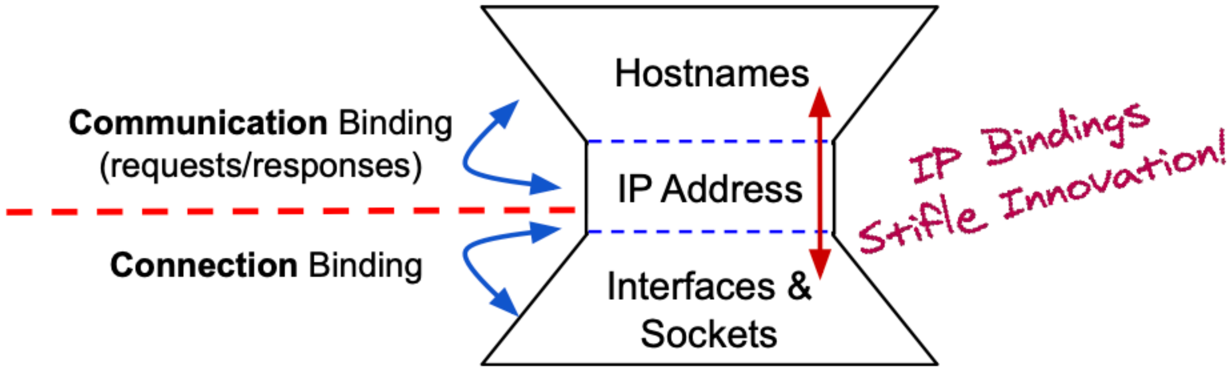

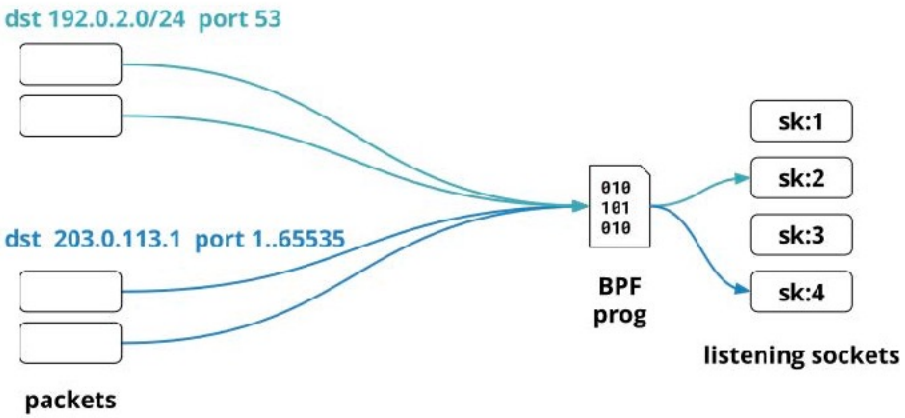

Recalling the relation between IP addresses, hostnames, and sockets, the IP address is bound to a connection because:

- It serves as parts of the connection identifier (as in TCP)
- System design strictly binds it to a socket

An **sk_lookup** example: Packets arriving on 192.0.2.0/24:53 to socket sk:2, while traffic on 203.0.113.1 to any port number lands in socket sk:4.

Image Credit: Fayed, Marwan, et al. "The Ties that un-Bind: Decoupling IP from web services and sockets for robust addressing agility at CDN-scale." SIGCOMM'21

P2: Proposal - Unbind IP (Both Sides)

Leverage the recently introduced IP-Socket unbinding technique. the IPs and ports are no longer associated with a “connection”, i.e., socket. We can maximize the use of socket mobility to fractionalize a connection. We assign a fixed connection ID/token in the packet header, e.g., TCP TS. Then the connection can be fractionalized into multiple links/paths at the packet level/

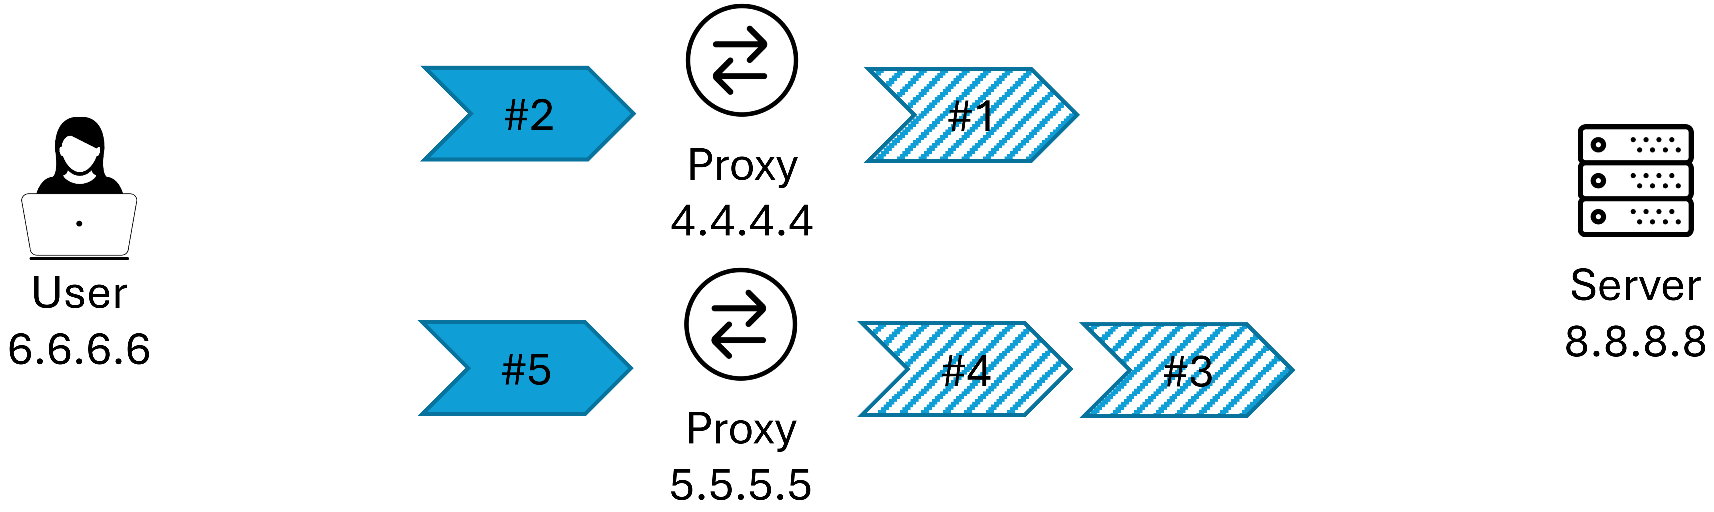

Yunming Xiao

yunming.xiao@u.northwestern.edu

Project 1: Eliminate the Trust: The Case of Ad Exchange

Background

- The best privacy-preserving system is the one that removes the dependency on the non-collusion agreement
- A good example is PDNS

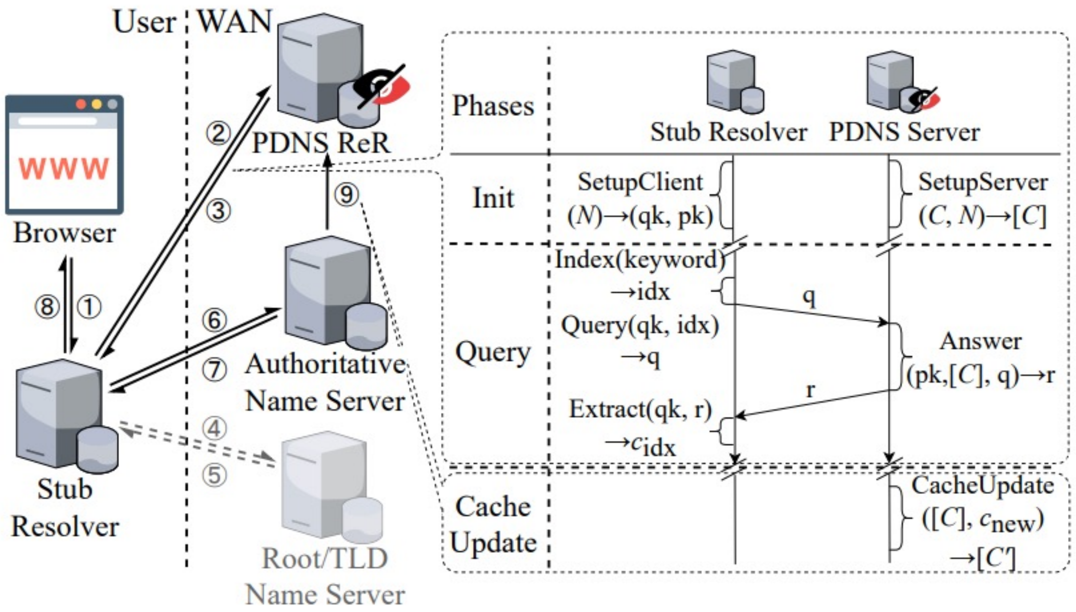

Figure 1: Visualization of PDNS and its workflow.

Background

- A good example is PDNS

Table 1: Comparison of privacy-preserving properties of current DNS solutions versus single and multi-server PIR.

Solution	Defend Pervasive Monitoring	Hide Individual Access Pattern	Hide Organizational or Regional Access Pattern	Survive Non-Collusion Agreement Violation
DoUDP [16] / DoTCP [43]	No	No	No	N/A
DoT [57] / DoH [54]	Yes	No	No	N/A
DoT/DoH + Resolver Rotation [56, 80]	Yes	Yes*	No	N/A
Oblivious DNS [82]	Yes	Yes	No	No
ODNS With Proxy Rotation [63]	Yes	Yes	Yes*	No
DoHoT [69, 70]	Yes	Yes	Yes*	Yes*
DNS with Multi-Server PIR	Yes	Yes	Yes	No
DNS with Single-Server PIR	Yes	Yes	Yes	Yes

Background

- Another example could be ad exchange

Figure 1: Online Advertising Auction

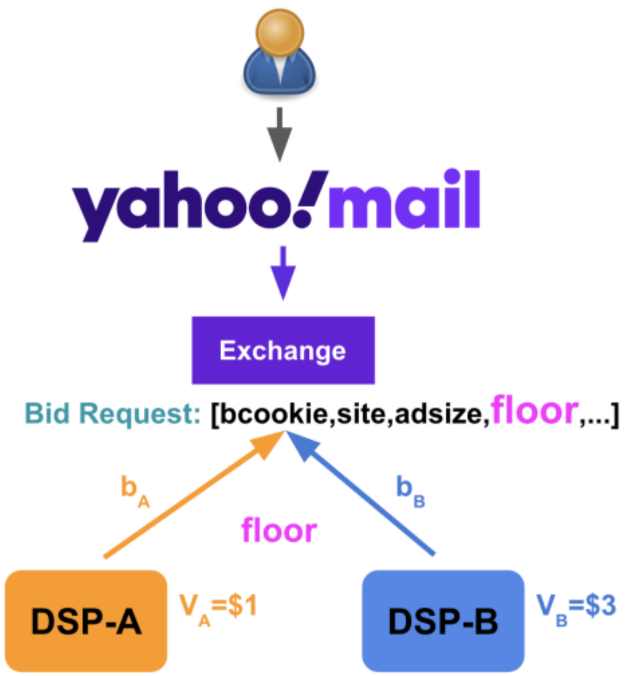

Figure helps to understand how bidding behavior of DSPs changes as a function of the auction mechanism and floors.

Background

- An improvement is to employ multi-party computation to hide the bidder's information
  - User privacy is not protected

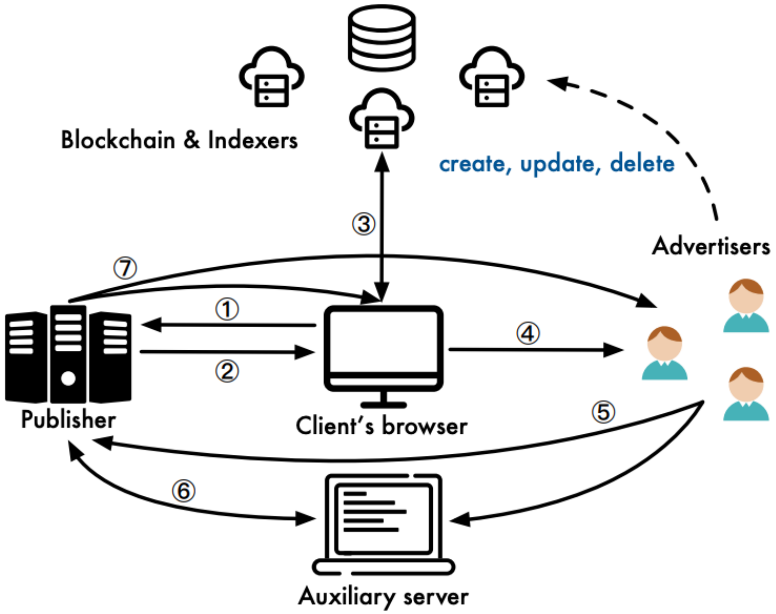

FIGURE 1—In Addax, the exchange's functionality is divided among the publisher, browser, an auxiliary server and a blockchain.

Background

- To protect user privacy, one solution is to not send it out to anyone but decides the ads to watch locally (Brave)
  - But here the ad provider cannot decide their auction schemes

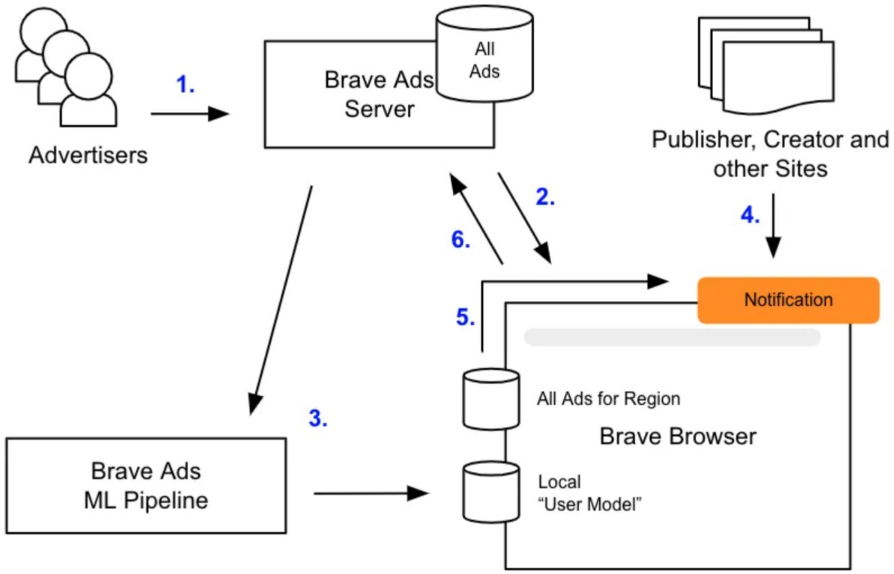

Proposal

- Our approach
  - The data shouldn’t leave the users’ devices
  - We ask the ad providers to send their auction algorithm to the user, which performs the computation and auction
  - Some cryptography tools can ensure that the user does not understand the ad providers’ algorithm while allowing the ad provider to verify that the user did not cheat
  - The cost is it incurs much higher computation requirements

- Goals of this study:
  - Find out whether the above can be done, and how it could be done
  - Evaluate the overhead of the above approach

Project 2: Optimal/Deterministic WAN Routing

Motivation

- Latency is critical for user experience in certain applications, e.g., online games (for both users and game developers).
  - For instance, I have worked on a traffic engineering project which aims to support fixed path routing across data centers. Here, a latency difference of 5ms means something
  - A->C->D is preferred

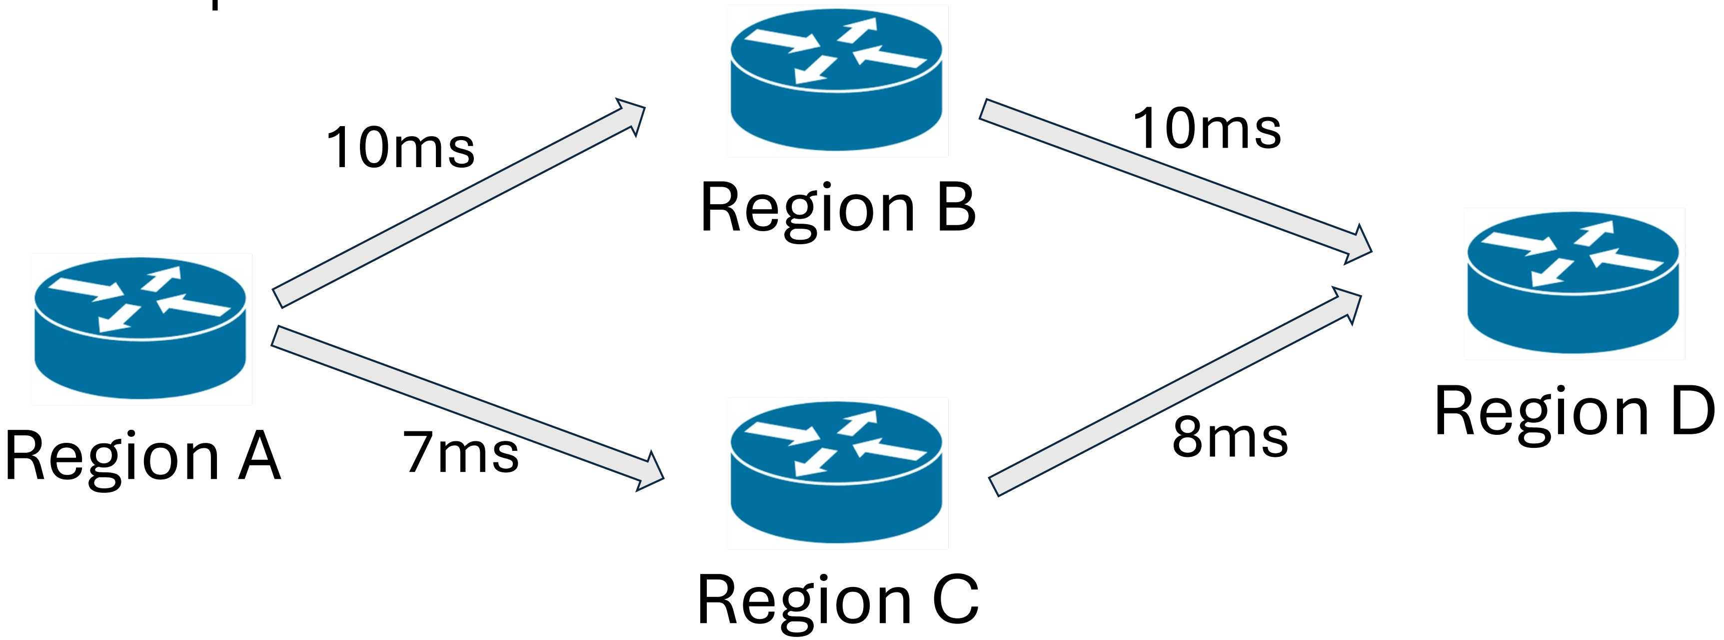

Motivation

- Latency is critical
- Inter-AS routing is complicated
  - Client-Edge + Edge-Cloud != Client-Cloud [Snatch, Eurosys'24]

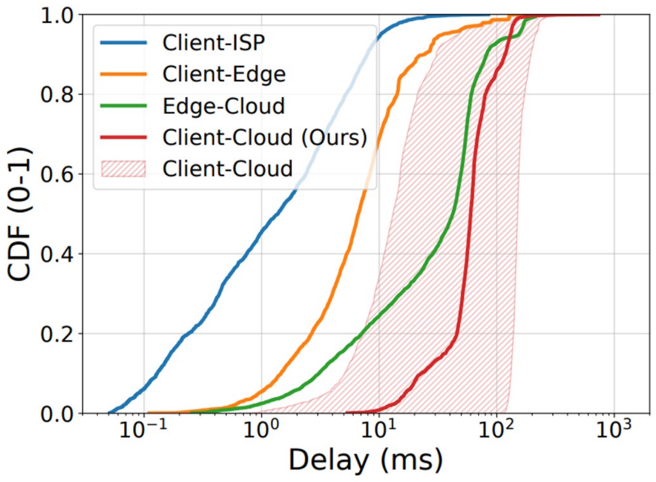

(a) Inter-component delays.

Motivation

- Latency is critical
- Inter-AS routing is complicated
- Intra-cloud and inter-cloud routing is also complicated
    - The default routing is not always optimal; instead, creating a private link is much faster and cost-efficient [Skyplane, NSDI'23 & SkyPilot, NSDI'23]

Insight #1: overlay routing to circumvent slow links

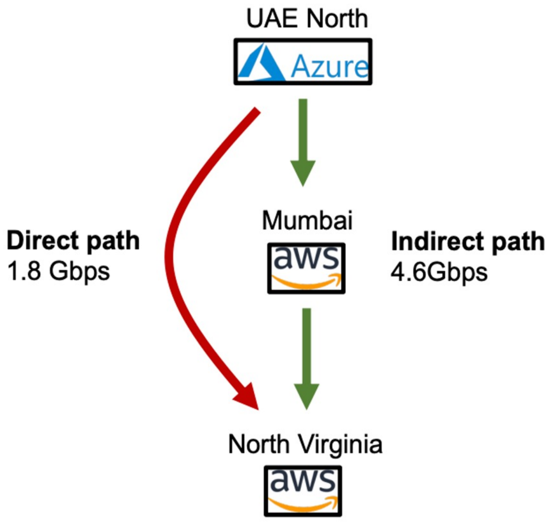

Proposal

- Create a deterministic private tunnel towards the target servers
  - minimize latency/loss rate/traffic velocity
  - with proxies (servers or switches) at edge networks as well as inside data centers
- Hypergiant’s off-nets work similarly, but we aim at making this faster, wider participation, and beyond content caching
  - Faster: more path choices & protocol innovation
  - Wider participation: involving more edge network as well as data centers
  - Beyond content caching: Support any network traffic with VPN

Illustration: Selecting the Best Path

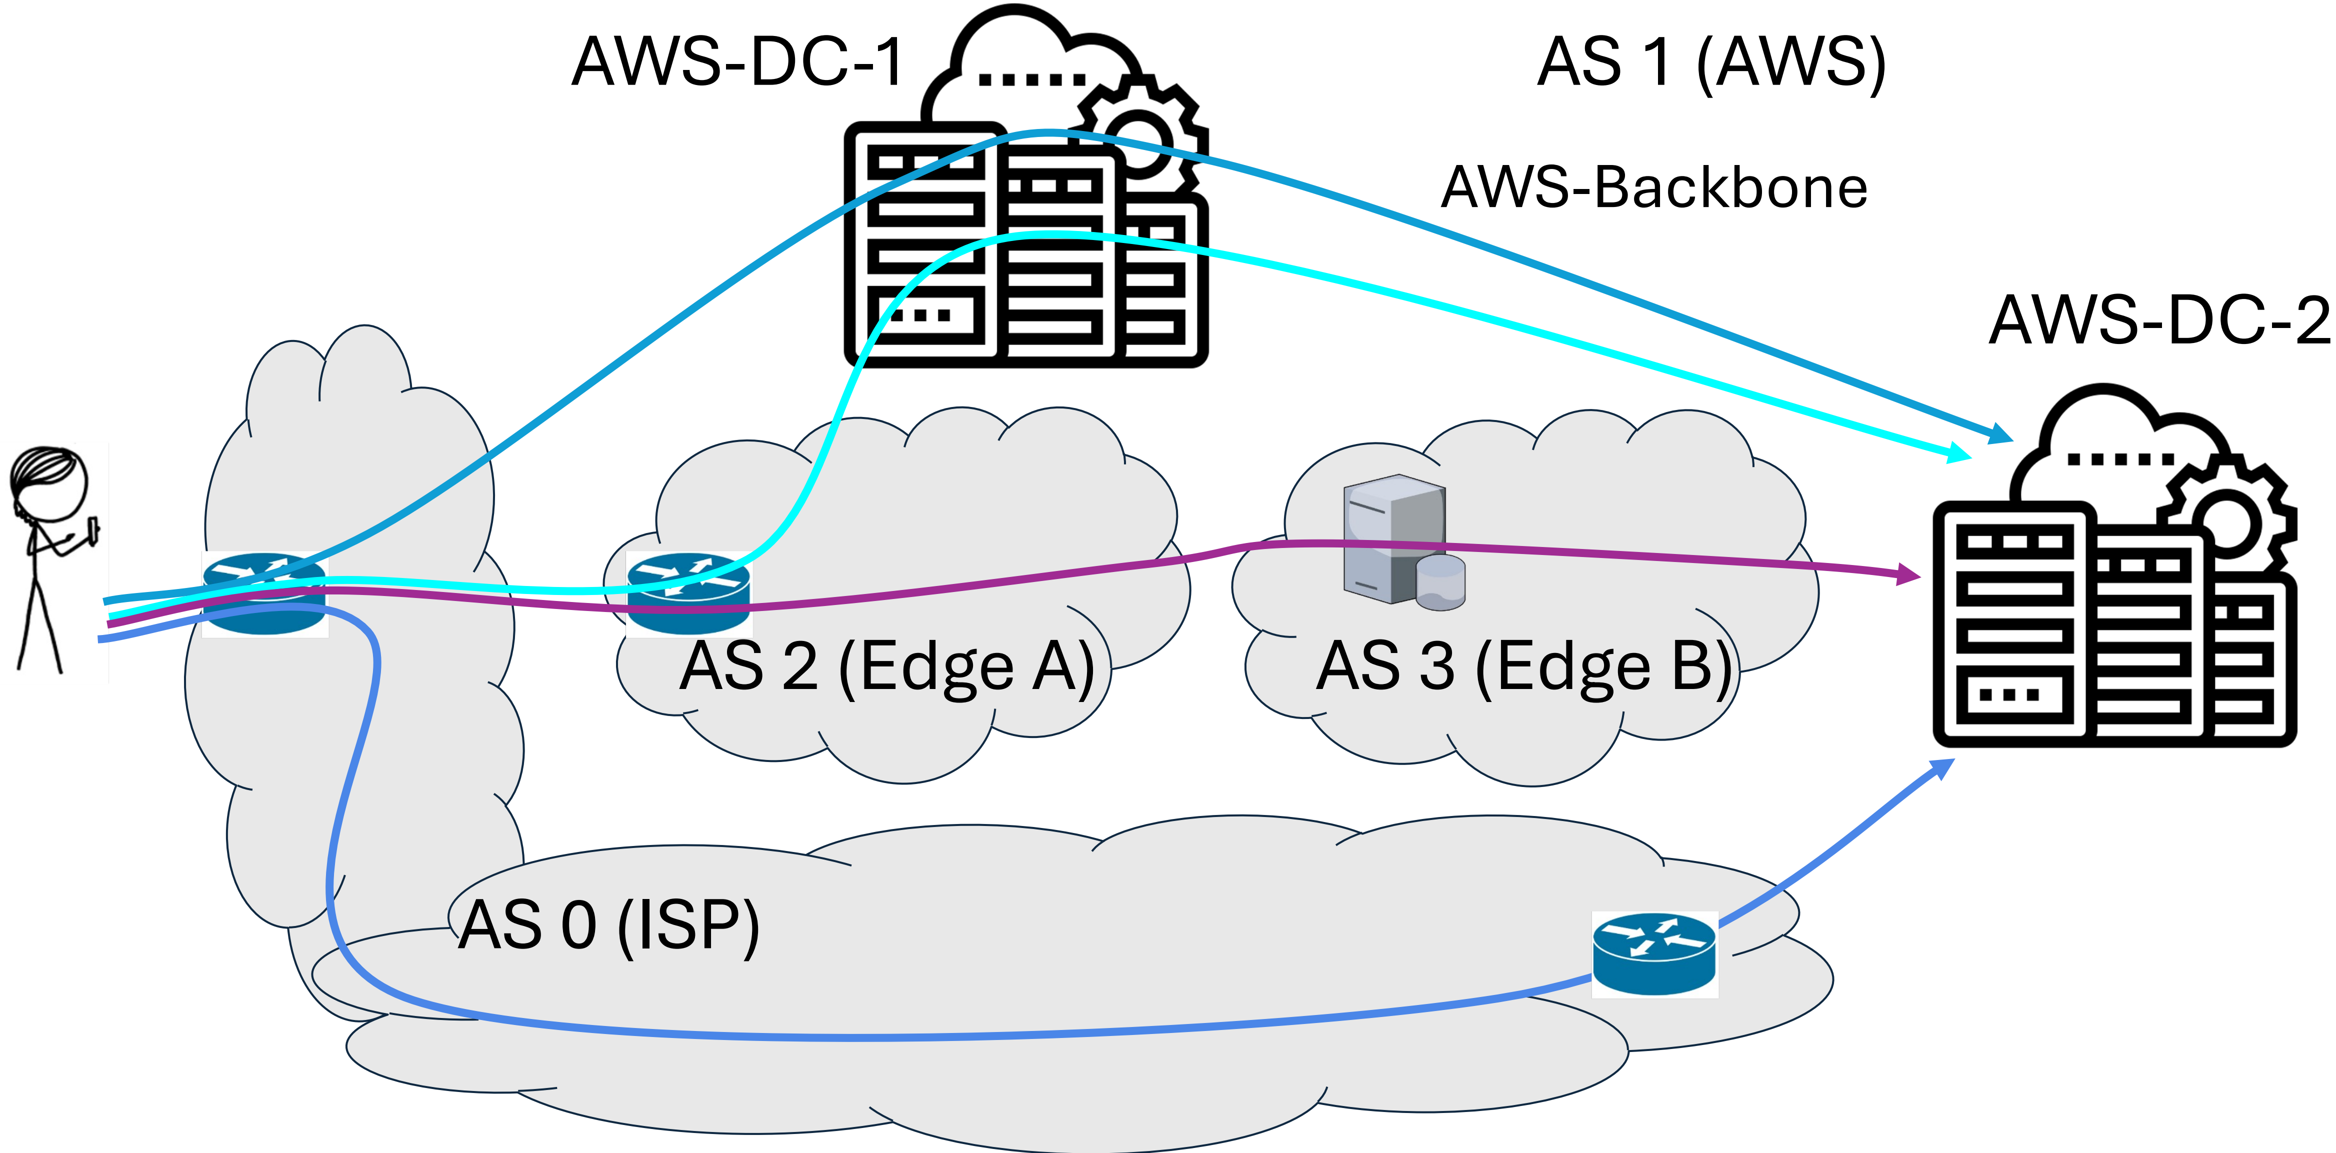

Protocol Innovation

- We can also advance the performance of lower-level protocols
- Example: instead of using existing VPN protocols, we propose to forward the traffic with much lighter overhead -- transparent QUIC proxy
  - No need to establish TCP/TLS connection between client and proxy, or between proxy and destination server
  - The proxy just needs to change the UDP source and destination IPs
  - Programmable switch or DPDK can help this
  - https://blog.cloudflare.com/unlocking-quic-proxying-potential

Yanzhi Li

yanzhili2026@u.northwestern.edu

Project 1: HTTPS Fallback Study

HTTPS Fallback

HTTPS Fallback is a mechanism used when a client and server are attempting to communicate over a more advanced or preferable protocol but encounter issues that prevent the communication from being established or maintained. If such issues occur, they "fall back" to using a older version protocol to ensure that the communication can still take place, albeit without the benefits of the newer protocol.

This fallback process is crucial for maintaining compatibility and accessibility, especially as the internet transitions to newer technologies.

HTTP3 is the newest protocol

Many websites currently support HTTP3

However, many traffic are still HTTP/2 even if both sides support HTTP3

Questions to answer

To understand, what triggers the mechanism

How fallback mechanism affect web performance

Approach

From a top list websites, retrieve 100 to 300 websites that support QUIC.

Then analysis HTTP3 OR 2 traffic

Find out when fallback happens

Measure the impact of fallback

Project 2: Spin Bit Accuracy

Concept

The Spin Bit is a feature introduced in the QUIC transport protocol, designed to provide a simple way for external observers (like network operators and diagnostic tools) to measure end-to-end latency between two points on a network without significantly compromising user privacy.

- The server changes the value of the Spin Bit each time it receives a packet from the client that has a different Spin Bit value than the last packet sent by the server.
- This behavior results in the Spin Bit value flipping once per round trip, allowing an on-path observer to measure the time between flips as an estimate of round trip time (RTT).

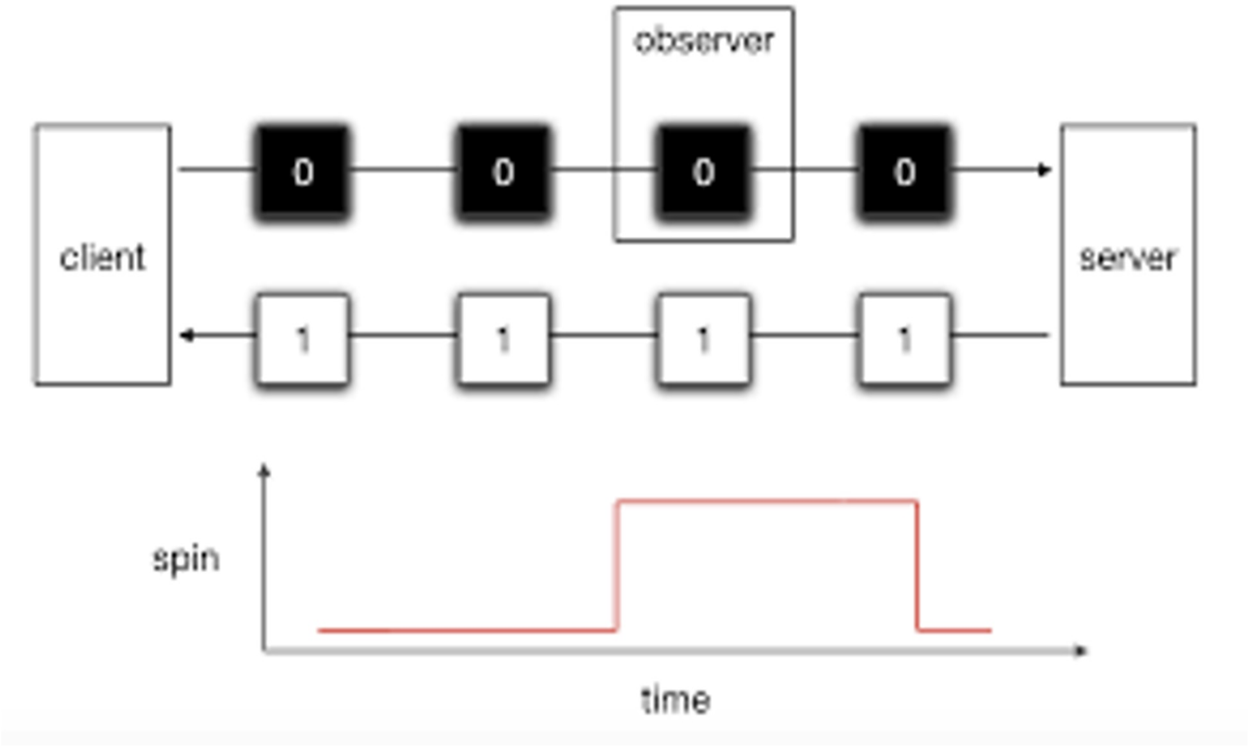

Why It Matters

1. Network Performance Measurement: The Spin Bit allows network operators to measure and monitor network latency in real-time, helping in traffic management, congestion control, and identifying network issues. It offers a balance between the need for privacy and the need for network performance data.

2. No Additional Overhead: Implementing the Spin Bit does not require additional packets or headers beyond what QUIC already uses, making it an efficient way to estimate latency without adding network load.

3. Privacy-Preserving: The Spin Bit provides limited information (essentially, just whether it's spinning or not), which helps preserve user privacy while still allowing useful network measurements.

Spin bit limitations

- Asymmetric Routes
- Reordering and Loss
- Limited Information
- Sampling Restrictions

Quantitative analysis

Measure the discrepancy between spin-bit estimation and real RTTs
Higher Goal,design a system to improve spin-bit estimation