Throughput-Delay-Reliability Tradeoffs in Multihop Networks with Random Access

Sunil Srinivasa and Martin Haenggi

Wireless Institute
Department of Electrical Engineering
University of Notre Dame

Sep 30, 2010

Introduction

- Ad hoc networks are intended to provide reliable broadband services across multiple hops, for example in mesh networks.
- Performance goals often conflict with one another.
  - Hardly possible to guarantee a high rate of transmission in conjunction with reliable packet delivery and low latency.
- In scenarios where reliable delivery is not critical, one can have the nodes forcibly drop a small fraction of packets.
- We characterize the throughput-delay-reliability (TDR) tradeoffs in multihop networks.

Prior work and its shortcomings:

- Focused on single-hop networks [Abouei '09].
- Provided scaling laws alone [Gamal '06], [Neely '05].
- Neglected dependence of dropping on success events [Xie '05].
- Assumed all nodes to be backlogged [Vaze '10].

Our contributions:

- Employ ideas from statistical mechanics to study TDR tradeoffs in ALOHA ad hoc networks.
- Present a simple framework to analyze ad hoc networks, which obviates the often-unwieldy queueing theory-based analysis.

System Model

- **Source nodes**: homogeneous PPP (δ).
- **Relays and destinations**: homogeneous PPP $(1-\delta)$.
- For each source node, the destination node is chosen at a random orientation, and at a random finite distance.

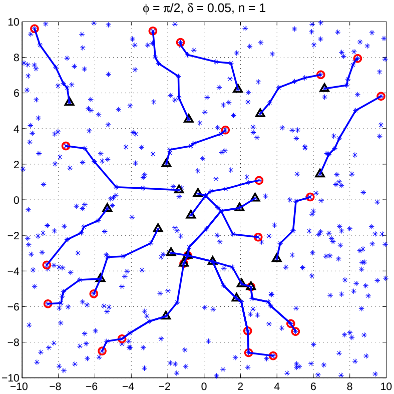

Each destination is assumed to be located 5 nearest-neighbor ($n = 1$) hops away from its source.

System Model (Contd.)

- **Routing**: each node that receives a packet relays it to its n^{th}-nearest-neighbor ($n \geqslant 1$) in a sector of angle $\phi \in [0, \pi]$ towards the destination.

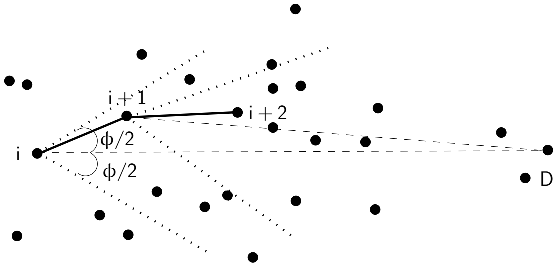

System Model (Contd.)

- All nodes use the same channel.
- Attenuation in the channel: modeled as the product of
  - Large-scale path loss with exponent γ.
  - Small-scale Rayleigh block fading.
- Interference: $I_\Phi(y) = \sum_{x \in \Phi} G_{xy} \|x - y\|^{-\gamma}$.

- Transmission success events are dictated by the SINR model.

$$p_s = \Pr\left(\frac{G_{xy}\|x-y\|^{-\gamma}}{N_0 + I_{\Phi\setminus\{x\}}(y)} > \Theta\right).$$

p_s: Success probability across each link.
Θ: SINR threshold. N_0: Noise (AWGN) variance.

A Revised Buffering and Transmission Policy

1. All the buffering is pushed back to the source, while relay nodes have buffer sizes of unity.
   Furthermore, the source node is always backlogged.
2. Nodes do not accept incoming packets if their buffer already has a packet.
3. Packets are retransmitted until they are successfully received.
   - A transmission is successful only if a node has a packet and its adjacent node has none.
   - Simple way to prevent packets from getting too close.
   - Self-organization: The exclusion principe regulates the traffic injected in a backpressure-like manner.

Advantages of the Single-Buffer Scheme

- Lowers average in-network delay.
  - Stacking-up of packets in buffers is minimal.
- Lessens the variance of the delay.
  - Packet delays are more tightly controlled.
  - Depending on the time a packet spends in its buffer, the source itself can judiciously decide whether to drop it or not.
- Reduces hardware cost and energy consumption.
- Minimizes end-to-end buffer usage [Venkataramanan '10], provides buffering gain [Bhadra '06], self-organizes network operation [Dousse '07].

A "Typical Flow"

- Since network is homogeneous, it is sufficient to consider a "typical" flow (across N relays).
- $\tau_i[t]$: configuration of site i, $0 \leqslant i \leqslant N$ in time slot t.
- $\tau_i[t] = 1$ if its buffer is occupied, otherwise $\tau_i[t] = 0$.
- $\Pr(\tau_i[t] = 1)$: occupancy of a node.
- A successful transmission occurs only if $\{\tau_i[t], \tau_{i+1}[t]\} = \{1, 0\}$.

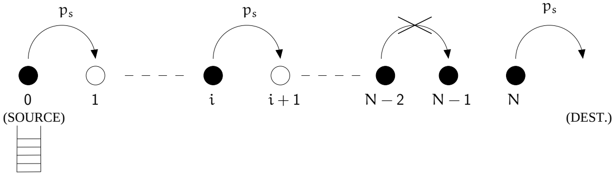

All the buffering occurs at the source; relays have unit-sized buffers.

MAC Scheme: **Slotted ALOHA**

- In each time slot, every node *having a packet* independently transmits w.p. q or remains idle w.p. $1 - q$.

Performance Metrics: **TDR**

- The per-flow **throughput** T, is defined as the average number of packets successfully delivered (to the destination) in unit time, along a typical flow in the network.
- The **mean end-to-end delay**, D, is defined as the average number of time slots it takes for the packet at the head of the source node[a] to successfully hop to the destination.
- The **end-to-end reliability** R is defined as the fraction of packets generated at the source that are eventually delivered.

[a] Note that we consider only the *in-network* delay since the source nodes are always backlogged.

The Regime $R = 1$

For an ALOHA-based line flow along N relays, the steady state throughput at full reliability ($R = 1$) is

$$T = \frac{qp_s B(N)}{B(N+1) + qp_s B(N)},$$

while the average end-to-end delay is given by

$$D = (1 + N/2)/T.$$

where $B(0) = 1$, and

$$B(k) = \sum_{j=0}^{k-1} \frac{1}{k} \binom{k}{j} \binom{k}{j+1} (1 - qp_s)^j, \quad k > 0.$$

The Regime R = 1 (Contd.)

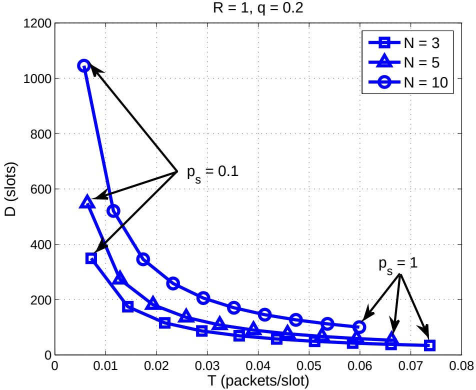

For each value of N, the TD curve is a hyperbola.

The Regime $R < 1$

- When $R = 1$, D and T performances are poor at small p_s.
- Nodes can choose to drop a small fraction of packets ($R < 1$).
  - Each node having a packet decides to drop the packet in its buffer or not stochastically w.p. ξ.

For clarity, we consider the following two regimes separately.

The Noise-Limited Regime: $p_s = \Pr(\mathsf{SNR} > \Theta)$

Noise power in the network is much stronger than the interference.

The Interference-Limited Regime: $p_s = \Pr(\mathsf{SIR} > \Theta)$

- Interference power in the network is much stronger than noise.
- Also covers the regime wherein the interference and noise powers are comparable.

The Regime $R < 1$ (Contd.)

- System Evolution: The following events affect τ_i:
  a) Node $i-1$ transmits its packet to node i.
  b) Node i transmits its packet to node $i+1$.
  c) Node i drops its packet.
- Employing mean-field theory, we obtain at steady state, $\mathbb{E}\lim_{t\to\infty}\Delta\tau_i[t] = 0$ for $1 \leqslant i \leqslant N$, i.e.,

$$p_s(1-\xi)q\Big[\underbrace{\mathbb{E}\tau_{i-1}(1-\mathbb{E}\tau_i)}_{a)} - \underbrace{\mathbb{E}\tau_i(1-\mathbb{E}\tau_{i+1})}_{b)}\Big] - \underbrace{\xi\mathbb{E}\tau_i}_{c)} = 0.$$

- The steady state occupancies $\mathbb{E}\tau_i$ may be obtained by numerically solving these N non-linear equations.

The Regime $R < 1$ (Contd.)

1. The steady-state throughput is

$$T = qp_s \mathbb{E}\tau_N.$$

2. The mean end-to-end delay is

$$D = \sum_{i=0}^{N} s_i^{-1},$$

where $s_i = qp_s(1 - \mathbb{E}\tau_{i+1})$.

3. The end-to-end reliability of the network is

$$R = \prod_{i=0}^{N} \frac{s_i(1-\xi)}{s_i + \xi - s_i\xi}.$$

The Noise-Limited Regime; $R < 1$

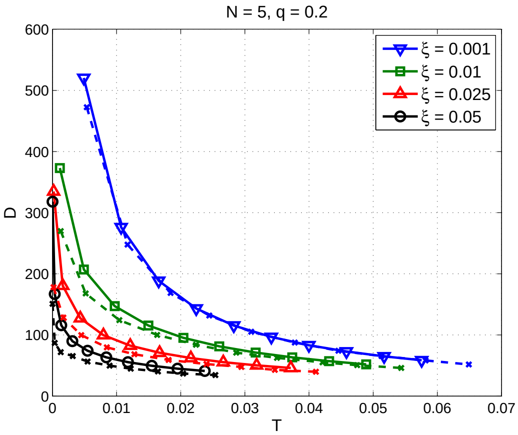

Increasing ξ helps reduce the end-to-end delay significantly

The Noise-Limited Regime; $R < 1$ (Contd.)

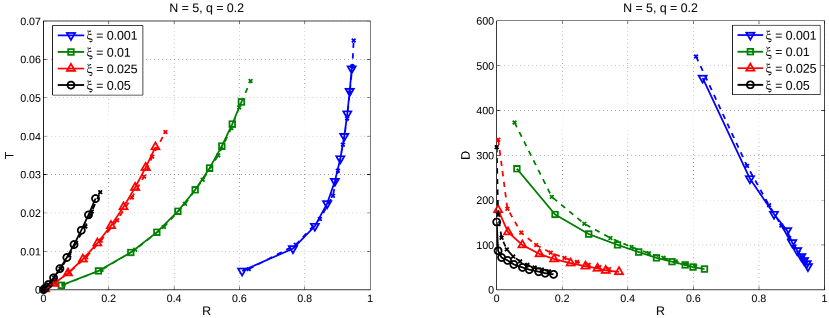

However, the throughput and reliability performances worsen too.

The Interference-Limited Regime; $R < 1$

- The average number of potential interferers in each flow is $1 + \sum_{i=1}^{N} \mathbb{E}\tau_i$.
  - The set of interferers (approximately) forms a PPP with density $\lambda_I = \delta q \left(1 + \sum_{i=1}^{N} \mathbb{E}\tau_i\right)$.
- The probability of a successful transmission for a typical link is

$$p_s \gtrsim \left(\frac{(1-\delta)\phi}{(1-\delta)\phi + 2\delta q \left(1 + \sum_{i=1}^{N} \mathbb{E}\tau_i\right) c} \right)^n .$$

The Interference-Limited Regime; $R < 1$ (Contd.)

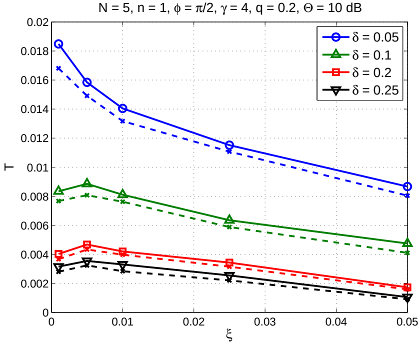

- When δ is small, increasing the packet dropping probability ξ reduces the system throughput.
- As δ gets larger, dropping a few packets helps mitigate the interference, and the throughput across a typical flow improves.

The Interference-Limited Regime; $R < 1$ (Contd.)

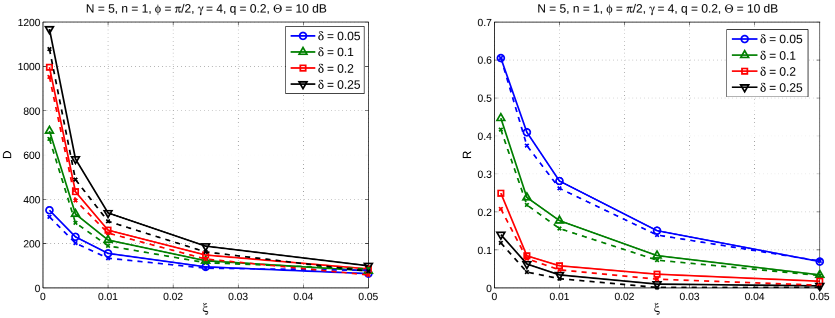

With increasing ξ or decreasing δ, the mean end-to-end delay decreases; the reliability also suffers.

Summary

- We considered a multihop wireless network consisting of several source-destination pairs communicating with each other in a completely uncoordinated manner.
- Employing the mean-field approximation, we presented a framework for computing the steady state mean node occupancies, and quantifying the network's TDR performance.
- In the noise-limited regime, dropping a small fraction of packets in the network leads to a smaller end-to-end delay at the cost of reduced throughput.
- In the interference-limited scenario, dropping a few packets in the network can help mitigate the interference in the network leading to an increased throughput.