

Course “Algorithmic Foundations of Sensor Networks”

Lecture 3: Energy-efficient and robust routing (single path)

Sotiris Nikolettseas
Professor

Department of Computer Engineering and Informatics
University of Patras, Greece

Spring Semester 2020-2021

Energy-Efficient and Robust Routing

Routing can be made robust and energy efficient by taking into account some pieces of *state information available locally* in the network, such as:

- *link quality*: e.g. packet reception rates obtained by periodic monitoring
- *link distance*: if link monitoring is problematic/expensive (e.g. in highly dynamic, rapidly fading environments) link distances can be alternatively used as *indicators* of link quality and energy consumption
- *residual energy*: avoid nodes with low energy, to prolong the network lifetime
- *location information*: if such information is available, geographic routing techniques can be used
- *mobility information* can also be exploited; sensor/sink mobility is both a challenge and an opportunity

- A. Link quality based methods (ETX, MOR)
- B. Link distance based methods
 - Routing with relay diversity (ExOR)
 - Greedy, local neighbor selection (LTP)

A. Link Quality Methods

- *Main goal:* Providing robustness via selecting routes that *minimize the number of end-to-end retransmissions or failure probabilities.*
- *Motivation:* If wireless links were (ideally) error-free, then a shortest hop-count path could be chosen. Such paths however require high distance links, which in realistic settings are likely to be error-prone. Thus, the link quality should be taken into account (via simple ACK signals for each successfully delivered packet).


A1. The ETX metric

- Let d_f be the packet reception rate (probability of successful delivery) over a link in the forward direction, and d_r the probability that the corresponding ACK is received in the reverse direction.
- Assuming each packet transmission can be abstracted as a Bernoulli trial, the expected number of retransmissions required for successful delivery of a packet on this link is:

$$ETX = \frac{1}{d_f \cdot d_r}$$

- This metric for a single link can then be incorporated into any routing protocol, so that *chosen end-to-end paths minimize the sum of ETX over all links on the path*, i.e. the total expected number of transmissions on the path.

ETX routing - example



- This example shows 3 routes from A to B, with the forward probabilities on each link (for simplicity, all reverse probabilities are taken $d_r = 1$).
- The direct A-B transmission would require an expected number of 10 retransmissions.
- The path via C,D,E would require $4(1/0.9) \simeq 4.44$ retransmissions (1.1 per link) on average
- The A-F-B path would require $2(1/0.8) \simeq 2.5$ retransmissions (1.25 per link) on average.

⇒ the ETX-minimizing path is A-F-B.

ETX routing - discussion

- As demonstrated in the example, ETX *neither favors long paths* with many short-distance, high-quality links, *nor very short paths* with a few long-distance, low-quality links. It actually takes an *“in-between” approach*.
- *ETX advantages*: It minimizes the number of transmissions required, as well as it improves energy efficiency. Also, it directly addresses the potential link asymmetries, by taking into account packet delivery probabilities in both directions.
- *ETX limitation*: It assumes knowledge of the d_r , d_f packet reception probabilities. In quite static networks these values may be periodically obtained via link monitoring, however in dynamic networks the obtained values may become obsolete very soon and the link monitoring overhead may be prohibitively high.

A2. The Minimum Outage Route (MOR) Metric (I)

- Under *high dynamics* the link quality fluctuates very rapidly and *ETX is not feasible*. In such cases, *analytic models* of link quality can be used. They explicitly model the wireless channel as having multi-path fading with Rayleigh statistics (fluctuations over time).
- Let d the distance between transmitter and receiver, h the path-loss exponent, SNR the signal-to-noise ratio without fading, f the fading state of the channel. Then, the *instantaneous capacity* of the channel is:

$$C = \log \left(1 + \frac{|f|^2}{d^h} \text{SNR} \right)$$

A2. The Minimum Outage Route (MOR) Metric (II)

- The *outage probability* P_{out} is defined as the probability that the instantaneous channel capacity falls below the transmission rate R . It is shown that:

$$P_{out} = 1 - \exp\left(\frac{-d^h}{\mu \cdot SNR^*}\right)$$

where $\mu = E[|f|^2]$ is the mean of the Rayleigh fading and SNR^* a normalized SNR.

where:

- Rayleigh fading: Rayleigh fading occurs when there are many objects in the environment that scatter the radio signal (especially in heavily built-up city centers where there is no line of sight between transmitter and receiver)
- f : measures the level of fading.
- h : typically ranges between 2 and 7.

A2. The Minimum Outage Route (MOR) Metric (III)

- Then, end-to-end reliability of a route is defined as the probability that *none of the intermediate links suffers outage*, and the most reliable route between two nodes is one that *minimizes the following path metric*:

$$\sum_i d_i^h$$

where d_i is the distance of the i^{th} hop in the path. The metric d^h for each link is called the minimum outage route (MOR) metric.


- The MOR approach does not require collection of link quality metrics like ETX, neither ACK messages are used. However, *its practicality is limited by the basic abstraction that the channel fading follows a certain distribution (the Rayleigh distribution)*.

B. Link Distance Methods

- In general, nodes *closer to the data destination* are chosen/favored.
- The relay diversity provided by the fact that wireless transmissions are *broadcast to multiple nodes* is exploited.

B1. Routing With Relay Diversity


- An example:



- In traditional routing, the reliability of the AC path depends on whether both AB and BC transmissions have been successful.
- However, if C is *also allowed to accept packets directly from A*, then the reliability can be further increased without much additional energy cost.
- Allowing such 2-hop packet reception, in the high SNR regime, *the end-to-end outage probability decays as $(SNR)^{-2}$* . When nodes within L hops can communicate with each other with high SNR , this probability can become as small as $(SNR)^{-L}$.
- A weakness of this method is that it requires *a larger number of receivers to be actively overhearing each message*, which may incur a *radio energy penalty*.

Extremely Opportunistic Routing (ExOR) (I)

- Main idea: the identity of the node which will eventually forward a packet, is not predetermined before the packet is transmitted. Instead, the method tends to ensure that the node closest to the destination that receives a given packet will forward the packet further.*



- The protocol has three stages:
 - 1. Priority ordering:** At each step, the transmitter includes in the packet a schedule with the *priority order* candidate receivers should forward the packet.

Extremely Opportunistic Routing (ExOR) (II)

2. *Transmission acknowledgements:* A MAC scheme is used so that each candidate receiver sends the ID of the highest-priority successful recipient known to it. All nodes listen to all ACKs, so they distributively determine *which node, among those who received the packet successfully, has highest priority.*
 3. *Forwarding decision:* After listening to all ACKs, the nodes that *have not* heard of any IDs with priorities greater than their own will transmit.
- Features of the ExOR protocol
 - *Nodes further away to the current node (yet closer to the destination)* are less likely to successfully receive the packet, but, whenever they receive it, they are favored to act as forwarders!
 - This *tends to make good progress towards the data destination*, without many transmissions and delays.
 - As with relay diversity, ExOR requires *a larger number of receivers to be active*. Also, the priority evaluation necessitates some inter-node packet delivery ratios to be tracked and maintained.

B2. Greedy, Local Neighbor Selection (LTP) (I)

We call network nodes “*particles*”.

- a) Each particle may have *two communication modes*: a *broadcast* (radio) *beacon mode* and a *directed to a point* transmission mode (laser beam). However, a radio broadcast suffices.
- b) Each particle may alternate between a **sleeping** and an **awake** mode. During sleeping periods particles cease any communication.
- c) Particles *do not move*.
- d) The particles are spread in a *two-dimensional* area (plane).

Greedy, Local Neighbor Selection (LTP) (II)

- e) A *receiving wall* $\mathcal{W}$ is a *line* in the plane. The wall represents the control center (multiple/mobile sinks).
- f) Each particle is *aware of the direction* toward $\mathcal{W}$.
- g) *No geolocation* abilities assumed

Definitions: Let d (in numbers of particles / m^2) be the **density** of the cloud.


Let $\mathcal{R}$ be the maximum (beacon/laser) **transmission range** of each particle.

The Local Target Protocol (LTP)


Let $d(p_i, p_j)$ the (vertical) distance of p_i, p_j and $d(p_i, \mathcal{W})$ the (vertical) distance of p_i from $\mathcal{W}$. Let $\text{info}(\mathcal{E})$ the info to be propagated. Each p' receiving $\text{info}(\mathcal{E})$ does the following:

- **Search Phase:** It uses a low energy broadcast of a beacon (angle α above and below the vertical line) to *discover* a particle closer to $\mathcal{W}$ (i.e. a p'' where $d(p'', \mathcal{W}) < d(p', \mathcal{W})$).
- **Direct Transmission Phase:** If found, p' *sends* $\text{info}(\mathcal{E})$ to p'' via a direct line (laser) transmission.
- **Backtrack Phase:** If repetitions of the search phase *fail* to discover a particle nearer to $\mathcal{W}$, then p' *sends* $\text{info}(\mathcal{E})$ *to the particle it received the information from*.

Example of the Search Phase



Example of Data Propagation



Definitions: Let h_{opt} the (*optimal*) number of “hops” (**vertical to $\mathcal{W}$** transmissions) needed to reach $\mathcal{W}$, if particles always exist in pair-wise distances $\mathcal{R}$ towards $\mathcal{W}$.

Let h the *actual* number of hops (transmissions) taken to reach $\mathcal{W}$.
The **“hops” efficiency** of the data propagation protocol is the ratio

$$C_h = \frac{h}{h_{opt}}$$

where $h_{opt} = \left\lceil \frac{d(p, \mathcal{W})}{\mathcal{R}} \right\rceil$

Why studying h , C_h ?

When a particle p “looks around” for a particle as close to $\mathcal{W}$ as possible to pass information, **it may not get any particle in the perfect direction** (on the line vertical to $\mathcal{W}$ passing from p), mainly because:

- a) There might never have been any particles in that direction.
- b) Particles of sufficient remaining battery power may not be available anymore.
- c) Particles available may temporarily “sleep” to save energy.

Simplifying Assumptions for a Rigorous Analysis

- The search phase **always finds a p''** in the semicircle of center p' and radius $\mathcal{R}$ towards $\mathcal{W}$.
This assumption can be *relaxed*: (a) by *repetitions* of the search phase (b) we may consider a *cyclic sector* defined by circles of radii $\mathcal{R} - \Delta\mathcal{R}$, $\mathcal{R}$ (c) if a search phase ultimately fails, the protocol **backtracks**.
- The position of p'' is **random uniform** in the arc of angle $2a$.
- Each target selection is stochastically **independent** of the others.

Lemma

The expected “hops” efficiency of LTP in the α -uniform case is $E(C_h) \simeq \frac{\alpha}{\sin \alpha}$, for large h_{opt} . Also, $1 \leq E(C_h) \leq \frac{\pi}{2} \simeq 1.57$.

Proof: A sequence of points is generated, $p_0 = p, p_1, p_2, \dots, p_{h-1}, p_h$ where p_{h-1} is the first particle found within $\mathcal{W}$'s range and p_h is beyond $\mathcal{W}$. Let α_i be the (positive or negative) angle of p_i w.r.t. p_{i-1} 's vertical line to $\mathcal{W}$. It is:

$$\sum_{i=1}^{h-1} d(p_{i-1}, p_i) < d(p, \mathcal{W}) \leq \sum_{i=1}^h d(p_{i-1}, p_i)$$

The (vertical) progress toward $\mathcal{W}$ is $\Delta_i = d(p_{i-1}, p_i) = \mathcal{R} \cos \alpha_i$. We get:

$$\sum_{i=1}^{h-1} \cos \alpha_i \leq h_{opt} \leq \sum_{i=1}^h \cos \alpha_i$$

From Wald's equation*, then

$$E(h-1) \cdot E(\cos \alpha_i) \leq E(h_{opt}) \leq E(h) \cdot E(\cos \alpha_i) \Rightarrow$$

$$\frac{\alpha}{\sin \alpha} \leq \frac{E(h)}{h_{opt}} = E(C_h) \leq \frac{\alpha}{\sin \alpha} \cdot \frac{h}{h-1}$$

since

$$E(\cos \alpha_i) = \int_{-\alpha}^{\alpha} \cos x \cdot \frac{1}{2\alpha} \cdot dx = \frac{\sin \alpha}{\alpha}$$

Assuming large values for h_{opt} and since for $0 \leq \alpha \leq \frac{\pi}{2}$ it is $1 \leq \frac{\alpha}{\sin \alpha} \leq \frac{\pi}{2}$ we get the result.

* Walds equation

- The linearity of expectation property: Let random variables $X_1, X_2, \dots, X_n$. Then:

$$E\left(\sum_{i=1}^n X_i\right) = \sum_{i=1}^n E(X_i)$$

even when X_i 's are not independent.

- Wald's equation generalizes this property to a random number of variables: Let $X_1, X_2, \dots, X_N$ random variables and let N a random number which is a stopping time for the X_i sequence (a random time when a random process exhibits a certain behavior). Then:

$$E\left(\sum_{i=1}^N X_i\right) = E(N) \cdot E(X_i)$$

The “min-two uniform targets” (M2TP) Protocol

We assume that the search returns **two points** p'' , p''' each uniform in $(-\alpha, \alpha)$ and that the protocol selects the best.

Let α_{i1} , α_{i2} **the angles** of the particles found and let $\alpha_i = \min \{|\alpha_{i1}|, |\alpha_{i2}|\}$. Then,

$$\mathbf{P}\{\alpha_i > \phi\} = \mathbf{P}\{|\alpha_{i1}| > \phi \cap |\alpha_{i2}| > \phi\} = \left(\frac{\alpha - \phi}{\alpha}\right)^2$$

Thus, the **distribution function** of α_i , is

$$F_{\alpha_i}(\phi) = \mathbf{P}\{\alpha_i \leq \phi\} = \frac{2\alpha\phi - \phi^2}{\alpha^2}$$

and the **probability density function** is,

$$f_{\alpha_i}(\phi) = \frac{d}{d\phi} \mathbf{P}\{\alpha_i \leq \phi\} = \frac{2}{\alpha} \left(1 - \frac{\phi}{\alpha}\right)$$

The **expected local progress** is:

$$E(\cos \alpha_i) = \int_0^\alpha \cos \phi \cdot f_{\alpha_i}(\phi) d\phi = \frac{2(1 - \cos \alpha)}{\alpha^2}$$

Lemma

The expected “hops” efficiency of the “min-two uniform targets” protocol in the α -uniform case is $E(C_h) \simeq \frac{\alpha^2}{2(1-\cos \alpha)}$, for $0 \leq \alpha \leq \frac{\pi}{2}$ and for large h .

We remark that

$$\lim_{\alpha \rightarrow 0} E(C_h) = \lim_{\alpha \rightarrow 0} \frac{2\alpha}{2 \sin \alpha} = 1$$

and

$$\lim_{\alpha \rightarrow \frac{\pi}{2}} E(C_h) = \frac{(\pi/2)^2}{2(1-0)} = \frac{\pi^2}{8} \simeq 1.24$$

Lemma

The expected “hops” efficiency of the min two uniform targets protocol is $1 \leq E(C_h) \leq \frac{\pi^2}{8} \simeq 1.24$ for large h and for $0 \leq \alpha \leq \frac{\pi}{2}$.

Note: sampling more than one targets may obviously incur additional energy dissipation.

Summary evaluation of LTP

- local, simple, greedy protocol
- no global structure (set of paths) maintained
- good for dense networks
- performance drops in sparse / faulty networks