Network Coding

IN 3300

Georg Carle

Outline

- Network Coding Taxonomy
- Reliable Multicast Transport
- Forward Error Correction at Layer 2
- Forward Error Correction at Application Layer
- Examples from Research

Classification of Network Coding

Starting with Forward Error Correction

Generalization to Network Coding

Packet-based Forward Error Correction (FEC)

- k original data packets form a **Transmission Group (TG)**
- h parity packets derived from the k data packets
- any k received out of k+h are sufficient
- Assessment
  + allows receiver (terminal) to recover lost packets
  - overhead at end-hosts
  - complex assessment of impact onto service quality
  - increased network load may increase loss probability

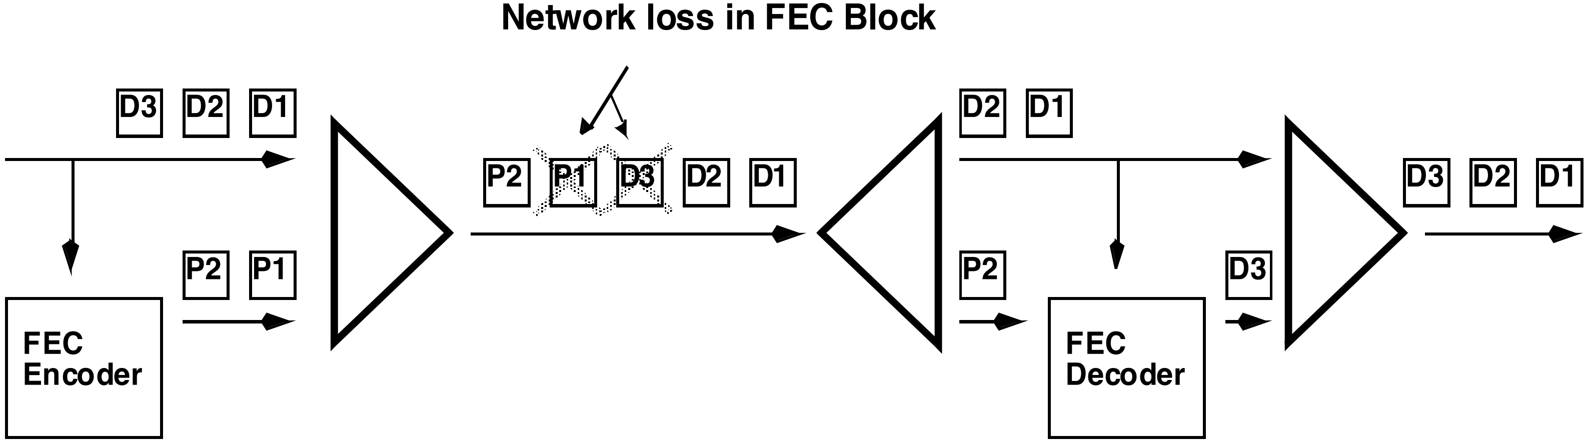

Potential Benefits of FEC

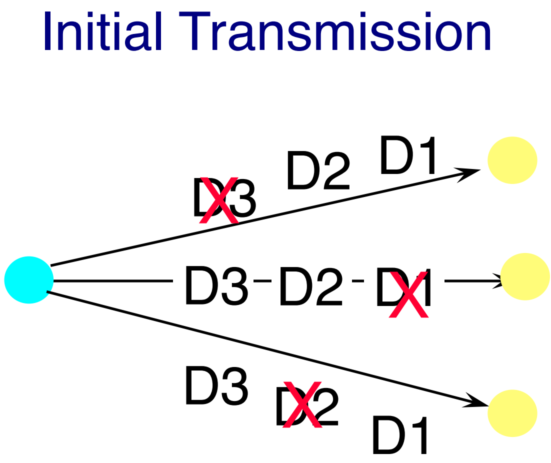

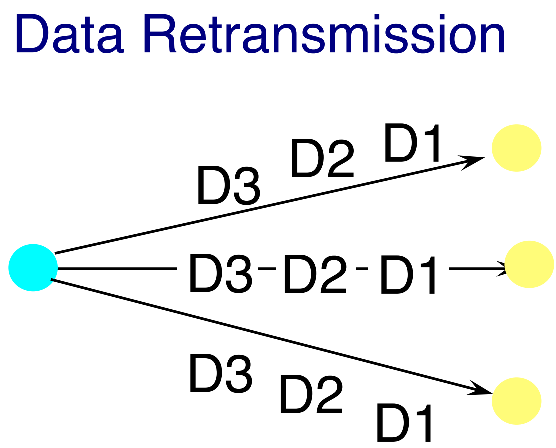

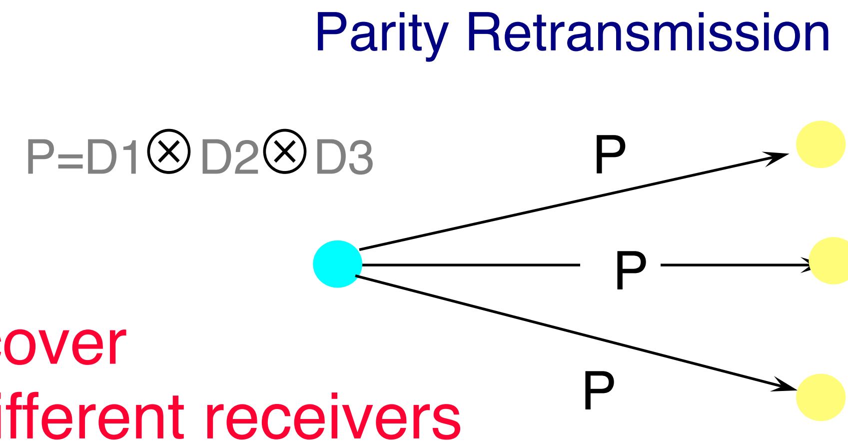

One parity packet can recover
different data packets at different receivers

Network Coding

- Assumption behind traditional network traffic
  - Information is separate, although it may share network resources. (say, cars in highways or fluids in pipes).
  - Network coding breaks this assumption.

- Network Coding
  - A technology to combine several data packets into one or several output packets

Network Coding

- Core ideas:
  - Coded packet mixing improves
    - bottleneck traffic
    - Broadcast
  - Fountain approach simplifies scheduling/coordination
  - In wireless networks, opportunistic listening allows benefits for packets distributed and mixing over the air

Network Coding

- Intermediate nodes transmit packets that are functions of the received packets.

- Potential Benefits
  - throughput improvements in bottleneck scenarios
  - robustness to link failures
  - energy savings
  - simplified operation
  - etc.

Network Coding

- Applications
  - Peer-to-peer multicast services
  - Wireless Networks
    - network types:
      sensor, adhoc, mobile, mesh, ...
    - functions:
      routing in wireless networks
      opportunistic listening
      opportunistic forwarding

Application of Network Coding

- **Multicast in Peer-to-Peer Networks**
  - Large application-level Service Data Unit (SDU) partitioned into blocks
  - Service types
    - File distribution
    - On demand media streaming
    - Real-time media streming
  - Peers exchange blocks according to certain policy
    - Important policy: random blocks („random gossiping)
  - Property of solution
    - Scalable and robust to peers joining and leaving the system

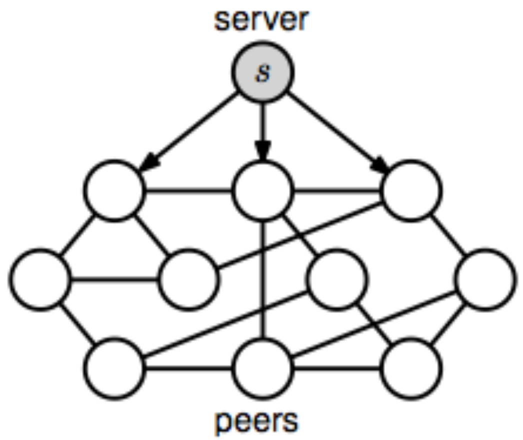

Assessment of Network Coding

- C. Gkantsidis and P. Rodriguez, “Network Coding for Large Scale Content Distribution”, IEEE INFOCOM 2005
  - “The performance benefits provided by network coding in terms of throughput can be more than 2-3 times better compared to transmitting unencoded blocks.”
  - http://research.microsoft.com/pubs/67246/tr-2004-80.pdf
- C. Gkantsidis, J. Miller, and P. Rodriguez, “Comprehensive view of a live network coding P2P system”, ACM IMC 2006
  - “Network coding incurs little overhead, both in terms of CPU processing and I/O activity, and it results in smooth, fast downloads, and efficient server utilization.”
  - http://research.microsoft.com/pubs/69452/imc06.pdf

Reliable Multicast Transport

Many Uses of Multicasting

- Teleconferencing
- Distributed Games
- Software/File Distribution
- Video Distribution
- Replicated Database Updates

⇨ multicast transport is done differently for each application

Multicast Application Modes

- Point-to-Multipoint:
  Single Source, Multiple Receivers

- Multipoint-to-Multipoint:
  Multiple Sources, Multiple Receivers

- Sources are receivers

- Sources are not receivers

Classification of Multicast Applications

Transport service type	***Fully reliable multicast***	***Real-time multicast***
Single source: ***1:N***	Multicast-FTP; Software update	Audio-visual conference; Continuous Media Dissemination
Multiple Sources ***M:N***	CSCW; Distributed computing	DIS; VR

- CSCW: Computer Supported Cooperative Work
- DIS: Distributed Interactive Simulation
- VR: Virtual Reality

Where Does Multicast Loss Occur

- Example measurements
  (April 96, Yajnik, Kurose, Towsely, Univ. Mass., Amherst)

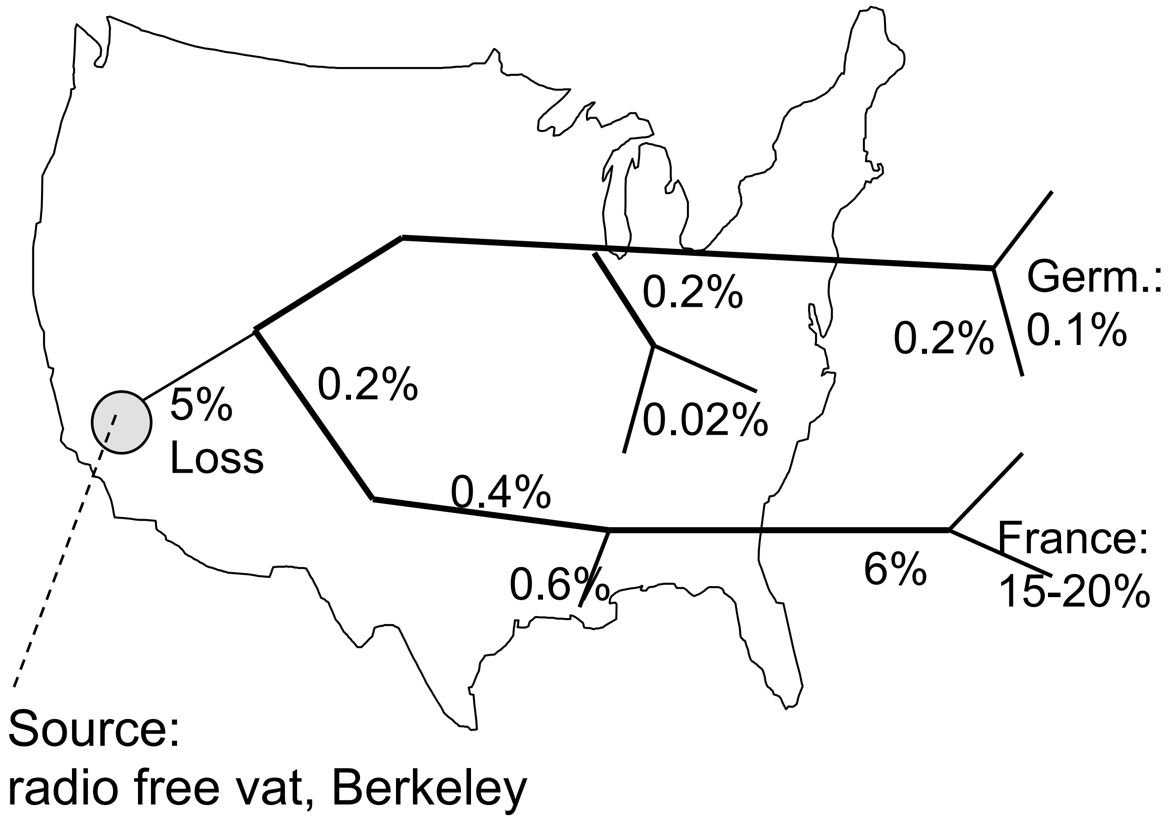

Simultaneous Packet Loss

- Q: distribution of number of receivers losing packet?

- Example dataset:
  47% packets lost somewhere
  5% shared loss
- Similar results across different datasets
- Models of packet loss (for protocol design, simulation, analysis):
  - star: end-end loss independently
  - full topology: measured per link loss independently
  - modified star: source-to-backbone plus star
    ⇨ good fit for example data set

Temporal Loss Correlation

Q: do losses occur individually or in "bursts"?

- occasional long periods of 100% loss
- generally isolated losses
- occasional longer bursts

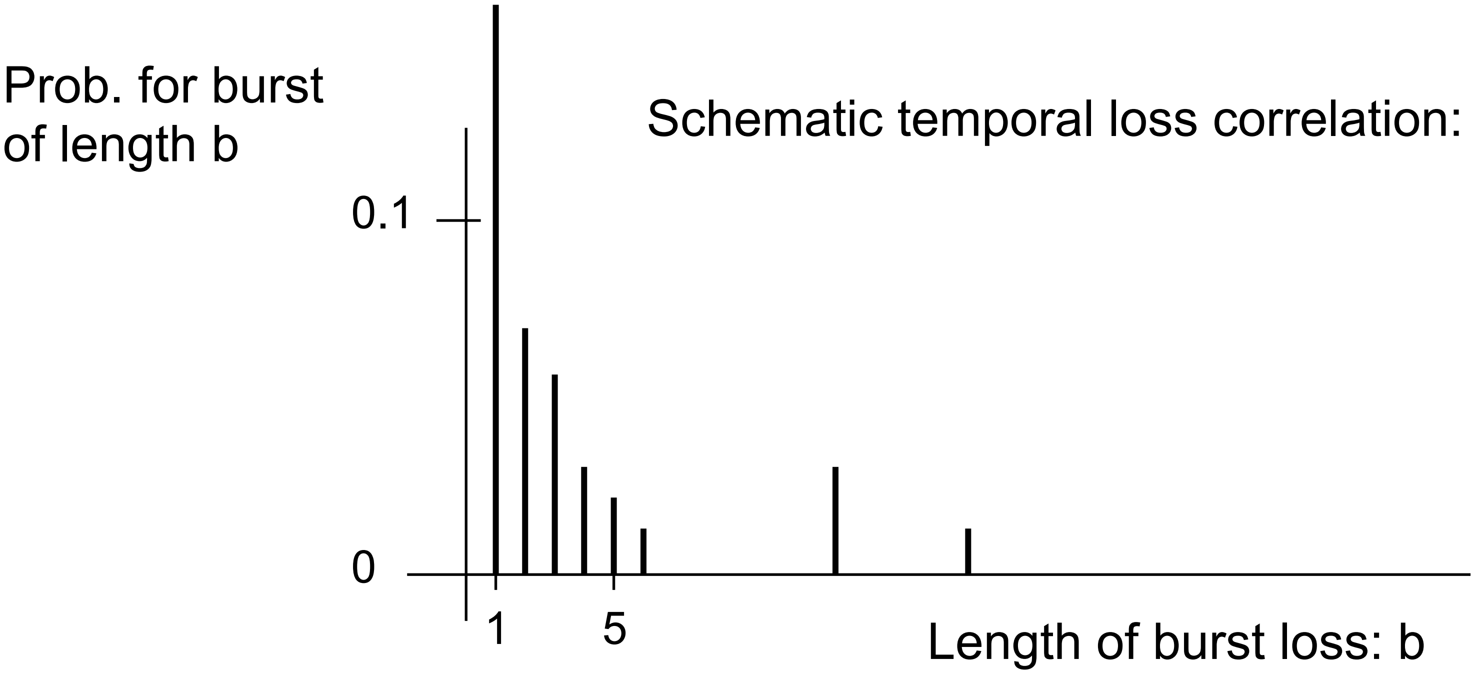

Reliable Multicast Challenge

- How to transfer data reliably from source to R receivers

- scalability: 10s - 100s - 1000s - 10000s - 100000s of receivers

- heterogeneity
  - different capabilities of receivers (processing power, buffer, protocol capabilities)
  - different network conditions for receivers (bottleneck bandwidths, loss rates, delay)

- feedback implosion problem

ARQ: Alternatives for Basic Mechanisms

- Who retransmits
  - source
  - network / servers
  - other group member.
- Who detects loss
  - sender based: waiting for all ACKs
  - receiver based:
    NAK, more receivers ⇨ faster loss detection.
- How to retransmit
  - Unicast
  - Multicast
  - Subgroup-multicast

Approaches

- shift responsibilities to receivers (in contrast to TCP: sender is responsible for large share of functionality)
- feedback suppression (some feedback is usually required)
- multiple multicast groups (e.g. for heterogeneity problems; can be used statically or dynamically)
- local recovery (can be used to reduce resource cost and latency)
- server-based recovery
- forward error correction (FEC)
  - FEC for unicast: frequently no particular gain
  - FEC for multicast: gain may be tremendous!

Multicast with FEC Error Correction

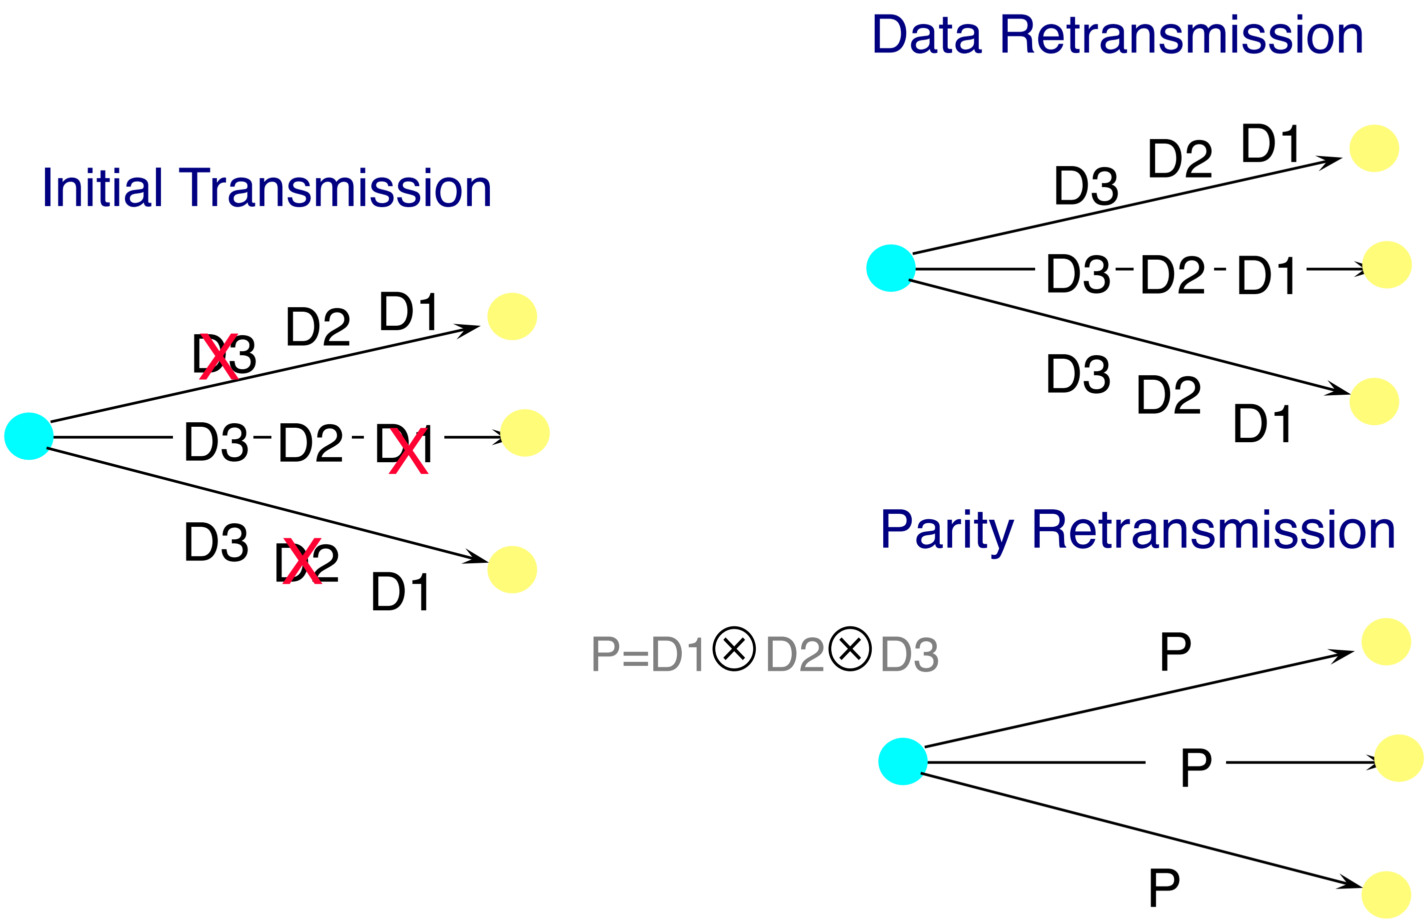

Classification of Multicast Error Control

Multicast Error Recovery

Centralized Error Recovery (CER):
Source retransmits

Distributed Error Recovery (DER): retransmission by server or receiver

grouped (local):
Multicast group is
partitioned into subgroups

ungrouped (global):
All group members participate in error recovery

Reliable Multicast: Building Blocks

- Elements from Unicast:
  - Loss detection
    - Sender-based (ACK): 1 ACK per receiver and per packet; Sender needs a table of per-receiver ACK
    - Receiver-based (NAK): distributed over receivers; potentially only 1 NAK per lost packet
  - Loss recovery: ARQ vs. FEC
- Additional new elements for Multicast:
  - Mechanisms for control message **Implosion Avoidance**
  - Mechanisms to deal with *heterogeneous receivers*