Performance Evaluation of Named Data Networking Forwarding Daemon (NFD)

Susmit Shannigrahi, Chengyu Fan, Christos Papadopoulos

Goal

- Evaluate NFD performance an application's perspective
  - Establish expectations for applications
  - Find the limitations and bottlenecks of the current NDN codebase

- Non-goals
  - Per-module performance profiling

Evaluation Topology

Baseline Network Performance

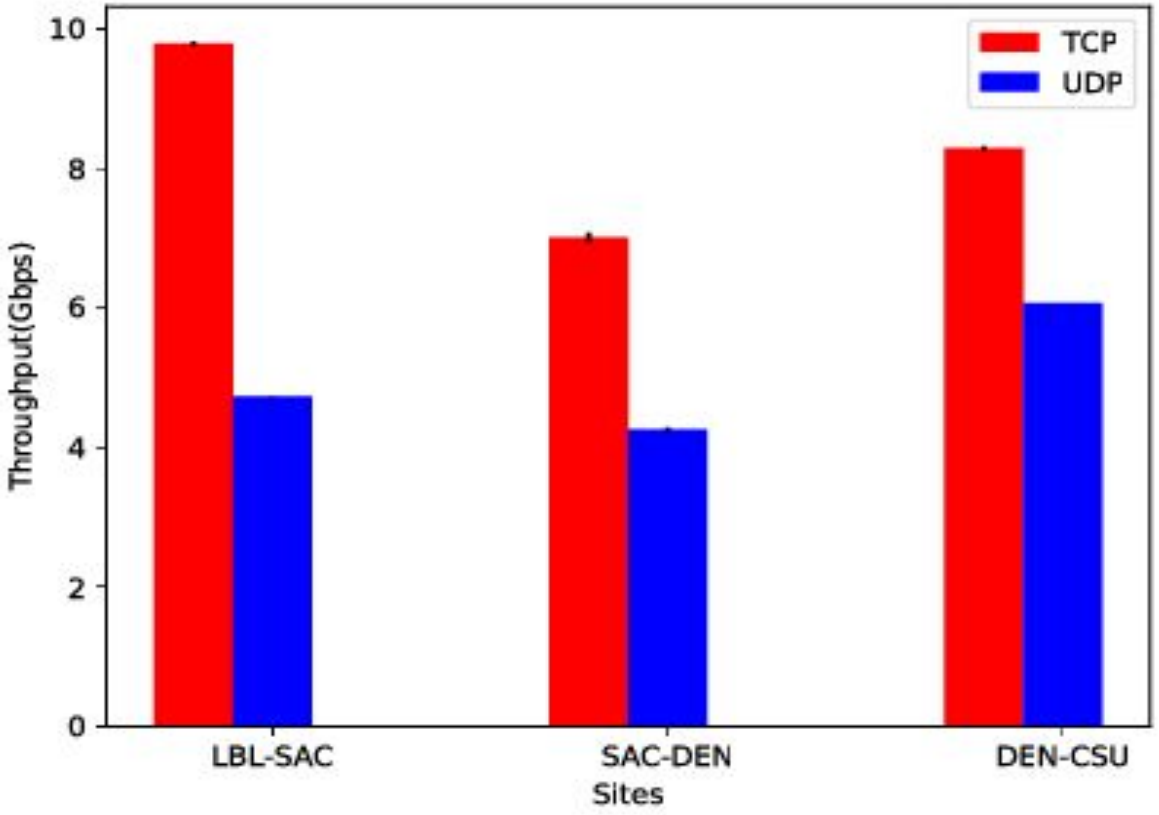

NFD Performance

- Single NFD
- ndncatchunks/putchunks
- Congestion control = off

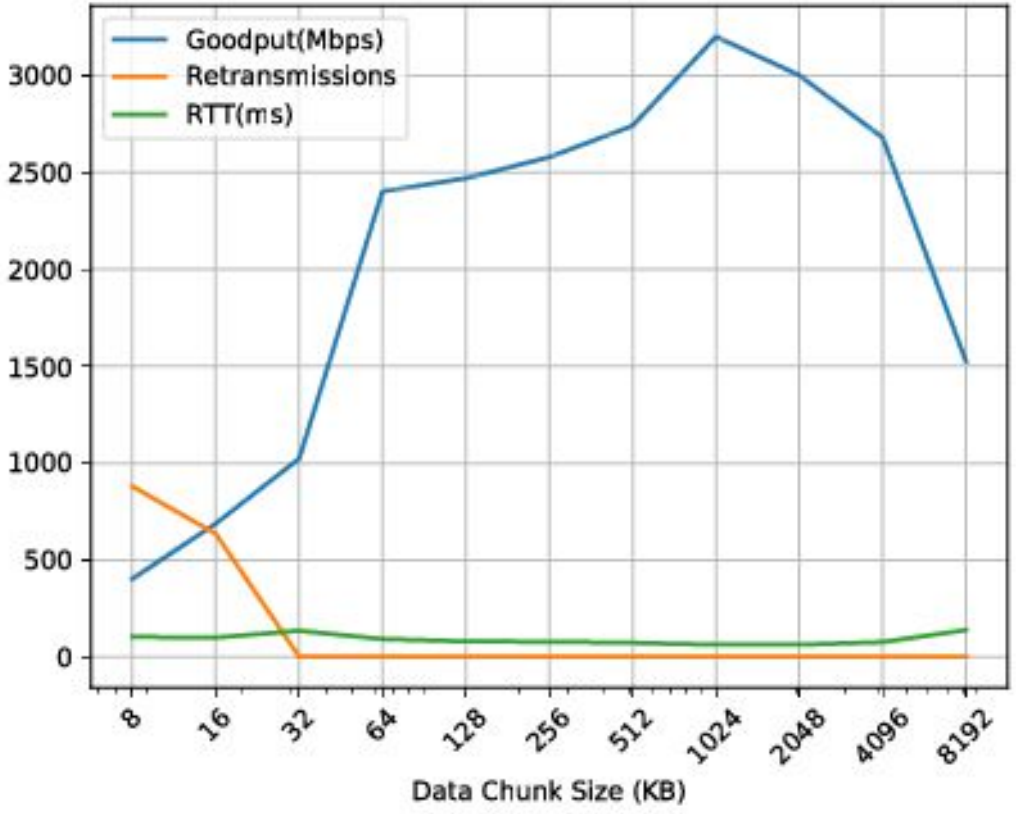

Better With Congestion Control

- Very few congestion markings, large improvements
- No packet loss

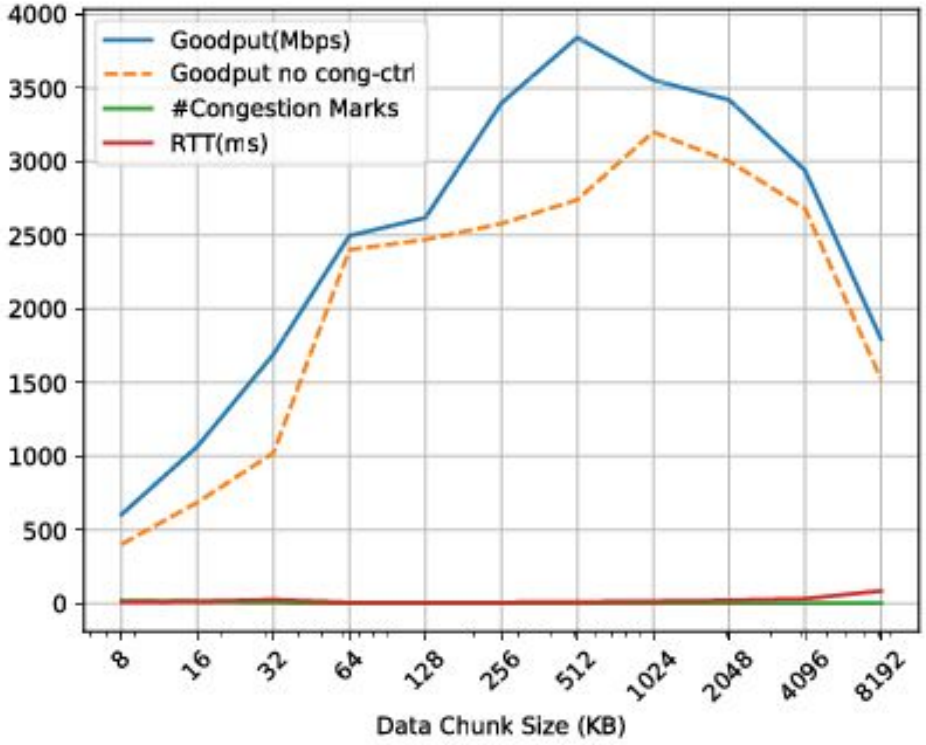

Caching slows down NFD significantly

- Candidate for a separate thread?

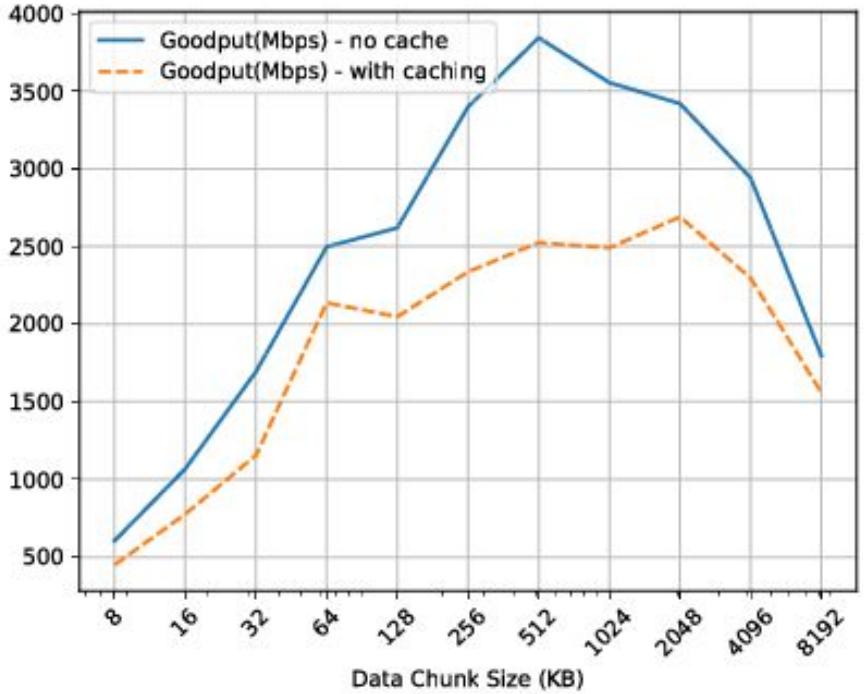

The results above are performance upper limits

- No network delay
- No caching

Performance in WAN

Performance over TCP tunnels

- Maxes out around 3Gbps
- Note that performance flattens for higher chunk sizes

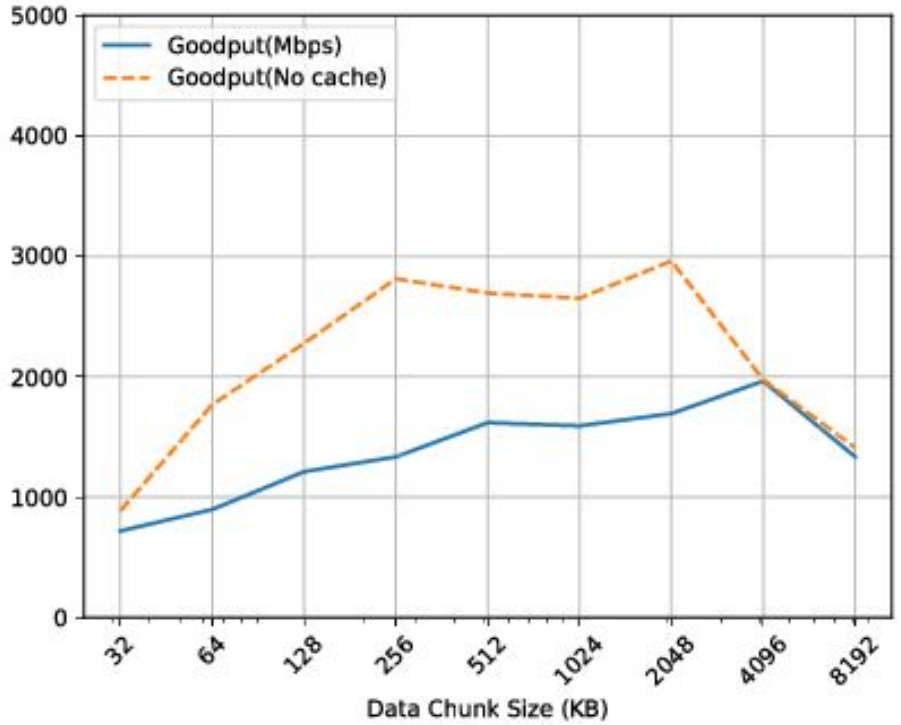

Performance over UDP tunnel

- Much lower than TCP
- Max segment size = 64K
- Chunk a Data Packet over multiple UDP packets?
  - Currently not possible

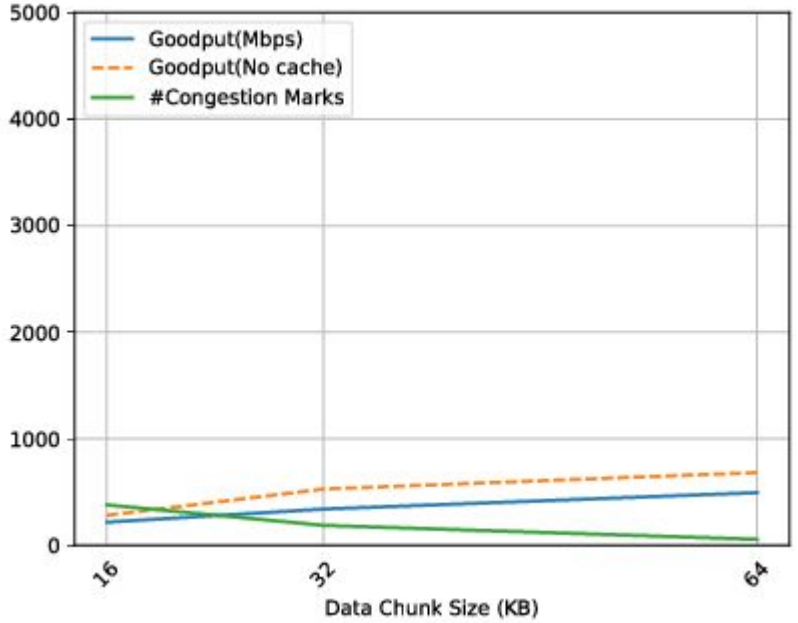

Performance over Ethernet Tunnel

- Possible to communicate without IP, but much slower than TCP/UDP
- ~150Mbps with 32K segments
- Possible bug: https://redmine.named-data.net/issues/4479

Throughput vs Number of Hops

- One, two, and three NFDs
- Decreases significantly with additional hops
- TCP tunnels can also attributes to performance loss

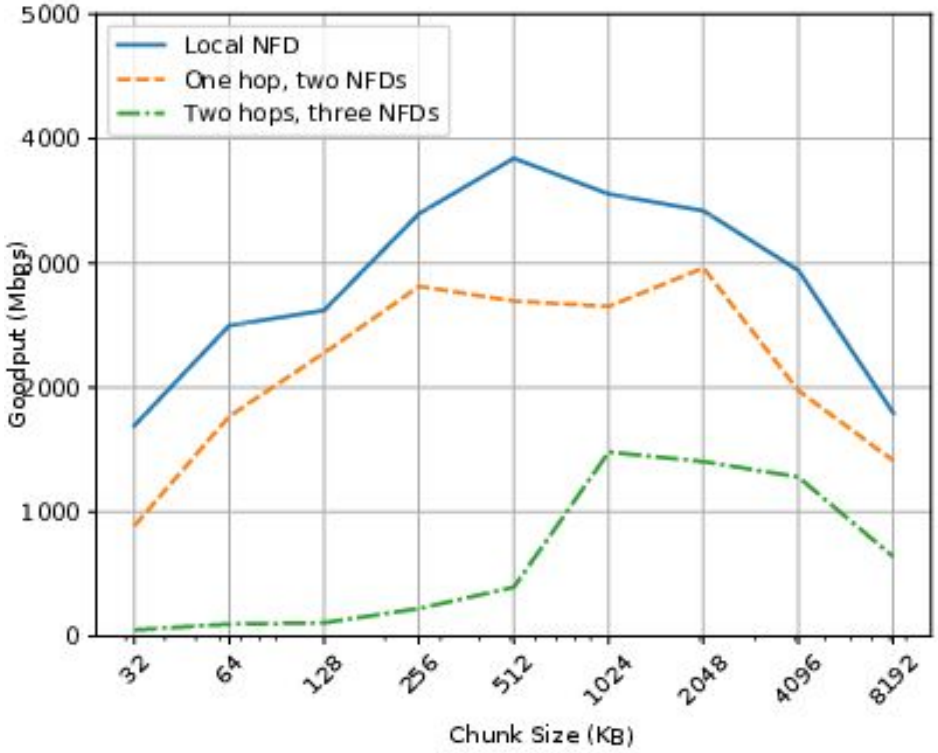

Receiver Side Processing Delay

- Higher with smaller chunk sizes
- Almost constant for higher chunk sizes
  - Queuing delay?

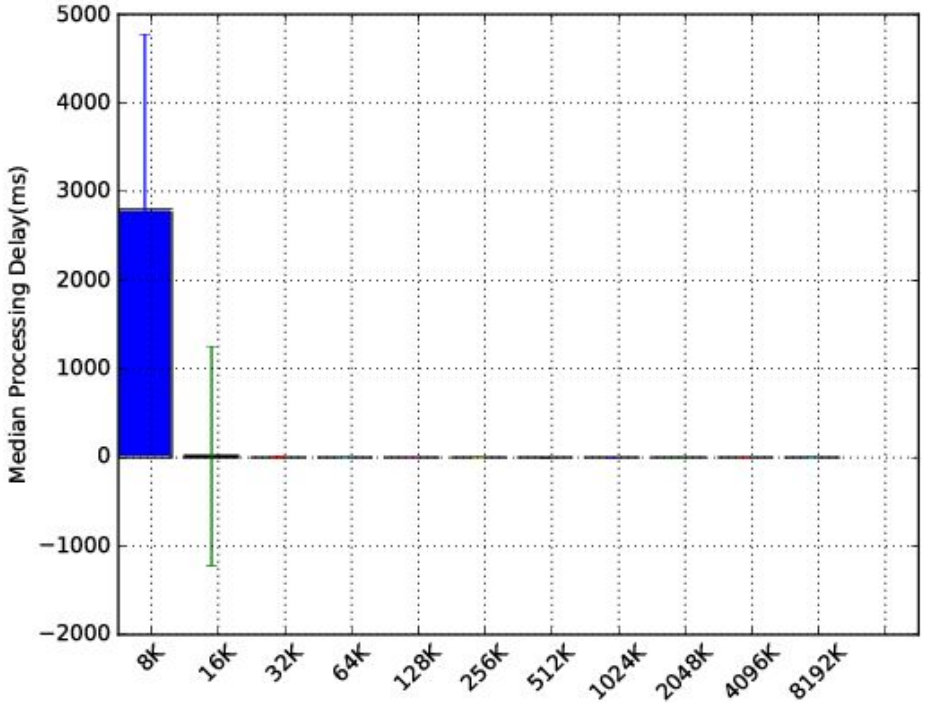

Server Side Processing Delay