

Large-scale Network Discovery: Design Tradeoffs in Wireless Sensor Systems

Deepak Ganesan[†], Bhaskar Krishnamachari[‡], Alec Woo[§],
David Culler[§], Deborah Estrin[†], Stephen Wicker[‡]

[†]Department of Computer Science, UCLA

[‡] School of Electrical and Computer Engineering, Cornell University

[§]Department of Electrical Engineering and Computer Sciences, UC Berkeley

We are investigating design tradeoffs relating to the structure of the network formed in large-scale self-organizing wireless sensor systems. In such systems, most traffic is in the form of many-to-one data flows, between all the networked sensors and a central monitoring node. We refer to the process of establishing a multi-hop routing path from every sensor node in the system to the monitoring node as *Network Discovery*.

Flooding, which is conventionally used for disseminating information, can also be used for network discovery. The central monitoring node initiates the flooding by forwarding a message to all its neighbors, each successive node forwarding it once. As the flooding proceeds, nodes keep track of who they received messages from. If a given node receives multiple copies of the flooding message from different neighboring senders, it may use a parent-selection mechanism to choose one of these as its parent. The nodes then send information to their parent, along along the reverse direction of the original flooding message, when they need to forward any data back to the monitoring node.

The network topology that is thus discovered is a spanning tree. It is of interest to know whether the discovered routing tree is *bushy*, with a few large clusters, or *sparse*, with many smaller clusters. Three factors affecting the structure of the discovered routing tree are the flooding mechanism, the transmission power, and the parent selection algorithm. The bushiness of the discovered tree structure will have direct implications in terms of application data aggregation, energy utilization, system throughput and robustness.

We would like to address a number of related questions in our investigations. What are the tradeoffs among these implications? Increasing the radio transmission power decreases the degree of spatial reuse in the network, but results in a bushier tree structure where greater application data aggregation can occur. How does this change the overall system throughput? Bushy trees also result in highly non-uniform energy utilization as opposed to sparse trees; but do they utilize less energy overall? How robust are bushy trees to different kinds of node failures? Finally, how can the use of different parent selection mechanisms help in controlling the bushiness of the discovered tree?

Our initial results in this study are obtained from a combination of theoretical models and empirical data from prototype wireless sensor networks consisting of over 150 nodes. These results suggest that when the transmission radius of nodes is large, and basic flooding with a simple parent-selection mechanism is used, the topology of the discovered routing tree is in fact bushy and follows a power law distribution. Most nodes are leaf nodes, and a small-yet-significant number of the nodes are cluster-heads with a large number of children. This happens because, in a given area, the first node which wins the medium-access contention and broadcasts the flood message first becomes the parent of all new receivers. This “winner gets all” scenario is

somewhat similar to the Internet, which shows a power law network topology as well because nodes attach themselves preferentially to bandwidth-rich nodes.

We intend to base our work on a combination of simulation and empirical measurements. To enhance our simulation environment, we are performing a series of experiments to create a statistical profile of connectivity among each node and run our simulations over these profiles. This allows us to compare and validate simulation results that are based on a simpler connectivity assumption such as circular cell coverage. We are also looking at some specific application scenarios for wireless sensor networks to study how these structure-related design tradeoffs are influenced by application-specific considerations.

Large-scale Network Discovery: Design Tradeoffs in Wireless Sensor Systems

Deepak Ganesan(UCLA), Bhaskar Krishnamachari(Cornell), Alec Woo(UCB), David Culler(UCB),
Deborah Estrin(UCLA), Stephen Wicker(Cornell)

Introduction: Challenge of Network Discovery

Network discovery is the mechanism of establishing a routing path from every networked sensor node to a centralized node. This routing path is used to send data to a central node for monitoring. We refer to these structures as **Discovery Trees**

Mechanisms

• Flooding: Is it efficient?

Flooding is a well accepted mechanism to route a packet to all nodes in the network. When the flood message is generated by a base-station, nodes can use the reverse path to return data. In densely deployed wireless networks, flooding causes massive redundancy, resulting in wasted energy and channel resources. Can this overhead be reduced?

• Network Self-Discovery

The network can self-organize to form discovery trees. What are the properties of these trees?

Tree structure

Are discovery trees **bushy** (few large clusters) or **sparse** (many small clusters). Three parameters affect their structure:

- Transmit power setting
- Parent selection criteria
- Tree construction mechanism

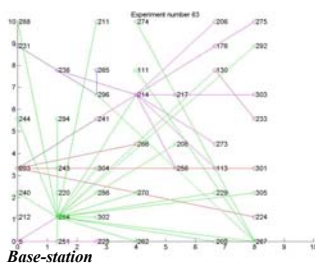
Tradeoffs of Discovery Trees

Assume that each node needs to send 1 sample per second to the base-station. Nodes with a large number of children expend **more energy** to receive and process data. On the other hand, nodes with larger number of children can perform significantly more **aggregation**. The **structure of discovery trees** determines the nature of these tradeoffs.

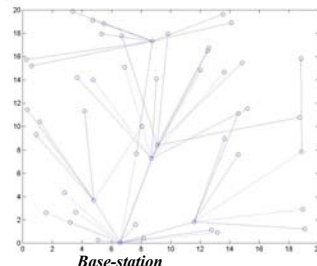
Power Laws in Sensor Networks: Initial Results and Methodology

Do Opportunistic Flooding Trees exhibit power laws?

Initial experiments and simulations suggest the existence of a power-law distribution of cluster sizes, where a cluster is defined as a node and its children on a discovery tree. The figures below represent sample flooding trees from experiments(left) and simulation(right).

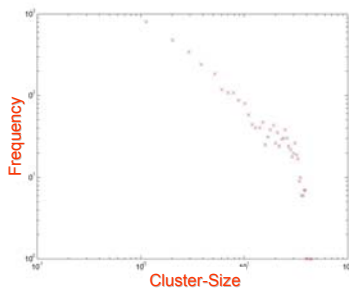


Flooding tree from 80 node experiment (8x10 foot grid)



Flooding tree from simulation (8x10 foot grid)

Simulation Results(below) show heavy-tailed distribution of cluster-sizes

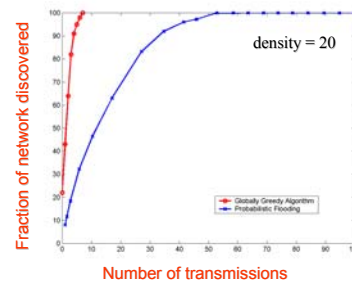


Intuition behind brute-force flooding: Greedy nature of flooding results in **early bird gets all** structure.

Similarity to internet power-law structure, where nodes attach themselves preferentially to **bandwidth-rich** nodes

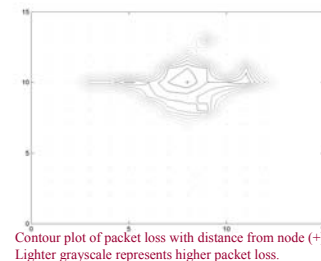
Reducing Flooding Overhead

- **Probabilistic flooding:** In dense sensor networks, high probability(>60%) is required to reach 95% of nodes. This results in low energy savings over opportunistic flooding.
- **Globally greedy algorithm** requires only 5-10% probability to reach all nodes.
- **Bridging the gap:** Flooding-optimized MAC



Simulating over Real Topologies

- Do theoretical results that assume circular cells hold in a "real topology"? Are there phase transitions in real topologies?
- Packet loss makes probabilistic flooding complex: send probability and receive probability need to be accounted for.
- Series of experiments to log connectivity of wireless cells (instance shown below) so that simulations can be run on experimentally generated topologies



Ongoing Work: Explore Tradeoffs

Aggregation:

- Aggregation near the leaves offers greater benefit towards reducing bandwidth requirement

Energy:

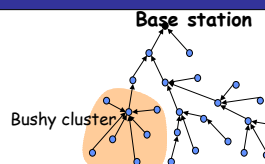
- What is a reasonable model for energy consumption?
- If each node listens only when its children are transmitting, cluster-heads of bushy clusters will have high energy consumption.

Bandwidth:

- Choice of transmit power setting may have implications on overall system throughput, as it changes the extent of spatial reuse possible.

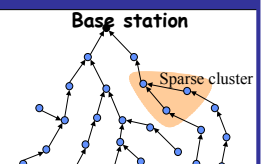
Robustness:

- Size of clusters may affect robustness to random and patterned failures.



Bushy Trees:

- High potential for aggregation close to leaves.
- Mechanisms for construction:
 - Opportunistic flooding with high transmit power setting.
 - Select parent preferentially based on aggregation criteria.
- High transmit radius may result in lower throughput.
- Non-uniform energy utilization among nodes
- Greater reliability to random failures. May be susceptible to catastrophic failures.



Sparse Trees:

- Low potential for aggregation close to leaves.
- Mechanisms for construction:
 - Opportunistic flooding with small transmit power setting
 - Randomized parent selection
- Greater spatial reuse may result in greater throughput.
- Uniform energy-consumption