

# From SRM to NDN: Three Decades of Lessons in Data-Centric Networking

Tianyuan Yu, Adam Thieme, Junxiao Shi, Lan Wang, Lixia Zhang  
UCLA • NIST • University of Memphis

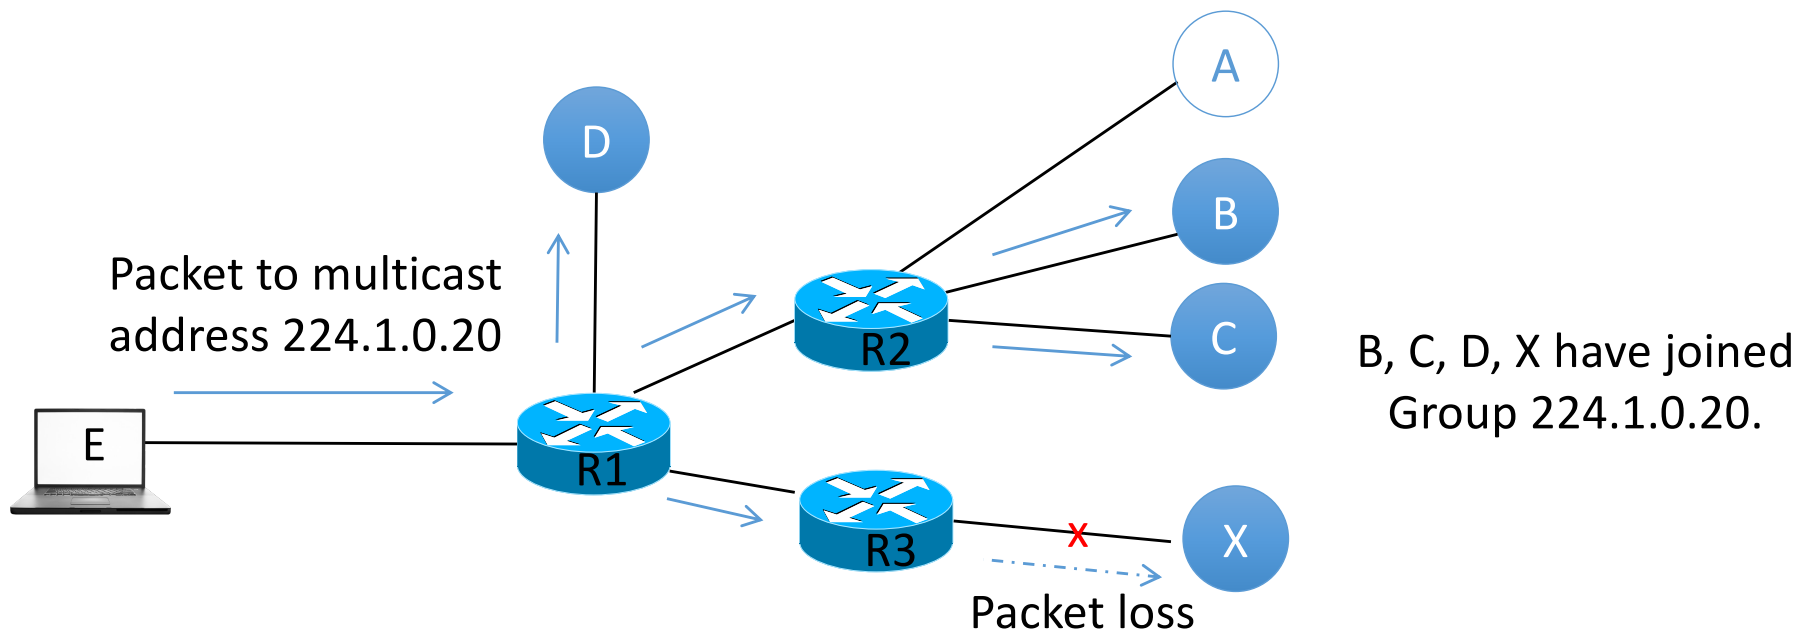
NDNComm 2026

# Motivation

- Multiparty communication is everywhere
  - Collaborative editing
  - Group messaging
  - Multiplayer games
  - AI agent communication
- Key challenge:
  - How to efficiently and reliably deliver shared data to many heterogenous and dynamic receivers?
  - Converting group delivery to individual TCP unicast pipes causes excessive traffic.

# IP Multicast [1]

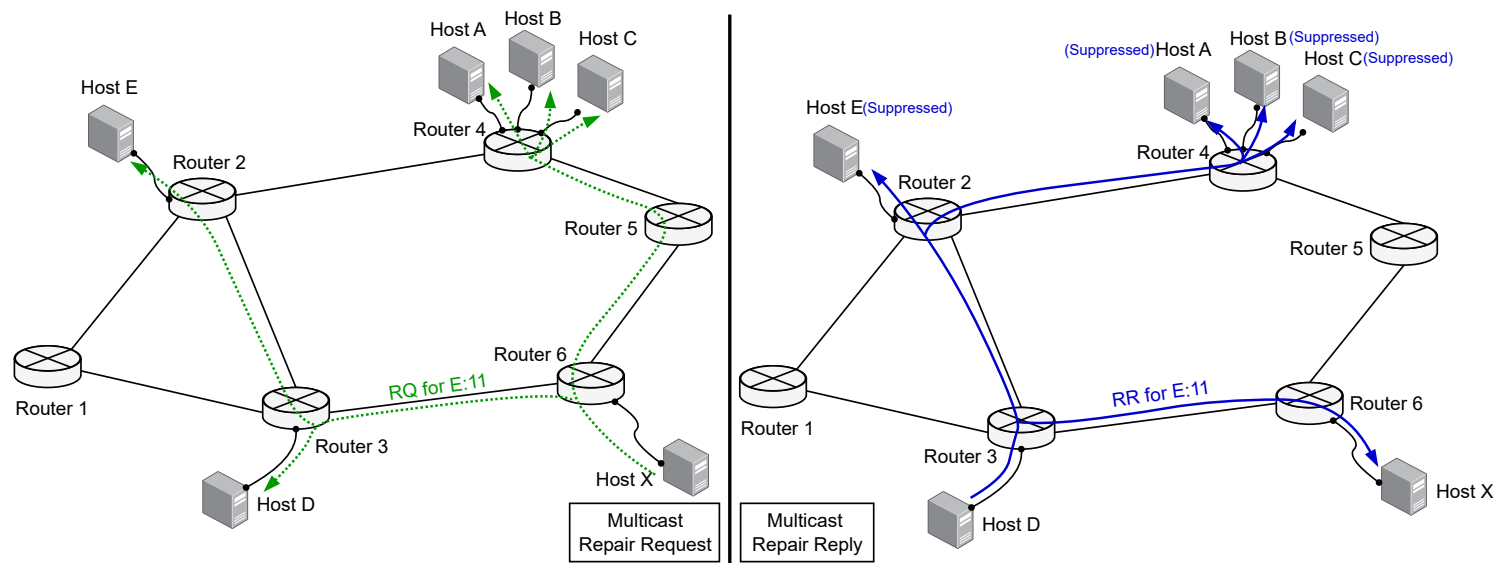
- Use logical IP addresses (e.g., 224.0.0.0 through 239.255.255.255) for multicast groups.
- Receivers join multicast group.
- Packets sent to the multicast address will be forwarded to all receivers in the group.



Source-based receiver tracking and retransmission do not scale.

# Scalable Reliable Multicast (SRM) [2]

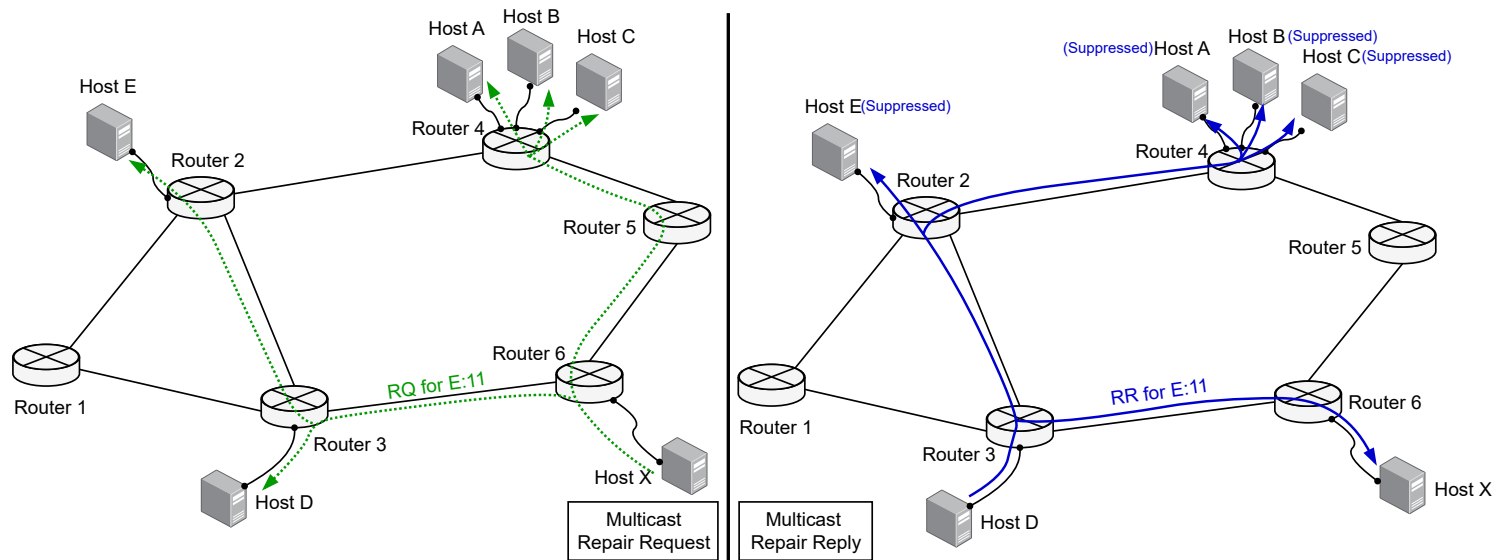
- Introduced in 1995
- Application-Level Framing [3]: naming data at the level of application data units, e.g., (E:11)
- Receiver-driven loss recovery
- SRM multicasts all communication, including (i) application data, (ii) session messages that convey production state, and (iii) repair requests (RQs) and repair replies (RRs) for collaborative loss recovery.



Insight: SRM already used a data-centric communication model.

# SRM Recovery Scales Poorly

- Repair requests multicast to entire group
- Recovery traffic spreads globally
- Duplicate suppression timers required



Even with ideal suppression, a single packet loss forces one RQ and one RR to be delivered to all multicast group members.

# Timer-Based Suppression

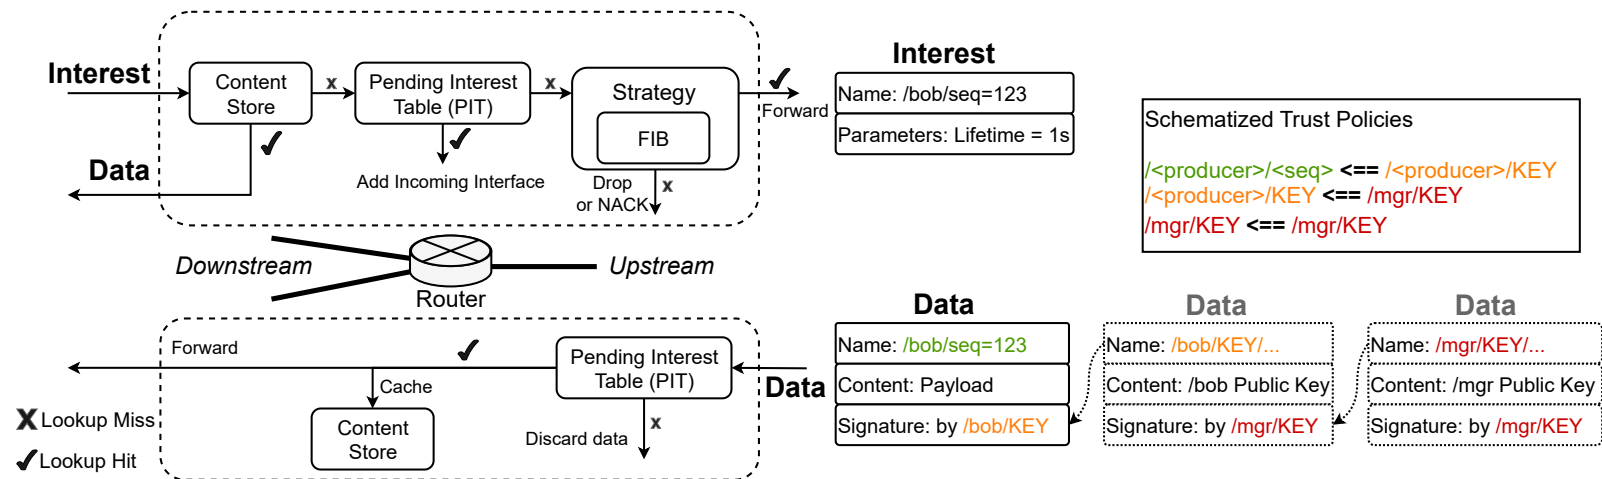
- Nodes estimate pairwise RTTs via background session messages.
- If a node hears a matching multicast RQ/RR from a peer while waiting, it suppresses its own pending traffic.
- Effectiveness depends strongly on network topology and member locations. [5]

# The Global Overhead of Local Loss

- Two key problems:
  - RQs and RRs identify data by name, but IP delivers them to multicast group addresses.
  - Routers are unaware of data names, preventing request aggregation and in-network data retrieval.
- Fundamental issue:
  - A mismatch between application's data-centric communication and IP's address-based multicast delivery model.

# Named Data Networking (NDN) [4]

- NDN integrates SRM/ALF's semantic naming directly into network routing and forwarding.
- Every Data packet is named, signed, and optionally encrypted, enabling receiver-driven retrieval and in-network caching.



Consumers: send Interests by data name  
 Network: forwards Interests toward data  
 Producers/caches: return matching data

# Built-In Data Multicast in NDN

- When multiple consumers request /E/seq=12
- Routers:
  - aggregate Interests
  - forward only one upstream request toward producer
- Result:
  - a multicast tree for data name during forwarding
  - no explicit multicast signaling


# Dataset State Synchronization (Sync) [6]

- NDN utilizes Sync to coordinate dataset states (similar to SRM session messages)
- The set of local data is encoded into an efficient representation, e.g., a digest, bloom filter, or state vector based on the data names.
- Nodes exchange their dataset state through Sync messages.
- If a node detects a missing data name, it sends an Interest to fetch the data



# Security

- SRM did not explicitly address communication security.
- In NDN, security becomes data-centric, instead of channel-centric.
  - Each entity undergoes a bootstrapping process to obtain its initial set of security parameters.
  - Certificates and trust schemas are named, secured data, and can be retrieved and updated.



# Summary

- Both SRM and NDN support reliable multiparty communication using receiver-driven loss recovery.
- Receivers may be in different reception states, so recovery must operate on missing data names rather than host addresses.
- NDN localizes recovery by forwarding and aggregating requests by data name and satisfying them from caches.
- In contrast, SRM relies on address-based IP multicast, forcing recovery traffic to be sent group-wide.
- NDN supports data-centric security. SRM doesn't address security.

# Discussion: CDN as a Workaround

- CDNs approximates data-centric delivery through application-layer overlay
  - caching
  - content retrieval by name
  
- But:
  - specialized infrastructures for paying customers
  - not general-purpose Internet support

# Discussion: Why HTTP is Not Enough

- “HTTP as the Narrow Waist”?
- HTTP names resources,
- BUT:
  - still fundamentally client-server
  - lacks native group communication
  - lacks receiver-driven repair
  - lacks network-layer data retrieval
- Result: application-level middleware complexity

# Discussion: Today's Multiparty Apps

- Typically implemented via:
  - full-mesh peer-to-peer connections
  - brokers: application-specific overlays that reimplement data dissemination, replication, and loss recovery on top of point-to-point transports
- Observation: **this repeated reinvention highlights the need for a general-purpose data-centric communication substrate**

# Looking forward

- Scalable multiparty communication requires aligning network delivery with data-centric semantics.
- NDN realizes this alignment
  - built-in multicast data delivery
  - localized loss recovery
  - secure named data

=> **scalable and secure multiparty communication**
- Future direction: data-centric networking as a general communication substrate.

# NDN Deployment

- NDN Testbed
  - operational for over a decade
  - multicontinental overlay
- Example Applications:
  - DRON (Docker Registry over NDN)
    - efficient Docker blob distribution
    - one retrieval → many consumers
  - Ownly collaborative editor
    - decentralized collaborative editing
    - secure synchronized updates

# References

1. Deering, S., and Cheriton, D. "Multicast Routing in Datagram Internetworks and Extended LANs." *ACM Transactions on Computer Systems (TOCS)*, 1990.
2. S. Floyd, V. Jacobson, S. McCanne, C.-G. Liu, and L. Zhang, "A reliable multicast framework for light-weight sessions and application level framing," in *Proceedings of SIGCOMM*, 1995.
3. D. D. Clark and D. L. Tennenhouse, "Architectural considerations for a new generation of protocols," in *Proceedings of SIGCOMM*, 1990.
4. L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, K. Claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang, "Named data networking," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 66–73, 2014.
5. C.-G. Liu, D. Estrin, S. Shenker, and L. Zhang, "Recovery Timer Adaptation in SRM," in *Proceedings of the International Conference on Computer Communication*, 1999.
6. P. Moll, V. Patil, L. Wang, and L. Zhang, "Sok: The evolution of distributed dataset synchronization solutions in ndn," in *Proceedings of ICN*, 2022

# Q&A