

Complex Behavior at Scale: An Experimental Study of Low-Power Wireless Sensor Networks

Deepak Ganesan, Bhaskar Krishnamachari, Alec Woo,
David Culler, Deborah Estrin, Stephen Wicker

Abstract A new class of networked systems is emerging that involve very large numbers of small, low-power, wireless devices. We present findings from a large scale empirical study involving over 150 such nodes operated at various transmission power settings. The instrumentation in our experiments permits us to separate effects at the various layers of the protocol stack. At the link layer, we present statistics on packet reception, effective communication range and link asymmetry; at the MAC layer, we measure contention, collision and latency; and at the application layer, we analyze the structure of trees constructed using flooding. The study reveals that even a simple protocol, flooding, can exhibit surprising complexity at scale. The data and analysis lay a foundation for a much wider set of algorithmic studies in this space.

I. INTRODUCTION

A new class of networked systems is emerging that involve very large numbers of small, low-power, wireless devices distributed over physical space. Today, numerous investigations in wireless sensor networks are utilizing hundreds of battery powered nodes that are perhaps a cubic inch in size [19], [28]. Laboratory studies have demonstrated nodes of a few cubic millimeters in volume, and we can certainly imagine many scenarios of computational fabrics, surfaces, and floating dust [5], [6]. The sheer number of devices involved in such networks and the resource constraints of the nodes – energy, storage, and processing – motivate us to explore extremely simple algorithms for discovery, routing, multicast, and aggregation. These algorithms should be localized, use minimal state, adapt to changes in structure, and have minimal communication cost. Our experience has been that these algorithms, while easy to build, often exhibit complex global behavior in real world settings. In particular, phenomena that may be side issues in typical wireless LAN environments become quite significant.

There are many instances of large-scale systems that exhibit unpredictable behavior when faced with unexpected operating conditions. Synchronization in periodic

signals in the internet leading to patterns of loss and decay (Floyd and Jacobson [39]); biological ecosystems exhibit remarkable robustness to variations in weather and predation, but catastrophic sensitivity to genetic mutation or a new virus (Carlson and Doyle ([38])); unexpected cascading failures in power grid systems; pathological routing oscillations resulting from route flapping at a single BGP router ([40]), are all well-known examples of such behavior. We show that as we scale wireless sensor networks, it becomes increasingly important to anticipate such behavior, and therefore attempt to understand them. We demonstrate this behavior with a simple protocol, flooding, that is often considered easy to use because of its simplicity. While trivial from the perspective of protocol design, complexity is introduced by interactions between different parts of the system: vagaries of radio propagation at very low power levels, typical non-idealities in hardware, radio processing and protocol choices. Coupling between these components is difficult to, and therefore not modeled, and can result in unanticipated behavior in large scale wireless sensor networks.

Many theoretical studies in the space assume a circular connectivity model ([12], [41]). Our data demonstrates that even a simple flooding algorithm entails complex large-scale behavior that such models fail to capture. To the best of our knowledge, there has been no systematic experimental study that analyzes how underlying network characteristics manifest themselves in protocol behavior at scale. The contribution of this paper is twofold. First, we provide a wealth of detailed empirical data from studies of relatively large scale, dense wireless network configurations. Our data suggests the need to revisit currently used models of connectivity in such networks. Second, we provide a systematic analysis of factors influencing the global behavior by separating them out each primary level: the link layer, medium access layer, and application layer. These experimental data and analysis lay a foundation for a much wider set of algorithmic studies.

The paper is organized as follows: in Section II we motivate complex behavior exhibited by flooding at scale.

We discuss the experimental platform in Section IV and describe our experiments, our methodology for examining the results and the need for useful metrics in Section V. This sets up our discussion of the experimental results showing link layer effects (Section VI), medium access layer effects (Section VII), and application layer effects (Section VIII). We highlight some of the related work in Section III. We discuss the implications of this study for the design of protocols for large scale wireless networks in Section IX.

II. MOTIVATING SCENARIO

Flooding is a one of the simplest, most widely used, and well studied protocol for disseminating data in many systems. A message initiated from a source is rebroadcasted by neighboring nodes and extends outward, hop by hop, until the entire network is reached. Such algorithms underly many sophisticated protocols, particularly in large scale wireless sensor networks, in which there may be a need for unattended operation. For example, they are used for single source-destination route discovery in reactive and hybrid ad hoc routing protocols [8], [9], [10]; for exploration in directed diffusion [37]; for multi-hop broadcast [2], [13]; for forming discovery trees [16]; for issuing network commands such as “sleep,” “wake up”; for changing network-wide parameters such as transmit power, and for multihop time synchronization [21].

Algorithm 1: Flooding-based Tree Construction

```

if message received for the first time then
  Set Parent on Tree = Source of message;
  Change Source field to MyId;
  Increment HopCount field;
  Rebroadcast Packet;

```

Algorithm 1 shows the algorithm for a flooding-based tree construction protocol. As seen in this schema, each node chooses the first node that it hears the message from, as its parent on a routing tree. The resulting tree connects all nodes that have participated in the flooding, and is rooted at the origin of the message. Under idealized settings, one would expect a flood to ripple outward from the source in an orderly, uniform, fashion. Our results from a large-scale implementation of this protocol on real hardware show, however, that the global behavior of this simple protocol can be surprisingly complex. We now present a motivating example to support this point.

Figure 1 shows a sequence of snapshots from traces of an experiment to illustrate how a flood propagated over time¹. 160 nodes are laid out in a 12x13 square grid on the

¹The unmarked nodes in Figure 1 correspond to failed nodes in the experiment

ground as indicated by dots and a flood originates from the node located at the coordinates (5,0). When a node receives the flood message, it immediately rebroadcasts once and squelches further retransmissions. Redundancy is expected as every node responds in this manner.

There are several noteworthy indicators of non-uniform flood propagation. Instead of extending outward step-by-step, some links show regions where the flood actually extends backward geographically towards the source. We call these *backward links*. An example of this can be seen in Figure 1(b) - the link between the node located at (6,3) and the node located at (6,2) is a backward link; Figure 1(c) also shows numerous backward links. In some instances, a message is received over a large distance and creates what we call a *long link*, such as the links from the source near (5,0) to nodes at (1,1) and (2,3) in Figure 1(a). Finally, some nodes are missed by the flood even though neighboring nodes transmit messages, such as the node near (3,4) in Figure 1(c). We refer to these nodes as *stragglers*. If we look at the tree structure that evolves as the flooding proceeds, we notice that it exhibits a high *clustering* behavior: most nodes in the tree have few or no descendants, while a significant few have many children.

Metric	Definition
<i>Straggler</i>	Node that misses a transmission, even though it would be expected to receive a packet with high probability
<i>Backward Link</i>	Link in which the recipient of the flood is closer from the base station than the transmitter
<i>Long link</i>	Link that is significantly longer than expected at given transmit power level
<i>Clustering</i>	Number of nodes attached to a single point on the data gathering tree

TABLE I
FLOODING METRICS

This simple example illustrates how flooding can exhibit complex behavior in a realistic setting. There are a number of factors across different layers that impact the dynamics of flooding. For example, the long links show that the cell region for each node is far from a simple disc (a link level effect), and the stragglers are very likely left behind due to MAC-level collisions. In the following sections, we dissect the contributions of different layers to the behavior of flooding at scale. Our experimental observations provide valuable input into the design of more sophisticated multicast mechanisms for large scale wireless networks.

III. RELATED WORK

There is currently a dearth of experimental measurements on the scaling of ad-hoc and sensor network proto-

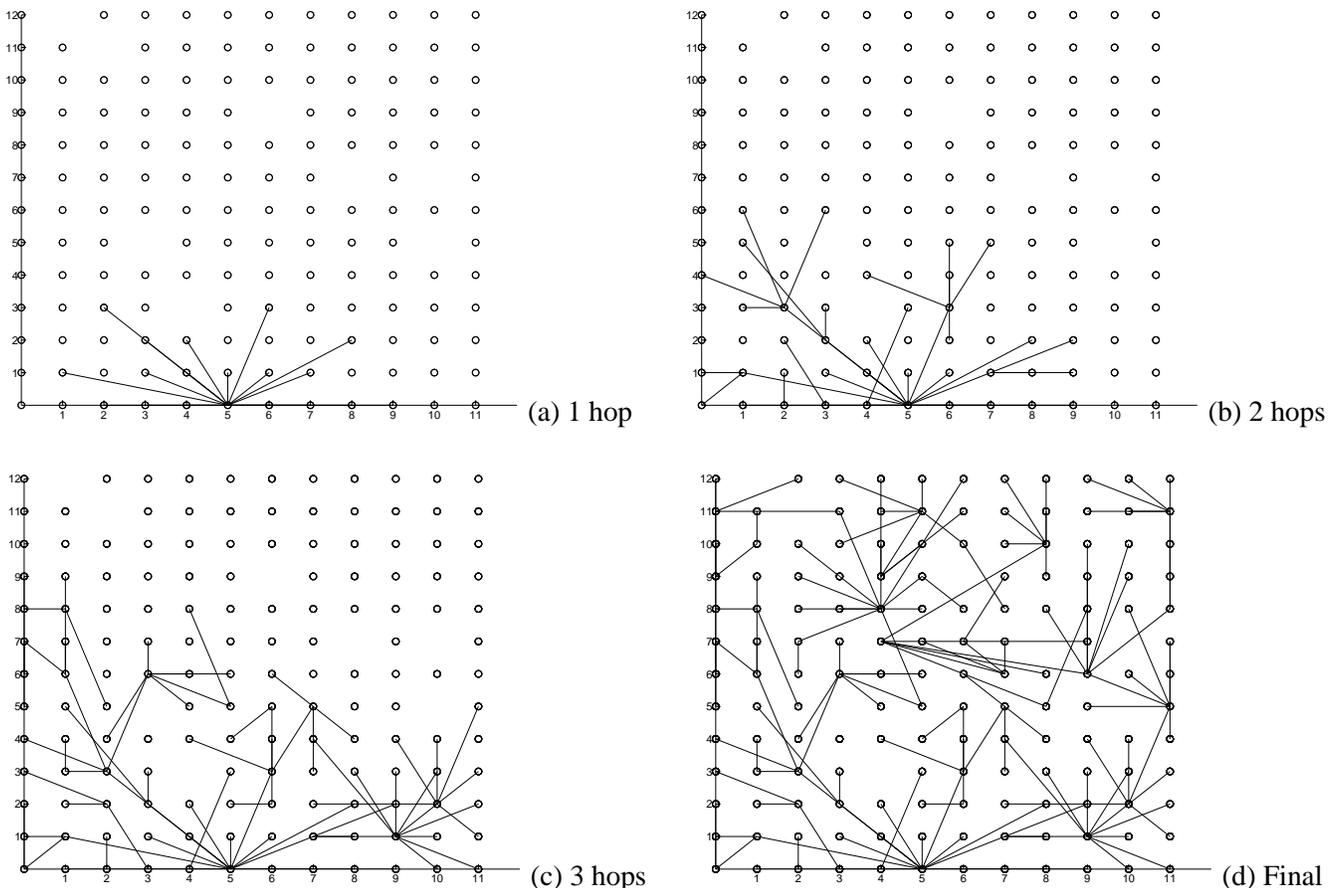


Fig. 1. Snapshots from a single run of flooding on the experimental testbed

cols. Prior experimental studies in this area have tended to focus on routing in wireless ad-hoc networks without addressing scaling for lack of large enough infrastructure. For example, [3] describes an experimental ad-hoc network with eight mobile nodes consisting of laptops and 802.11 cards driven around in a 300 x 700 m area. This test-bed is used to provide some results on the performance of dynamic source routing (DSR). Similarly, [4] describes an experimental testbed involving one desktop and five laptops with 802.11 cards, used to test the performance of the ad hoc on-demand distance vector (AODV) routing protocol. In [22], the performance of data aggregation in directed diffusion is tested on a sensor network consisting of 14 PC/104 nodes equipped with Radiometrix RPC modems. Some experimental work has sought to validate and test medium access protocols. An 11-node experimental setup using Berkeley notes is used to analyze the performance of an adaptive rate control mechanism for medium access in sensor networks in [24]. In [23], a small experimental setup consisting of 5 Berkeley notes is used to validate the performance of the proposed S-MAC protocol. Another small-scale experiment involving Berkeley notes for signal strength measurements is

described in [25].

Most previous work on analyzing the behavior of routing protocols in large-scale wireless sensor networks has been done in simulation [26][37]. These studies are not entirely satisfactory because the realistic modelling of link-layer characteristics in simulation settings is a very challenging problem, and the final validation of the protocol performance has to be in real settings.

Flooding is also the subject of [2] which investigates the “broadcast storm” problem associated with flooding, showing both analytically and through simulations inefficiencies such as redundancy, MAC-level contention and collision. In contrast, our work examines the impact of link layer non-idealities on the dynamic behavior of this protocol.

Significant research studies have addressed various aspects of the MAC layer: improving channel utilization, proportional fairness, energy efficiency, and collision avoidance. While these are critical issues, little published work has systematically evaluated the coupling between subtle effects at the link layer and MAC layer behavior, in dense, large-scale low-power wireless networks.

IV. EXPERIMENTAL PLATFORM

The nodes used in these experiments are shown in Figure 3. Each node has a 4 MHz Atmel [17] processor with 8 kB of programming memory, and 512 B of data memory. The node is equipped with a 916 MHz, single channel, low power radio from RFM [18], capable of delivering 10 kbps of raw bandwidth using on off keying (OOK) modulation. The transmission power of the radio is dynamically tunable with different potentiometer (*pot*) settings as shown in Figure 2. For the rest of this paper, we use the mapping in Table II to refer to the potentiometer setting.

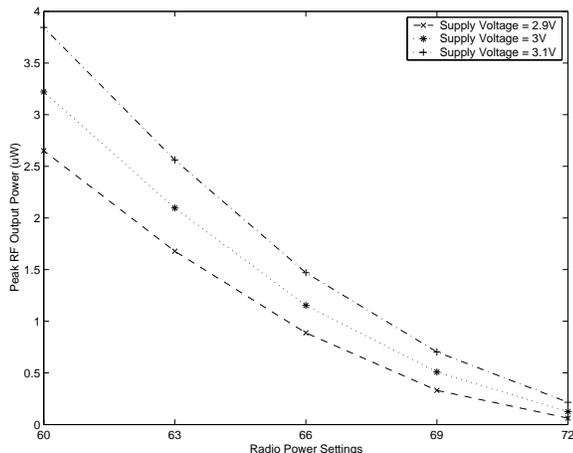


Fig. 2. RF output power at different power settings.

Potentiometer Setting	Strength
60	Very High
63	High
66	Medium
69	Low
72	Very Low

TABLE II

MAPPING BETWEEN HARDWARE POTENTIOMETER SETTING AND LEGEND

To avoid undesirable effects on transmission power due to unregulated voltage supply from the batteries, all nodes are equipped with fresh AA batteries before conducting all experiments. Furthermore, all nodes have the same antenna length with uniform vertical orientation. We should note, however, that perfect calibration of the radio hardware on our experimental nodes is difficult and thus even with the same nominal hardware settings, actual transmit power on different nodes can vary [32].

The TinyOS [19] platform provides the essential runtime system support. It includes a complete network stack



Fig. 3. The Rene Mote: a hardware platform for wireless networking

with bit-level forward error correction, 16-bit CRC error checking, medium access control, network messaging layer, non-volatile storage, and timing capability. The default packet size is 38 bytes long, with a payload of 30 bytes. The medium access control protocol [24] is a variant of the simple carrier sense multiple access (CSMA) protocol [33]. It waits a random duration before each transmission and goes into random backoff if the channel is busy. The delay and backoff durations are randomly picked from a fixed interval between 6ms and 100ms. During backoff, the radio is actually powered off to save energy, but the tradeoff is that no communication is possible during that period. Unlike many MAC protocols such as IEEE 802.11 [35] which will drop packet transmission after a maximum number of backoffs, this MAC protocol keeps trying until it finds a clear channel.

V. DESCRIPTION OF EXPERIMENTS & METHODOLOGY

To understand the dynamics of flooding, we conducted two separate sets of experiments. The first set focused on understanding the characteristics of links among all nodes in a large test bed. The second set focused on studying the dynamics of flooding over a similar test bed. Table III summarizes these two sets of experiments.

A. Experiment Set 1

For these experiments, 185 nodes were laid out over an flat, open parking structure in a regular grid, with a grid spacing of 2 feet. The goal was to map the connectivity characteristics between all nodes at 16 different radio transmit power settings, in which nodes transmitted in sequence in response to commands sent by a base-station.

Experiment Set	Network Size	Number of Transmit Power Settings	Comments
1	169	16	Packet reception statistics at different power levels over a grid
2	150	8	Flooding at different power levels over a grid

TABLE III
SUMMARY OF THE TWO SETS OF EXPERIMENT

The base-station issued commands to all nodes to control the experiment periodically and ensured that only one node would transmit at a time to eliminate collisions. For each transmit power setting, each node transmitted twenty packets, one node at a time. All packets were sent in sequence, 100ms apart. Receivers logged the transmitter's ID, sequence number, and transmit power setting, which were embedded in the packet payload, into their local EEPROM storage.

Prior to the experiment, nodes were subject to a diagnostic test to detect unresponsive and failed nodes, and broken antennas. Sixteen of the nodes were removed, bringing the number down to 169, which was arranged in a 13×13 grid. A total of about 54000 ($20 \times 16 \times 169$) messages were transmitted in the system, allowing us to construct a map of packet reception statistics at each power level. Some of these power levels (pot setting < 60) were beyond the useful extent of the map, and results from these settings are omitted from the analysis. This entire set of experiments was conducted over a four-hour period.

B. Experiment Set 2

The second set of experiments involved 156 nodes over an open parking structure, under identical settings as the first. No obstacles were present in the immediate vicinity. The nodes were laid out in a 13×12 grid, again with a 2 ft. separation. The base-station was placed in the middle of the base of the grid. The base-station initiated flooding periodically, with the period long enough to let the flood settle. Each node would rebroadcast a message only once upon first reception of a new flood. Eight different transmit power settings were chosen and 10 non-overlapping floods were issued at each setting.

Both the application and MAC layers logged necessary information to reconstruct the flooding message propagation. At the application layer, the identifier of the node from which a message had been received was logged. Since we used globally unique identifiers for each node, this gave us a causal ordering of message propagation [7], which was used to reconstruct the propagation tree. At the MAC layer, timing information was crucial for us to extract metrics such as backoff time and collisions. While

absolute time-synchronization [20] was an option, this proved to be unnecessary for our needs. To obtain timing information, the MAC layer stored two locally generated timestamps, with granularity $16 \mu s$. The first timestamp recorded the total amount of time that a message was stored on a node before being retransmitted. The second timestamp recorded the interval for which the node was in backoff mode. The fact that flood propagation through a large network occurs quite quickly is our ally, since clock skew and drift is small during the flooding period. However, we still had to contend with receiver delay (as noted in [20]), which we reduced to a minimum by recording timestamps at the link layer. Thus, we restricted reconstruction errors to under a bit-time per hop, which is $100 \mu s$ at 10 kbps .

C. Analysis of Experiments

Our methodology in analyzing the vast quantity of data collected during the experiments is to decompose the behavior into layers, analyze them independently with different metrics, and combine the analysis as a composite to explain the global behavior.

At the link layers, we attempt to quantitatively define and measure the effective communication radius at a given transmit power in a real setting. We explore packet reception statistics over distance, define what constitutes a bidirectional link and an asymmetric link, and measure these effects. At the medium access control level, we use timing information to identify metrics that capture both end-to-end properties of the flood propagation, and local properties such as contention and collision. At the application layer, we analyze the resulting structure of the flood. As a composite of this analysis, we reconstruct the process of the message propagation and explain how the interactions across levels lead to the final global behavior. A comprehensive understanding of the characteristics of the radio, effective communication range, packet reception behavior, extent of asymmetry, and MAC layer behavior provides guidance for algorithms designed to work under similar large scale wireless sensor networks.

VI. LINK LAYER ANALYSIS

Our first step towards analyzing the data is to develop a set of metrics that will help us understand the basic link characteristics of the testbed. These metrics include node-to-node packet reception rate with respect to distance, cell radius coverage with respect to different radio transmission power setting, and the degree of link asymmetry.

A. Packet Reception Statistics

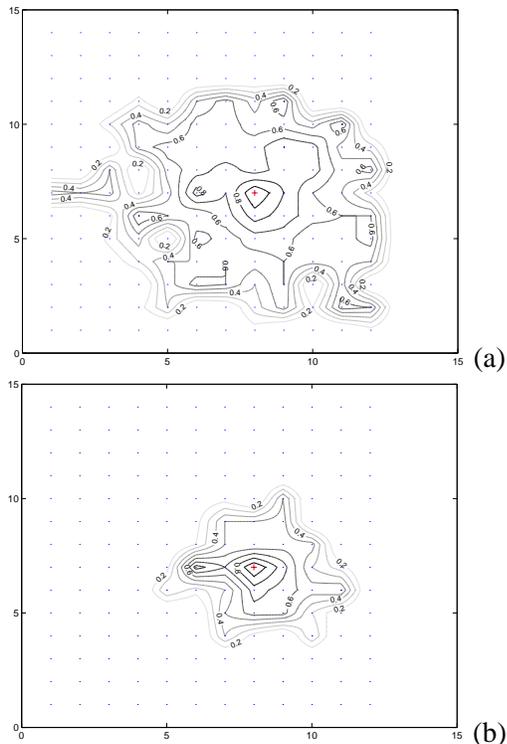


Fig. 4. Contour of probability of packet reception from a central node at two different transmit power settings

A fundamental metric while evaluating link-layer connectivity is packet throughput. In our experiments, packets that fail to pass CRC checking are considered lost. The distribution of packet reception rate over distance is quite non-uniform, as Figure 4 shows. In fact, individual contours clearly exhibit *directionality*, with propagation being better at some directions than others. This non-uniformity is also observed in Figure 5, which shows the variation of packet reception rate with distance. Although packet reception falls off quite rapidly with distance, the plots have a heavy tail i.e. there is a *non-zero probability* of receiving a packet at a long distance from the transmitter. Both these observations clearly indicate the presence of long-links (Table I), that result in greater propagation in some directions than others.

Another observation from these curves is that the throughput is lower than 100% even at short distances

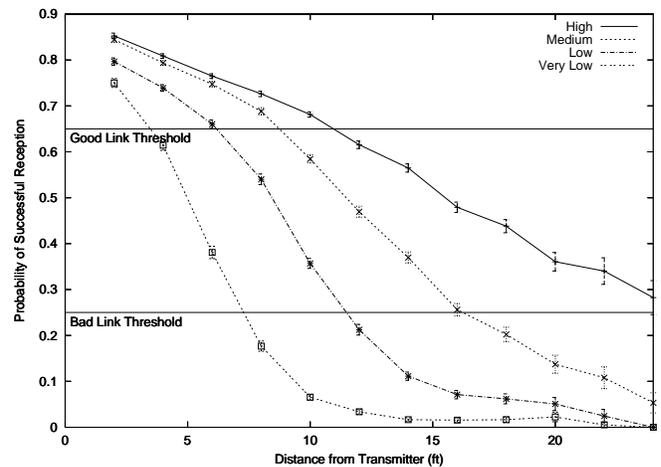


Fig. 5. Probability of packet reception over distance with different transmission power settings.

from the transmitter. This is due to two factors: increased fading rate due to deployment on the ground [36], and insufficient signal processing and forward error correction due to the limited computational and energy resources available on this platform.

B. Measuring the Connectivity Radius

Algorithm designers often conceptualize these systems in terms of the connectivity radius and the notion of a circular connectivity cell. Many analytical results involve working with circular cells, since this simplifies analysis and allows a geometric approach.

Our definition of connectivity radius is based on a packet-reception threshold. A guiding argument for choosing this threshold would be to treat a link as a “good link” if we can use forward error correction(FEC) and other techniques to improve the raw packet throughput to adequate levels. Correspondingly, a “bad link” would be regarded as one which cannot possibly be salvaged by such means, as it offers very poor throughput. Using these criteria, and based on Figure 5 we take the threshold for a “good link” to be 65% and the corresponding threshold for a bad link to be 25%

C. Asymmetric Links

Asymmetric links arise relatively infrequently in sparse wireless networks, such as typical 802.11 LAN and ad-hoc configurations, and are often filtered out by protocol levels [1], [31]. However, with a large field of low-power wireless nodes such asymmetric links are very common, even if all nodes are set to have the same transmit power. Our experiments allow us to quantify asymmetric links and understand their behavior. The definitions used for

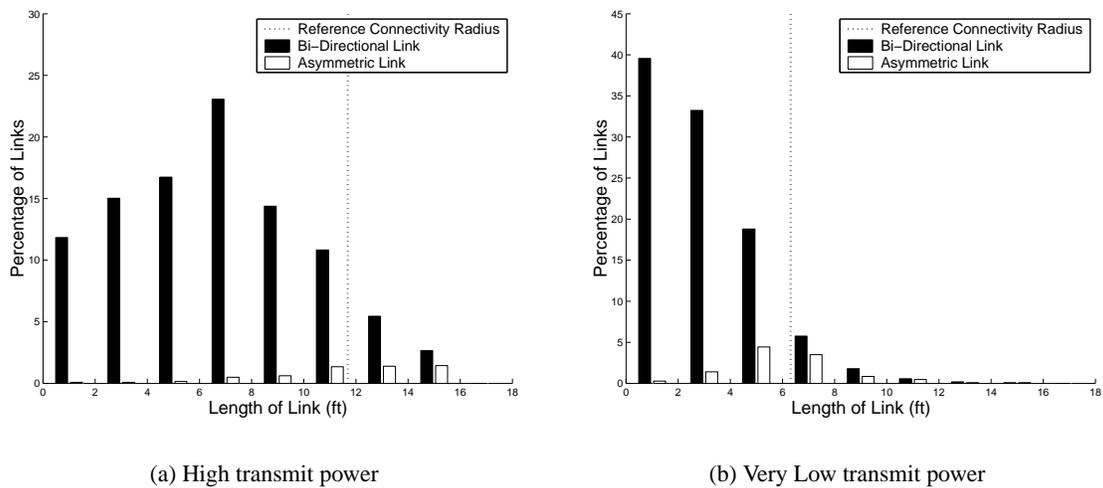


Fig. 7. Distribution of bidirectional and asymmetric links over distance.

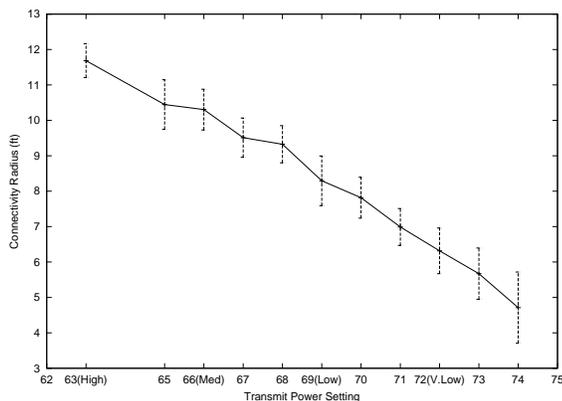


Fig. 6. Connectivity radius at different radio transmission power.

a “good” and “bad” link in Section VI-B is used in developing this metric. An *asymmetric link* is defined as one which has a “good” link in one direction and a “bad” link in the other. A *bidirectional link* is one which has a good link in both directions. An analysis over the data reveals that for the range of transmit power settings studied, approximately 5-15% of all links are asymmetric, the percentage increasing with decreasing transmit power setting.

Figure 7 shows the distribution of bi-directional and asymmetric links over distance. At short distances from the transmitter, a negligible percentage of links are asymmetric, but this percentage grows significantly with increasing distance, especially at lower power settings. The dotted vertical line shows the connectivity radius calculated as per section VI-B for the particular transmit power.

At the fading edge of a connectivity cell, small differences between the nodes in transmit power and reception sensitivity become significant, resulting in asymme-

try. The aggregate effect of the small differences in the radios and hardware, as mentioned in Section IV, and slight differences in energy levels of the nodes contribute significantly to link asymmetries in this regime (discussed further in Section IX).

VII. MEDIUM ACCESS LAYER ANALYSIS

We now turn to the dynamics and MAC layer effects during message propagation. We examine three metrics that capture different aspects of the propagation: maximum backoff interval, reception latency and settling time.

Maximum backoff interval captures the interference within each cell, since the node the transmits last has backed off multiple times as a result of contention. Reception latency and settling time are end-to-end metrics: reception latency captures the time taken for all nodes in the network to receive the flood, and settling time is the amount of time taken by all nodes to transmit the packet, and therefore settle down. Figures 9(a) and 9(b) illustrate the two metrics at different transmit power settings. The reception latency is typically lower than the settling time, the difference between them being dependent on the transmit power level of the radio. The three metrics can be related as follows:

$$\max(\text{MaxBackoffTime}, \text{ReceptionLatency}) \leq \text{SettlingTime} \leq \text{MaxBackoffTime} + \text{ReceptionLatency}$$

The above relation has a simple explanation: the settling time is at least as long as the time until the node with the longest backoff time has retransmitted a packet or the last node receives a packet. It is bounded above by the case that the last node to receive the flood propagation also chooses the maximum backoff interval.

Power	#Expts	MaxBackoff Interval(s)	Reception Latency (s)		Settling Time (s)	
			95%	100%	95%	100%
Very High	7	3.265 ± 0.169	1.285 ± 0.302	2.958 ± 0.767	3.842 ± 0.330	5.145 ± 1.171
High	8	2.773 ± 0.099	1.569 ± 0.278	2.558 ± 0.393	3.663 ± 0.125	4.553 ± 0.384
Medium	7	2.587 ± 0.075	1.688 ± 0.307	2.870 ± 0.712	3.469 ± 0.172	4.786 ± 0.74
Low	4	2.202 ± 0.073	1.861 ± 0.476	2.853 ± 0.814	3.258 ± 0.185	4.056 ± 0.826
Very Low	4	1.302 ± 0.057	2.174 ± 0.235	2.192 ± 1.139	2.985 ± 0.161	3.611 ± 0.473

TABLE IV

MAXIMUM BACKOFF INTERVAL, RECEPTION LATENCY, SETTILING TIME, WITH CORRESPONDING 95% AND 100% CONFIDENCE INTERVALS AT DIFFERENT TRANSMIT POWER SETTINGS.

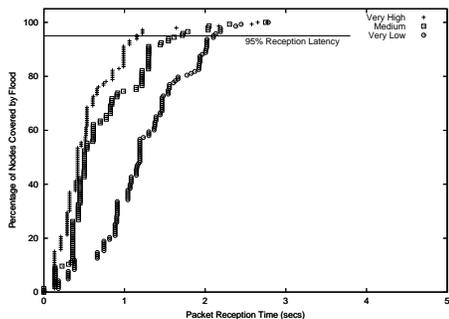
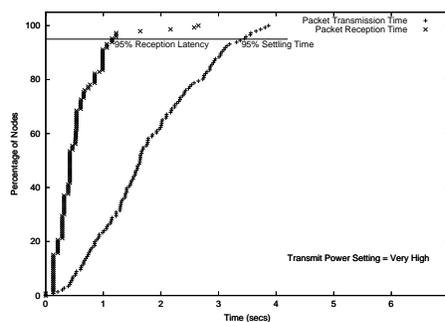
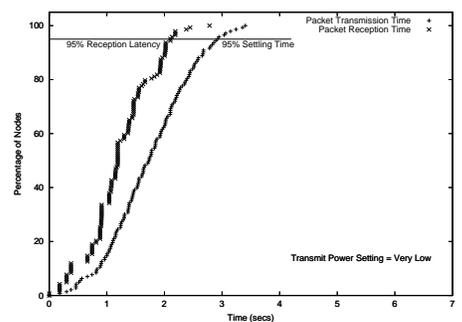


Fig. 8. Reception Latency grows with Increasing Transmit Power Setting.



(a) Very High transmit power setting.



(b) Very Low transmit power setting.

Fig. 9. Timeseries of packet transmission and reception

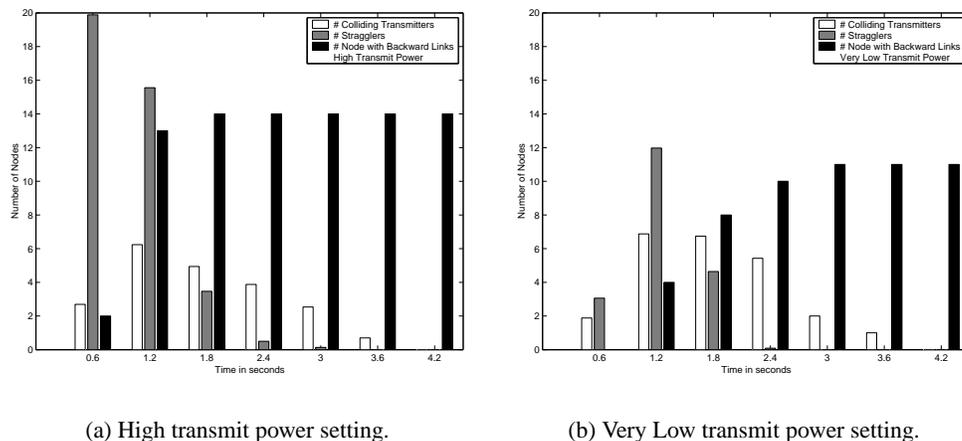
The 95% values of the parameters (Table IV) follow Equation 1 as expected and vary predictably with the transmit power setting. Interestingly, however, the 100% values of the metrics show significant deviation from the norm. In fact, Figure VI-C shows that at different transmit power levels, the time taken for 100% of the nodes to receive the flood is almost equal, although this hardly reflects the significant differences between the curves. Similarly, in Figure 9(b), the last 5% of the nodes take as much time to receive their packets as the first 95%. We hypothesize that these nodes typically constitute backward links that are observed in flooding behavior. The deviation in the last 5% results from a combination of nodes being missed by the flood on its outward propagation, and the flood rebounding off the boundary of the network (edge effect). To further validate our claim, we look at collisions.

A. Collision

Measuring collisions is a difficult problem, since a packet could be dropped due to packet-loss rather than collisions. To distinguish collisions from other possibilities, we use an approximate measure combining link layer and MAC layer metrics. The connectivity radius in Fig-

ure 6, gives us a rough idea of the size of a communication cell. We use this metric to estimate the number of nodes that are within communication range of each transmitter, and therefore should receive the packet. Timing information at the MAC level provides additional information about nodes that were transmitting within each others transmission period, potentially causing collisions. These metrics give us an estimate of the number of nodes that should receive the packet, but do not since they lie within overlapping cells of colliding transmitters.

Figure 10 shows the relation between number of colliding transmitters, stragglers and backward links. Collisions are frequent at the initial stages of the flood, and leave behind many stragglers. These stragglers miss the propagation in the early stage of the flood and form backward links from a later reception. At higher transmit power setting, each node has a larger communication cell, and the number of hidden terminals (reflected by stragglers) is larger. This results in a larger number of backward links being generated, as the flood rebounds to capture stragglers.



(a) High transmit power setting.

(b) Very Low transmit power setting.

Fig. 10. Histogram showing distribution of colliding transmitters, distribution of stragglers and cumulative distribution of backward links over time

VIII. APPLICATION LAYER ANALYSIS

We use two metrics to characterize the application level characteristics of the tree constructed from flooding: Node Level and Cluster-Size.

Node Level is defined as the number of hops between the base-station and a node on the tree. Figure 11 shows that the distribution varies widely around the mean. At the upper left, nodes that were physically close to the base-station were many more hops away from it than their peers and in the lower right, nodes far away from the base-station were few hops from it. Both these observations can be easily explained given observations at link and MAC layers; the former by backward links, and the latter by long links.

Cluster-Size is defined as the number of children attached to a particular parent on the tree, and is a property of the parent selection algorithm that we use. Our simple parent selection mechanism is for each receiver to select the source of the first message it received as its parent. Figure 12 shows a histogram (in logscale) of the cluster sizes for various runs. Evidently, the distribution shows that there is a significant likelihood of large clusters at various transmit power settings. This observation can be explained by a combination of link and MAC layer behavior. Long links cause the flood to propagate further along certain directions than others. Nodes at the end of long links see a clear channel, and more importantly a clear field of nodes that have not been reached by the flood. The former results in low MAC layer backoff intervals, and correspondingly faster flood propagation along long links, and the latter results in the transmission reaching a large number of uncovered nodes. The combined behavior is observed at the application layer as high clustering.

High clustering can be both a blessing and a curse. On

one hand, it implies that most nodes in the network are fewer hops away from the base station, which saves energy in routing data. On the other hand, Section VI-C shows that these long links are more likely to be asymmetric than others, resulting in a fragile data-collection tree.

IX. DISCUSSION AND CONCLUSION

Our experiments have numerous implications, from modeling dense large scale networks, to vertically integrated experimental methodology.

A. Flooding: the bigger picture

We have identified the complexities in simple flooding by examining contributions from the various layers. Recall the scenario presented in Section II where we discussed four notable effects: long links, backward links, stragglers and clustering. The incidence of long links is explained in our link-layer discussion (Section VI-A). Figure 5 shows that while packet reception decreases with distance, it has a fairly long tail. With many nodes, it is likely that a given transmission will reach some nodes far away. Stragglers can be explained by collision effects caused at the MAC layer (Section VII-A). Backward links can be explained as a combination of the effect of long links and collisions. Long links resulted in the flood propagating faster in certain directions, and rebounding to fill areas where the propagation was slower or where stragglers remained, forming backward links. The opportunistic, earliest-first, parent selection mechanism at the application layer, results in highly clustered trees (Section VIII).

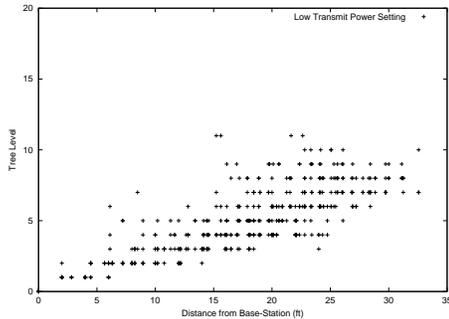


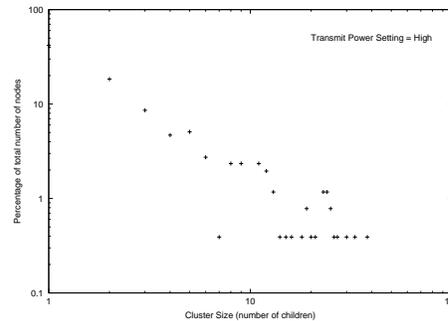
Fig. 11. Level of the nodes in the constructed tree against node distance from the root of the tree.

B. Modeling Link Layer characteristics for Sensor Networks

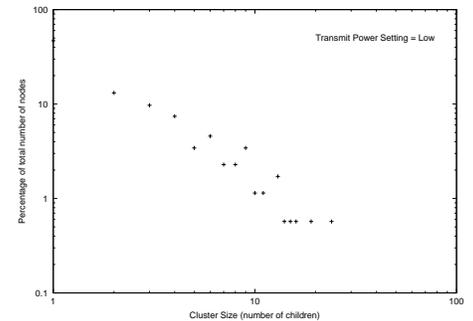
Our experiments reveal several interesting effects at the link-layer: highly irregular packet reception contours, nodes in deep fades (Figure 4), directionality in transmission resulting in long links (Figure 5) and greater asymmetry in long links (Figure 7). Some of these can be modeled by well-known RF-propagation models. For instance, a shadowing model with high noise variance can generate a non-uniform packet reception plot. Other features such as directionality in transmission, result from spatially correlated behavior, where nodes along certain directions consistently have better throughput to a transmitter. Such features cannot be captured by gaussian models of noise. Current models, while sufficient for protocol designs in sparse, mobile networks, might need to be expanded to include a wider range of parameters while testing sensor network protocols, where the density of deployment could result in spatially correlated RF characteristics. Parameters that might be irrelevant in high cost, high power systems have relevance in our setting. Small differences in radio hardware, synchronization at different nodes, and energy levels often cause dramatic variations that are not captured in current models. At scale, however, these small differences manifest themselves and impact protocol behaviors in various ways that are currently not tested in simulations. We believe that simulators to study scaling behavior of sensor network protocols should include a wider range of parameters and operating scenarios, thereby fully testing protocol behavior.

C. Impact on Protocol Design in low-power, large-scale systems

Sensor network deployments are expected to be ad-hoc and in potentially difficult environments, precluding



(a) High Transmit Power Setting



(b) Low transmit power setting.

Fig. 12. Cluster size histogram in log log scale.

highly engineered designs such as cellular networks. Instead, protocols should be designed to be self-organizing and robust to widely varying deployment scenarios: environmental effects, density of deployment, antenna orientation, radio calibration etc. Our empirical study points out that in a large scale wireless system operating at very low power, small environmental effects, and deviations from the norm, can have a huge impact on protocol behavior. Protocol designers should, therefore, remember that circular or probabilistic models are insufficient for fully understanding complex system interactions.

Our experimental study suggests that asymmetric links are indeed likely to be significant in large scale, sensor networks and that robust protocols must deal appropriately with asymmetric links through mechanisms such as the use of a sub-routing layer to provide bidirectional abstraction [31]. Some of the ad hoc routing protocols proposed in the literature, such as DSR [8] and ZRP [9], can route using asymmetric links; however there are others, such as AODV [10] and TORA [11], that either assume that all links are symmetric or filter out asymmetric links.

REFERENCES

- [1] D-K. Kim, C-K. Toh, and Y-H. Choi, "On Supporting Link Asymmetry in Mobile Ad Hoc Networks," *Proceedings of IEEE GLOBECOM 2001*, San Antonio, Texas, November 2001.
- [2] S.-Y. Ni, Y.-C. Tseng, Y.-S. Chen, J.-P. Sheu, "The Broadcast Storm Problem in a Mobile Ad Hoc Network," *Mobicom '99*, Seattle, Washington, USA.
- [3] D.A. Maltz, J. Broch, D.B. Johnson, "Experiences Designing and Building a Multi-Hop Wireless Ad Hoc Network Testbed," *CMU School of Computer Science Technical Report CMU-CS-99-116*, March 1999.
- [4] S. Desilva and S.R. Das, "Experimental Evaluation of a Wireless Ad Hoc Network," *Proceedings of the 9th Int. Conf. on Computer Communications and Networks (IC3N)*, Las Vegas, October 2000.
- [5] J. M. Kahn, R. H. Katz and K. S. J. Pister, "Mobile Networking for Smart Dust," *ACM/IEEE Intl. Conf. on Mobile Computing and Networking (MobiCom 99)*, Seattle, WA, August 17-19, 1999

- [6] Harold Abelson et al, "Amorphous Computing," *Communications of the ACM*, Volume 43, Number 5, May 2001.
- [7] Leslie Lamport, "Time, Clocks, and the Ordering of Events in a Distributed System," *Communications of the ACM* vol 21, no 7, p. 558-565, 1978
- [8] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networking," *Mobile Computing*, T. Imielinski and H. Korth, Eds., Kluwer, 1996.
- [9] M. R. Pearlman and Z. J. Haas, "Determining the Optimal Configuration for the Zone Routing Protocol," *IEEE Journal on Selected Areas in Communications: Wireless Ad-Hoc Networks*, vol. 17, no. 8, p. 1395-1414, August 1999.
- [10] C. E. Perkins and E. M. Royer, "Ad hoc on-demand distance vector routing," *IEEE WMCSA*, vol. 3, p. 90-100, 1999.
- [11] V. Park and M.S. Corson, "A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks," Proc. IEEE INFOCOM '97, Kobe, Japan, Apr 1997.
- [12] B. Das and V. Bharghavan, "Routing in ad-hoc networks using minimum connected dominating sets," *IEEE International Conference on Communications*, p. 376-380, June 1997.
- [13] C. Ho, K. Obraczka, G. Tsudik, K. Viswanath, "Flooding for Reliable Multicast in Multi-Hop Ad Hoc Networks," *ACM DIAL M '99*, Seattle, Washington, USA, 1999.
- [14] R. Chandra, V. Ramasubramanian, and K. P. Birman, "Anonymous Gossip: Improving Multicast Reliability in Mobile Ad-Hoc Networks", International Conference on Distributed Computing Systems, 2001.
- [15] L. Li, J. Halpern, Z. J. Haas, "Gossip-based Ad Hoc Routing," unpublished.
- [16] D. Ganesan et al. "Large-scale Network Discovery: Design Tradeoffs in Wireless Sensor Systems," poster presented at the 18th ACM Symposium on Operating System Principles, Banff, Canada, October 2001. <http://lecs.cs.ucla.edu/~estrin/>.
- [17] ATMEL 8-bit RISC Processor <http://www.atmel.com/atmel/products/prod23.htm>
- [18] RF Monolithics <http://www.rfm.com/products/data/tr1000.pdf>
- [19] J. Hill et al., "System architecture directions for network sensors," in *ASPLOS 2000*
- [20] J. Elson and D. Estrin, "Time Synchronization for Wireless Sensor Networks," *Proceedings of the 2001 International Parallel and Distributed Processing Symposium (IPDPS), Workshop on Parallel and Distributed Computing Issues in Wireless and Mobile Computing*, San Francisco, California, USA. April 2001.
- [21] Jeremy Elson, Lewis Girod and Deborah Estrin, "Fine-Grained Network Time Synchronization using Reference Broadcasts". *Submitted for review*, February 2002.
- [22] J. Heidemann et al. "Building Efficient Wireless Sensor Networks with Low-Level Naming," in *Proceedings of the Symposium on Operating Systems Principles*, pp. 146-159. Banff, Alberta, Canada, October, 2001
- [23] W. Ye, J. Heidemann, and D. Estrin, "An Energy-Efficient MAC Protocol for Wireless Sensor Networks, IEEE INFOCOM, New York, NY, USA, June, 2002.
- [24] A. Woo and D. Culler, "A Transmission Control Scheme for Media Access in Sensor Networks," *Mobicom 2001*.
- [25] S. Klemmer, S. Waterson, and K. Whitehouse, "Towards a location-based context-aware sensor infrastructure," *unpublished*, 2000. <http://guir.berkeley.edu/projects/location/Location.pdf>
- [26] J. Broch et al. A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols *Mobicom '98*, ACM, Dallas, TX, October 1998.
- [27] B. Krishnamachari, D. Estrin, and S. Wicker, "Impact of Data Aggregation in Wireless Sensor Networks," *DEBS'02*.
- [28] D. Estrin et al. "Next Century Challenges: Scalable Coordination in Sensor Networks," *ACM/IEEE International Conference on Mobile Computing and Networks (MobiCom '99)*, Seattle, Washington, August 1999.
- [29] C. Intanagonwivat et al., "Impact of network density on data aggregation in wireless sensor networks," *submitted for publication to International Conference on Distributed Computing Systems (ICDCS-22)*, November 2001.
- [30] W. Heinzelman, J. Kulik, and H. Balakrishnan, "Adaptive Protocols for Information Dissemination in Wireless Sensor Networks," *Proc. 5th ACM/IEEE Mobicom Conference (MobiCom '99)*, Seattle, WA, August, 1999.
- [31] V. Ramasubramanian, R. Chandra and D. Mosse, "SRL: A Bidirectional Abstraction for Unidirectional Ad-Hoc Networks," *INFOCOM 2002*, June 23-27, New York, 2002.
- [32] J. Hightower, C. Vakili, G. Borriello, and R. Want, "Design and Calibration of the SpotON Ad-Hoc Location Sensing System", *unpublished*, August 2001.
- [33] S.S. Lam. "A carrier sense multiple access protocol for local networks," *In Computer Networks*, volume 4, pages 21-32, 1980.
- [34] D. Allen. "Hidden terminal problems in Wireless LAN's," IEEE 802.11 Working Group Papers, 1993.
- [35] ANSI/IEEE Std 802.11 1999 Edition.
- [36] K. Sohrabi, B. Manriquez, and G. Pottie, "Near Ground Wideband Channel Measurement", *Vehicular Technology Conference IEEE*, volume 1, pages 571-574, 1999.
- [37] C. Intanagonwivat, R. Govindan and D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," *ACM/IEEE International Conference on Mobile Computing and Networks (MobiCom 2000)*, August 2000, Boston, Massachusetts
- [38] J. M. Carlson and J. Doyle. "Highly optimized tolerance: a mechanism for power laws in designed systems", *Physics Review E*, 60(2):1412-1427, 1999.
- [39] S. Floyd and V. Jacobson. "The Synchronization of Periodic Routing Messages" *IEEE/ACM Transactions on Networking*, volume 2, pages 122-136, 1994.
- [40] Timothy G. Griffin and Gordon Wilfong "An Analysis of BGP Convergence Properties." ACM SIGCOMM, 1999
- [41] S. Basagni, "Distributed Clustering for Ad Hoc Networks", in *Proceedings of International Symposium on Parallel Architectures, Algorithms and Networks*, June 1999