Network Architectures
Sensor Networks
Lecture 2:
Sensor Networks Routing

Dr. Cecilia Mascolo

UNIVERSITY OF CAMBRIDGE

What is covered in this lecture

- We will first cover two example of routing protocols for sensor networks
  - Directed diffusion
  - MintRoute

Network Protocols

- Can we apply ad hoc networks protocols?
- Yes protocols like epidemic can be applied but overhead is an issue
- Aims are usually different: not communication but data reporting to single or multiple source

- Specific protocols have been devised
- Specific nodes are interested in specific events
  - Sink interested in all results
  - Sink interested in a sensor reading change

UNIVERSITY OF CAMBRIDGE

Routing in Sensor Networks

- In the previous lecture have seen an example of static tree routing

- In general sensors are producers and sinks are consumers
  - A bit like a publish/subscribe system
- Information tends to flow one way
  - Exceptions are reprogramming and actuation

Directed Diffusion: A scalable and robust communication paradigm for sensor networks

Intanagonwiwat, C., Govindan, R., and Estrin, D.
Mobicom 2000.

UNIVERSITY OF CAMBRIDGE

Directed Diffusion

- Data-centric approach
- Nodes send “interests” for data which are diffused in the network
- Sensors produce data which is routed according to interests
- Intermediate nodes can filter/aggregate data

UNIVERSITY OF CAMBRIDGE

Interest Propagation

- Each sink sends expression of interests to neighbours
- Each node will store interests and disseminate those further to their neighbours.
  - Cache of interest is checked not to repeat disseminations
- Interests need refreshing from the sink [they time out]
- Interests have a “rate of events” which is defined as “**gradient**”

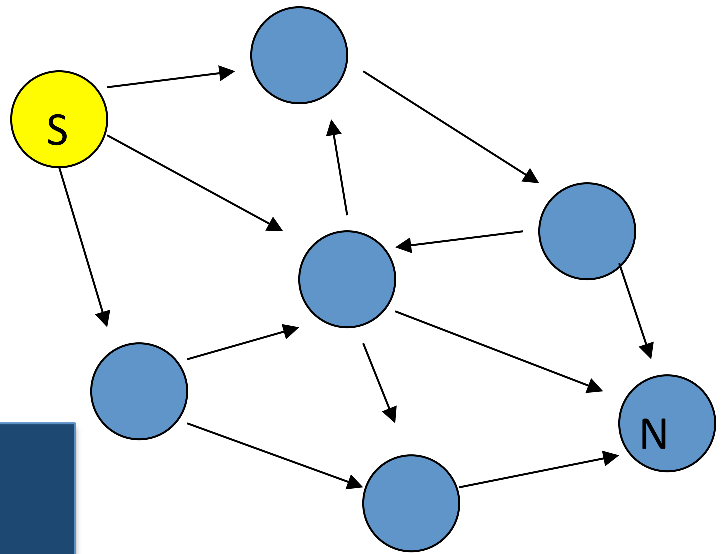

Interest:
Type: animal
Interval= 20 msec
Rect=[-100,100,200,400]
Duration: 10 min

UNIVERSITY OF CAMBRIDGE

Data delivery

- Sensor data sources emit events which are sent to neighbours according to interest [ie if there is a gradient]
- Each intermediate node sends back data at a rate which depends on the gradient
  - Ie if gradient is 1 event per second and 2 events per second are received send either the first or a combination of the two [aggregation]
- Events are stored to avoid cycles [check if same event received before]
- Data can reach a node through different paths. Gradient enforcement needed

UNIVERSITY OF CAMBRIDGE

Gradients Reinforcement

- When gradients are established the rate is defined provisionally [usually low]
- Sinks will ‘reinforce’ good paths which will be followed with higher rate
- A path expires after a timeout so if not reinforced it will cease to exist
  - This allows adaptation to changes and failures

Directed diffusion Two-phase pull

- **Phase 1**: nodes distribute ***interests*** in certain kinds of named data
  - Specified as attribute-value pairs
- Interests are flooded in the network
  - Apparently obvious solution: remember from where interests came, set up a convergecast tree
  - Problem: Node X cannot distinguish, in absence of unique identifiers, between the two situations on the right – set up only one or three convergecast trees?

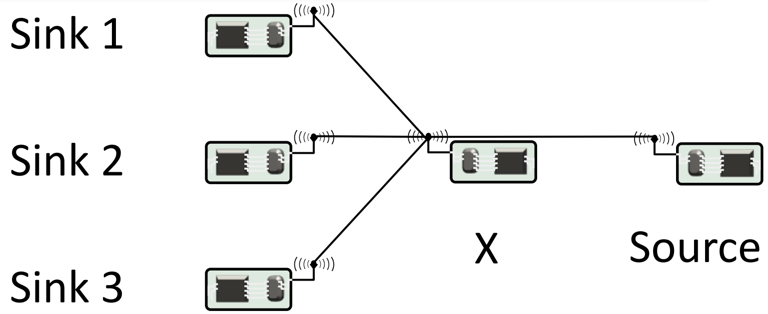

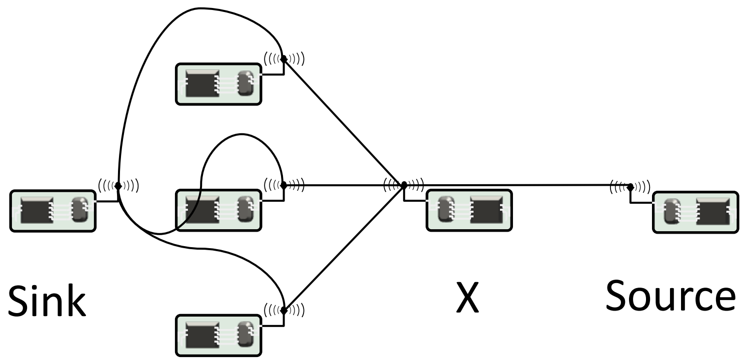

UNIVERSITY OF CAMBRIDGE

Direction diffusion: Gradients in two-phase pull

- Option 1: Node X forwarding received data to all “parents” in a “convergecast tree”
  - Not attractive, many needless packet repetitions over multiple routes
- Option 2: node X only forwards to one parent
  - Not acceptable, data sinks might miss events
- Option 3: Only provisionally send data to all parents, but ask data sinks to help in selecting which paths are redundant, which are needed
  - Information from where an interest came is called ***gradient***
  - Forward all published data along all existing gradients

Evaluation

- ns2 simulation
- Modified 802.11 MAC for energy use calculation
- Comparison against flooding and omniscient multicast
- Experiment with node failure
- Did not overload system
- Standard random node placement (but only 3 hops across entire topology)

UNIVERSITY OF CAMBRIDGE

Metrics

- Average dissipated energy
  - Ratio of total energy expended per node to number of distinct events received at sink
- Average delay
  - Average one-way latency between event transmission and reception at sink
- Both measured as functions of network size

Topology

- 50-250 nodes in 50 node increments
- Avg. Node density constant with network size
- Square of 160m, radio range of 40m
- 5 sources, 5 sinks uniformly distributed
- 1.6Mbps 802.11 MAC
  - reliable transmission, RTS/CTS, high power, idle power ~ receive power
  - Set idle power to 10% of receive power, 5% of transmit power
  - Tried also with comparable power consumption for idle/receiving

Sim: Average energy and delay

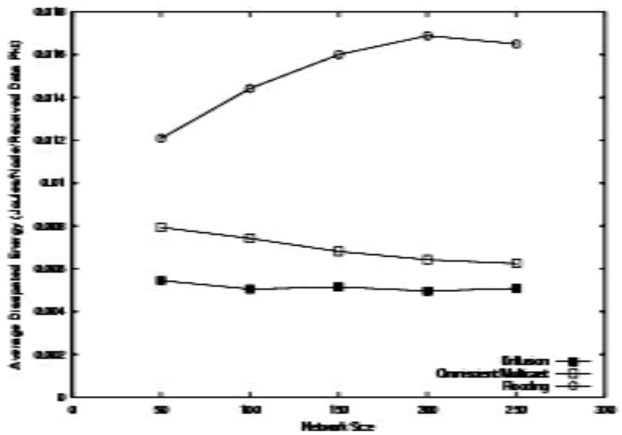

(a) Average dissipated energy

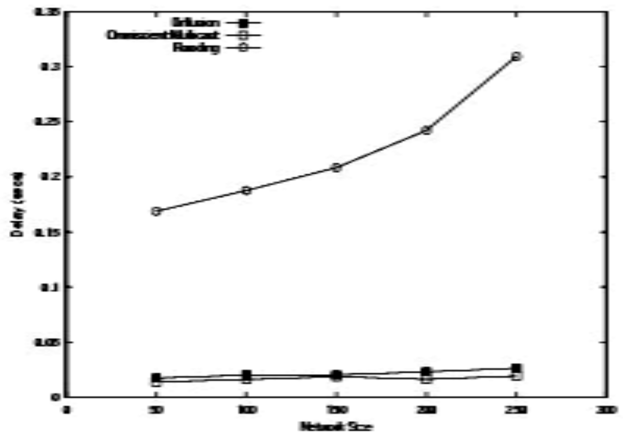

(b) Average delay

Figure 4: Directed diffusion compared to flooding and omniscient multicast.

UNIVERSITY OF CAMBRIDGE

Sim: Average energy and delay

- Directed Diffusion is better than Omniscient Multicast…
  - Omniscient multicast sends duplicate messages over the same paths
    - Why not suppress messages with Omniscient Multicast just as in Directed Diffusion?

Sim: Failures

- Dynamic failures (no settling time), adverse network conditions (10-20% failure at any time)
- Each source sends different signals
- <20% delay increase, fairly robust
- Energy efficiency *improves*:
  - Reinforcement maintains adequate number of high quality paths

UNIVERSITY OF CAMBRIDGE

Directed diffusion – extensions

- Problem: Interests are flooded through the network

- Geographic scoping & directed diffusion
  - Interest in data from specific areas should be sent to sources in specific geo locations only
- Push diffusion – few senders, many receivers
  - Same interface/naming concept, but different routing protocol
  - Here: do not flood interests, but flood the (relatively few) data
  - Interested nodes will start reinforcing the gradients

Issues

- Purely theoretical work
- A part from the flooding of the interests…

- MAC Layer issues (assume nodes are awake…
- or does not discuss it)

- More recent approaches have considered more directly link capabilities as part of the routing decision making

UNIVERSITY OF CAMBRIDGE

Taming the underlying challeges of reliable multihop routing in sensor networks

A. Woo, T. Tong, D. Culler. ACM Sensys 2003.

UNIVERSITY OF CAMBRIDGE

Link Estimation

- Routing algorithms should take into account underlying network factors and under realistic loads.
- Link connectivity in reality is not spherical as often assumed

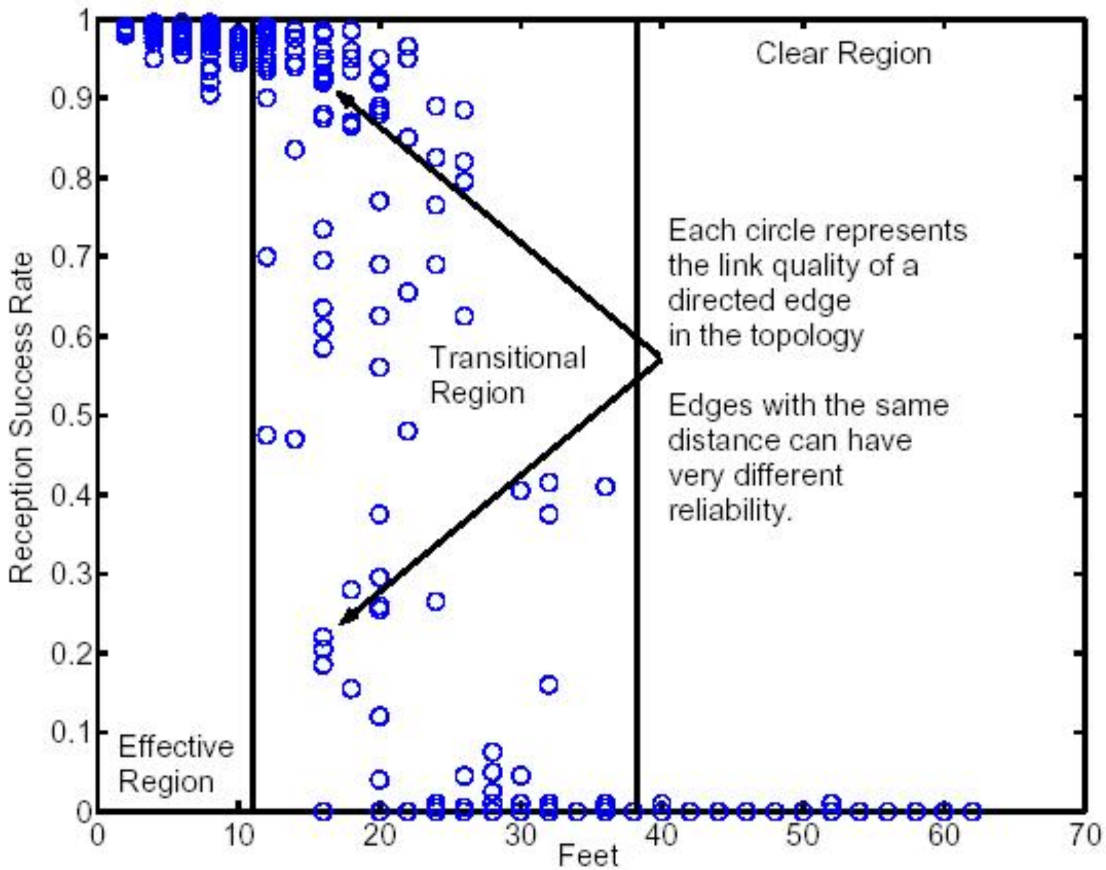

UNIVERSITY OF CAMBRIDGE

Link Estimation

- A good estimator in this setting must
  - Be stable
  - Be simple to compute and have a low memory footprint
  - React quickly to large changes in quality

  - Neighbour broadcast can be used to passively estimate

WMEWMA

- Snooping
  - Track the sequence numbers of the packets from each source to infer losses

- Window mean with EWMA
  - MA(t) = (#packets received in t) / max(#packets expected in t, packets received in t)
  - $EWMA(t_x)=a\ (MA(t_x)) + (a-1)EWMA(t_{(x-1)})$
  - t_x : last time interval; a: weight

UNIVERSITY OF CAMBRIDGE

WMEWA (t =30, a =0.6)

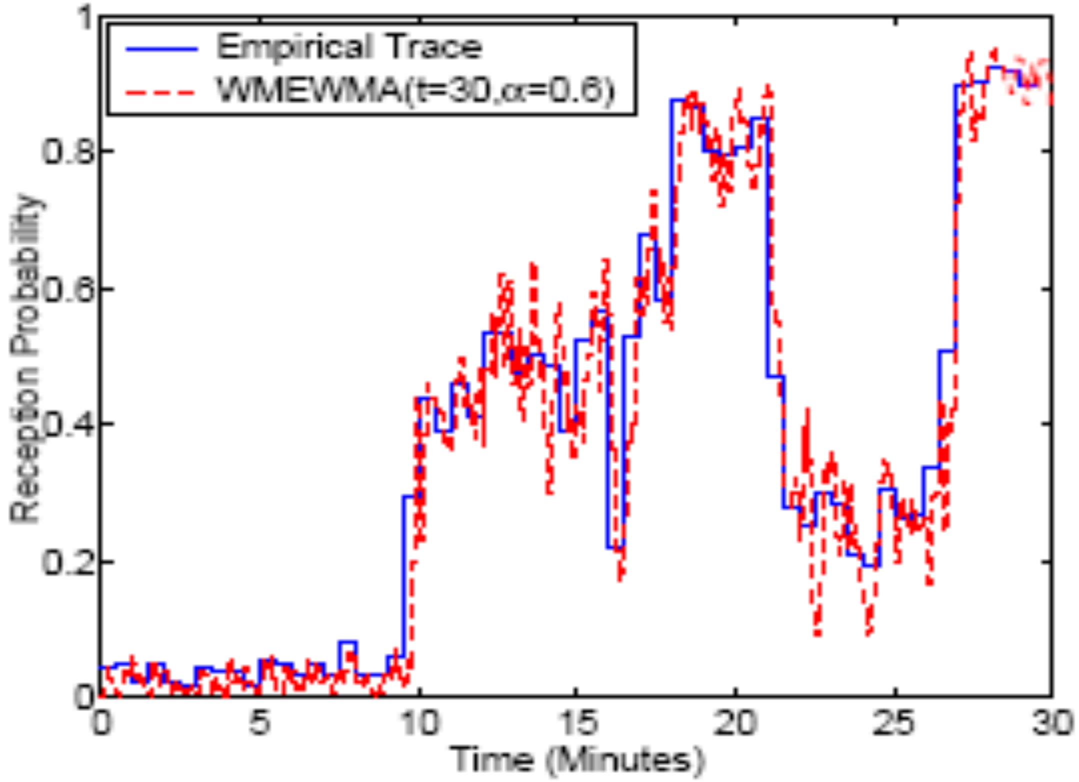

UNIVERSITY OF CAMBRIDGE

Neighborhood Management

- Neighborhood table
  - Record information about nodes from which it receives packets (also through snooping)
- If network is dense, how does a node determine which nodes it should keep in the table?
- Keep a sufficient number of good neighbours in the table
- Similar to cache management for packet classes

UNIVERSITY OF CAMBRIDGE

Link Estimation based Routing

- Focus on “many to one” routing model
  - Information flows one way
- Estimates of inbound links are maintained, however outbound links need to be used!
  - Propagation back to neighbours
- Each node selects a parent [using the link estimation table]
  - Changes when link deteriorates (periodically)

UNIVERSITY OF CAMBRIDGE

Distance vector routing: cost metrics

- Routing works as a standard distance vector routing
- The DVR cost metric is usually the hop count

- In lossy networks hop count might underestimate costs
  - Retransmissions on bad links: shortest path with bad links might be worse than longer path with good links

  - Solution: consider the cost of retransmission on the whole path

UNIVERSITY OF CAMBRIDGE

MIN-T

- MT (Minimum Transmission) metric:
  - Expected number of transmissions along the path
  - For each link, MT cost is estimated by 1/(Forward link quality) * 1/(Backward link quality)
    - backward links are important for acks
- Use DVR with the usual hop counts and MT weights on links

En Example

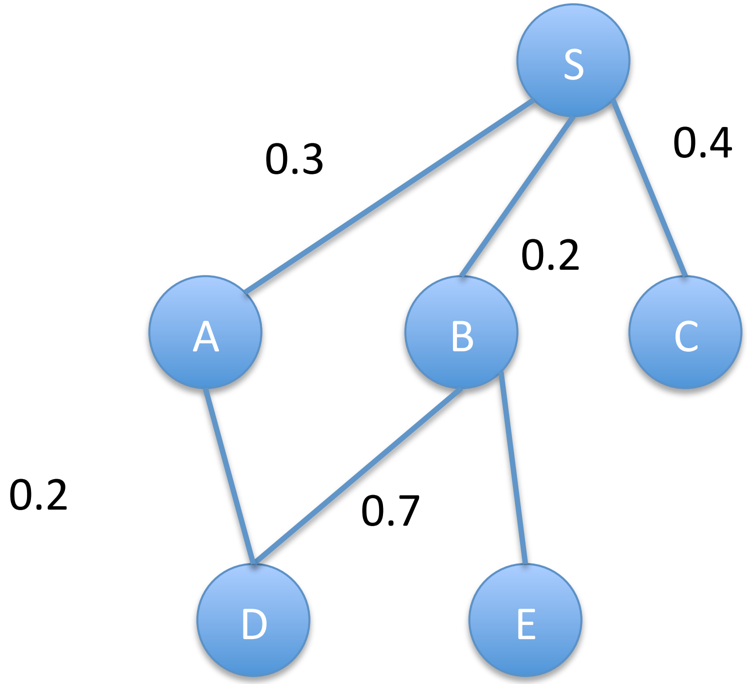

Routing Table on D:

Id	Cost	NextHop
A	0.2	A
B	0.7	B
S	0.5	A

Performance Evaluation: Tested Routing Algorithms

- Shortest Path
  - Pick a minimum hop count neighbour
  - SP: A node is a neighbour if a packet is received from it
  - SP(t): A node is a neighbour if its link quality exceeds the threshold t
    - t = 70%: only consider the links in the effective region
    - t = 40%: also consider good links in the transitional region

UNIVERSITY OF CAMBRIDGE

continued

- Minimum Transmission (MT)
- Broadcast
  - Periodic flooding: Choose a parent based on the source address of the 1st flooding message in each epoch
- Destination Sequence Distance Vector (DSDV)
  - Choose a parent based on the freshest sequence number from the root
  - Maintain a minimum hop count when possible
  - Ignore link quality – consider a node a neighbour once heard from it
  - Periodically reevaluate

UNIVERSITY OF CAMBRIDGE

Packet level simulations

- Built a discrete time, event-driven simulator in Matlab

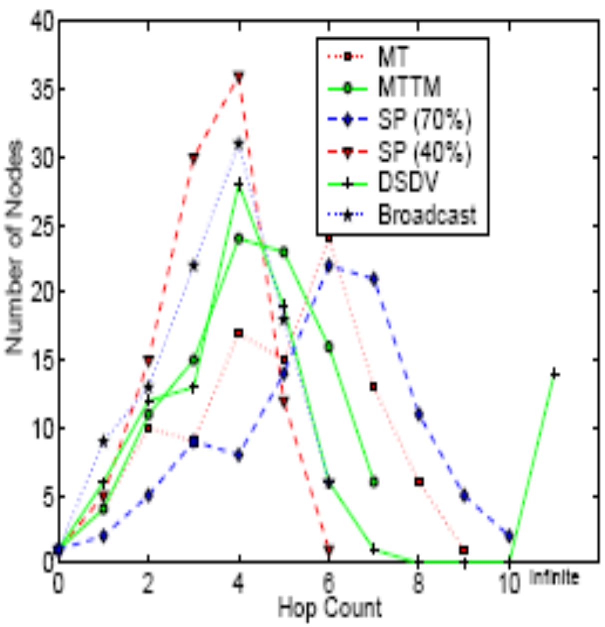

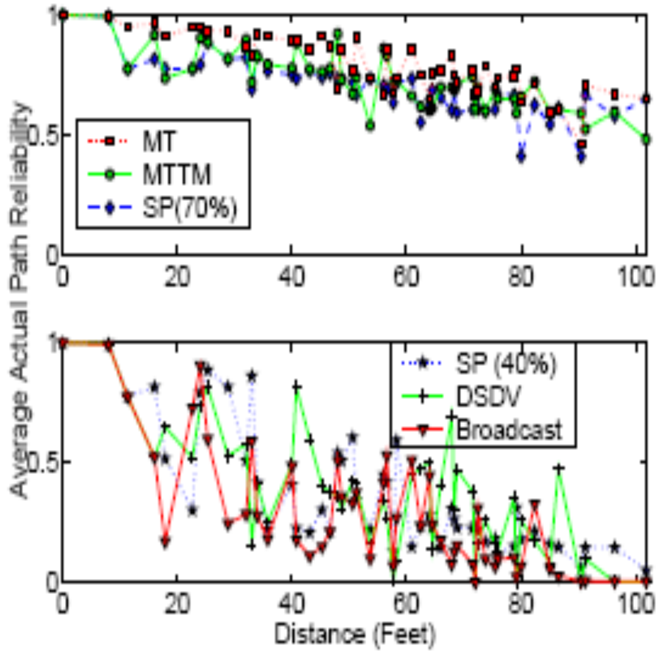

UNIVERSITY OF CAMBRIDGE

Empirical study of a sensor field

- Evaluate SP(40%), SP(70%), MT
- 50 Berkeley motes
- 5 * 10 grid w/ 8 foot spacing
  - 90% link quality in 8 feet
- 3 inches above the ground

Results

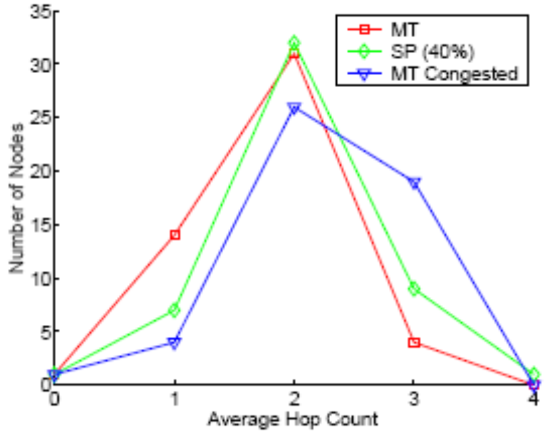

Hop Distribution

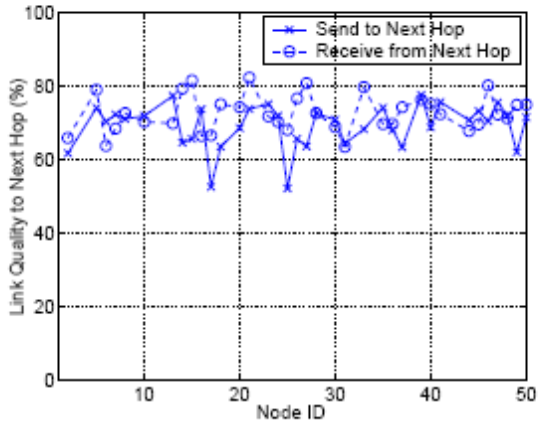

Link Quality of MT

UNIVERSITY OF CAMBRIDGE

Results

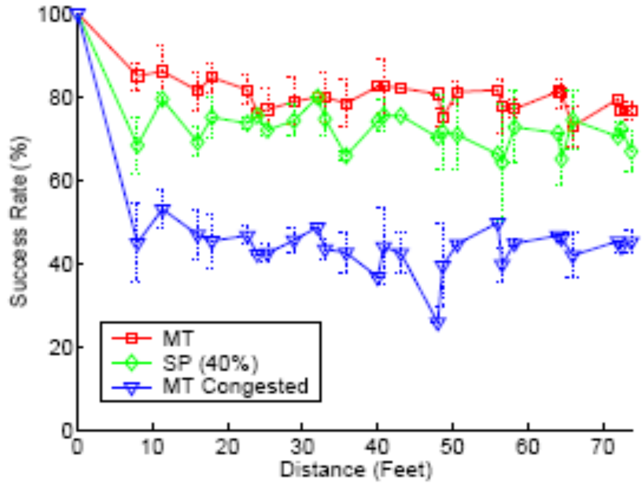

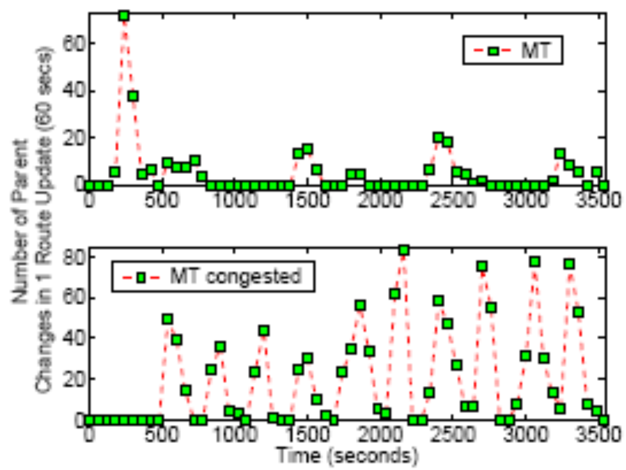

E2E success rate

Stability

UNIVERSITY OF CAMBRIDGE

Irregular Indoor Network

- 30 nodes scattered around an indoor office of 1000ft^2

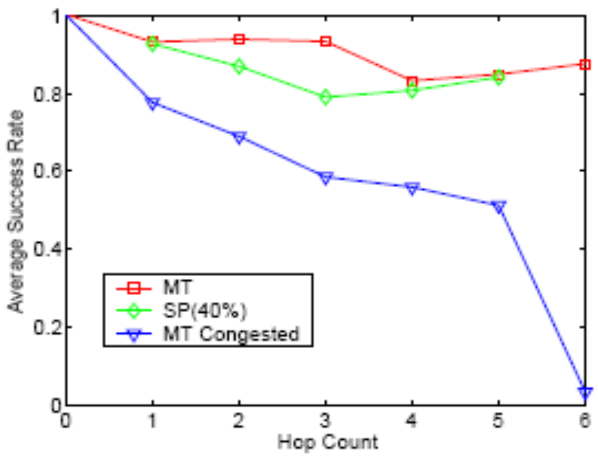

E2E Success Rate

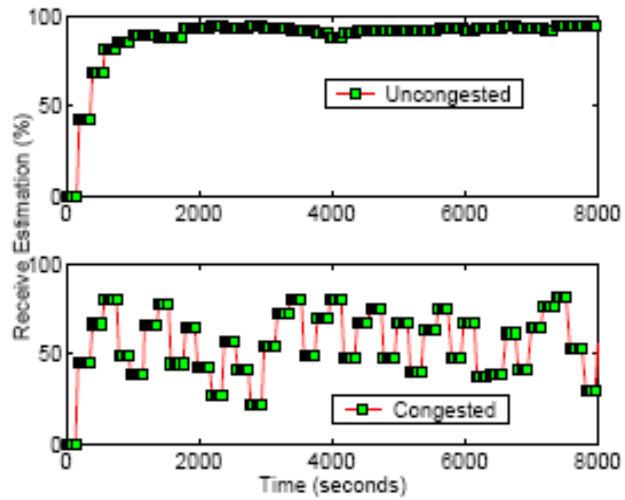

Link Estimation

UNIVERSITY OF CAMBRIDGE

Conclusions

- Link quality estimation and neighborhood management are essential to reliable routing
  - WMEWMA is a simple, memory efficient estimator that reacts quickly yet relatively stable

- A follow up paper:

Gnawali, O., Fonseca, R., Jamieson, K., Moss, D., and Levis, P. 2009. Collection tree protocol. In *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems* (Berkeley, California, November 04 - 06, 2009). SenSys '09. ACM, New York, NY, 1-14.

UNIVERSITY OF CAMBRIDGE

References

- Intanagonwiwat, C., Govindan, R., and Estrin, D. 2000. Directed diffusion: a scalable and robust communication paradigm for sensor networks. In Proceedings of the 6th Annual international Conference on Mobile Computing and Networking (Boston, Massachusetts, United States, August 06 - 11, 2000). MobiCom '00. ACM, New York, NY, 56-67.

- Woo, A., Tong, T., and Culler, D. 2003. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Proceedings of the 1st international Conference on Embedded Networked Sensor Systems* (Los Angeles, California, USA, November 05 - 07, 2003). SenSys '03. ACM, New York, NY, 14-27.

UNIVERSITY OF CAMBRIDGE