Postcards from the Edge: A Cache and Forward Architecture for the Future Internet

WINLAB
WIRELESS INFORMATION NETWORK LABORATORY

Sanjoy Paul (Presenter), Roy Yates, Dipankar Raychoudhuri

Wireless Information Network Laboratory (WINLAB), Rutgers University

And

Jim Kurose

University of Massachusetts, Amherst

Wireless and Video will be the *key* drivers for Future Internet

- Wireless has overtaken wired as the primary mode of connectivity to Internet
  - **500M Internet-connected server/PC's vs. >2B cell-phones; 400M Internet capable**
  - **New types of wireless data devices: Blackberry, PDA, iPoD**
  - **Sensor deployment just starting: ~5-10B units by 2015**

- Variety of Wireless Network Usage Scenarios would impact Future Internet design
  - **Mobile Data applications**
  - **Multi-Hop Wireless Mesh networks**
  - **Pervasive Sensor networks**
  - **Vehicular networks**
  - **Cognitive Radio networks**
  - **Info-stations**

- Content (video) distribution and sharing is the most dominant application in terms of bw consumption on the Internet today and may be for the foreseeable future
  - **YouTube hosts over 6M videos, growing at about 20% per month**
  - **Requires 45 TB of storage, several million $ worth of bw/month to transmit**
  - **Total time spent watching YouTube videos since it started last year is 9,305 yrs!**
  - **Most popular items get an especially large percentage of the traffic**
  - **Cameras are everywhere (cellphones, video surveillance, handycams etc.)**
  - **Hundreds of TV channels on the Internet and growing**

Drivers for Future Internet design

Wireless Network Usage Scenarios	Characteristics	Impact/Requirements
Mobile Data Wireless Mesh	Time-varying link (capacity, error rate)	Opportunistic transport Cross-layer protocol design
Vehicular Network Info-station	Intermittent connectivity, High bw	Opportunistic transport Cache and Carry
Sensor Networks	Low CPU/Memory/Power Intermittent connectivity, Low bw	Opportunistic transport, Computation at nodes, Data driven networking

Other Trends	Characteristics	Impact/Requirements
Proliferation of video (YouTube)	High bandwidth and storage Popularity of some clips	Efficient Transport (Caching)
Sharing of content (Myspace, Facebook)	Content (music, photos, video) is bulk of traffic	Efficient Sharing (Caching)

Caching and Opportunistic Transport are emerging as key themes

WINLAB WIRELESS INFORMATION NETWORK LABORATORY

Basic Technology Trend:
Moore's Law Enablers of New Network Architecture

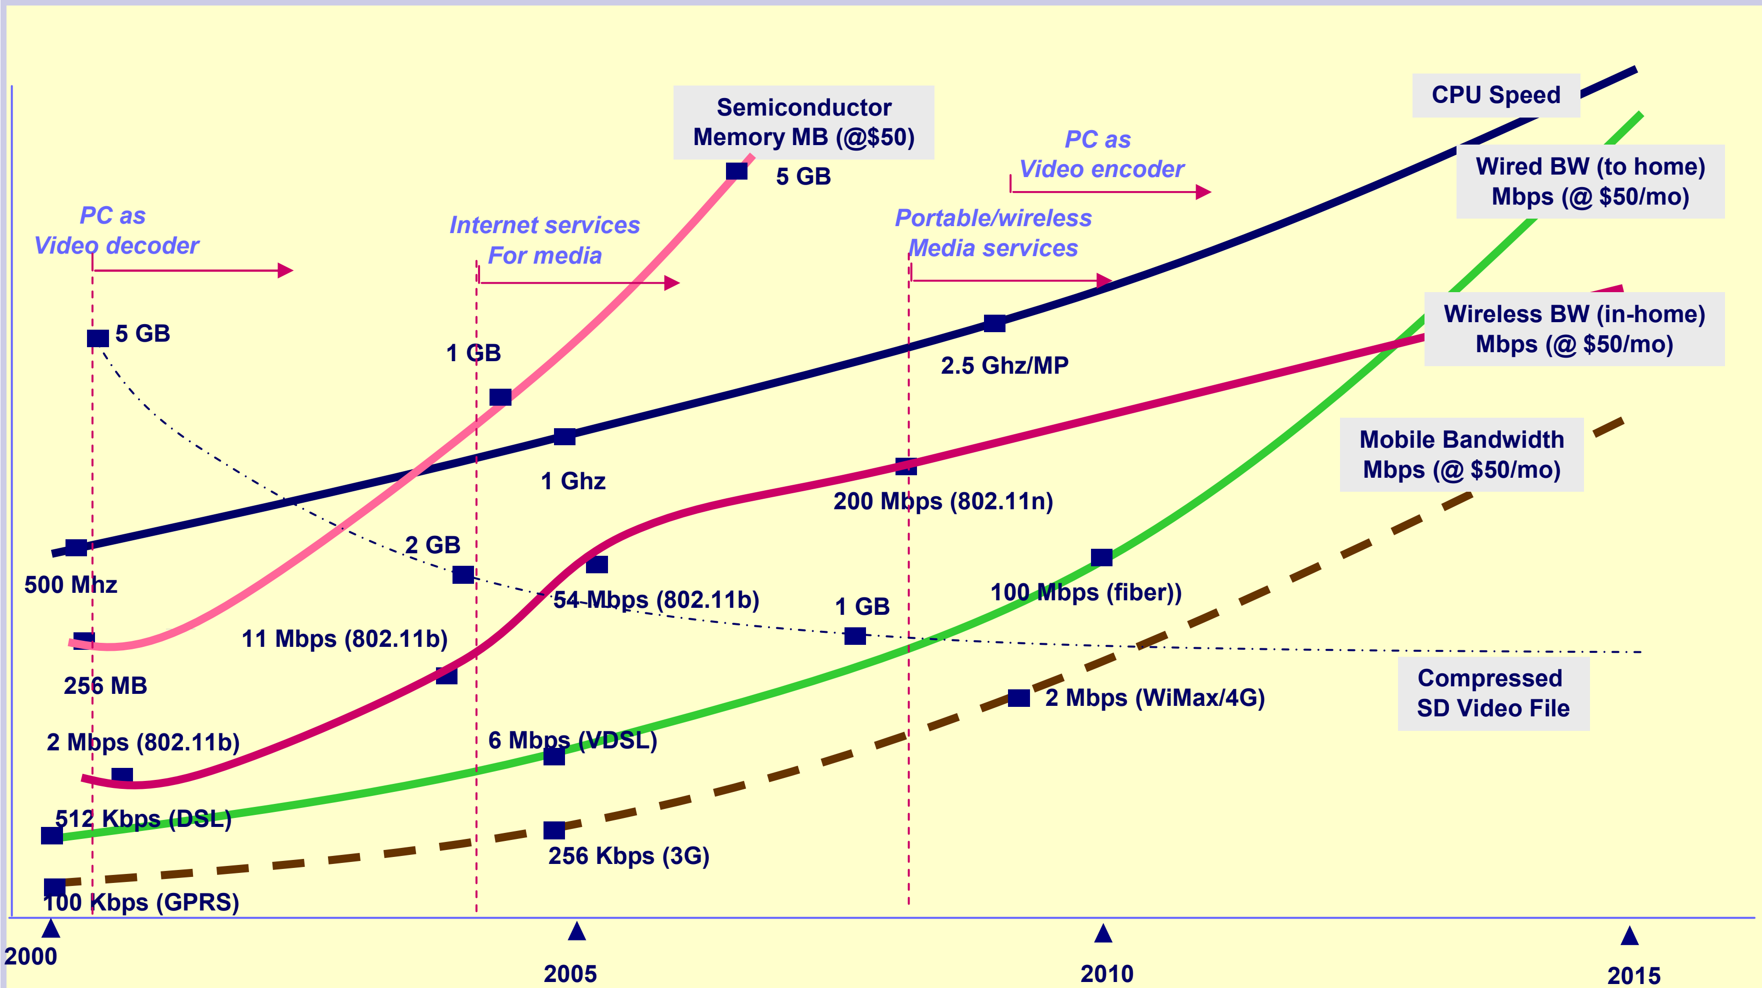

Semiconductor Memory for $50: 256MB in 2000; 1GB in 2004; 5GB in 2006

CPU: 500Mhz (2000); 2.5 GHz (2006); Cost of storage: $5/GB (2000); $0.50/GB (2006);

WINLAB WIRELESS INFORMATION NETWORK LABORATORY

Storage is cheap and plenty of Processing available at low cost

Cache aNd Forward (CNF) Architecture

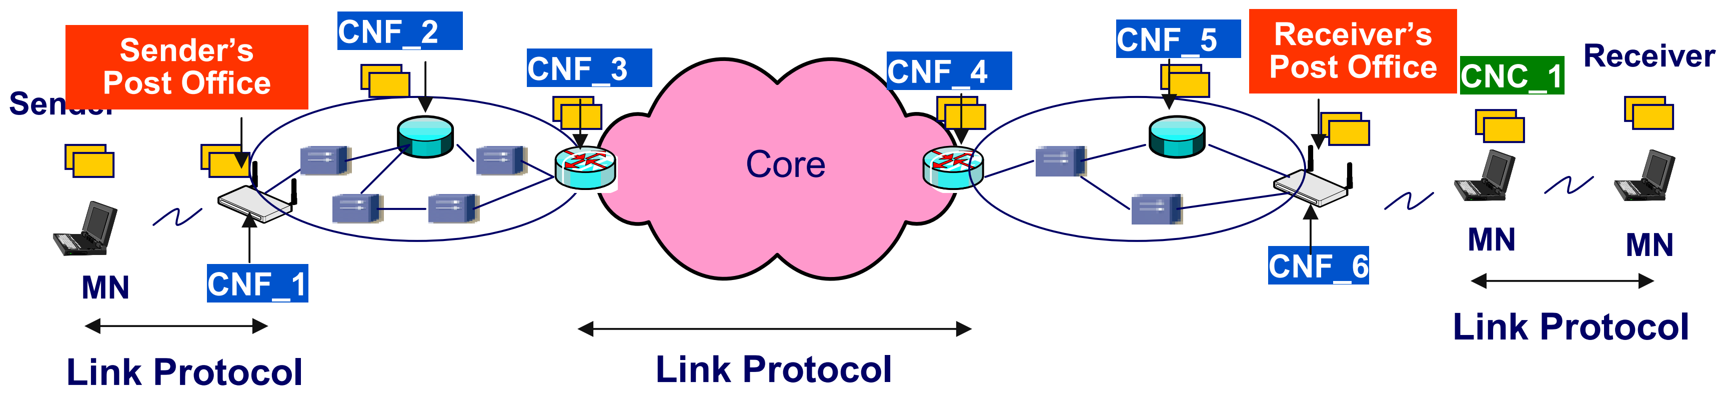

Push some complexity into the network and make the end-to-end transport protocol simpler.

This is justified because of plummeting cost of storage and the exponential increase in CPU speed and processing capacity

✓Network of CNF r... ward of Large files

✓Link Protocol be...

✓Routing Protoco... ", "content", "location" etc.

✓Two step Forwar...

- ✓Forwarding of query (based on content "handle");
- ✓Forwarding of response (based on destination's "handle"/address)

✓End-to-end connectivity not needed: only CNF-to-CNF (Opportunistic Transport)

✓File stored in Post Office until retrieved by mobile (Disconnected Operation)

✓In-network caching handles distribution of Popular Files very efficiently

Cache aNd Forward (CNF) Architecture: Overlay on IP?

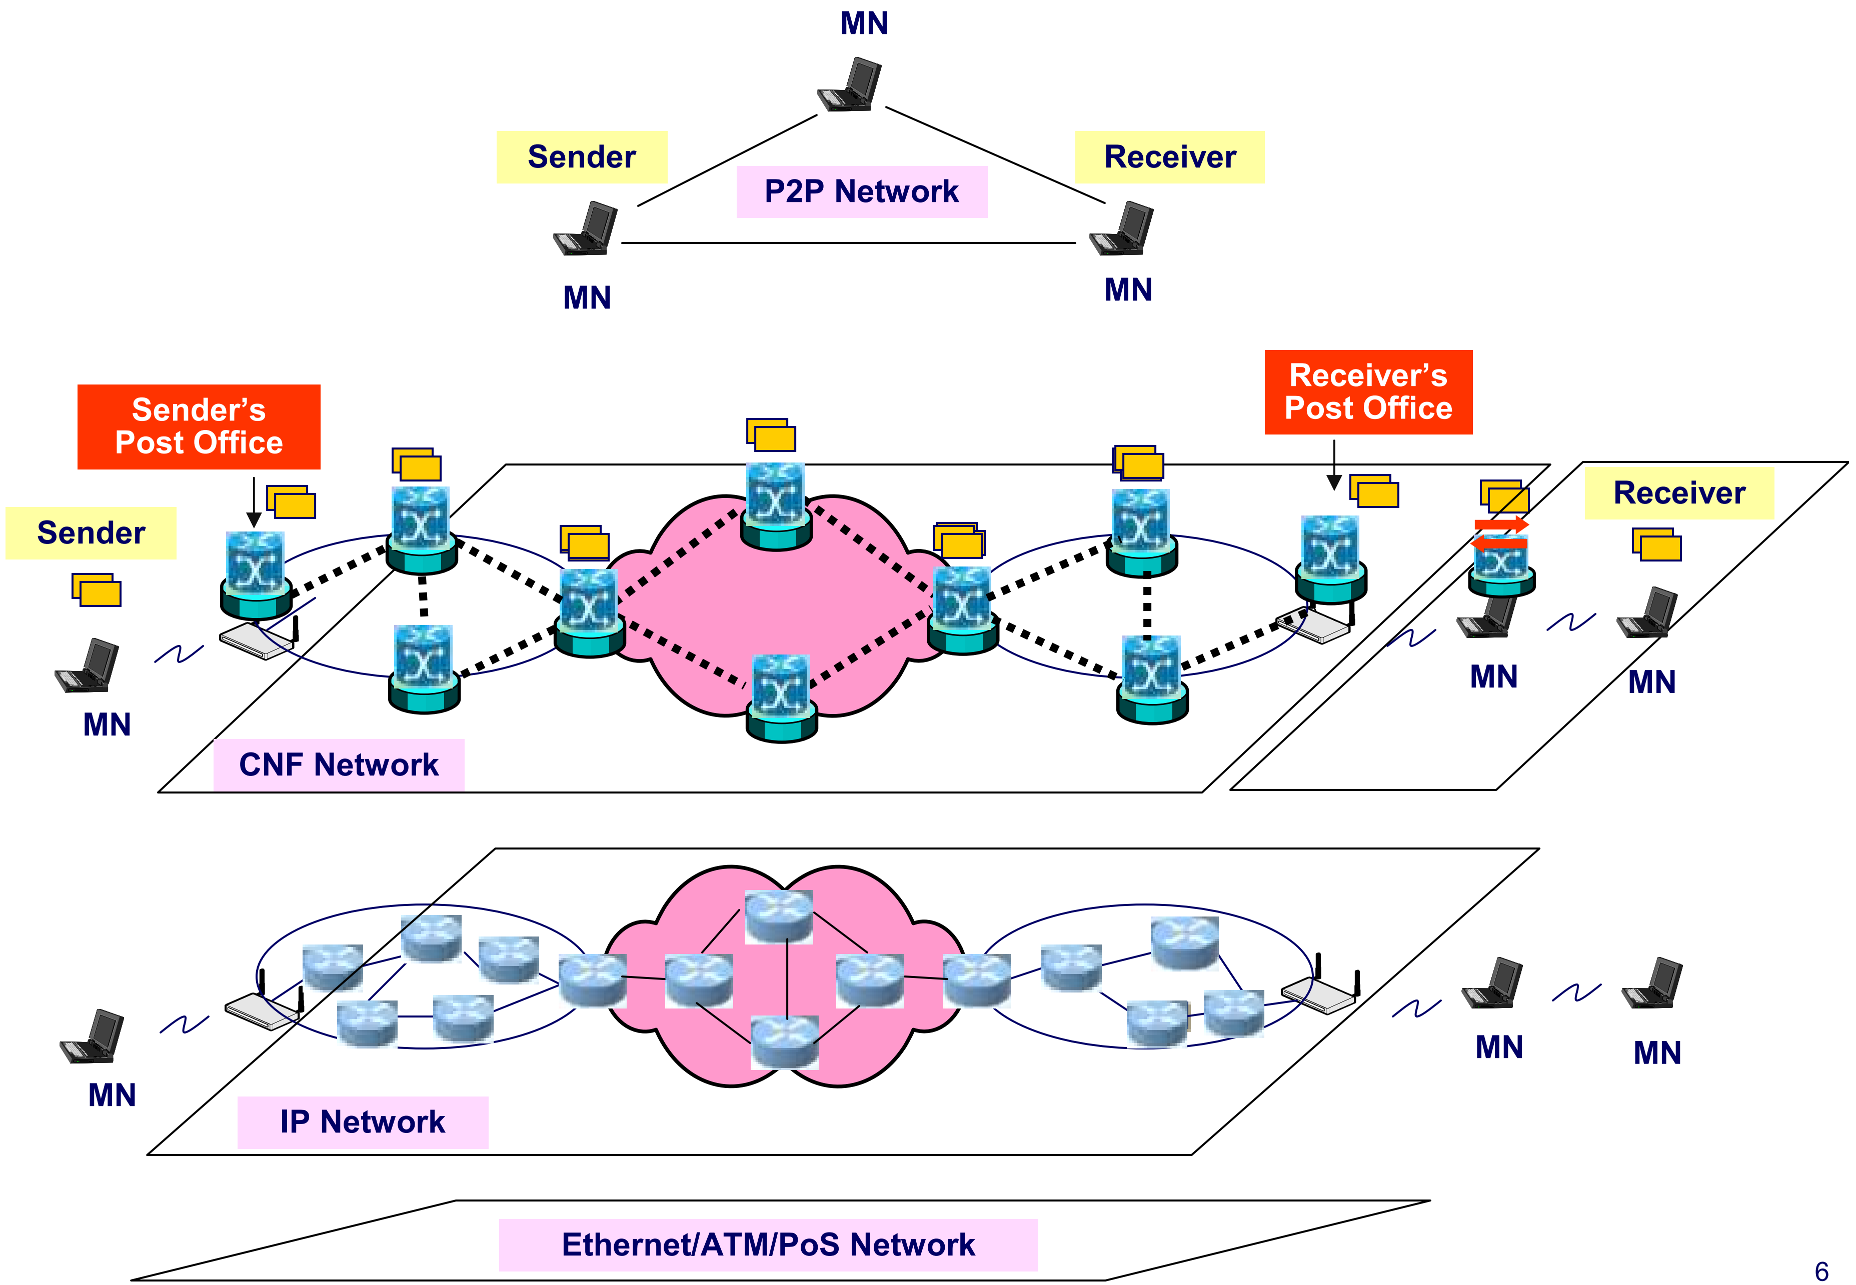

WINLAB
WIRELESS INFORMATION NETWORK LABORATORY

Postcards Architecture: Protocol Stack

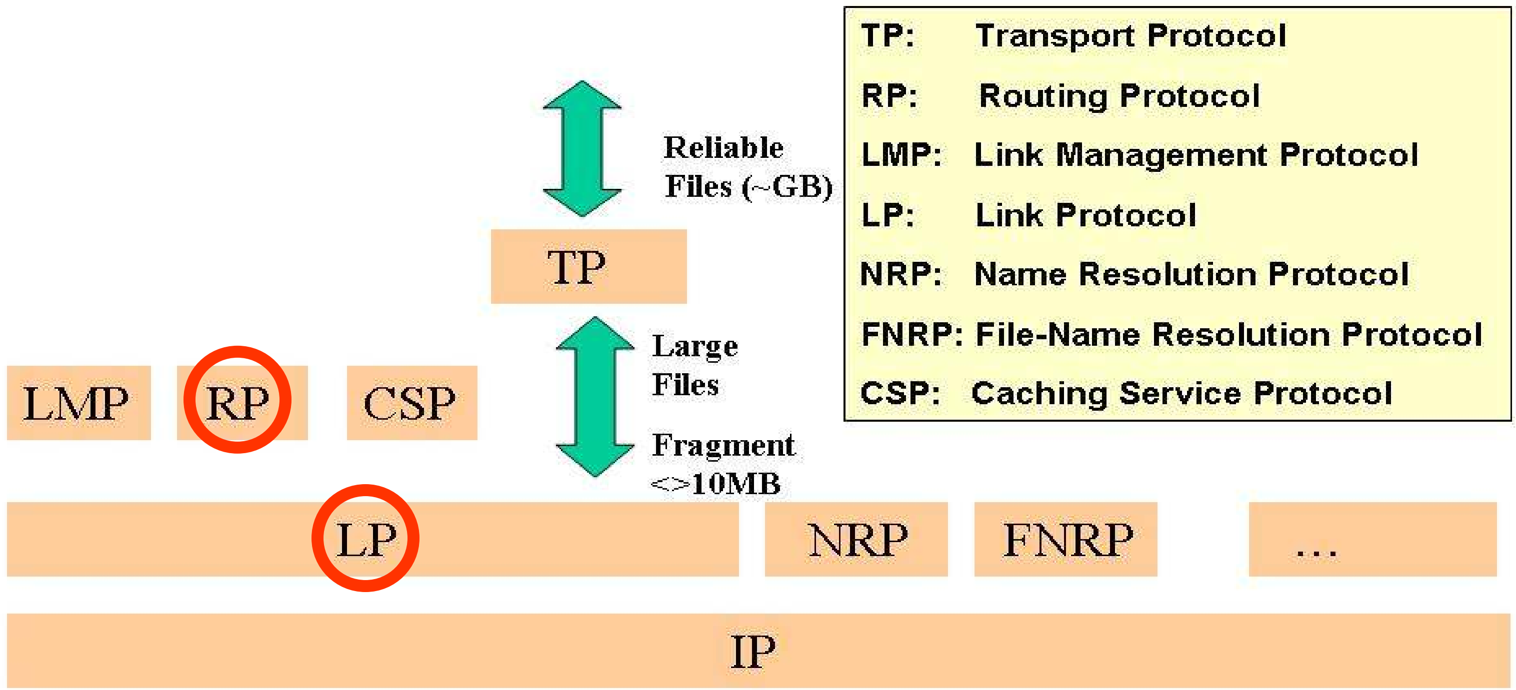

CNF Architecture: Link Protocol

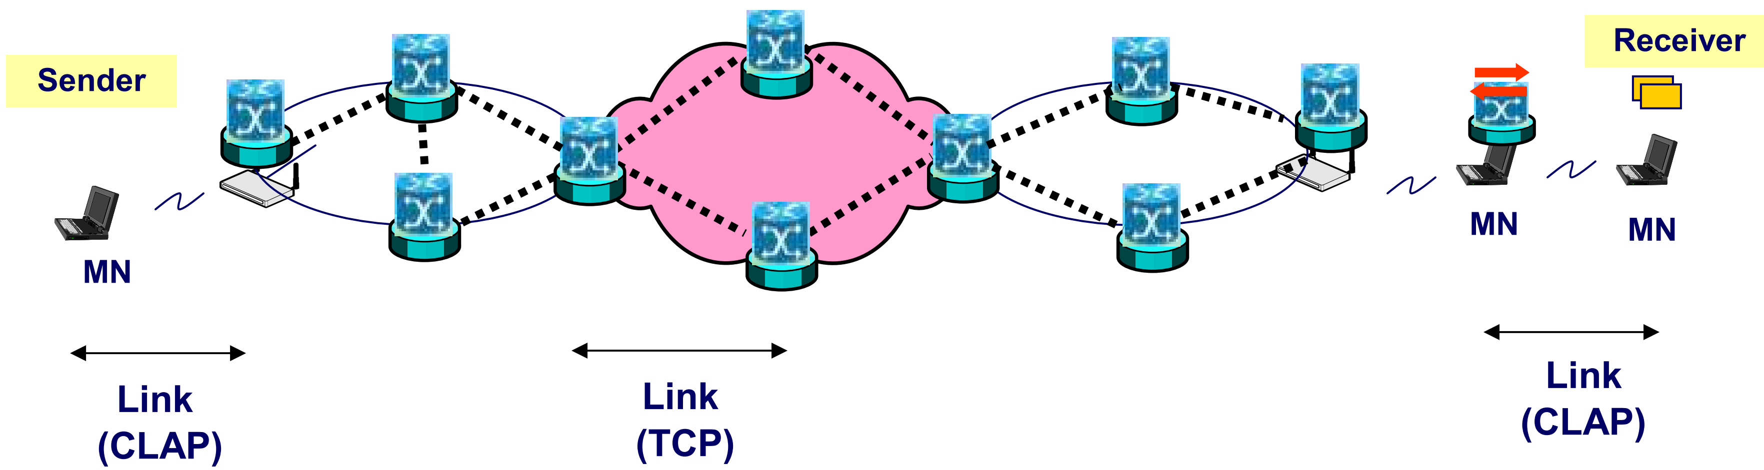

- Link session Protocol to exchange link characteristics and decide on choice of parameters
- Link Protocol
  - TCP variants
  - Cross-Layer Aware Transport Protocol (CLAP)
  - etc.

Cross Layer Aware Transport Protocol (CLAP) vs. TCP: Single hop, single flow

Thanks to Sumathi Gopal

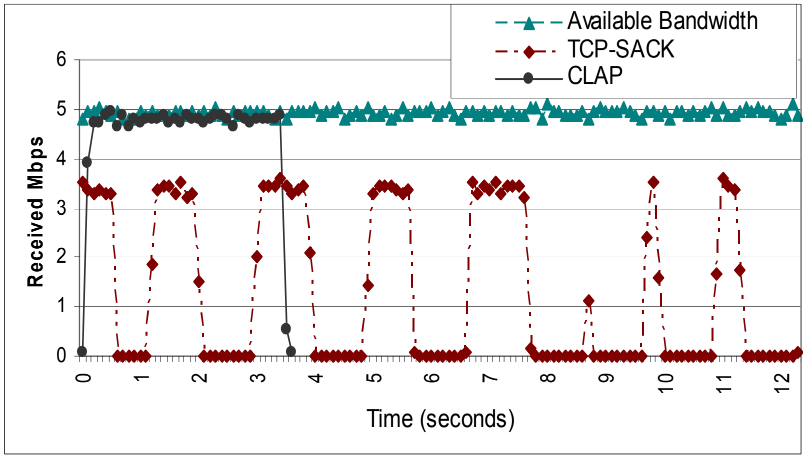

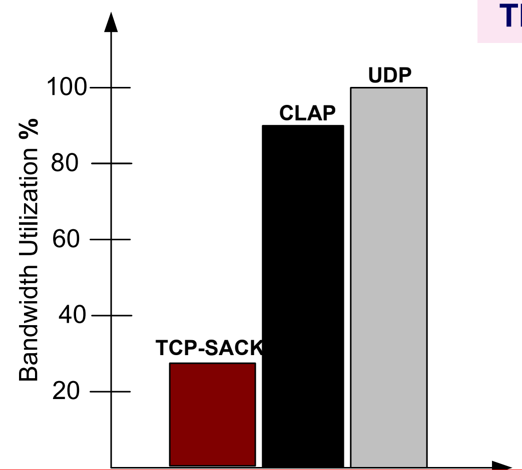

No Error Case

When no errors TCP suffers timeouts due to *self-interference*.

CLAP gains 300% over TCP-SACK, making far better utilization of link bandwidth

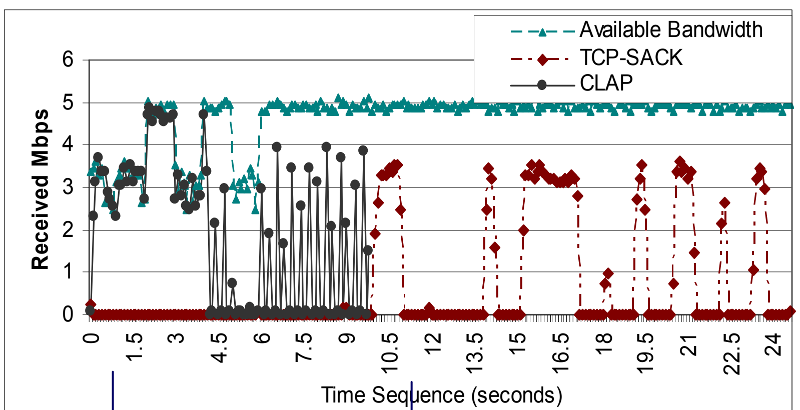

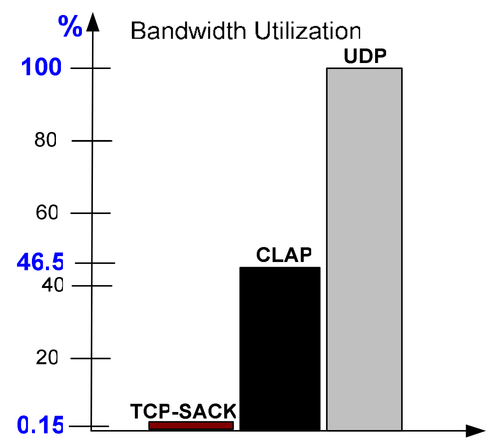

Time-varying Error Case

TCP shuts down operation in the presence of time-varying link errors.

TCP is slow to adapt to available bandwidth due to exponential backoff

CLAP operates despite errors and bw fluctuations; achieves immense (4500%) gains over TCP-SACK

CNF Architecture: Routing

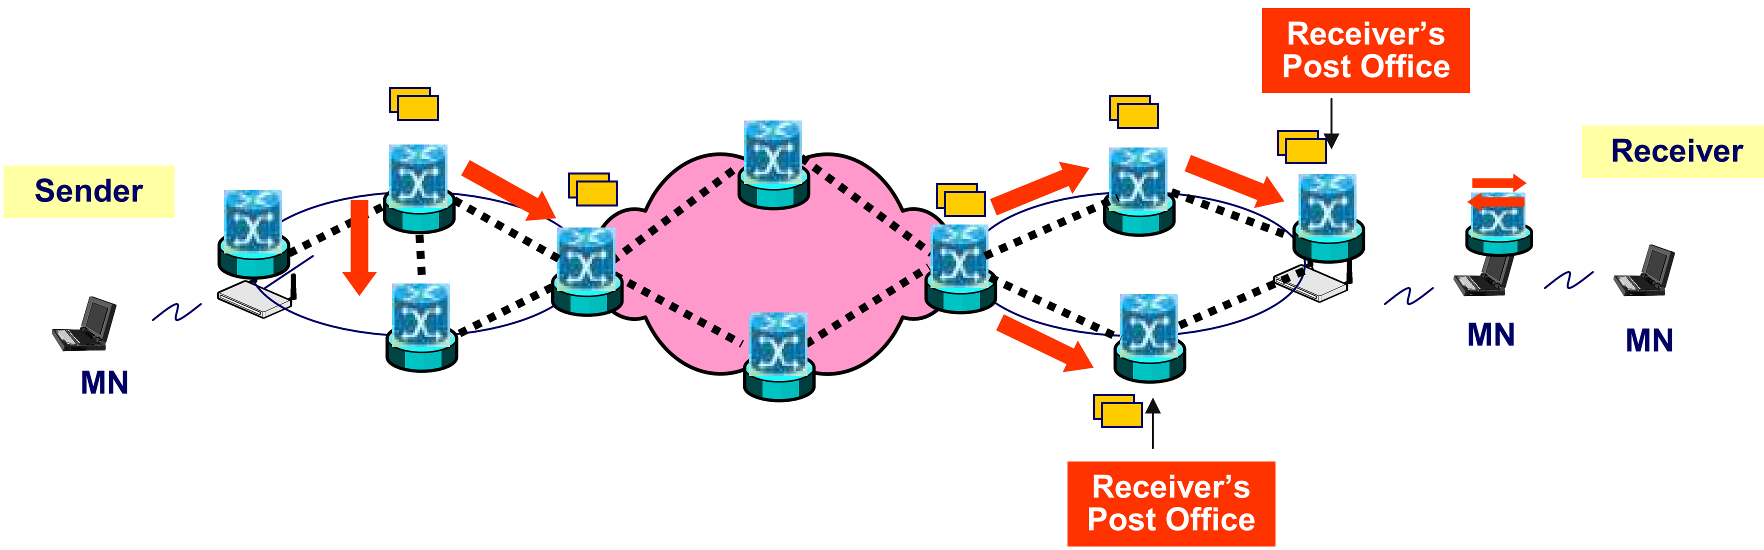

- Routing is done hop-by-hop towards the Receiver's Post Office(s)
- A given file may be sent to multiple Post Offices (PO)
- Routing from the PO to a Mobile Node (Receiver) can happen in many ways:
  - PO notifying the MN
  - MN checking with the PO (Polling)
  - Multi-hop delivery of the file from PO to MN
  - Can leverage vehicles on the highway for delivery

CNF Architecture: Query Routing and Content Delivery

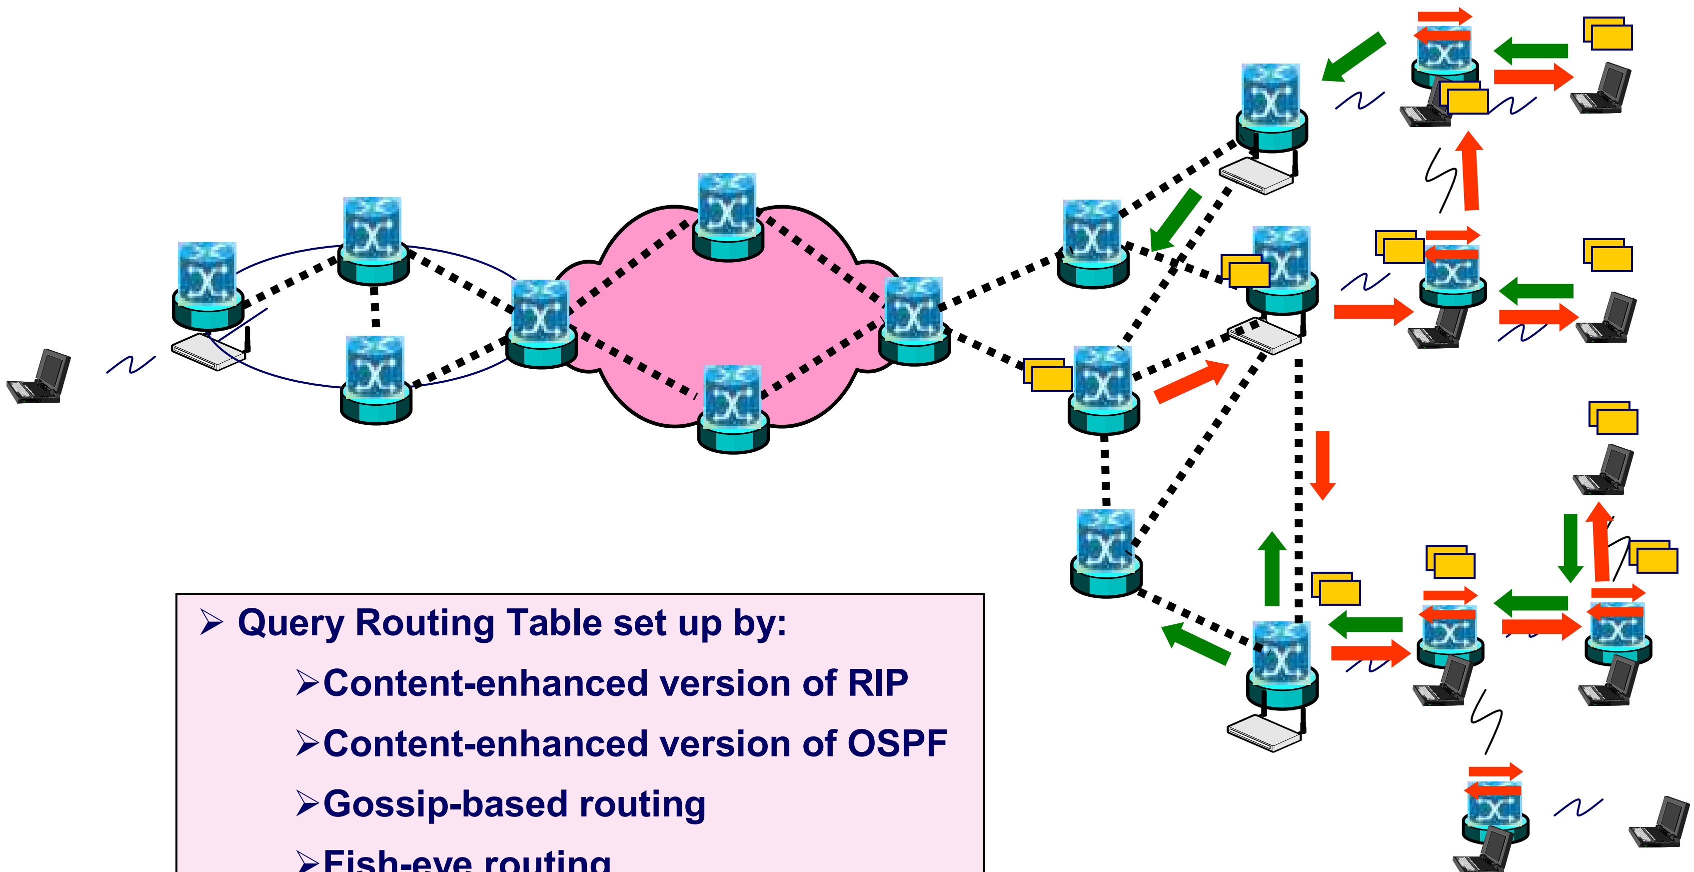

- Query Routing Table set up by:
  - Content-enhanced version of RIP
  - Content-enhanced version of OSPF
  - Gossip-based routing
  - Fish-eye routing
- Content-enhanced version of BGP

WINLAB

Preliminary Evaluation: Simulation Results

* Thanks to Umut

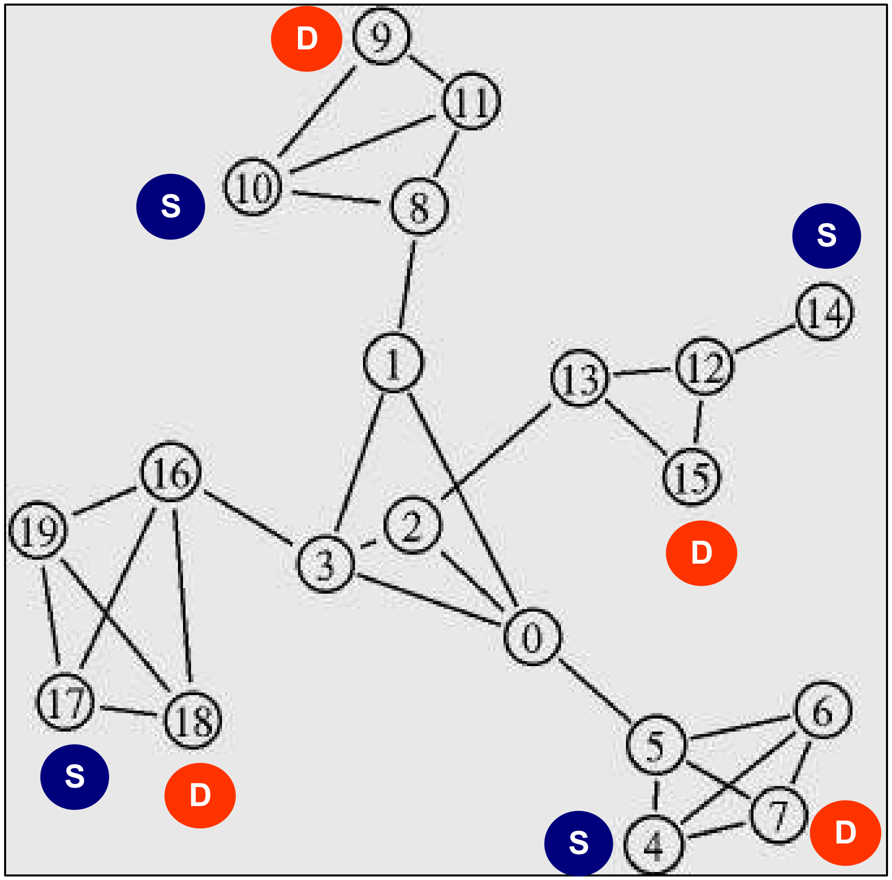

- GT-ITM's transit-stub model
- Transit-transit links
  - 10Gbps bandwidth
  - Uniform (20,23)msec delay
- Transit-stub links
  - 155Mbps bandwidth
  - Uniform (20,23)msec delay
- Stub-stub links
  - 1Gbps bandwidth
  - Uniform (2,11)msec delay
- 15 source-destination pairs

- ON State: Exponentially distributed with mean 100sec
- OFF State: Exponentially distributed with 3600sec
- An ON state is always followed by an OFF state (vice versa)
- Initial state of the source is ON
- During ON State: Fixed file size (50MB), Poisson arrivals

Throughput, Delay, Loss in CNF Architecture

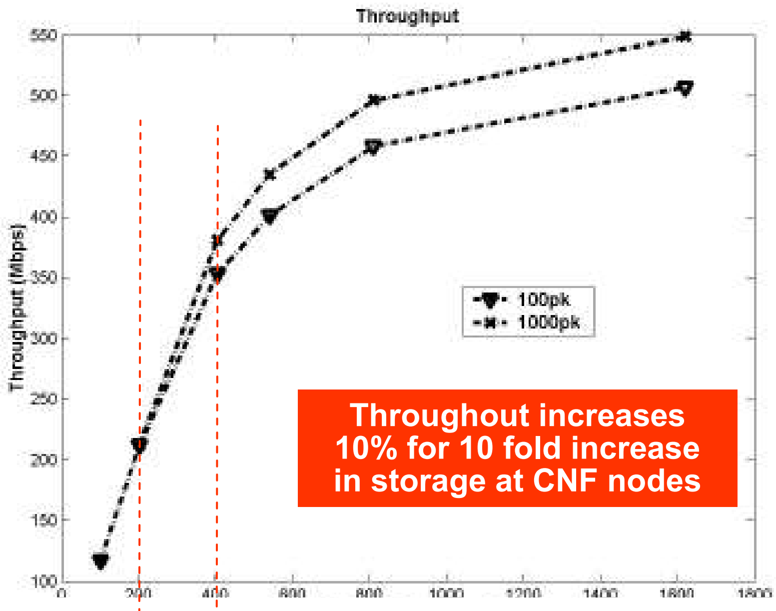

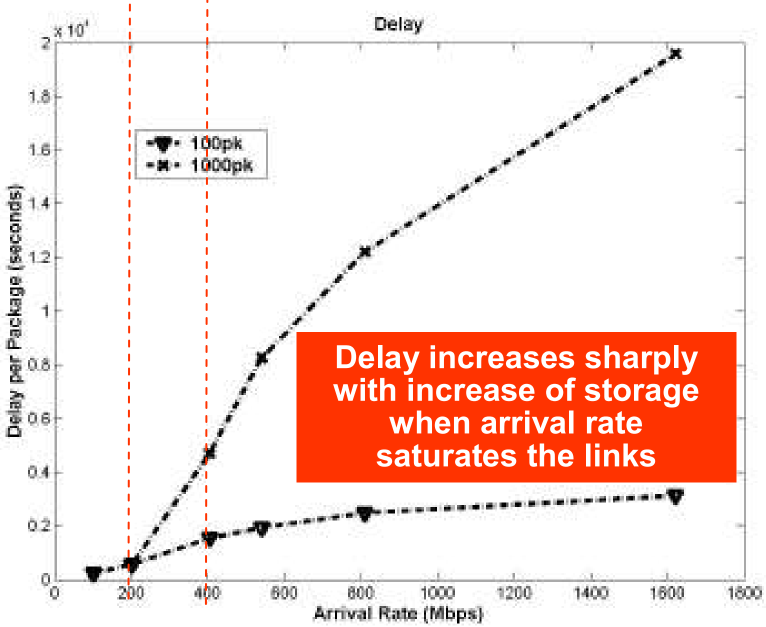

$$\frac{\text{Packages Dropped}}{\text{Packages Sent}}$$

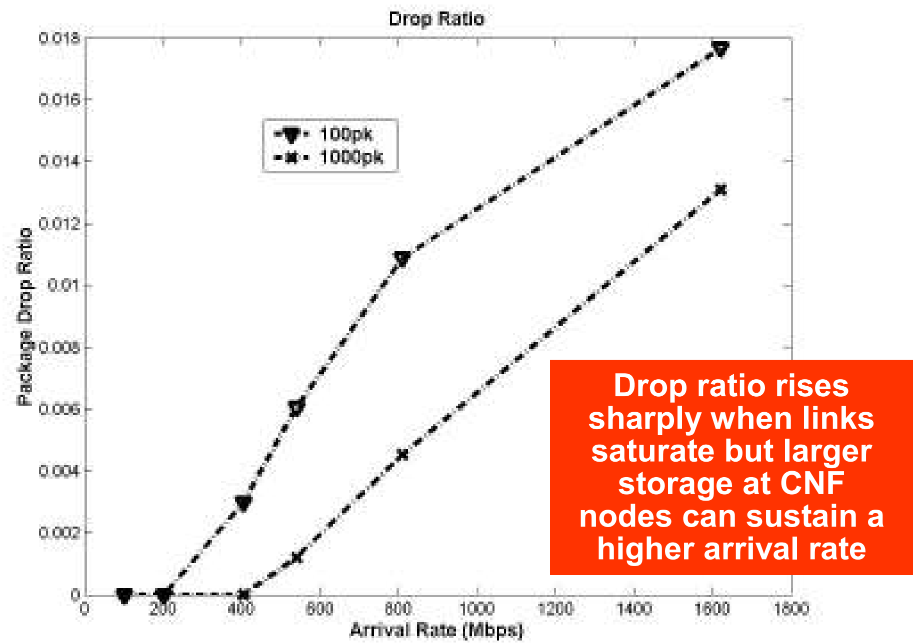

CNF vs. TCP

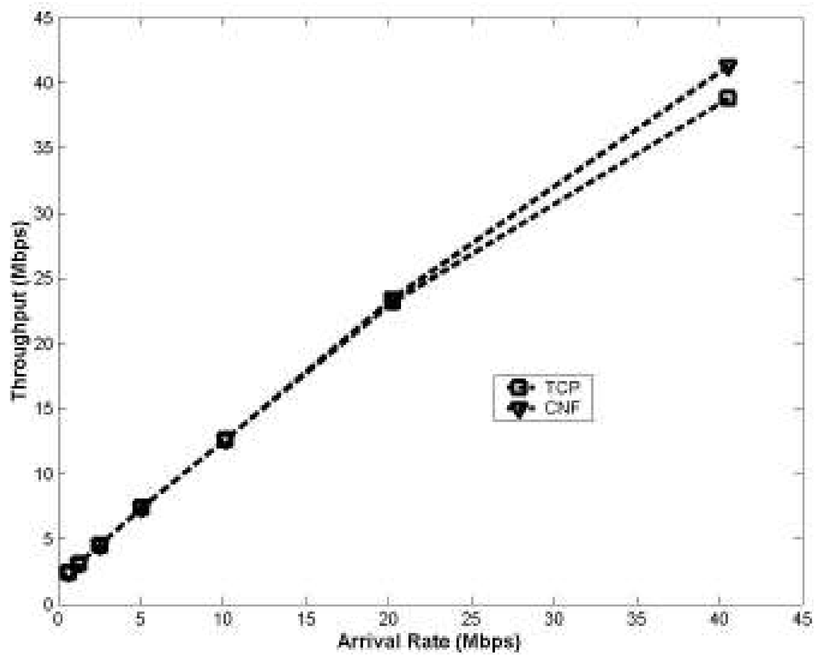

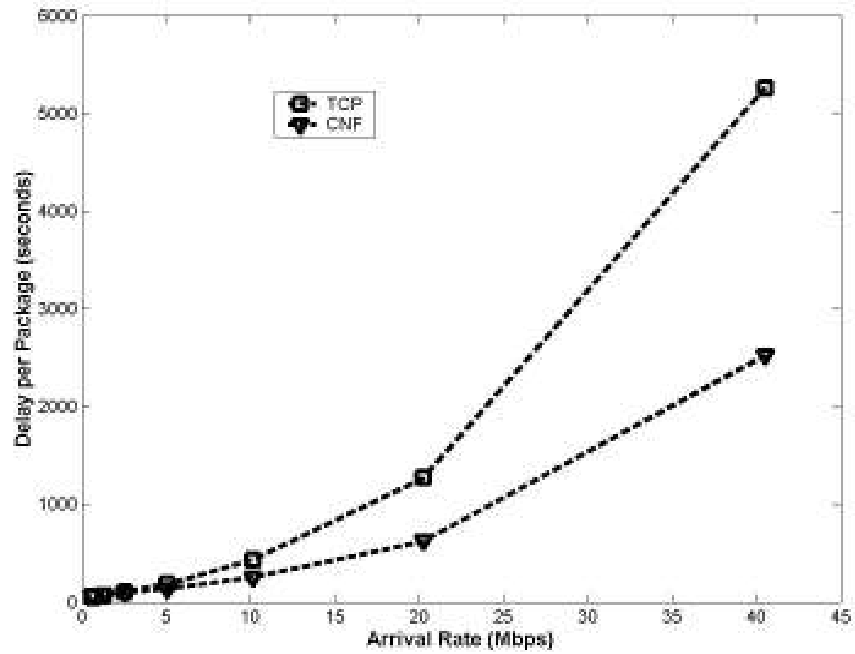

CNF starts to outperform TCP in terms of throughput for On-OFF traffic as the overall arrival rate increases

End-to-end delay for CNF is lower than that for TCP for On-OFF traffic and the difference increases as the overall arrival rate increases

WINLAB WIRELESS INFORMATION NETWORK LABORATORY

Postcards Experiments in GENI (Experiments at scale)

- **Compare Hop-by-Hop file transfer vs. TCP's pipelining**
- **Compare Routing protocols**
- **Compare Uniform vs. Heterogeneous Link Protocols**
- **…**

Summary of Postcards Architecture

- **Cache and Forward Architecture**
  - Exploits decreasing cost and increasing capacity of storage devices
  - Provides a unified and efficient transport service to the end hosts
  - **End-hosts may be wired, wireless, static, mobile, and/or intermittently connected, and either resource rich or poor**

- **Key Architectural Concepts:**
  - **Reliable hop-by-hop transport of large files**
    - Classical store and forward of large files with in-network storage + reliable link layer
  - **Push-Pull architecture**
    - For mobile nodes the architecture enables “opportunistic” push-pull delivery of files, both to and from wired network
  - **Enhanced Naming**
    - Name Resolution Protocol: Resolve a destination’s name into destination’s Post Office (s) by leveraging location information of mobile nodes
    - File Name Resolution Protocol: Resolve a file’s name into a unique “handle”
  - **In-network Distributed Caching**
    - Distributed caching of popular content will occur throughout the network, thus making peer-2-peer file sharing a first-class service

Summary of Benefits of Postcards Architecture

- **Efficient multi-hop wireless transmission:**
  - Enables "opportunistic" transmission
  - Avoids "self interference"
- **Facilitates cache-and-carry to increase capacity in mobile scenarios:**
  - Mobile info-station (bus/train/airplanes) can use physical motion to greatly reduce distances over which radio transmission must take place
- **Unified routing solution for wired and wireless networks**
  - Potential diversity of routes is increased because an end-to-end real-time connection is no longer maintained
- **Makes file sharing a first-class service**
  - Can avoid multiplicity of protocols and architectures to achieve the same goal: P2P file transfer and can provide a TCP-like interface for P2P

Perspective on CNF: CNF adopts just the right concepts from P2P, Overlay and DTN Networks

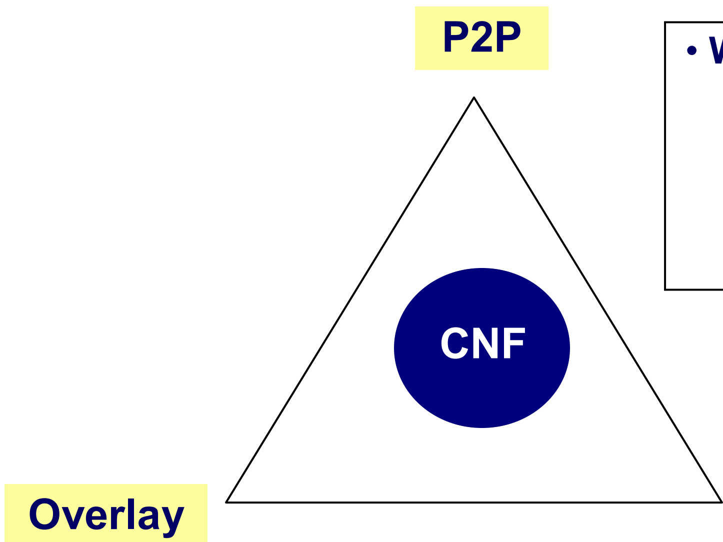

- **Without P2P features in CNF**
  - **CNF will perform no worse than current Internet**
  - **Content cannot be queried using a location-independent content "handle" and retrieved from a nearby CNF node rather content will have to retrieved based on "location-dependent" url that identifies the desired content**

- **Without Overlay features in CNF**
  - **CNF will perform no worse than current Internet**
  - **Content cannot be retrieved from a nearby CNF node, and opportunistic transport by alternative routing cannot be leveraged**
  - **Without Post Office, delivery from/to disconnected nodes will be inefficient**

- **Without DTN features in CNF**
  - **CNF will perform no worse than current Internet**
  - **Content cannot be delivered when end-host is not connected**