Evaluating Sensor Network Protocols

Phil Levis

ICIR Lunch, 22.i.2003

Protocol Design and Evaluation

- Opportunities abound
- Evaluation needs
  - Controlled experiments
  - Comparative results
- Real world for sensor networks
  - Uncontrollable, noisy
  - Time consuming (deployments)
- Simulation!

But What Do We Simulate?

Not The Internet…

- Application/protocol interactions
- Tiny resources
- Network conditions
- *Simulation must capture as much of this as possible*
  - Emulate empirical results

Outline

- Simulation
- Empirical network observations
- TinyOS approach
- Conclusion

Simulation

- Approximates the real world
- What approximations are needed?
  - Resources
  - Empirical data
    - Link asymmetry
    - Flooding is harmful
    - Long-term behavior

Resource Limits

- 512 bytes to 4K of RAM
  - All state and the stack
  - Major determining factor on use
  - Dynamic allocation problematic
  - No large routing tables…
- 40 packets/sec (30 byte payload)

Small Scale Experiments

- Packet loss rates over distance
  - Alec Woo
- Effects of flooding
  - Alec Woo, Deepak Ganesan, et al.

Link Asymmetry

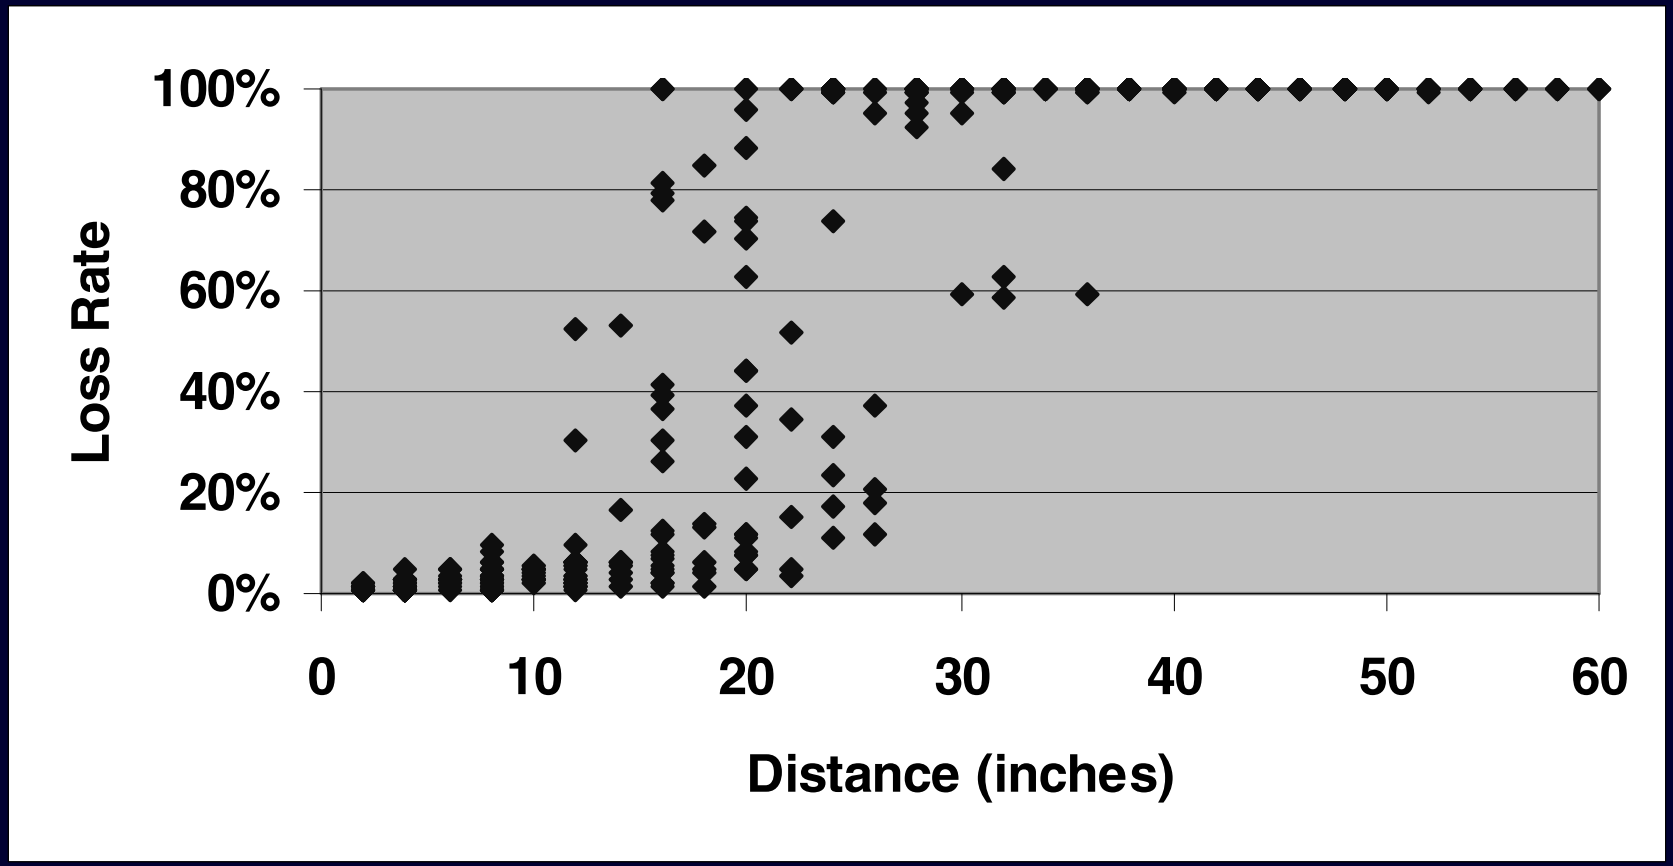

The Flooding Effect

- Overlapping cells
  - Hidden terminal effect
- Collisions ➡ stragglers
- Tree propagation
  - Folds back on itself
  - Rebounds from the edge
  - Picks up stragglers.

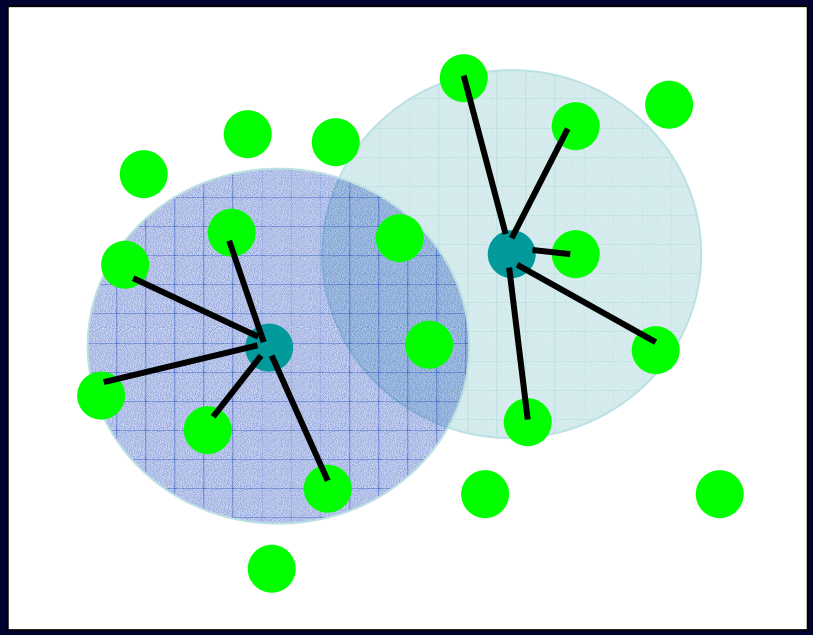

Flooding Is Harmful

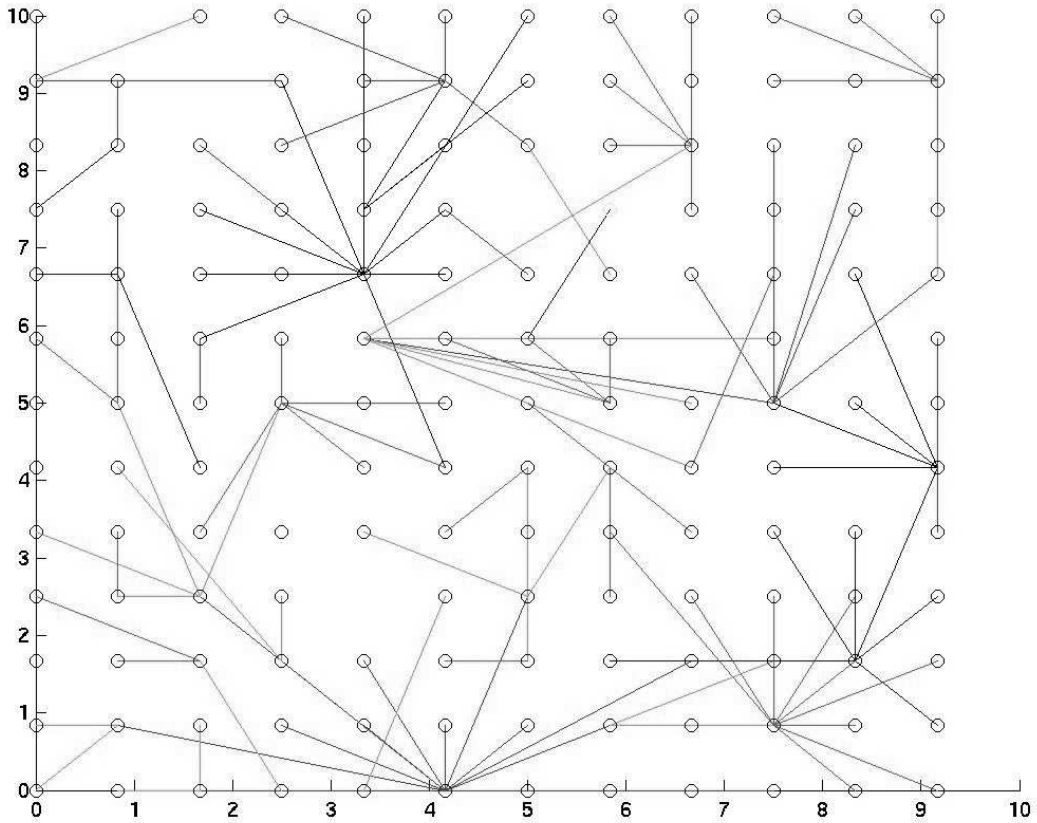

Data from Great Duck Island

- Alan Mainwaring, Joe Polastre, Rob Szewczyk
- WSNA '02
- Long-term deployment
- Packet loss rates over time
- Mote synchronization
  - due to sleep cycle

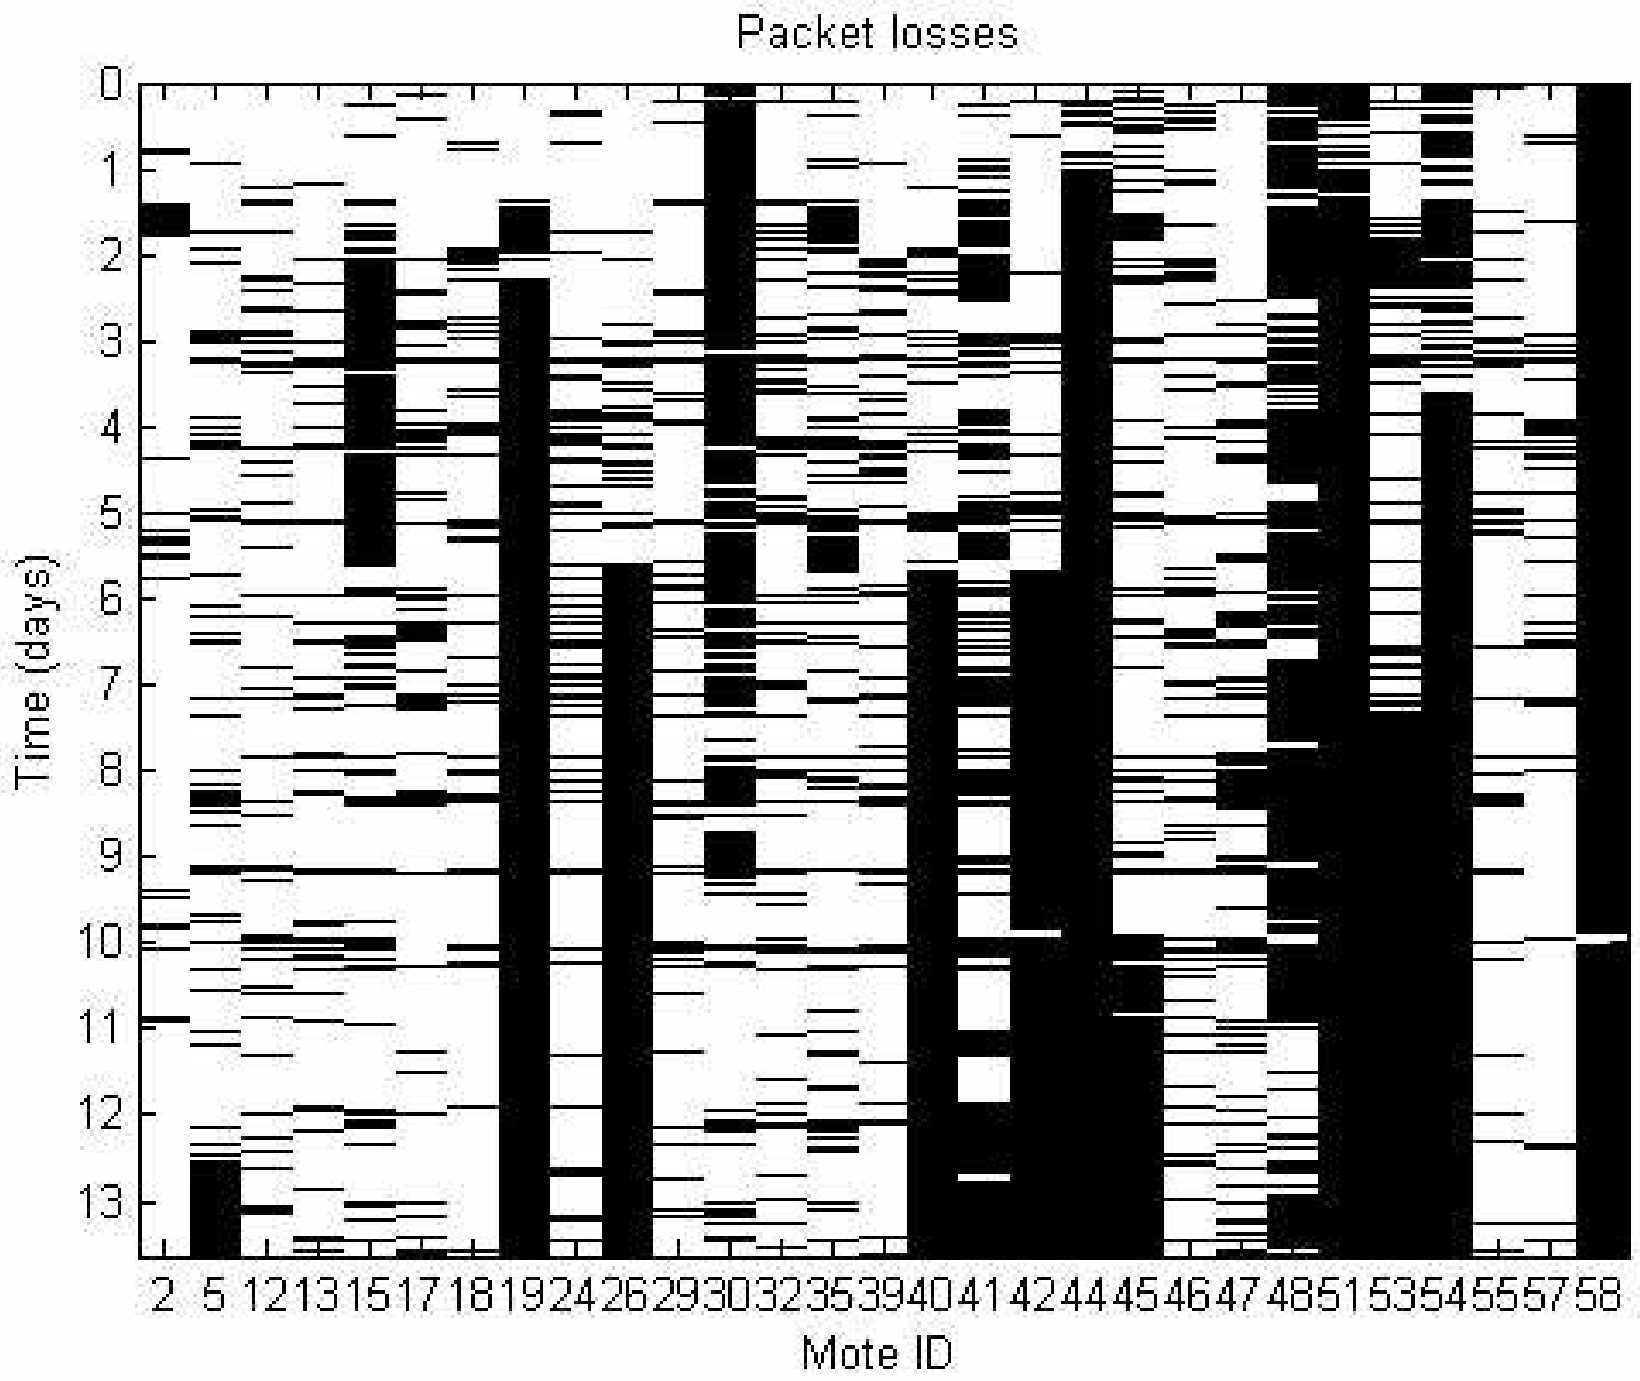
Packet losses
Time (days)
0
1
2
3
4
5
6
7
8
9
10
11
12
13
2 5 12 13 15 17 18 19 24 26 29 30 32 35 39 40 41 42 44 45 46 47 48 51 53 54 55 57 58
Mote ID

Mote Synchronization

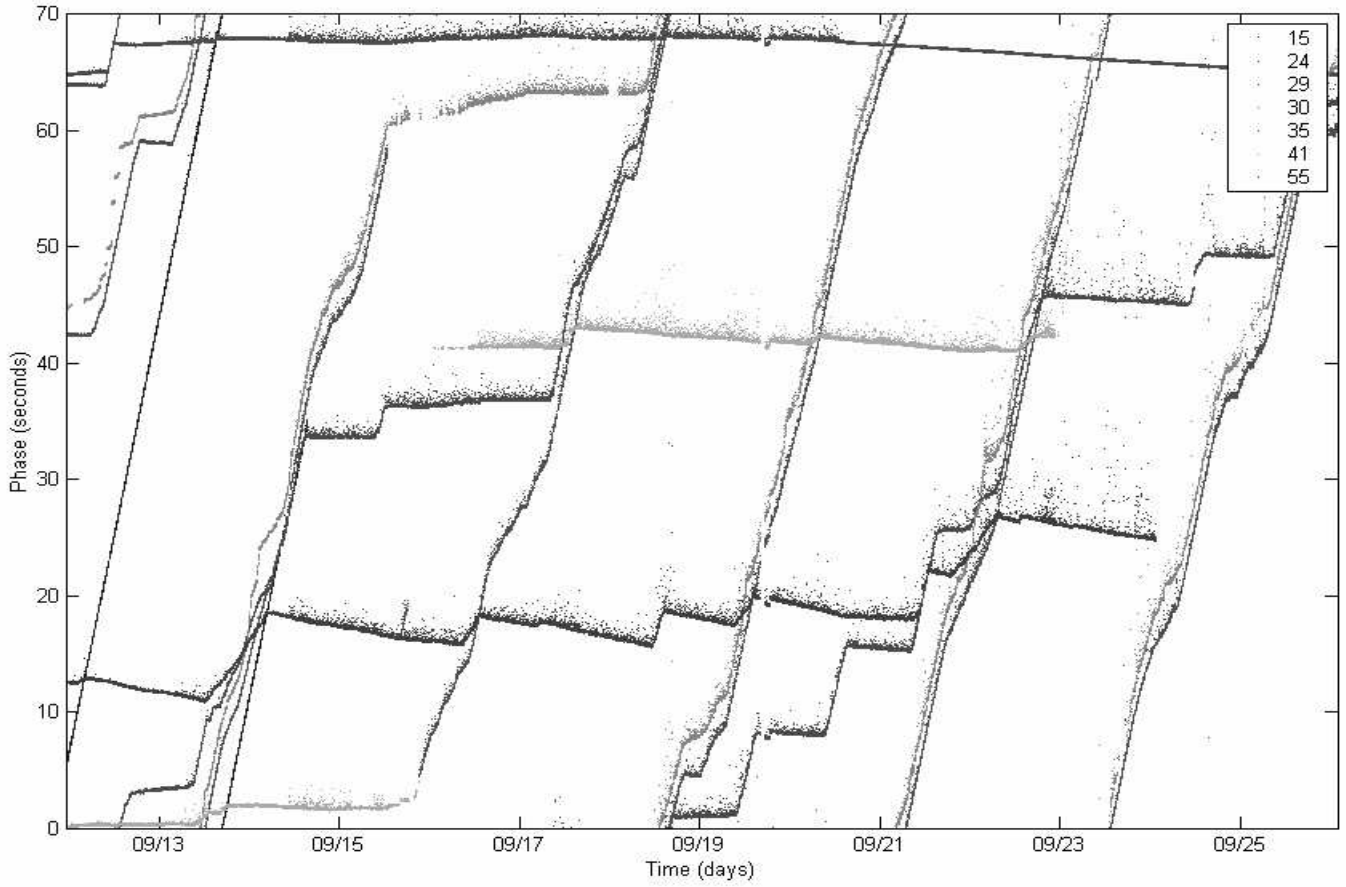

ns + 802.11 + AODV?

(Not the Internet, revisited…)

- ns
  - No application-protocol interactions
- 802.11
  - Power-hungry
  - High reliability
- AODV
  - Route discovery (flooding)
  - Do we need any-to-any?

TinyOS Approach: TOSSIM

- Simulates application + protocol
  - Complete TinyOS applications
  - Per-bit radio simulation
- Top-down statistical models
  - Models observed behavior
  - Ambivalent to causes
  - Directed per-link error rates
- Not perfect (transience, time)

Validated Simulation

- Simulation for evaluation
  - Controlled, reproducible
  - Comparative evaluation possible
  - Real world approximation (not truth)
- Deployment for validation
  - Resources, noise, time, everything we forgot
  - Our protocol really is better

Conclusion

- Simulation for comparative analysis
- Subtle interactions between applications and protocols
- Top-down models
- Validate simulations with deployments

Questions