Routing

Daniele Puccinelli

daniele.puccinelli@supsi.ch

http://web.dti.supsi.ch/~puccinelli

Routing in WSNs

- Collection
  - Get the data from the nodes to the sink
  - Many-to-one traffic pattern (convergecast towards sink)
  - Sample usage scenario: monitoring
- Dissemination
  - Get the data into the network
  - One-to-many traffic pattern
  - Sample usage scenario: event detection
- Point-to-point routing
  - Get the data from one specific node to another node
  - More limited applicability: data-centric queries

Our focus: **collection routing**

Why do we need to route?

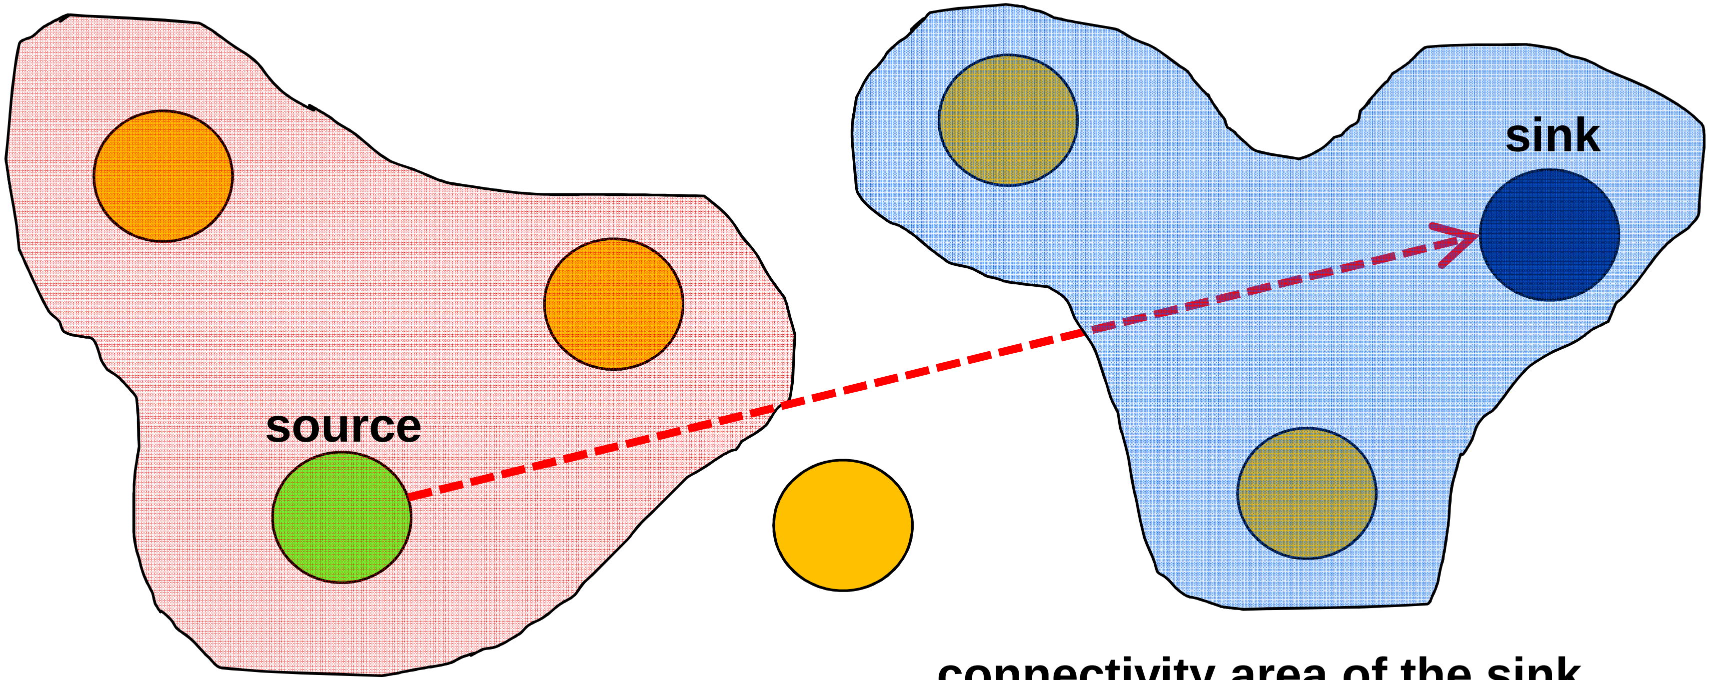

You typically cannot hit the sink over one hop!

Flooding

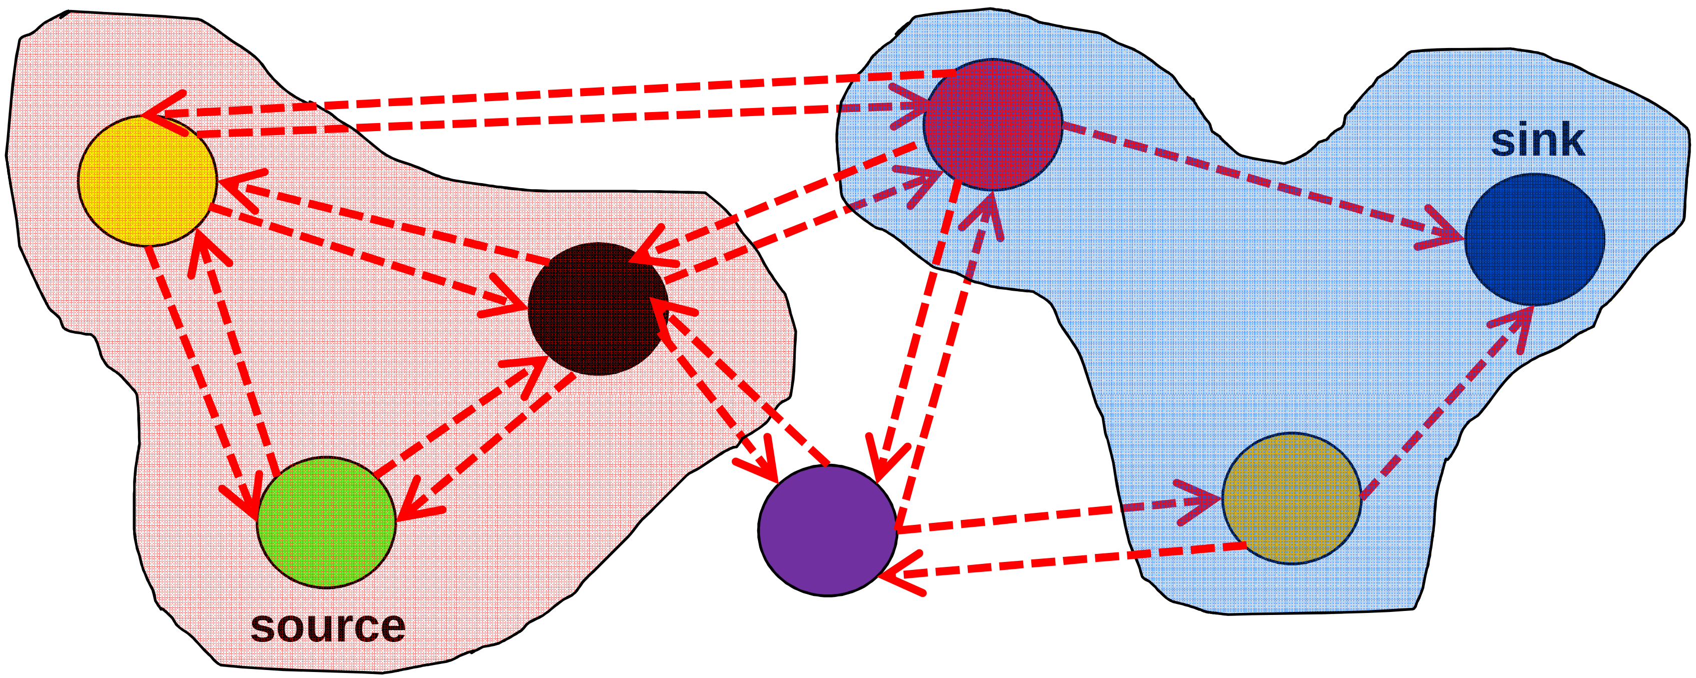

A Flooding Experiment

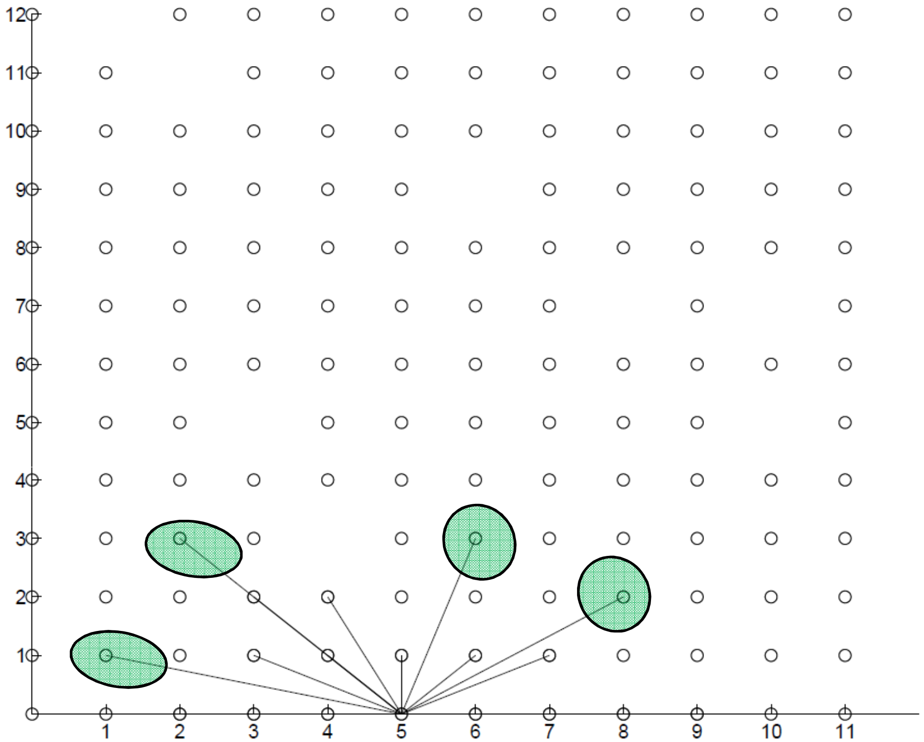

After one hop

Some long links (great)

D. Ganesan, D. Estrin, A. Woo, D. Culler, B. Krishnamachari, S. Wicker
Complex Behavior at Scale: An Experimental Study of Low-Power WSNs
UCLA Technical Report, 2002

A Flooding Experiment

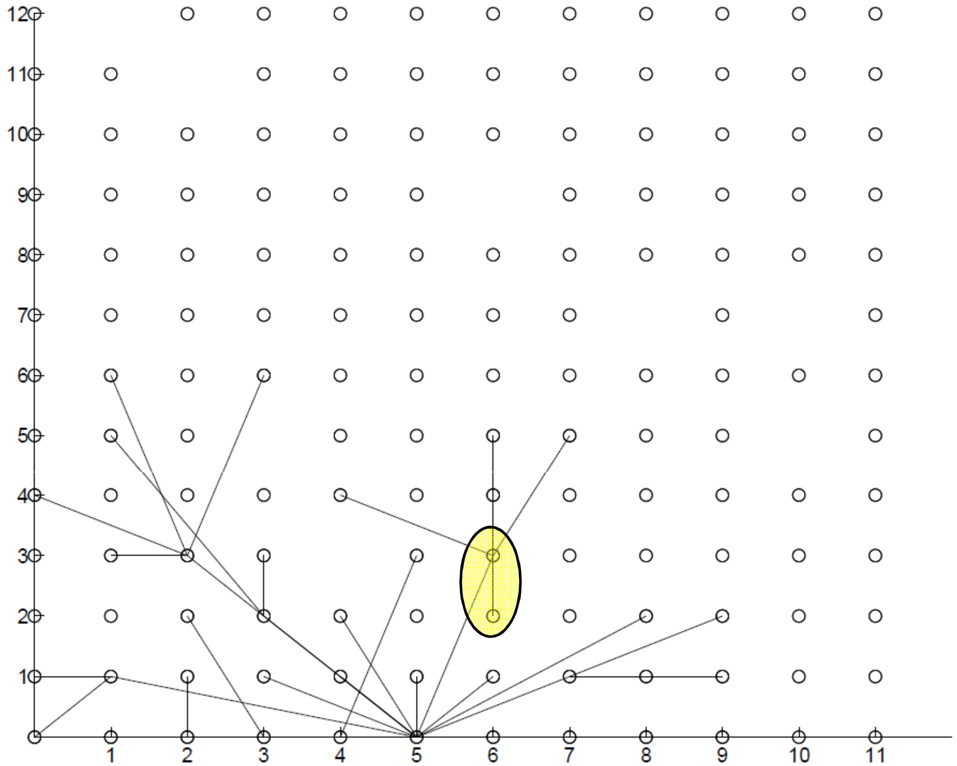

After two hops

Already one backward link

D. Ganesan, D. Estrin, A. Woo, D. Culler, B. Krishnamachari, S. Wicker
Complex Behavior at Scale: An Experimental Study of Low-Power WSNs
UCLA Technical Report, 2002

A Flooding Experiment

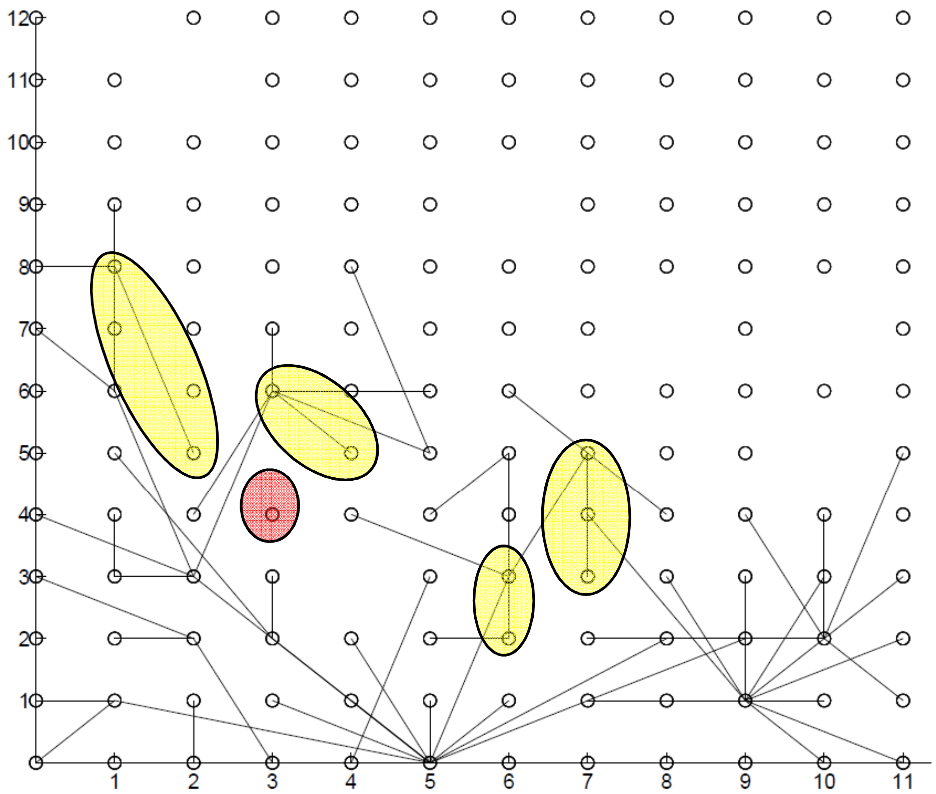

After three hops

Emergence of backward links

Presence of stragglers

D. Ganesan, D. Estrin, A. Woo, D. Culler, B. Krishnamachari, S. Wicker
Complex Behavior at Scale: An Experimental Study of Low-Power WSNs
UCLA Technical Report, 2002

Flooding: What's Wrong?

- Flooding is not suitable for WSNs
  - o requires too many resources!
  - o interference
  - o loops
  - o useless transmissions

- Gossiping
  - o Recipients act as relays with a given probability
  - o Less resources required
  - o More delay, likely disconnection

- Coordination among relays?

Traditional approaches

• Proactive approaches (routes computed up front)

• Distance-vector routing

oCompute distance to all nodes

oPeriodically tell your neighbors

Network-wide connectivity information to a few others

• Link-state routing

oPeriodically tell everyone about your neighbors

oCompute shortest path

Local connectivity information to everyone else

Challenges of WSN Routing

Programming a large network...

...of highly resource-constrained nodes...

...to self-organize...

...into some globally consistent and robust behavior...

...using only simple local rules...

...over a noisy and dynamically changing environment!

Centralized or Distributed?

Centralized

- One powerful node (gateway) has a global picture of the network
- It computes optimal routes and directs the nodes
- Requires hierarchical organization

Distributed

- Each node takes local decisions based on partial state knowledge
- Generally applicable

Sender-Based vs. Receiver-Based

Sender-based

- The sender knows where to unicast
- No effort required from the receiver

Receiver-based

- The sender broadcasts its packets
- The receiver must figure out whether to forward

Single-Path vs. Multi-Path

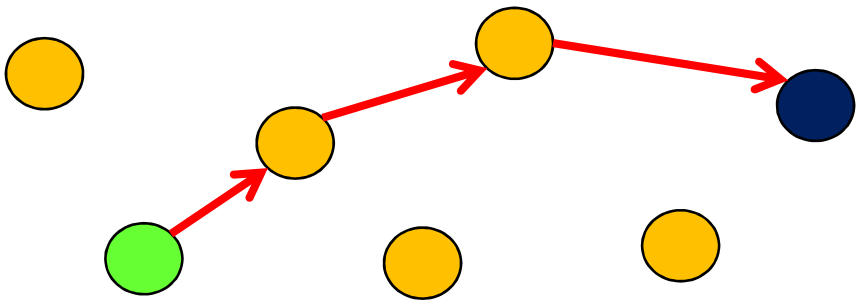

Single-path
Only one route between 2 given nodes at any given time

Multi-path
Multiple routes between 2 given nodes at a given time
- Potentially more robust
- Definitely more wasteful

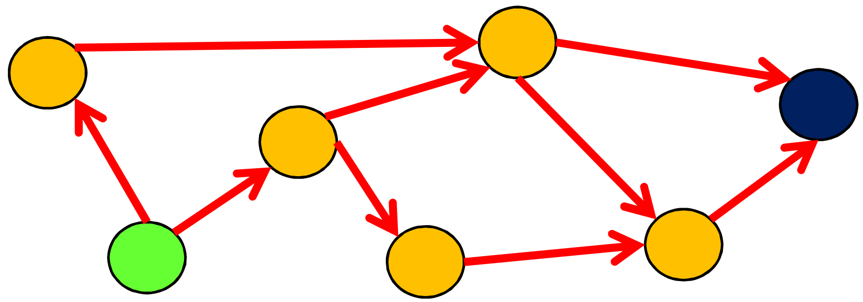

Geographic

- Nodes know where they are and where the sink is
- Routing decisions are based on geographic information

Cost-based

- Nodes have no sense of Euclidean distance
- They don’t know where they are or where anyone is
- Each node assesses the cost of reaching the sink
- Cost is the node’s equivalent of distance

WSN Routing Design Space

- Centralized vs. distributed
Centralized requires heterogeneous architecture

- Geographic vs. cost-based
Geographic requires localization

- Reactive vs. proactive
Depends on application

- Single-path vs. multipath
Multipath has a huge overhead

- Sender-based vs. receiver-based
Receiver-based is simpler, but it comes with redundancy

Address-Centric vs. Data-Centric

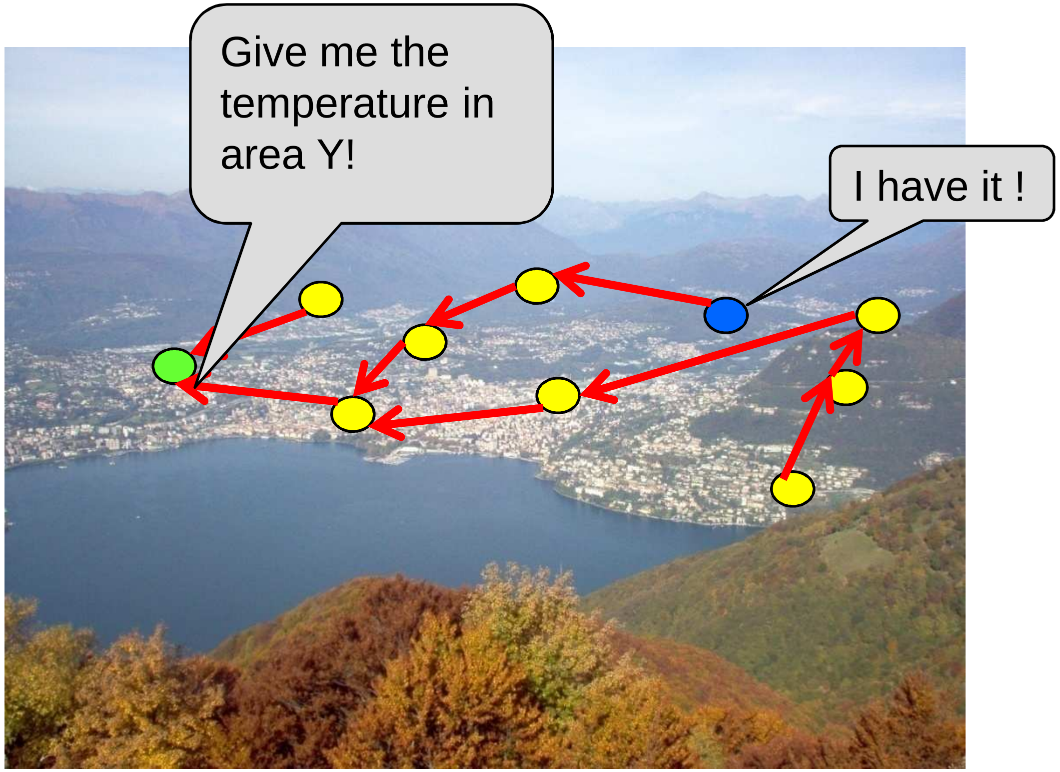

- Oftentimes in WSNs addresses don’t really matter
- The sink may not necessarily want data from a given node
- Most likely, the sink wants a given kind of data

- Interest dissemination
- Data diffusion

Cost-based routing

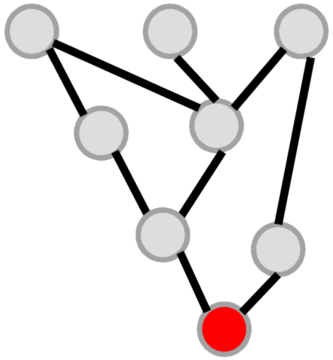

Cost=0 hops

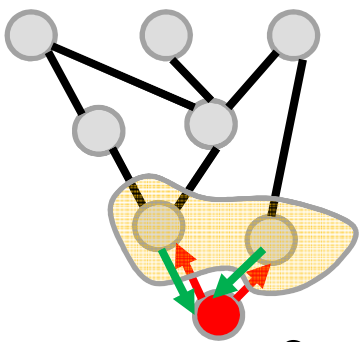

Cost=1 hop

- Basic principle: propagation of a cost field through control traffic
- Cost: hop distance from sink
- Upstream nodes send data down the field

Cost=2 hops

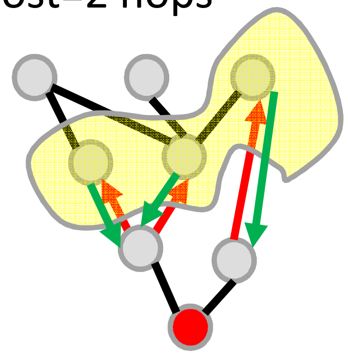

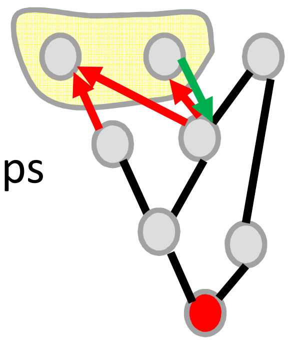

Cost=3 hops

Cost Field Diffusion

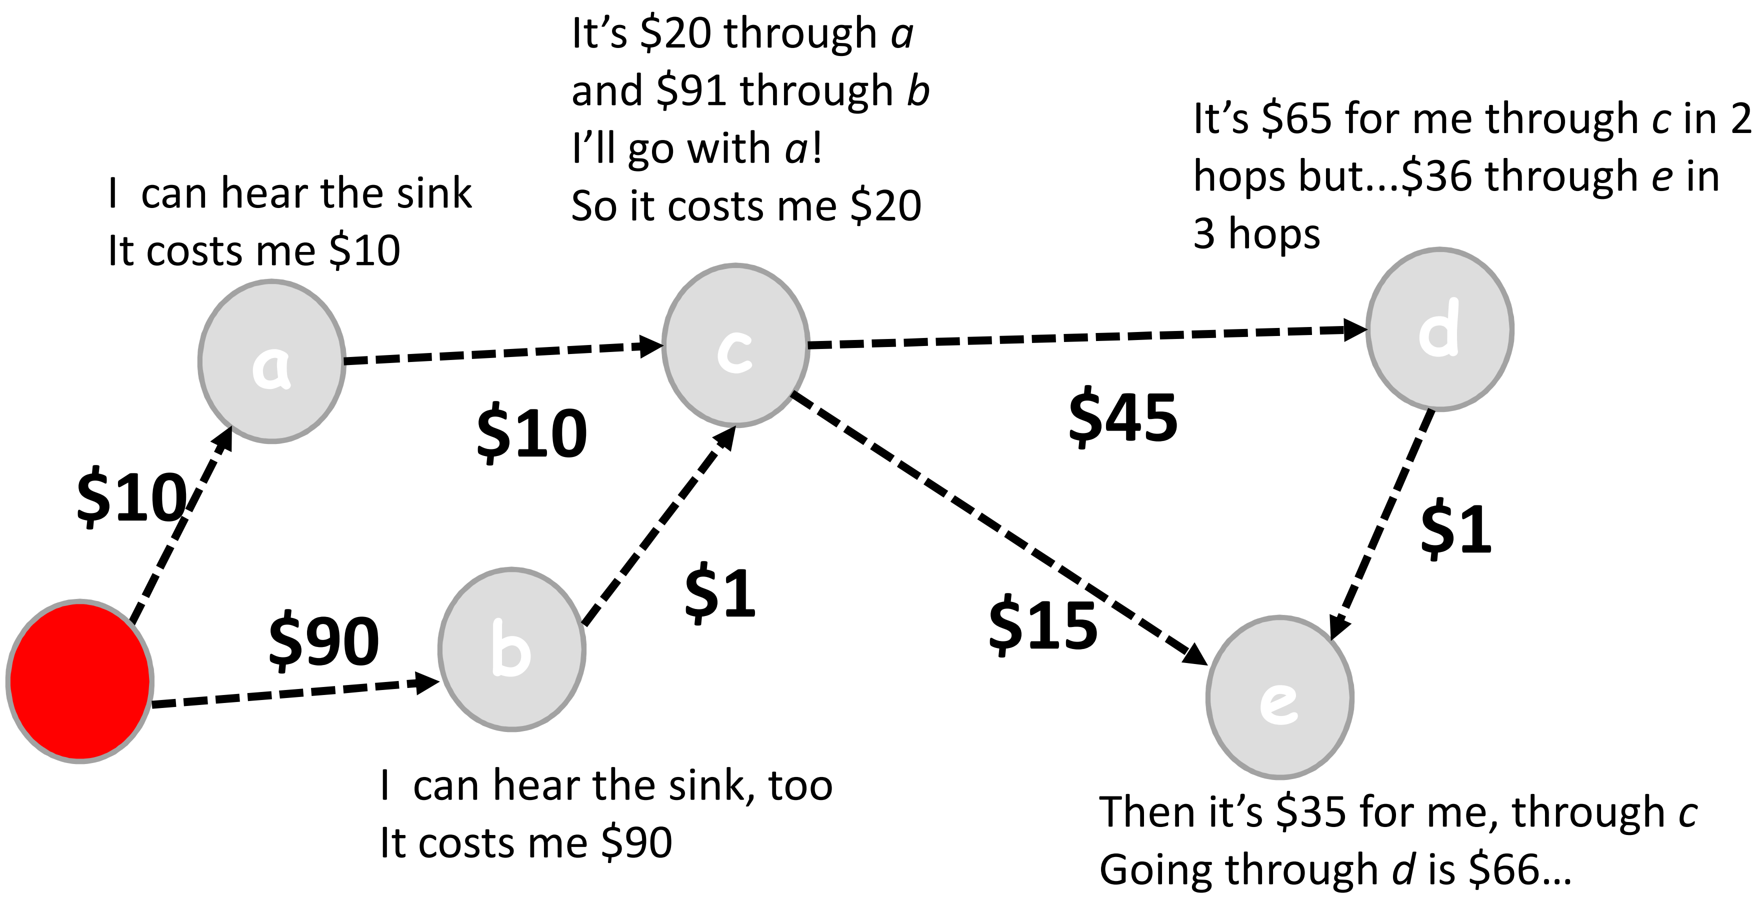

Routing Decisions

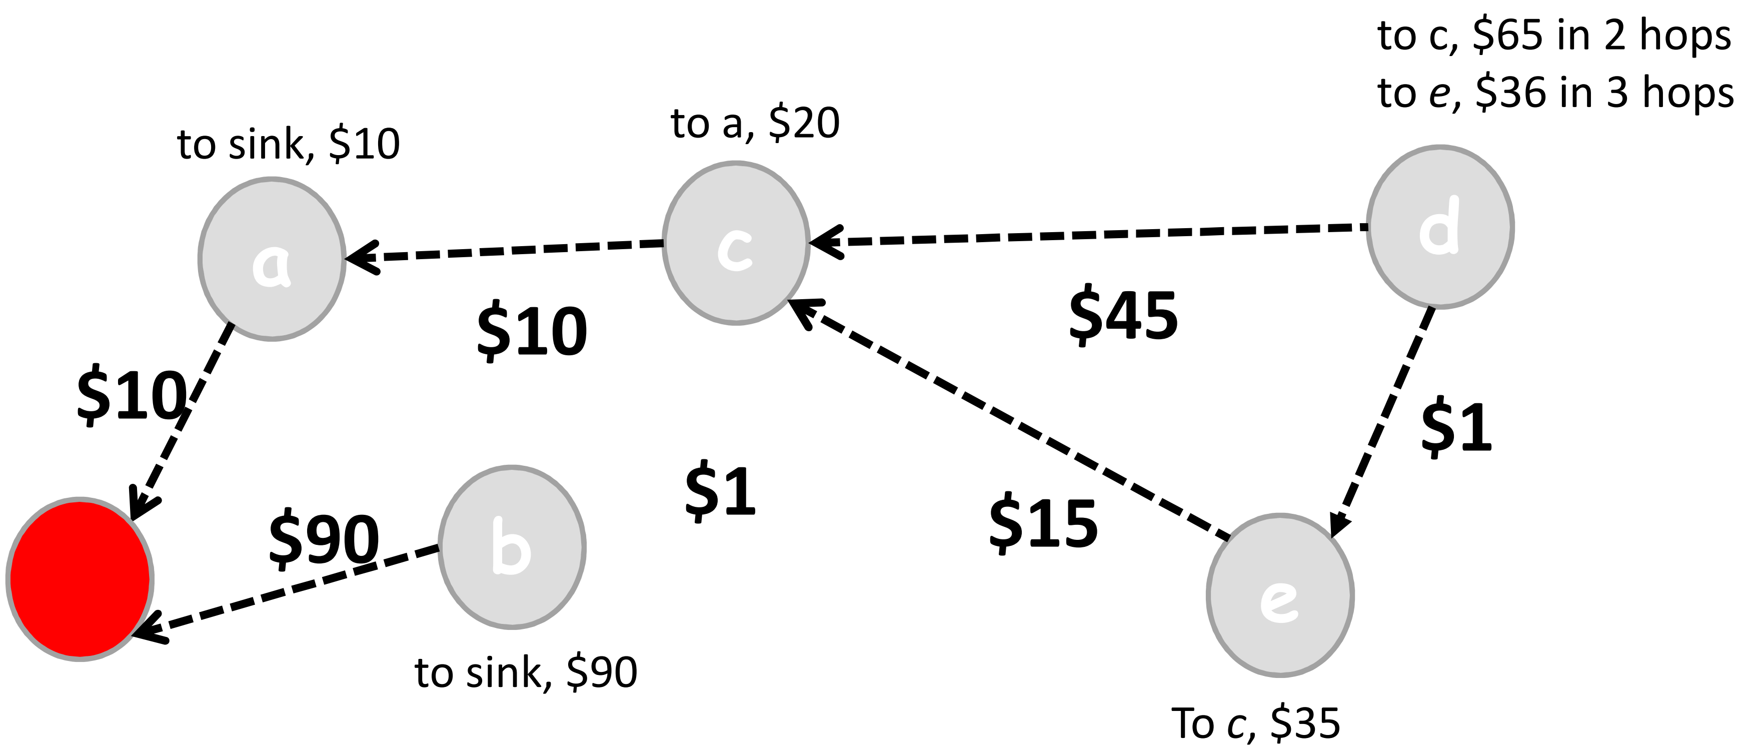

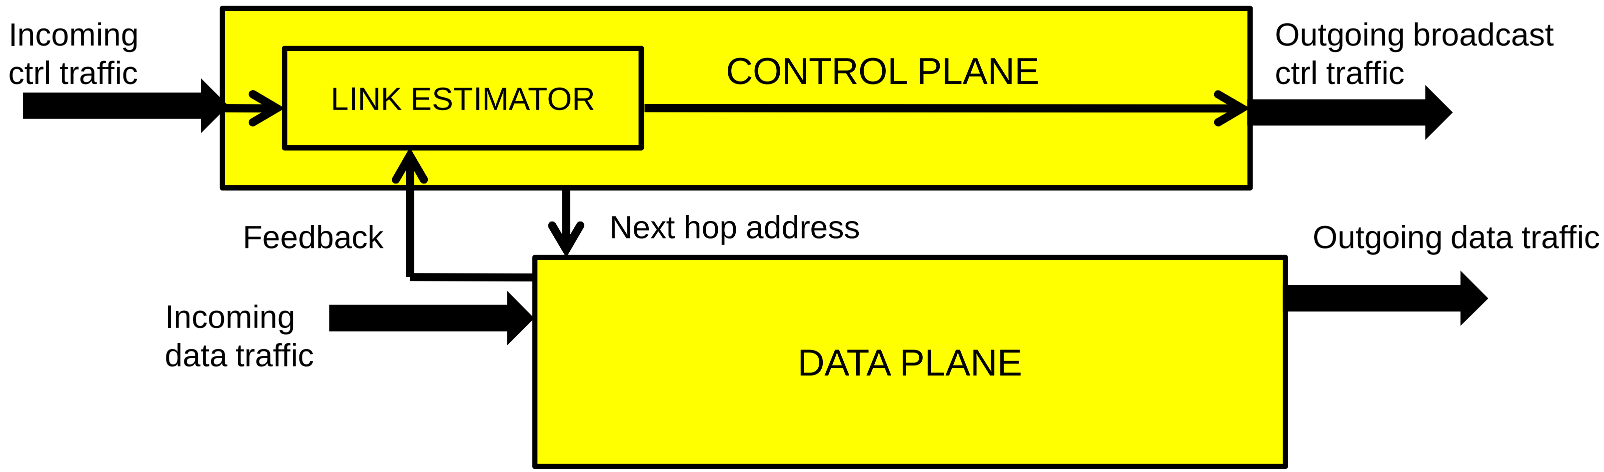

Link estimator: builds distributed cost field
Control plane: cost-based parent selection, cost field maintenance
Data plane: unconstrained retx, congestion control

Remember the RNP?

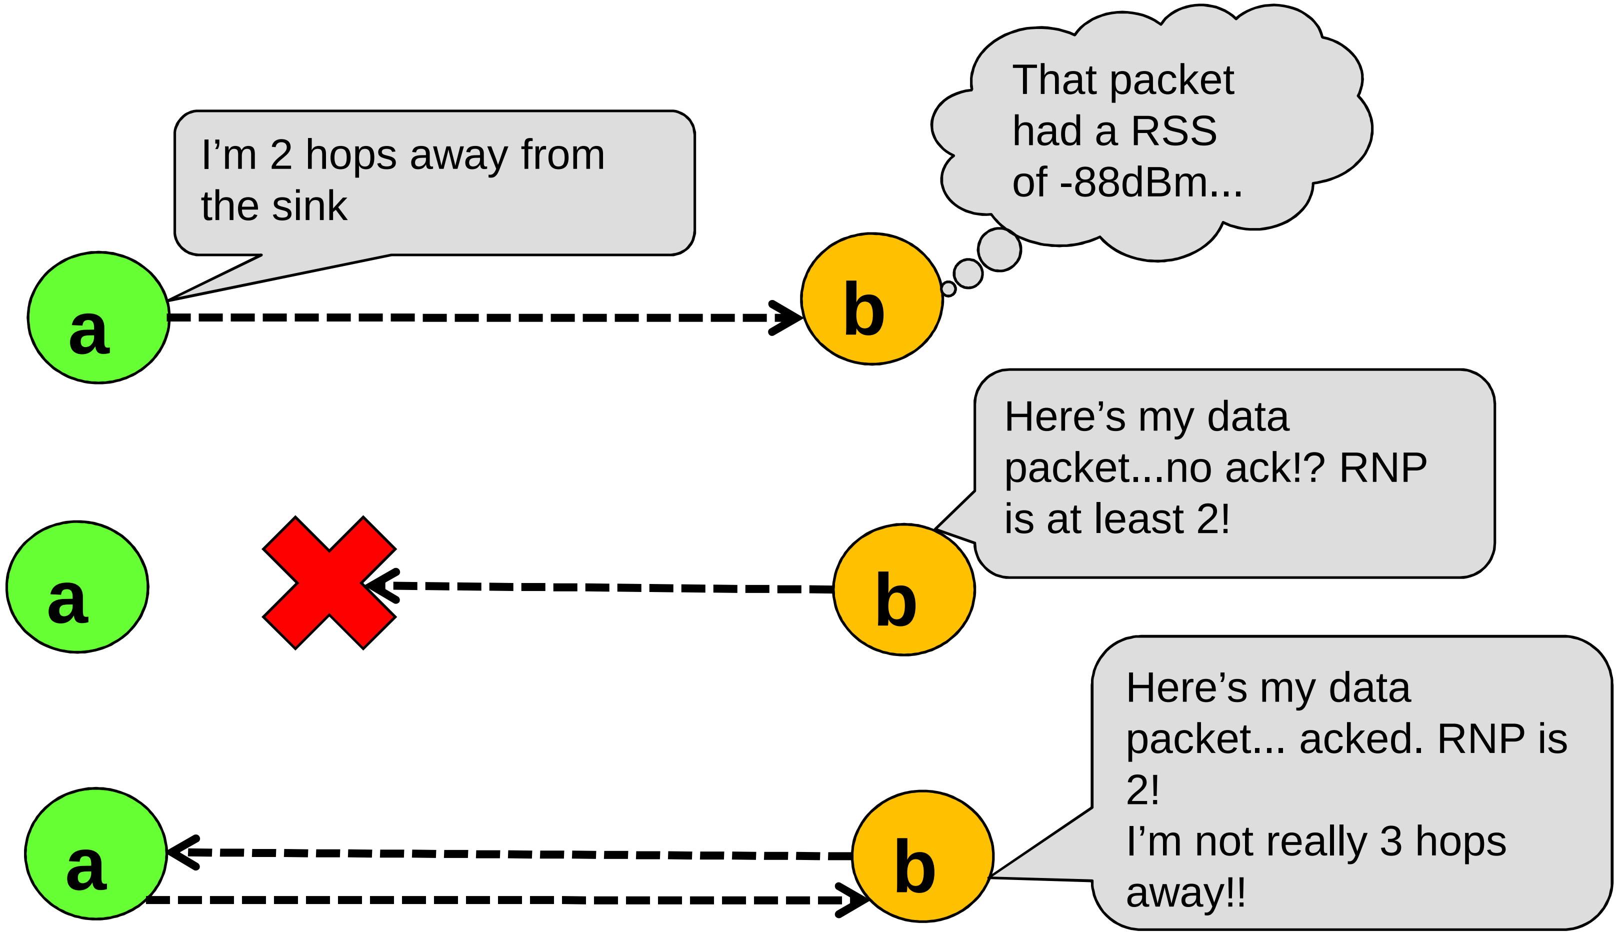

But why do we need to have this?

Asymmetric Links

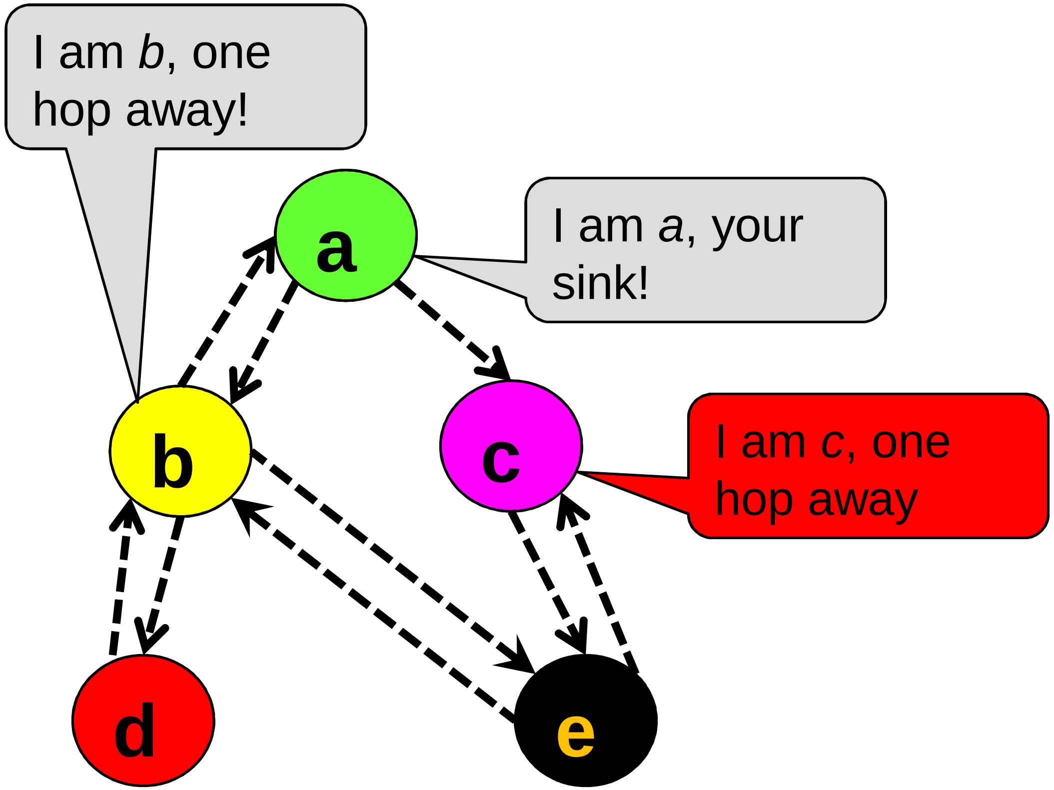

- Node *a* advertises its status as the sink
- Nodes *b* and *c* think they're one hop away
- Node *b* is, but *c* isn't!!!

- Links are NOT boolean
- Hop count doesn't work as a cost metric

Routing Decisions

- Decisions are made based on control beacons
- Control beacons enable connectivity discovery
- Control beacons provide link quality estimates (RSS, LQI, PDR)

- Control beacons cost extra energy
- Control beacons travel on a different timescale than data traffic
- Control beacons estimate the reverse link

Performance Metrics

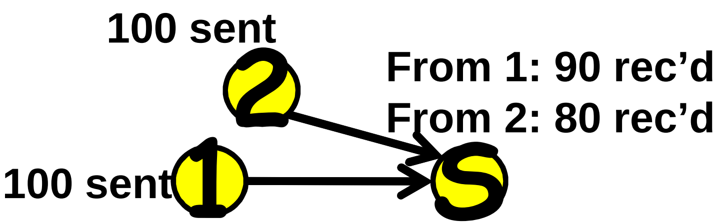

RELIABILITY

Delivery ratio at the sink

On average, 85% here

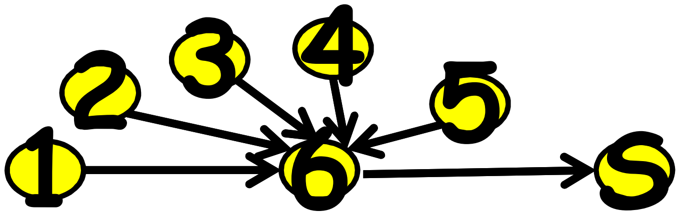

GOODPUT

Number of delivered packets per time unit

With an offered load of 1 pkt/sec/node, 6 pkts/sec

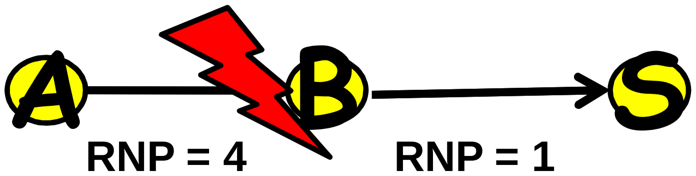

ROUTING COST

To get a packet to S, a total of 5 transmissions over 2 hops are needed

Link Metrics

What needs to be captured?
- Wide range of link dynamics
- Asymmetric links
- Inter-hop interference

Possible Metrics

Hop count?

No, there may be lossy links

Product of Path PRRs?

Doesn't account for inter-hop interference

1 hop at 90% is better than 2 hops at 100%

End-to-End Delay

Depends on the load

Expected Transmission Count

- Each transmission attempt is viewed as a Bernoulli trial
- Success probability: delivery prob. times ACK prob.
- Expected transmission count: inverse of success prob.
- ETX is a landmark metric

but...

- Delivery and ACK are not independent
- Real tx attempts experience time correlation
  Bernoulli trials are independent
- The wireless channel changes over time
  Bernoulli success probability is the same across trials

D. De Couto, D. Aguayo, J. Bicket, and R. Morris.
A High-Throughput Path Metric for Multi-Hop Wireless Routing, MobiCom'03.

More on the ETX

• Each transmission attempt is viewed as a Bernoulli trial

• Use geometric probability model to determine the expected number of tx attempts to achieve one successful delivery

• Let **X**=*x* be the number of tx attempts that you'll need
• You must have had *x-1* failures
• So, $P(\mathbf{X}=x)=(1-p)^{x-1}p$

• Now compute the expected value of **X**

$$E[X] = \sum_{x=0}^{\infty} x(1-p)^{x-1}p = \frac{1}{p}$$

MintRoute

• Goal: to get up-to-date connectivity information

MintRoute (MIninum Number of Transmission Routing)

• Use beacons to estimate the PRR
• Use the PRR to estimate the Expected Number of Transmissions (ETX)
• Filter it out to get a good tradeoff between agility and stability
• Keep a neighbor table and manage it so that it doesn't grow too big (with dense connectivity)

A. Woo, L. Tong, and D. Culler,
Taming the underlying challenges of reliable multihop routing in sensor networks, SenSys'03.

Accounting for Channel Information

• Since the vagaries of RF have such a huge impact, use channel state information in your link metric

• Radios typically make RSS available
• Use it to get soft information about links
• Can tell noise from in-network interference

• RSS-based blacklisting?
  Possible network partitioning
  Unclear where to place the threshold

• With 802.15.4, you also get LQI (*a.k.a.* CCI)

LQI = Link Quality Indicator
CCI = Chip Correlation Indicator

Basic PHY Information

- The *de facto* standard PHY for WSNs follows 802.15.4
- Uses (a variant of) Direct Sequence Spread Spectrum (DSSS)

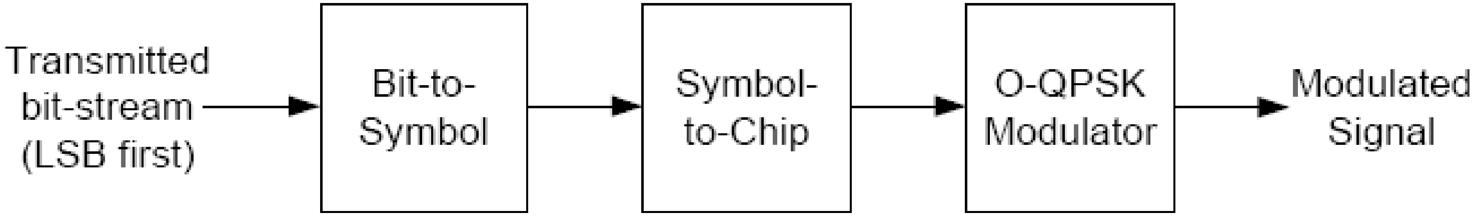

- Break each byte into two 4-bit symbols
- Map each symbol to a predetermined 32-bit chip sequence

Symbol	Chip sequence (C_0, C_1, C_2, … , C_{31})
0	1 1 0 1 1 0 0 1 1 1 0 0 0 0 1 1 0 1 0 1 0 0 1 0 0 0 1 0 1 1 1 0
1	1 1 1 0 1 1 0 1 1 0 0 1 1 1 0 0 0 0 1 1 0 1 0 1 0 0 1 0 0 0 1 0
2	0 0 1 0 1 1 1 0 1 1 0 1 1 0 0 1 1 1 0 0 0 0 1 1 0 1 0 1 0 0 1 0
3	0 0 1 0 0 0 1 0 1 1 1 0 1 1 0 1 1 0 0 1 1 1 0 0 0 0 1 1 0 1 0 1

Basic PHY Information (cont’d)

- O-QPSK = Offset Quadrature Phase Shift Keying
- Modulate even-indexed chips onto the In-Phase carrier
- Modulate odd-indexed chips onto the In-Quadrature carrier

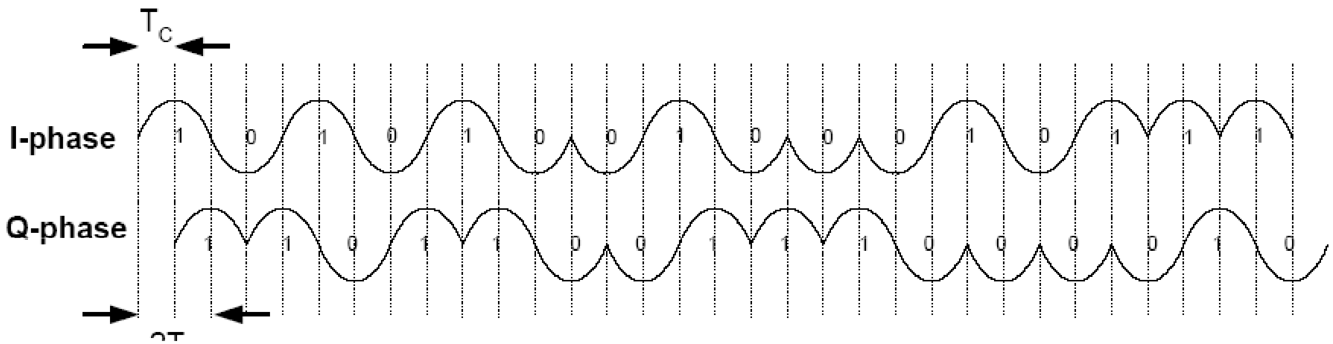

- Transmit at 2MChips/sec
- At the receiver, demodulate
- Map received chip sequence to known chip sequences
- CCI = correlation between received chip sequence and known chip sequence that the received sequence gets mapped to

LQI (CCI) and RSS

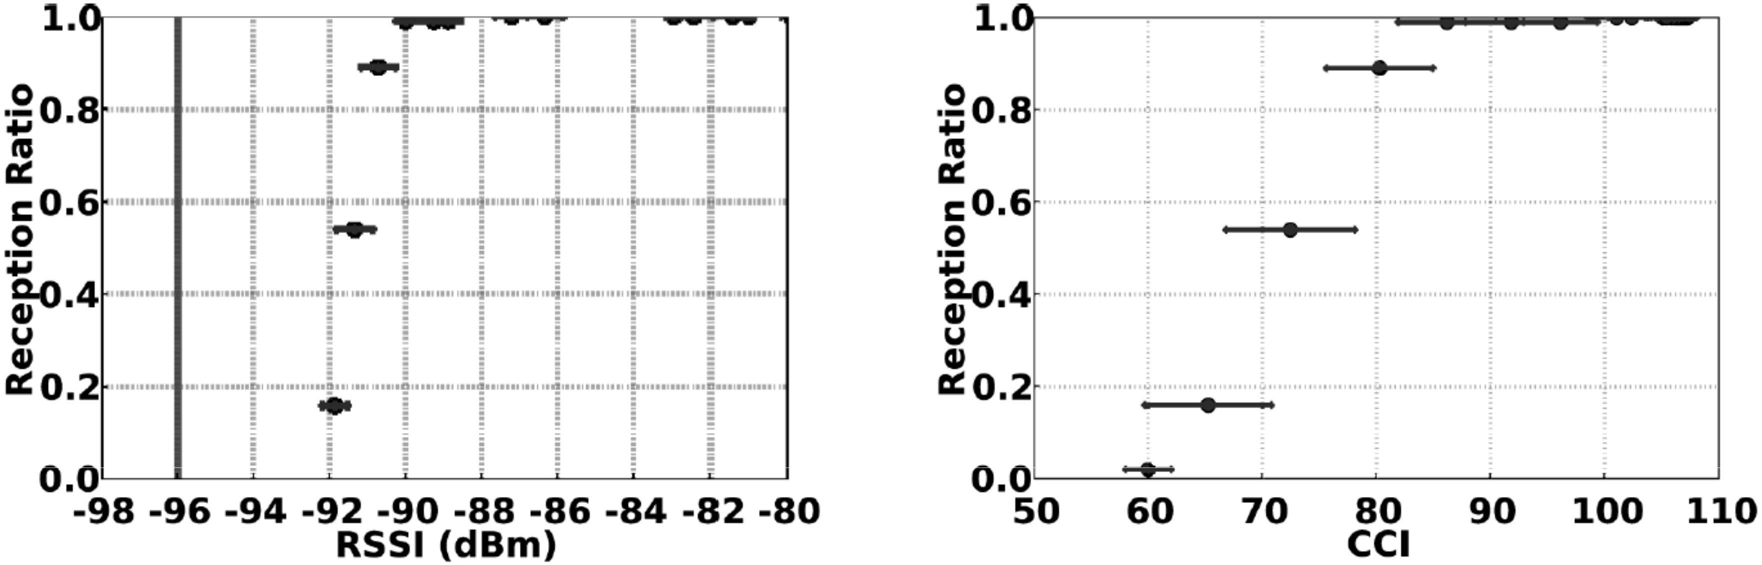

Clean PRR/RSS and PRR/LQI curves obtained using wired MICAz motes and controlled attenuation

K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis, "**An Empirical Study of Low-Power Wireless**", ACM Transactions on Sensor Networks, 2010

Link Quality Indicator

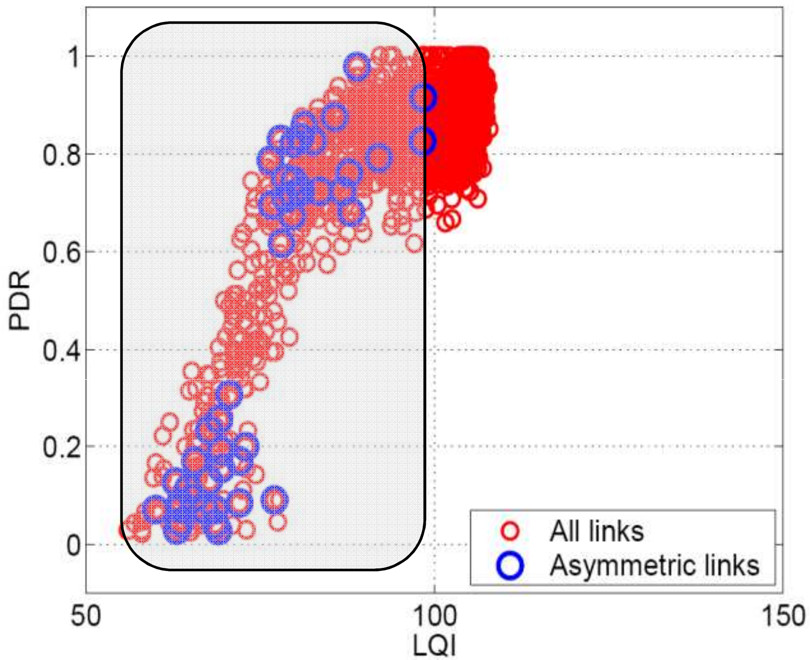

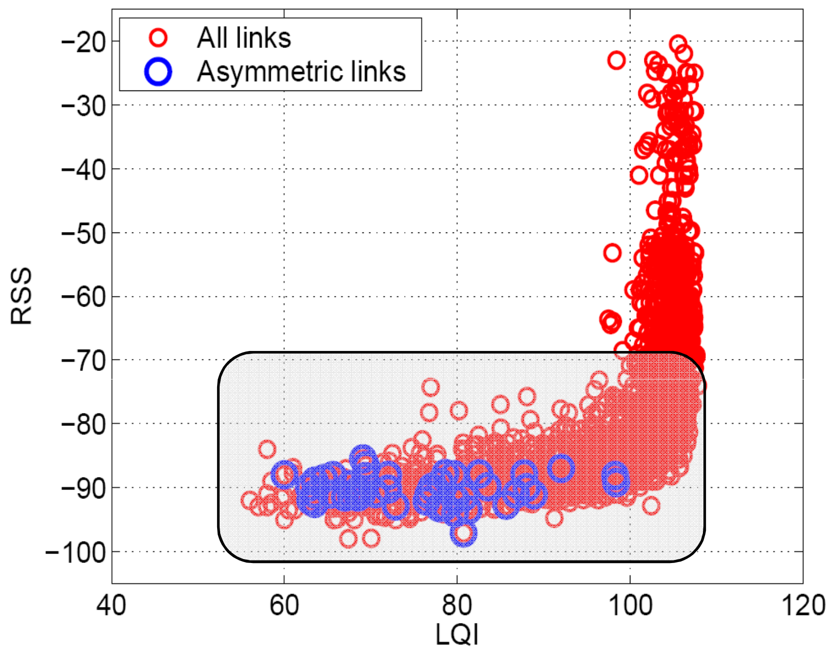

- LQI: soft information about bad links
- RSS: soft information about good link

Link Quality Indicator

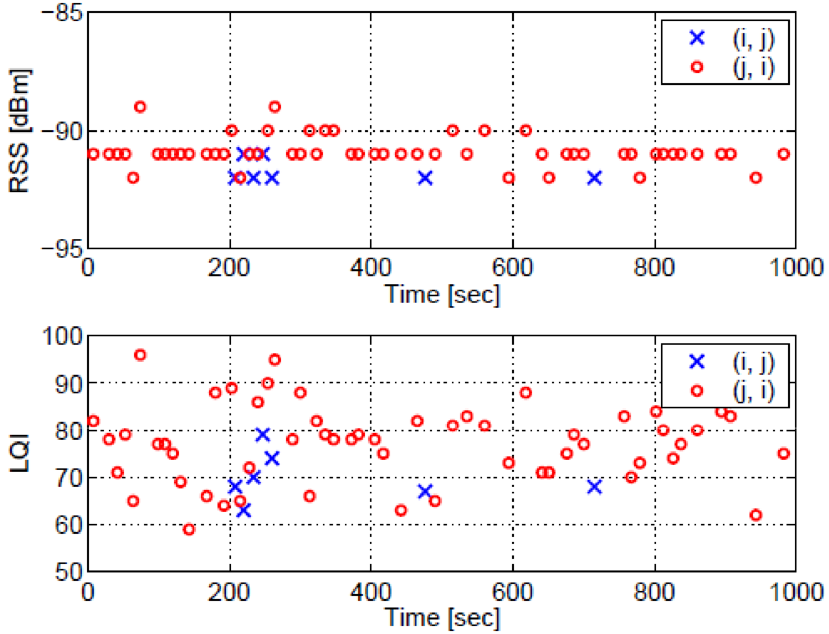

D. Puccinelli and M. Haenggi **DUCHY: Double Cost Field Hybrid Link Estimation**, HotEmNets'08

- LQI: soft information about asymmetric links

Channel Information is Not Enough

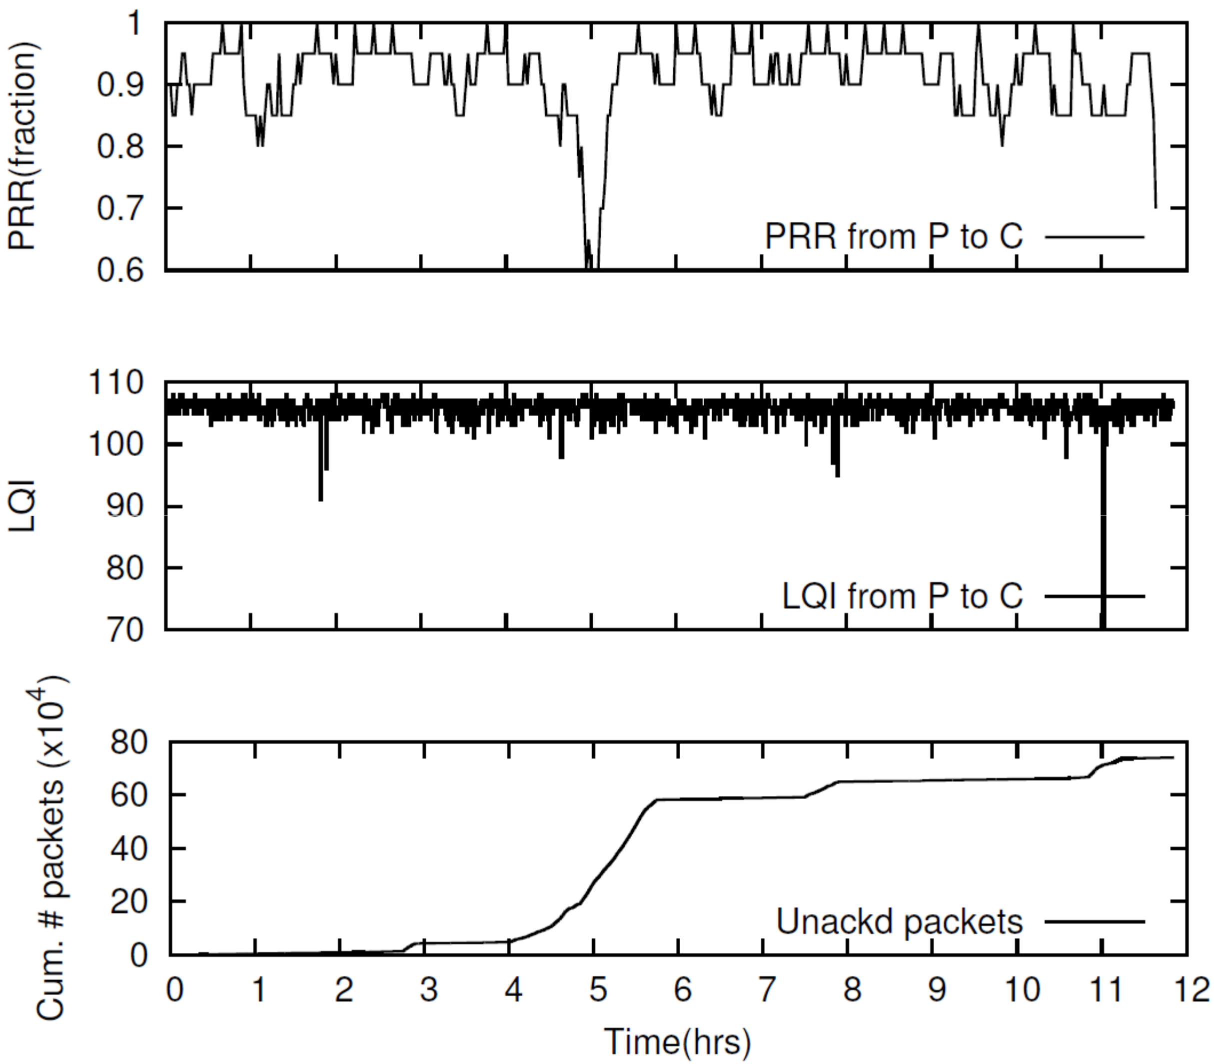

R. Fonseca,
O. Gnawali,
K. Jamieson,
P. Levis
Four-Bit Link Estimation,
HotNets'07

Four-Bit Link Estimator

• A hybrid estimator using control and data traffic

Four Bit Link Estimation

• Fuse information from each layer
• PHY: was it easy to decode the packet over this link?
• MAC: did you get an ACK over this link?
• NET: does this link matter?

Four Bits: white, ack, compare, pin

Four-Bit gets to measure the RNP with link-layer ACKs

Assumption: each node keeps a neighbor table

R. Fonseca, O. Gnawali, K. Jamieson, P. Levis, , **Four-Bit Link Estimation**, HotNets'07

Collection Tree Protocol

- Control Plane injects broadcast beacons
- Link Estimator (Four-Bit) manages the neighbor table
- Distributed cost field gets set up
- Control Plane chooses the best route

- Data Plane forwards to the address indicated by the Control Plane
- Link-layer acknowledgments and retransmissions
- No congestion control (upper layer)

O. Gnawali et al., **Collection Tree Protocol**,SenSys'09

Congestion Control

Source rate-limiting

- Each node adapts its offered load to the parent's load
- Self-regulatory approach

Hop-by-Hop Backpressure

- Parent node asks its child nodes to slow down
- Queue-based backpressure: act before your queue fills up
- If the medium is congested the parent may not get heard

Congestion-aware MAC

- Make it easier for congested nodes to access the medium

B. Hull, K. Jamieson, H. Balakrishnan
Mitigating Congestion in Wireless Sensor Networks
SenSys'04

Arbutus

Control Plane
- Data plane feedback for link estimates (Four-Bit)
- Novel link estimator to actively enforce long-hop routing
- No neighbor tables

Data Plane
- Deferred, unconstrained retransmissions
- Congestion control services: backpressure + rate-throttling
- Congestion control management
- Loop recovery decoupled from cost field

D. Puccinelli and M. Haenggi
Reliable Data Delivery in Large-Scale Low-Power Sensor Networks
ACM Transactions on Sensor Networks, 2010

Performance Variations

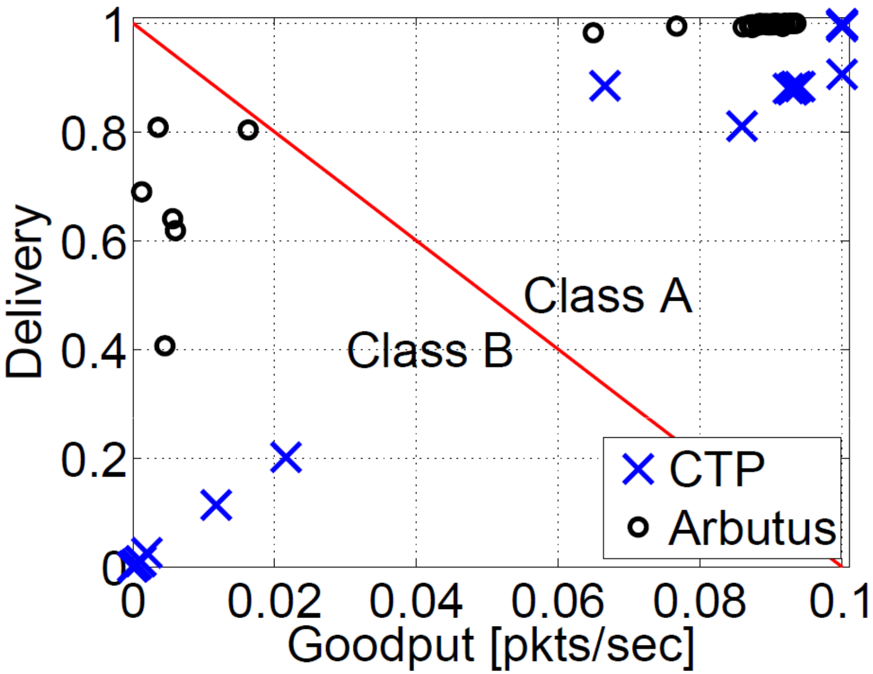

D. Puccinelli, O. Gnawali, S. Yoon, S. Santini, U. Colesanti, S. Giordano, and L. Guibas
The Impact of Network Topology on Collection Performance
EWSN 2011

- All points come from the same network
- Each point is a run with a different sink
- Huge performance variations with both protocols

- Link dynamics change over time (transitional links)
- Different sinks see different networks

Fair Performance Comparisons

- Open research question
- Critical and timley: most experimental WSN research papers are centered around a comparison among protocols

- I run protocol A on network N using sink s from 12pm to 1pm
- You run protocol B on network N using sink s from 2pm to 3pm
- A delivers 99% of the injected traffic
- B only delivers 80%
- Can we conclude that A is better than B?

No! It could have been the network's fault

- A fair comparison requires similar network conditions

Similar Network Conditions?

- Even if we use the same network, we cannot control the link dynamics
- Run multiple protocols concurrently [1]
  - Interplay between protocols may affect performance
- Explicitly capture the state of your network as each protocol is run [2]
  - Distil the network down to a protocol-independent topology metric
  - Show both the performance results and the topology metric
  - Only leverage the protocol's traffic to measure the network state

[1] O. Gnawali, L. Guibas, and P. Levis, **A Case for Evaluating Sensor Network Protocols Concurrently**, WINTECH 2010

[2] D. Puccinelli, O. Gnawali, S. Yoon, S. Santini, U. Colesanti, S. Giordano, and L. Guibas, **The Impact of Network Topology on Collection Performance,** EWSN 2011

Expected Network Delivery (END)

• Use the protocol's traffic to estimate the link PRRs
• Compute the optimal routes (with respect to delivery) offline using Dijkstra or Bellman-Ford
• Compute the Expected Path Delivery from each node to the sink
• Average out and get your END (Expected Network Delivery)

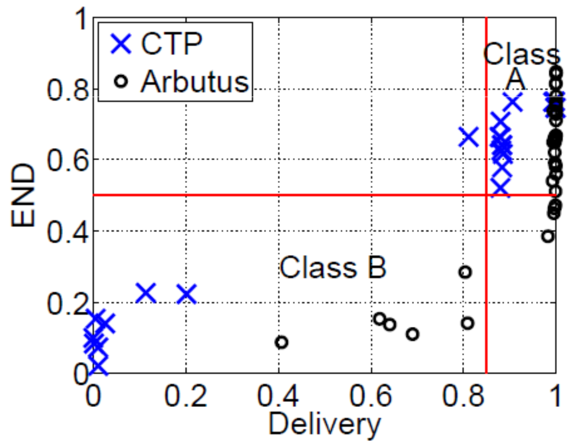

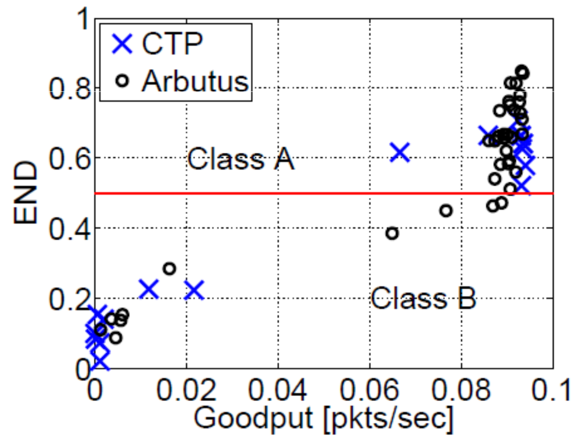

D. Puccinelli, O. Gnawali, S. Yoon, S. Santini, U. Colesanti, S. Giordano, and L. Guibas,
The Impact of Network Topology on Collection Performance, EWSN 2011

The END is not where it ends

- Not completely protocol-independent
- The END only views the network through the protocol's eyes

- The measured network state cannot be ground-truthed without interfering with the protocol under test

- Works well for short experiments, but for longer ones it should be a function of time (remember bimodal links)

- Does not completely account for inter-link interplay

D. Puccinelli, O. Gnawali, S. Yoon, S. Santini, U. Colesanti, S. Giordano, and L. Guibas,
The Impact of Network Topology on Collection Performance, EWSN 2011