Reliable Data Multicast Across Inter-Domain Multipoint Layer-2 Networks

Yuanlong Tan, Michele Co, Malathi Veeraraghavan
University of Virginia

Presented by
Yuanlong Tan
yt4xb@virginia.edu
Feb. 6, 2019

Outline

- Background
- Cross-layer architecture & LDM7
- Evaluation of LDM7 on Chameleon
  - VLAN/Sender-tc rate
  - Packet loss rate
  - Feedtype
  - Round-trip time (RTT)
- Chameleon experience

Background

- UCAR Unidata IDD program collects and distributes real-time meteorology data to 576 sites at 230 domains
  - 30 data feeds, referred as feedtypes
  - Receives 20 GB/hr from various sources (e.g., radar)
  - Sends 1 TB/hr out to subscribers
- Software: Local Data Manager (LDM)
  - LDM6: push mode; uses Application-Layer Multicast
    - As number of receivers ↑, traffic volume ↑, UCAR needs higher-speed WAN access, and larger sender clusters
  - LDM7: cross-layer architecture; uses Network Multicast

Our solution

Comparison between LDM6 & LDM7, please see our previous work:

S. Chen, et al, "A cross-layer multicast-push unicast-pull (MPUP) architecture for reliable file-stream distribution," in Computer Software and Applications Conference (COMPSAC), 2016 IEEE 40th Annual. IEEE, 2016.

UCAR: University Corporation for Atmospheric Research

IDD: Internet Data Distribution

Outline

- Background
- Cross-layer architecture & LDM7
- Evaluation of LDM7 on Chameleon
  - VLAN/Sender-tc rate
  - Packet loss rate
  - Feedtype
  - Round-trip time (RTT)
- Chameleon experience

Cross-layer architecture & LDM7

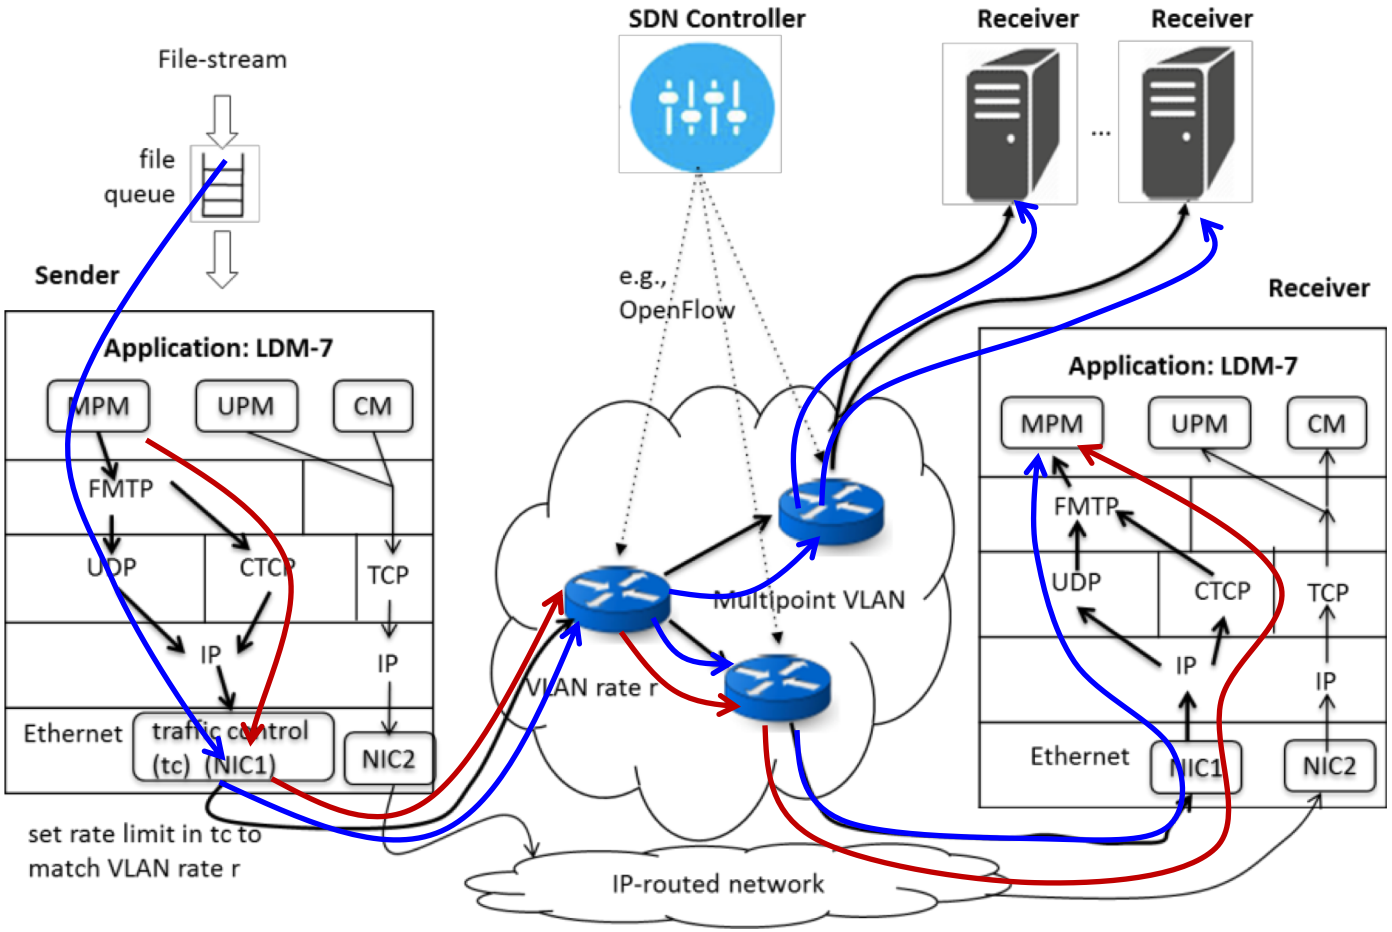

MPM: Multicast-Push Module
UPM: Unicast-Pull Module
FMTP: File Multicast Transport Protocol

CM: Control Module
CTCP: Circuit-TCP

Outline

- Background
- Cross-layer Architecture & LDM7
- Evaluation of LDM7 on Chameleon
  - VLAN/Sender-tc rate
  - Packet loss rate
  - Feedtype
  - Round-trip time (RTT)
- Chameleon experience

Evaluation of LDM7 on Chameleon

Experimental Setup:

- Haswell bare-metal nodes reserved in advance at TACC.
- Isolated network VLAN
- Linux network traffic control utility, tc, was used to set rate limit and emulate RTT (with netnem).
- Linux iptables was used to inject packet loss rate.

Input parameters:

- Feed: metadata for products in 1-hr file-streams
- Num. of receivers varied from 1 to 16
- Loss rate: 0%, 5%
- Rate of the multipoint VLAN = sender-tc rate: 20 Mbps – 1000 Mbps
- RTT: 0.1 ms – 50 ms

Output metrics:

- Throughput (LDM log): Dividing aggregated file sizes by aggregated latency.

Impact of r, m and p on throughput

- r, VLAN/sender-tc rate; m, number of receivers; and p, injected loss rate.
- Feedtype: NGRID; RTT: 0.1ms.

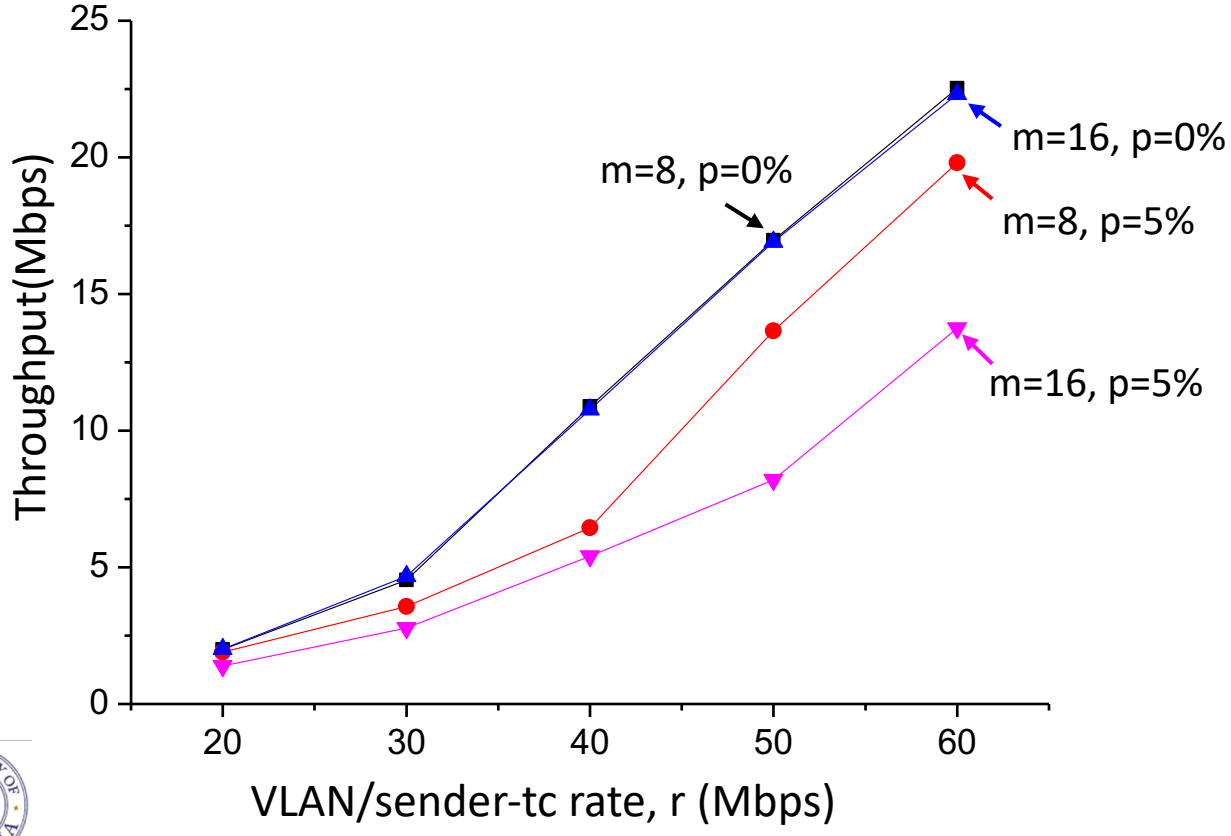

- as VLAN rate increases, throughput goes up
- m doesn't matter when loss rate is 0
- when loss rate is 5%, the value of throughput depends on m.
- Reliability --> Throughput goes down

Impact of loss rate on throughput

- Feedtypes; packet loss rate.
- Num. of receivers: 8; VLAN/sender-tc rate: 50 Mbps; RTT: 0.1 ms.

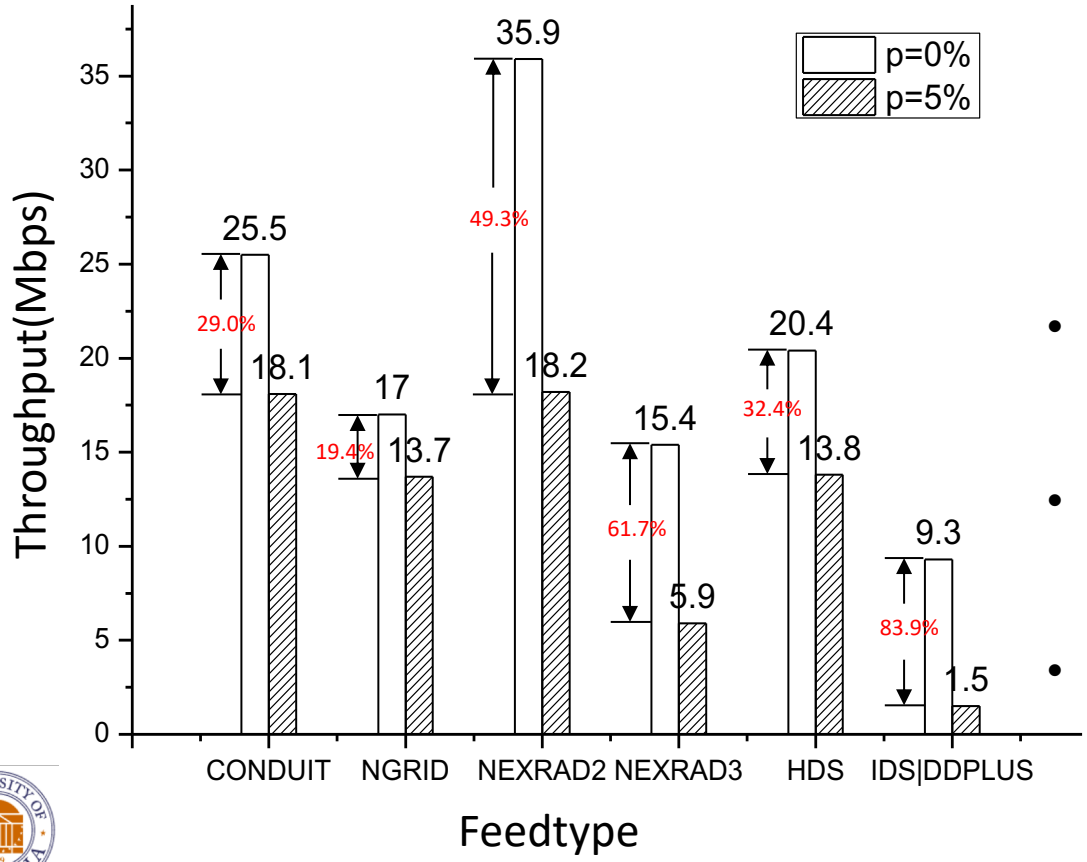

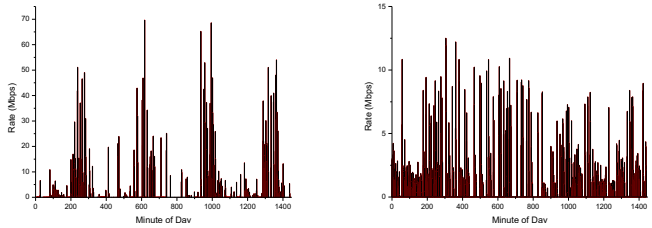

CONDUIT NEXRAD2

- All of feedtypes suffer a throughput drop with injected 5% packet loss.
- Different feedtypes have different drop rate, depending on its characteristics.
- Taking CONDUIT and NEXRAD2 as examples, we have a close look the buffer delay vs. transmission delay.

Impact of feedtype on throughput

- Feedtypes; VLAN/sender-tc rate.
- Num. of receivers: 8; Packet loss rate: 0; RTT: 0.1 ms.

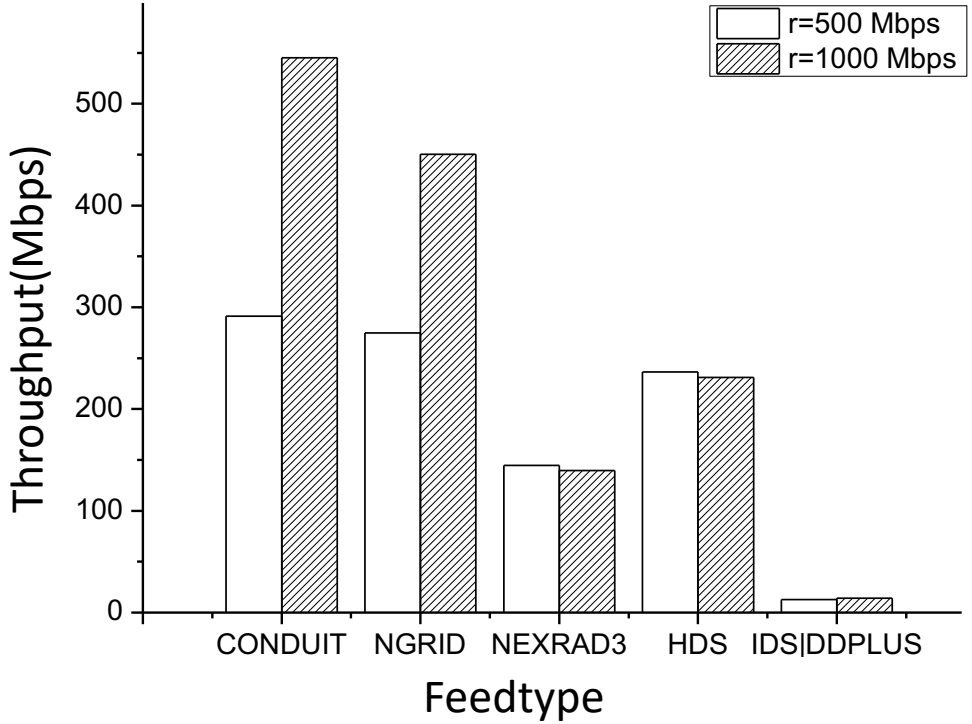

- Different delay play an important role in latency.
- CONDUIT and NGRID, buffer delay and transmission delay dominates the latency.
- NEXRAD3, HDS, and IDS|DDPLUS other factors influence the latency more

Impact of RTT on throughput

- Parameters: RTT, round trip time; and r, VLAN/sender-tc rate.
- Num. of receivers: 1; packet loss rate: 0; and feedtype: NGRID.

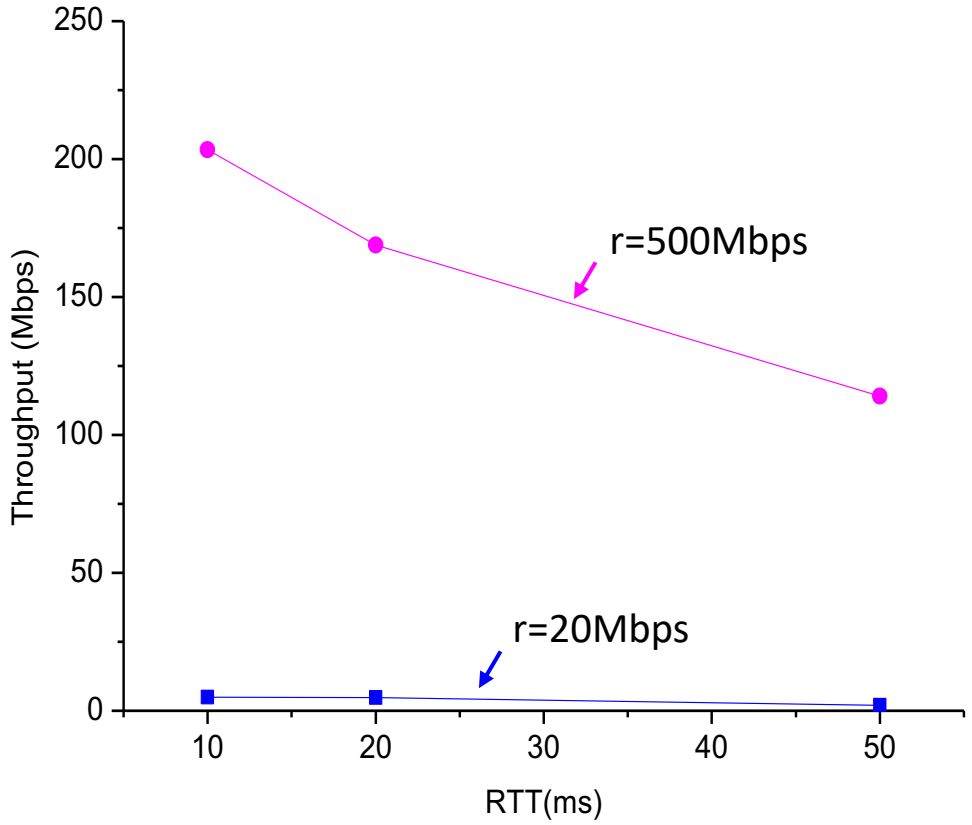

- When VLAN/sender-tc rate is high, propagation delay is the main factor contributing latency
- When VLAN/sender-tc rate is low, buffer delay will deeply influence throughput.

Conclusions

- Several factors are investigated which may have an impact on throughput.
  - Transmission delay
  - Sender tc buffering delay
  - Propagation delay
  - Processing delay
  - Clock synchronization

Outline

- Background
- Cross-layer Architecture & LDM7
- Evaluation of LDM7 on Chameleon
  - VLAN/Sender-tc rate
  - Packet loss rate
  - Feedtype
  - Round-trip time (RTT)
- Chameleon experience

Chameleon experience

- Pros:
  - Allocating resources conveniently, like:
    - reserving the resources in advance
    - discovering resources with lease calendar
  - Isolated network VLAN
  - Multiple network interfaces (specifically, 4 network interfaces in Haswell bare-metal nodes)

- Our concerns:
  - Isolated network VLAN sharing the interface with ssh.
  - Assigning multiple networks, more than 1, to the instances.
  - Scaling up experiments to a larger size, such as over 50 instances.

Thanks!

Q&A