

Real-Time Deployment of Multihop Relays for Range Extension

Michael R. Souryal, Johannes Geissbuehler,^{*} Leonard E. Miller,[†] Nader Moayeri
Wireless Communication Technologies Group
National Institute of Standards and Technology
Gaithersburg, Maryland, USA
souryal@nist.gov, hannes.geissbuehler@gmail.com, moayeri@nist.gov

ABSTRACT

When the range of single-hop wireless communication is limited by distance or harsh radio propagation conditions, relays can be used to extend the communication range through multihop relaying. This paper targets the need in certain scenarios for rapid deployment of these relays when little or nothing is known in advance about a given environment and its propagation characteristics. Applications include first responders entering a large building during an emergency, search and rescue robots maneuvering a disaster sight, and coal miners working underground. The common element motivating this work is the need to maintain communications in an environment where single-hop communication is typically inadequate. This paper investigates the feasibility of the automated deployment of a multihop network. A deployment procedure is proposed that employs real-time link measurements and takes into account the physical layer characteristics of a mobile multipath fading environment and the radio in use. A prototype system is implemented based on 900 MHz TinyOS motes supporting low-speed data applications including text messaging, sensor data and Radio Frequency Identification (RFID)-assisted localization. Results of deployments in a hi-rise office building are presented.

Categories and Subject Descriptors

C.3 [Computer Systems Organization]: Special-Purpose and Application-Based Systems—*real-time and embedded systems*; C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*wireless communication*

General Terms

Algorithms, Design, Experimentation, Measurement, Performance, Reliability

^{*}J. Geissbuehler was a guest researcher at NIST and is now with AdNovum, Zurich, Switzerland.

[†]L. E. Miller was with NIST until his death in June 2006.

Keywords

Deployment, multihop wireless networks, first responders, breadcrumbs, sensors, RFID

1. INTRODUCTION

Reliable communication is a critical resource for public safety operations including emergency response. A common scenario for responding to indoor emergencies, such as a building fire, is for an incident commander to be located outside of the building and for first responders to enter the building to conduct search and rescue. When the building is large, incident command can lose radio communications with first responders due to severe attenuation of the signal,¹ thereby losing the ability to track the status of personnel and to coordinate units.

One approach to alleviating the risk of communication loss (or outage) is to use multihop wireless communications to extend radio coverage. In the event that the signal strength on a single hop is too weak to sustain a desired quality of service, relays deployed between the end points can serve to extend the range of communication by repeating transmissions in a coordinated manner. The deployed relays have been referred to by some as “breadcrumbs,” conjuring the fairy tale of Hänsel and Gretel who left a trail of breadcrumbs leading back home.

The applications for rapid deployment of a multihop network extend beyond that of firefighters entering a large building to any application in which the rapid establishment of reliable communications is needed in an environment in which direct radio communication is difficult or impossible. For example, robots searching an earthquake or chemical disaster sight, or rescue teams trying to reach trapped coal miners, could use deployed breadcrumbs, or wireless relays, to maintain contact with a command and control station.

A key challenge in deployments of this nature is determining where the relays should be deployed. Existing work suggests a distance-based deployment rule. However, distance is only one factor in link quality. Non-line-of-sight propagation, shadowing and multipath fading lead to a non-deterministic relationship between signal attenuation and distance. Therefore, a simple distance-based rule is unable to capture the rich and varied channel propagation dynamics of a given environment, not to mention generalize to other environments.

¹A radio propagation study of a 14-story apartment complex showed losses of up to 50 dB going into the building [5].

This paper investigates the feasibility of a more sophisticated deployment protocol that employs real-time assessment of the physical layer of the wireless link. We propose a relay deployment algorithm that automatically adapts to the environment encountered and determines when a new relay needs to be deployed as a user moves to the edge of the existing network. The objective is to minimally impact the mission by not requiring the user to manually determine where to place relays or backtrack to repair a broken link. Furthermore, the deployment function must operate concurrently with and not interfere with the medium access control (MAC) and routing protocols that facilitate reliable delivery of application traffic. While the strategy is generic to any radio platform, we specify the algorithm and implement a prototype on a platform of 900 MHz TinyOS motes. Results are provided of the detailed link studies used in the design of the deployment algorithm, as well as of the end-to-end trials of the system in hi-rise office buildings. The prototype system provides delivery of end-user sensor measurements (e.g., first responder vital signs), two-way text messages, and RFID tags read by the mobile user and used for localization purposes.

2. RELATED WORK

Researchers at UC Berkeley in collaboration with the Chicago Fire Department have developed a wireless sensor network, SmokeNet, composed of a pre-installed backbone of nodes located at each smoke detector and doorway in a building [13]. Firefighters equipped with a MICA2 mote, a small wearable computer, and a head-mounted display are provided with their locations and messages from incident command using this network. While the intention is for the network to be pre-installed in buildings, a simple intuitive protocol for dropping small radio beacons has been suggested, such as placing them every 15–20 feet [3]. Rajant Corporation provides the BreadCrumb system for creating an IEEE 802.11b-based network on the fly. As a guideline for deployment, a table is given of suggested distances to leave between relays as a function of device power, and assuming line-of-sight propagation [10]. By contrast, our work assumes no pre-existing network and aims to automate the deployment of nodes in the network using rapid link assessment to account for the environment-specific and even radio-specific characteristics of the channel.

Researchers at Virginia Tech, in collaboration with SAIC, are also investigating rapidly-deployable wireless communications for disaster and emergency response, but their focus is on the deployment of a broadband wireless backbone network for the wider incident area, on the order of several kilometers, using, for example, Local Multipoint Distribution Service (LMDS) [8]. Their work includes a Geographic Information System (GIS) tool to generate coverage estimates for preliminary planning of the backbone hub and field unit sites. Similar to our work, this effort also includes a broadband channel sounder used to optimize the final placement of the hub and field units and an adaptive data link protocol to adjust to changing environment conditions. Our work is a natural complement to this effort in that while we focus on rapid deployment of a local area network inside buildings or tunnels, this effort would provide the backbone services to connect to a wider area network.

In developing our deployment algorithm, we undertook a number of experiments to characterize the radio channel

and investigate techniques for rapid link assessment. The topics of link characterization and assessment have received considerable attention in work on mesh and sensor networks. For example, the Roofnet study of an outdoor IEEE 802.11b mesh network concluded that links of marginal quality are the norm and that the signal-to-noise ratio (SNR) reported by the physical layer is not predictive of link reliability [1]. However, a subsequent study demonstrated that in the absence of co-channel interference not captured by the 802.11b cards' SNR measurements (Roofnet nodes experienced severe inter-symbol interference), SNR is indeed, as expected, predictive of link reliability [12].

In several other studies, a communication “gray zone” or “transitional region” was observed whereby links at some intermediate range of distances experience a large variability in reliability [7, 14, 15], largely due to small-scale multipath fading [16]. Other studies have investigated the temporal characteristics and asymmetry of wireless links (e.g., [2, 4, 15]). We cite these works for comparison with our results in the relevant sections describing our link studies below.

Using the results of link characterization studies, we design a deployment algorithm that utilizes a channel probe-and-response protocol to assess link quality. A related approach is used during the initial deployment phase of FireWxNet, a wireless sensor system for monitoring weather conditions in wildland fire environments. Whereas the primary objective of our deployment algorithm is to reduce or eliminate manual user involvement in the deployment decision, the objective of the FireWxNet deployment mechanism is simply to verify connectivity between nodes and permits active user involvement. This mechanism is more closely related to our optional *local placement assistance* feature, described in Section 5.3, which can be used in conjunction with the basic automated deployment algorithm to improve link robustness.

3. RELAY DEPLOYMENT STRATEGY

As noted earlier, one approach to relay deployment is to adopt a simple rule such as “deploy a relay every 10 m,” or “every floor in a staircase.” The assumption behind such static rules is that the transmission range of the relays is known, and the rule is to simply ensure that each relay is within range of other relays. However, simple static rules do not capture the wide variety of potential radio implementations, including different radio types, antenna types, and transmission power levels, all of which affect transmission range. More importantly, static deployment rules do not adapt to different channel propagation environments. For example, the range in an office corridor might be very different from that on a factory floor.

An adaptive deployment strategy, on the other hand, would attempt to account for differences in radio implementation and propagation environment to consistently provide connectivity across a variety of scenarios. Ideally, relays would be deployed in a manner that maximizes some performance or quality-of-service metric, such as end-to-end transmission reliability or end-to-end delay, usually subject to resource constraints, such as the number of relays deployed. In practice, however, a global optimum may be very difficult to achieve with a real-time deployment strategy, especially when link reliability is not monotonic with distance, as in the case of channels with shadowing or multipath fading.

In the absence of a constraint on the number of relays,

another simple rule might be to intentionally overbuild the network, deploying relays very close² to one another. However, even if an unlimited number of relays are available, deploying too many relays can adversely affect performance.

To understand why overbuilding the network can be disadvantageous, we first consider a simple example consisting of a linear topology with static nearest-neighbor routing, a simplified channel model, and a performance objective of minimizing the end-to-end delay of a message transmission. Let the packet success probability p_s on a link in this simplified channel model be related to the link distance d through

$$p_s(d) = e^{-Kd^n} \quad (1)$$

where n is the path loss exponent and K is a constant. The relationship given by (1) can be viewed as the average probability that the SNR exceeds a threshold in a Rayleigh fading channel, averaged over all fading realizations and conditioned on the link distance, where the average SNR is proportional to $1/d^n$.

Let us assume that delay is dominated by the number of transmissions required to send a message end-to-end. Furthermore, on each link a packet is retransmitted, if necessary, until it is successfully received. Therefore, the optimum relay spacing is that which minimizes the number of transmissions, including retransmissions. For simplicity, we assume an unlimited number of retransmissions on a link, each occurring with probability $1 - p_s$ independently of the others. The average number of transmissions per packet on a link of distance d , then, is $1/p_s(d)$. Assuming the routing protocol routes messages from one relay to the next toward the destination, then for an end-to-end transmission over m equally spaced hops³ covering a total distance of D , the average number of transmissions required to send a packet end-to-end is

$$N_m = \frac{m}{p_s(D/m)} = m e^{K(D/m)^n}. \quad (2)$$

Minimizing (2) with respect to the number of hops m yields optimum values for m and the probability of success per link of

$$m^* = D(nK)^{1/n} \quad (3)$$

$$p_s^*(D/m^*) = e^{-1/n}. \quad (4)$$

(In fact, m is an integer, but (3) ignores this for simplicity.)

We observe, first, that an optimum number of hops exists in this case, below which there would be too many retransmissions due to errors on average, and above which there would be too many retransmissions due to multiple hops. The optimum number of hops balances these two opposing factors. Second, the optimum probability of success per link transmission (4) is a function of neither the optimum number of hops nor of the end-to-end link distance, suggesting that relay deployment can be triggered based on current (last hop) link state information, without *a priori* knowledge of the final end-to-end parameters.

The example given above is artificial in that, in practice, the success probability of a deployed link would be a function of the sample SNR, not averaged over the ensemble of

²“Very close,” here, is relative to the transmission range of a relay which, as noted above, is environment-specific.

³It can be shown that uniformly spaced hops are optimum when, as in this case, the link cost is a convex function of distance [6].

possible fading realizations. In any case, it illustrates that minimizing the end-to-end delay (or the number of retransmissions, in this case), a desirable objective for supporting real-time services (e.g., two-way voice, video), cannot be achieved by overbuilding the network.

Even in other cases where more intelligent routing is used and where end-to-end reliability rather than delay is the performance metric of interest, overbuilding the network can result in excessive interference. As the number of a node’s neighbors increases, contention for the channel and collisions can also increase resulting in a decrease in reliability, long queues, as well as increased delay. Practically, there is also a limit on the number of relays available, so that efficient use of relays is often desirable.

Given the observation, albeit from the simple example above, that relays can be deployed using local link state information, we propose a deployment strategy in which the mobile node continuously measures the state of its link(s) to the network and indicates when a deploy-criterion is satisfied. Upon such indication, a new relay is activated and introduced into the network at that point. The new node may then start relaying traffic between the mobile node and the rest of the network.

The nature of the problem requires that the measurement procedure have low latency, in order to avoid disconnection of the mobile node resulting from a late deployment indication. Measuring link quality based on link layer information (e.g., percentage of frames successfully received) as described in [14], for example, incurs high latency because of the need for a sufficient number of messages to obtain a reliable statistic. Link assessment based on physical layer information (e.g., signal strength, or signal-to-noise ratio), on the other hand, can be made on a message-by-message basis and provide more timely information on the state of the link. The next section discusses the feasibility of physical-layer-based link assessment for real-time deployment, using a specific radio platform for illustration.

4. REAL-TIME LINK ASSESSMENT

To assess the feasibility of deploying wireless relays in real time, we conducted a series of experiments using 900 MHz TinyOS Crossbow MICA2 Motes (MPR400CB).[‡] Among the questions to be answered are whether physical layer measurements such as received signal strength are predictive of link reliability in practice; how useful are these measurements in a mobile environment with multipath fading; what is the impact of the displacement between the point at which measurements are made and the point at which a node is deployed; and whether measurements made in one direction are indicative of the link state in the other direction of a bidirectional link. Answers to these questions are used to guide the development and specification of a deployment algorithm.

4.1 RSS-based Link Assessment

The purpose of the first experiment was to determine whether the received signal strength (RSS) measurement of the motes is a useful predictor of link quality and, if so, to characterize the relationship between link reliability and RSS. Results were collected using a layout of 12 transmitter locations and 10 receiver locations on a single floor of an office building (see Fig. 1). At each transmitter location, a batch of 200 packets was transmitted, and the receivers

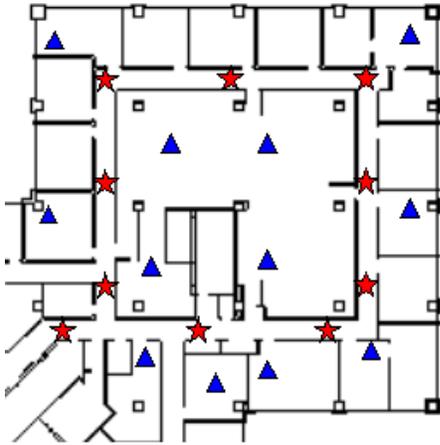


Figure 1: Layout for RSS-based link assessment experiment. Transmitter locations are triangles. Receiver locations are stars.

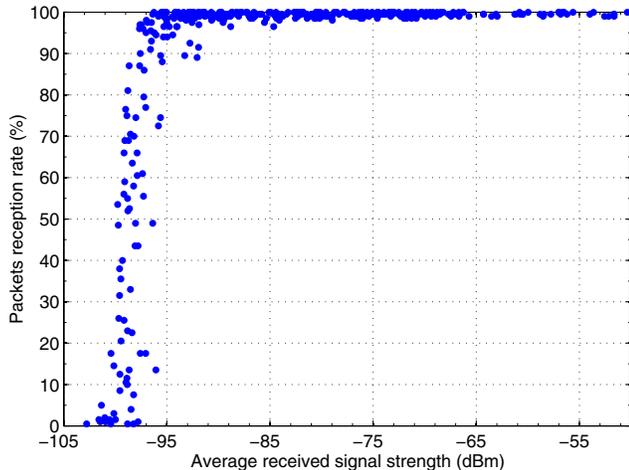


Figure 2: Packet success rate vs. received signal strength

recorded the sequence number, CRC result and RSS (in dBm) of each detected packet. The transmitter repeated the transmission batch at six different transmission power levels, (-20 , -15 , -10 , -5 , 0 and 5) dBm, in order to obtain a finer range of RSS data points.

Fig. 2 plots the percentage of correctly decoded packets as a function of the average RSS (one data point for each batch received by a mote). The results demonstrate a clear drop-off in the packet success rate around -95 dBm. This RSS threshold for link reliability has been consistently observed across multiple transceivers and is close to the receiver sensitivity of the radio, which is listed as -98 dBm.

Two important conclusions can be drawn from the results in Fig. 2. First, the sharp decrease in packet success rate with RSS implies that monitoring the number of dropped packets is not a feasible approach for determining when a new relay needs to be deployed. Assuming the average RSS decreases with distance, it is likely that by the time packet losses are encountered the link is already unsatisfactorily unreliable. Second, because of the clear relationship between

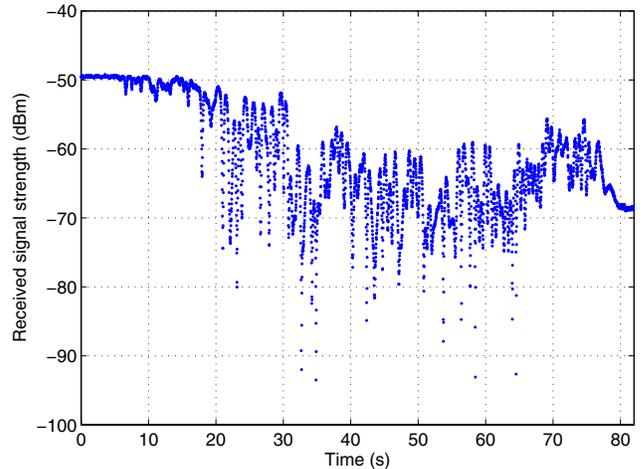


Figure 3: Received signal strength vs. time at a mobile receiver

measured RSS and packet success rate, RSS measurements are indeed predictive of link reliability and can be used to trigger deployment. A cautionary note, however, is that these measurements were made in an interference-free environment. In the presence of interference, a high RSS on a successfully received packet may not be indicative of a reliable link [12]. In such cases, a separate measure of the interference is needed and the relevant metric becomes the signal-to-interference-and-noise ratio (SINR).

4.2 Temporal Variability of a Mobile Link

The results of the first experiment measuring packet success rate as a function of RSS were gathered from a fixed topology. If the measured RSS is to be used as a metric for triggering deployment, an important question is how the RSS behaves on a mobile link. One would expect that, due to multipath reflections and scattering, RSS measurements will not be monotonic with distance.

To gain an appreciation for the extent of fluctuations in RSS at pedestrian speeds in a typical building environment, we placed a receiver on a small vehicle moving down an office corridor away from a fixed transmitter at a speed of approximately 0.3 m/s. The total distance covered was 20 m. Packets were transmitted at a rate of 50 packets/s, and the mobile receiver recorded the RSS of each detected packet. The measured RSS as a function of time is illustrated in Fig. 3. These results demonstrate fluctuations due to multipath fading of up to 20 dB over distances on the order of 10 cm, which is expected as this distance is between one-quarter and one-third of the carrier wavelength. One also observes general trends in the RSS measurements, first decreasing as the node moves away from the transmitter, then increasing, possibly due to reception of a strong signal component at the end of the corridor.

The results of the mobility experiment indicate that fast fading causes significant variation in the RSS and can make determination of the deployment point difficult. If deployment is triggered by comparing RSS measurements with a threshold, then a single measurement taken during a deep fade of the channel can prematurely trigger deployment. One can average several consecutive RSS measurements to

try to capture the large scale propagation effects, but then there is no guarantee that the final placement of the relay will not result in a worse fade than the measured average. Therefore, additional measures must be taken to reduce the probability of deploying a disconnected relay.

Two commonly used strategies to protect against deep fades are (i) building a fade margin into the link budget, and (ii) utilizing diversity. First, a fade margin can be included in our application by choosing a deployment RSS threshold that is higher than the minimum RSS needed for reliable transmission (found to be -95 dBm in the previous experiment). The current version of our prototype utilizes a fade margin of 15 dB. Second, diversity can be introduced either in frequency or in space to mitigate the effects of multipath fading.⁴ Frequency diversity can be obtained through the use of wideband physical layers, such as direct sequence or frequency hopping spread spectrum, or multicarrier systems such as orthogonal frequency division multiplexing (OFDM). Spatial diversity is often obtained through the use of multiple antennas, which may be challenging on small devices because of the spatial separation needed to guarantee low fading correlation across the antenna elements. Another approach for spatial diversity that is particular to this application is to attempt to maintain connectivity with at least two other nodes in the network. The probability that a new relay is deployed in a deep fade with respect to both nodes is significantly lower than it is with respect to either node alone.

4.3 Receiver Height

A practical consideration in many applications of real-time relay deployment is the effect of a receiver’s height on the quality of the link. For example, if the transceiver monitoring link quality is positioned at a given height, we wish to know if there is a consistent degradation in link quality if the new relay is deployed at a different height, say on the floor. A scenario of this type might be a first responder with a monitoring radio strapped to his/her belt along with a canister ejecting relays onto the floor as needed.

A fixed transmitter positioned at 38 cm above the floor of an office corridor transmitted packets to a fixed receiver positioned at one of three heights above the floor: 120 cm, 38 cm, and directly on the floor. The experiment was repeated at several transmitter-receiver separation distances. At each distance and height, 250 packets were transmitted, and the receiver logged the RSS of each detected packet. Fig. 4 plots the average RSS as a function of transmitter-receiver distance for each of the three receiver heights. The first six distance measurements were made with line-of-sight (LOS) links, while the last two (beyond 23 m) were non-LOS.

Although we observe differences in average RSS across the receiver heights of up to 10 dB for LOS links and as much as 17 dB for non-LOS, one receiver height is not consistently better or worse than the others. The differences in RSS are on the order of those observed in the preceding mobile experiment, leading us to believe that the variability is largely due to differing multipath fading realizations from one height position to another. Thus, the same precautions taken to protect against the effects of fading (margin and diversity) can be used to increase the likelihood of main-

⁴Temporal diversity is discounted because of the low time-selectivity of the channel on fixed links.

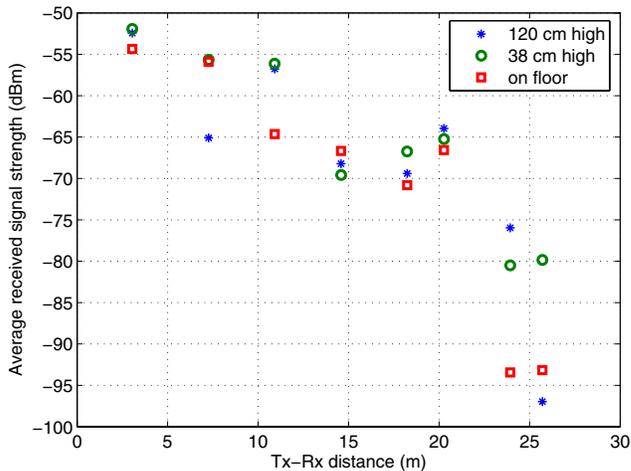


Figure 4: Received signal strength vs. tx-rx distance and receiver height

taining connectivity when relays are deployed at a different height from which the link monitoring occurs.

4.4 Link Symmetry

When using network protocols that require bidirectional links, it is important to assess the quality of each direction of the link. Bidirectional links are useful for duplex communications as well as for transmitting acknowledgments of successfully received transmissions. If the quality of a bidirectional link is symmetric—that is, transmissions in both directions of a link have similar reliability—then the link assessment function can be simplified to measure only one direction; otherwise, we need to assess both directions of a link.

In general, link reliability cannot be assumed to be symmetric. The success of a transmission depends on the received signal strength and the interference-plus-noise at the receiver at the time of the transmission. Provided each transceiver uses the same transmission power and carrier frequency, and by the antenna reciprocity property, the received signal strength in each direction is theoretically equal *at a given point in time*. In practice, same-frequency transmissions are not simultaneous and therefore temporal changes in the channel propagation profile with mobility or the changing environment result in different RSS. However, if the channel is static, or highly correlated in time, the RSS can be approximated as being symmetric. Aside from time-varying propagation conditions, differences in RSS measurements can also result from measurement error or differently calibrated receivers.

The primary source of asymmetry in static environments with equal power transmitters is the presence of differing levels of interference at the receivers. For example, one receiver may be closer to a co-channel emitter than the other, resulting in higher interference power. Other sources of asymmetry could be differences in the transmit power and receiver sensitivity of the radios [4].

To test RSS symmetry, we compared the RSS measured at each end of a fixed point-to-point link.⁵ One transceiver

⁵While previous studies (e.g., [4]) investigated symmetry in terms of packet delivery rate, we are more interested in RSS

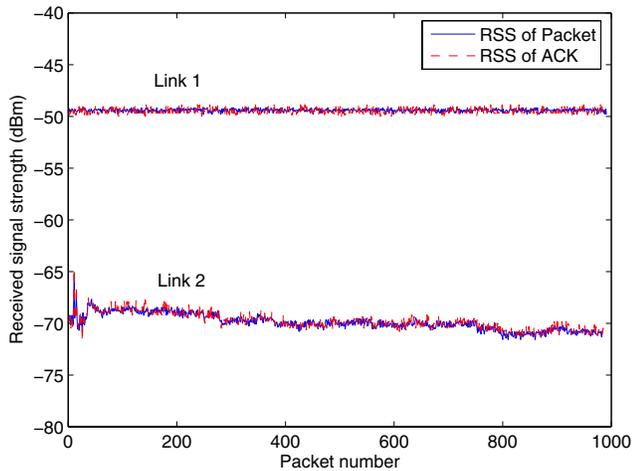


Figure 5: Instantaneous received signal strength of received packets and acknowledgments

transmitted 1000 packets at a rate of 4 packets/s, and the other transceiver immediately replied with an acknowledgment (ACK) for each packet it successfully received using the built-in ACK of the MAC. The second transceiver recorded the RSS of each packet it detected, while the first transceiver recorded the RSS of each ACK it detected. In this way, we were able to make nearly simultaneous measurements of the RSS in both directions of the link. Measurements were taken for a relatively strong link (approximately 2 m, LOS) and a second link roughly 20 dB weaker (approximately 7 m, non-LOS). Using a spectrum analyzer, we observed no other emissions on the same 900 MHz channel (i.e., an interference-free environment).

Fig. 5 plots the instantaneous RSS measurements made on receipt of the packets and the ACKs for both links over a total duration of 250 s. We observe almost identical measurements in both cases. The data labeled “Link 2,” the weaker of the two links, exhibit some variability due to slight changes in the environment during the data collection, but nevertheless the measured RSS of both directions track each other closely. Based on these results, we observe that an RSS measurement made on a reception in one direction is indicative of the RSS in the opposite direction, at the point in time at which the measurement was made.⁶

4.5 Temporal Variability of Fixed Links

An ancillary observation of the results of Fig. 5 is that the RSS measurements on these fixed links tend to be quite stable, moreso for the stronger Link 1. While Fig. 5 shows results for a period of about 4 min, our related work demonstrates stability of RSS measurements on fixed links over much longer time periods (hours to days) [12]. Other work that has investigated the temporal variability of fixed links has observed significant variability, mostly in terms of the short-term packet reception rate, especially on weaker links [2, 14, 15]. The difference in variability between weaker and

symmetry, which need not be the same for the reasons outlined above.

⁶The emphasis on the time of the measurement is important. Beyond the *coherence time* [11] of the channel, the measurement is obsolete.

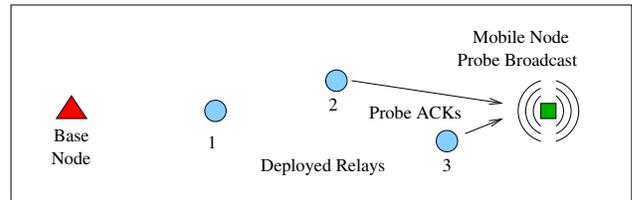


Figure 6: Channel probing

stronger links can be attributed to the fact that links operating close to the critical transition in packet reception rate (PRR) with signal strength shown in Fig. 2 are sensitive to small changes in RSS. Thus, changes in the environment of these links (e.g., moving people, doors opening or closing) can cause large changes in PRR.

The deployment algorithm described in the following section aims to deploy links operating in the region to the right of the transition of PRR with RSS in Fig. 2. Deployed links in this RSS region tend to be more stable over time in terms of packet delivery. In the event that a deployed link ends up close to the transition, a proposed power control algorithm (described at the end of Section 6.4) would increase the transmission power of each node on the link with the objective of moving the operating point well to the right of the transition where packet delivery is more stable over time.

5. RELAY DEPLOYMENT ALGORITHM

Based on the results of the preceding section, we devised a parameterized deployment algorithm based on rapid probing by the mobile node. Every Δ seconds, the mobile node broadcasts a channel probe. Any relay in the network that hears the probe responds with a probe acknowledgment packet (Fig. 6). The mobile node utilizes channel measurements made on reception of the probe ACKs to determine when a new relay needs to be deployed. Because of the expected symmetry of RSS measurements for short probe-ACK turnaround times (see Section 4.4), the mobile node only uses the RSS measurements made on receipt of the ACKs. If bi-directional measurements were needed (for example, of the SINR in each direction), a node receiving a probe could include that measurement in a data field of its replying ACK packet.

To mitigate collisions of probe ACKs from possibly multiple responding relays, all transmissions utilize Carrier Sense Multiple Access (CSMA) with random backoff. Through simulation and experimentation, we found that initial and congestion backoff windows of 128 bytes (53 ms) and 64 bytes (27 ms), respectively, provide an adequate tradeoff between delay and collision avoidance for a probe period of $\Delta = 100$ ms.

The mobile node measures the RSS of each probe ACK it receives and stores measurements for the most recent N probe periods from each responding relay. The mobile uses these previous N measurements to maintain a running average of the RSS of each link. The question arises as to how to treat missing probe ACKs. A probe ACK might be missed either as the result of weak signal strength or because of a collision with an interfering transmission. The former event should somehow be factored into the computation of average RSS, to avoid biasing the computed average higher than

the true average. In the latter event, that of a collision, it might be better to ignore the missing ACK, since what we are interested in measuring is the average RSS, not collision rate. In the interest of erring towards a more conservative deployment protocol, we elected the first alternative, that is, to record a low default value, S_0 , of the RSS for each missing ACK from a given node, where S_0 is less than the receiver sensitivity of the radio (observed to be between -95 and -100 dBm in Fig. 2).

A visual indication is given to the user to deploy a new relay when the highest running RSS average goes below a threshold, S_{th} . In other words, letting $S_{i,j}$ be the RSS measurement from the i^{th} replying node during the j^{th} probe period, deployment is triggered when

$$\max_i \left\{ \frac{1}{N} \sum_{j=1}^N S_{i,j} \right\} < S_{th}.$$

As discussed in Section 4.2, protection against link outage from deploying into an unfavorable fading position can be given in two ways. First, fading margin can be provided by choosing a deployment threshold $S_{th} > -95$ dBm. For example, a choice of $S_{th} = -80$ dBm provides a 15 dB margin and would account for most of the small-scale fades observed in the mobility experiment of Section 4.2 and height experiment of Section 4.3. Second, by modifying the deployment criterion to trigger deployment when the L th strongest link crosses the threshold—rather than the strongest link, as stated above—the algorithm aims to maintain multiple links with the network, providing link diversity.

5.1 Parameter Selection

Implementation of the deployment algorithm requires selecting values for the following parameters: the probe period (Δ), the RSS averaging filter length (N), the threshold for triggering deployment (S_{th}), the default RSS value for missed probe ACKs (S_0), and the link diversity order (L).

Values for the probe period Δ and averaging filter length N were selected after studying the performance of the algorithm for different (Δ, N) pairs. The product ΔN represents the duration of the observation window over which the RSS average is computed. First, we tested values for (Δ, N) corresponding to a fixed observation window of $\Delta N = 4$ seconds. For each trial, the measuring node executing the real-time link assessment algorithm was carried away from a fixed relay in an office building environment at walking speed. When the measuring node gave the indication to deploy, the node was placed on the floor and a long sequence of packet transmissions was initiated over the fixed link to measure the steady-state RSS.

Fig. 7 depicts the results of four separate trials for each of three pairs of (Δ, N) . The trials consisted of two different paths, and two trials of each path. Results are given in terms of the difference between the steady-state RSS and the deployment threshold, chosen here to be $S_{th} = -80$ dBm. In most cases, the four-second filter latency results in a post-deployment RSS that is lower than the deployment threshold RSS. We observe, though, that a probe period of $\Delta = 100$ ms results in more consistent results for the post-deployment RSS than probe periods of 200 and 500 ms, due to the higher resolution sampling of the observation window. However, due to the latency of the averaging filter, the choice of $\Delta = 100$ ms and $N = 40$ results in a post-deployment RSS

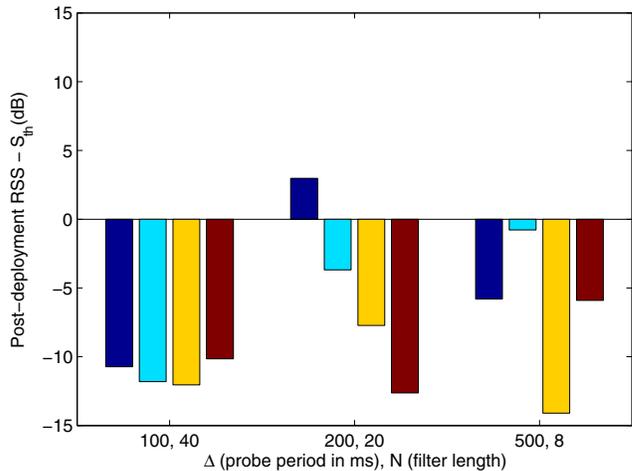


Figure 7: Deviation from deployment threshold for three different choices of probe period Δ and RSS filter length N

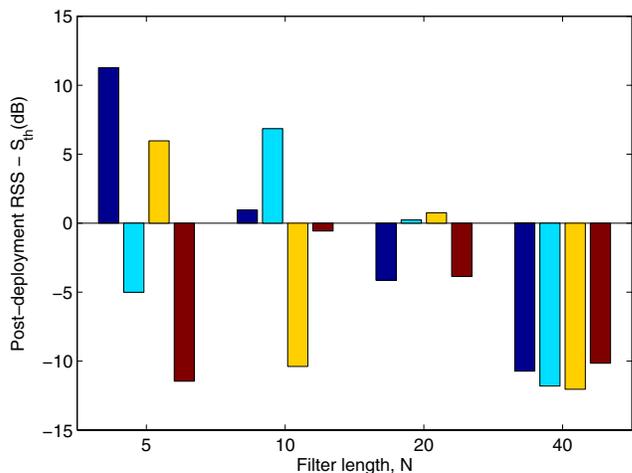


Figure 8: Deviation from deployment threshold vs. RSS filter length N ; probe period $\Delta = 100$ ms

roughly 10 dB below the deployment threshold, leaving little margin for additional fading attenuation.

We then examined different choices of the RSS filter length, N . Fig. 8 illustrates results of trials for four different values of N and a fixed probe period of $\Delta = 100$ ms. A clear trade-off is apparent between the sensitivity to fading fluctuations of a short filter and the improved averaging but higher latency of a long filter. The choices of $N = 20$ and $\Delta = 100$ ms appears to strike a balance between consistency and latency. Five additional trials of this choice of parameters were done, and the results are shown in Fig. 9. They indicate that the deployed link is usually within 5 dB of the threshold, but that there still exists the probability of a wider deviation, as evidenced by trial number 6.

Table 1 summarizes the selected values of the deployment algorithm parameters which are used in the end-to-end implementation and tests described later.

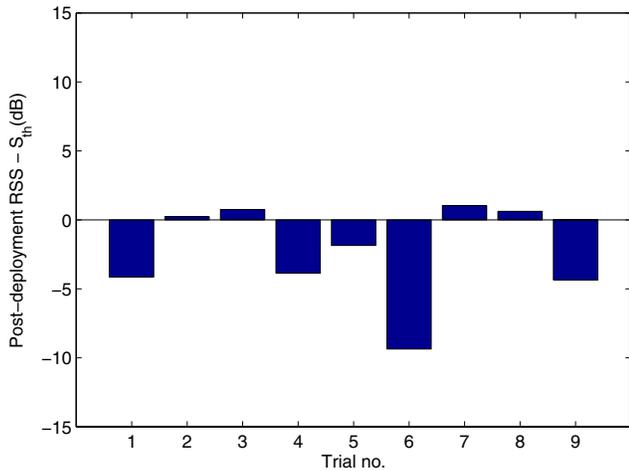


Figure 9: Deviation from deployment threshold with probe period $\Delta = 100$ ms and RSS filter length $N = 20$

Table 1: Deployment Algorithm Parameters

Δ	Mobile probe period	100 ms
N	Averaging filter length	20
S_{th}	Threshold signal strength	-80 dBm
S_0	Default RSS measurement value	-100 dBm
L	Link diversity order	1

5.2 Adaptive Probing

Rapid probing of the channel is particularly necessary when the channel characteristics are varying in time, such as when the monitoring node is moving. However, 100 ms probing is overkill if the node is stationary, unnecessarily occupying the channel that could otherwise be used for transmitting user traffic. Therefore, we implemented an algorithm to adapt the probe period based on the motion of the monitoring node. Using accelerometer readings of an accompanying sensor board, the link assessment function switches to a longer probe period $\Delta' > \Delta$ when the node is detected to be stationary (i.e., when differential accelerometer readings are below a threshold). As soon as the accelerometer registers changes above the threshold, rapid probing with shorter probe period Δ is restarted. For our implementation, we chose the slower probe period to be $\Delta' = 500$ ms.

5.3 Local Placement Assistance

The mobile link experiment in Section 4.2 confirmed that small changes in position of the relay can lead to substantial changes in the received signal strength, due to the effects of multipath fading. *Local placement assistance* is an optional feature that lowers the probability of deploying a new relay in a deep fade. With this feature, a visual indication is provided by the relay node when its strongest link to the network is below some threshold. The deployment mechanism, whether it be a human operator or a robot, can then choose to displace the relay a small distance, on the order of one-quarter of the carrier wavelength (e.g., 8 cm at 916 MHz). In many cases, such a displacement will move the relay out of an unfavorable fade.

The measurement of link quality is made by overhearing

the probe ACKs of other nodes that are still responding to the mobile node’s channel probes. If the average RSS of the probe ACKs is below a critical threshold (-90 dBm in our implementation), the newly deployed relay activates a red LED; otherwise, it activates a green LED. Deployment experience with and without local placement assistance is described in the following section.

Local placement assistance can be viewed as a limited form of backtracking for use by applications in which the time to adjust the position of the relay is available, such as search operations using robots. When relay adjustment is not possible (e.g., when used by firefighters in an emergency), local placement assistance can be ignored or turned off.

Recall that alternative means of avoiding deep fades are provided by other physical layer technologies. While this study and prototype used narrowband radios, wideband technologies such as direct sequence or frequency hopping spread spectrum and OFDM are inherently more robust in the face of multipath fading. A final solution for actual use in real-time deployment scenarios would likely use such wideband physical layer alternatives, obviating the need for small-scale local placement assistance and relying solely on the macro-scale assessment of the basic, automated deployment algorithm.

6. REAL-TIME DEPLOYMENT PROTOTYPE

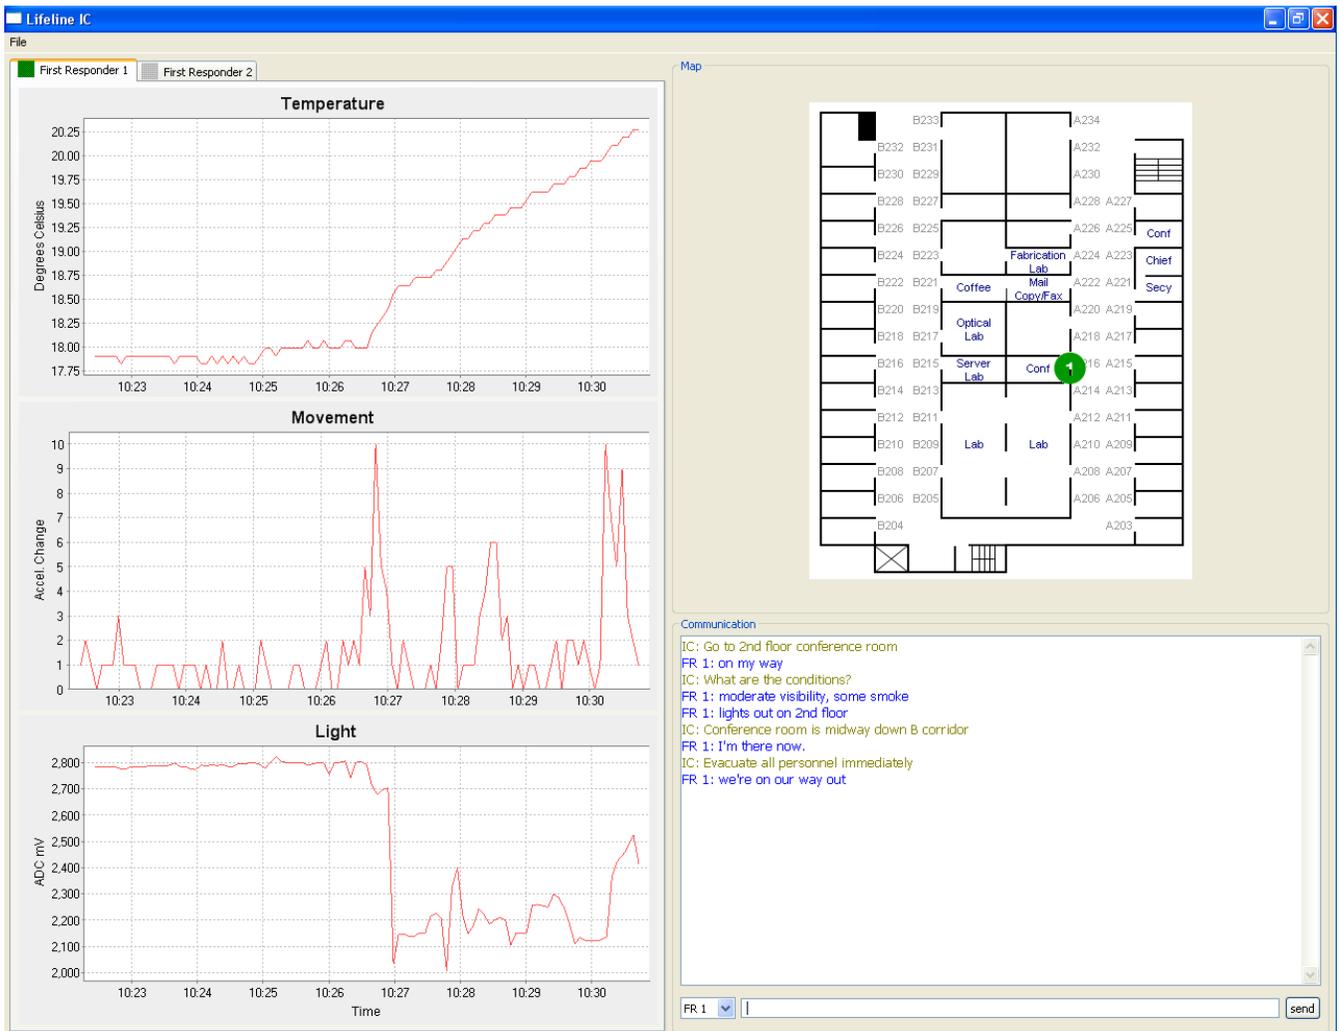
This section describes a prototype system we implemented for testing and evaluating the end-to-end performance of a multihop communications network deployed in real time. First, an overview is given of the application, followed by a system-level description of the hardware components. We then describe the routing protocol we developed for forwarding messages between the base and mobile nodes. The section ends with results of experimental trials emulating a first responder ascending an office building.

6.1 Application Overview

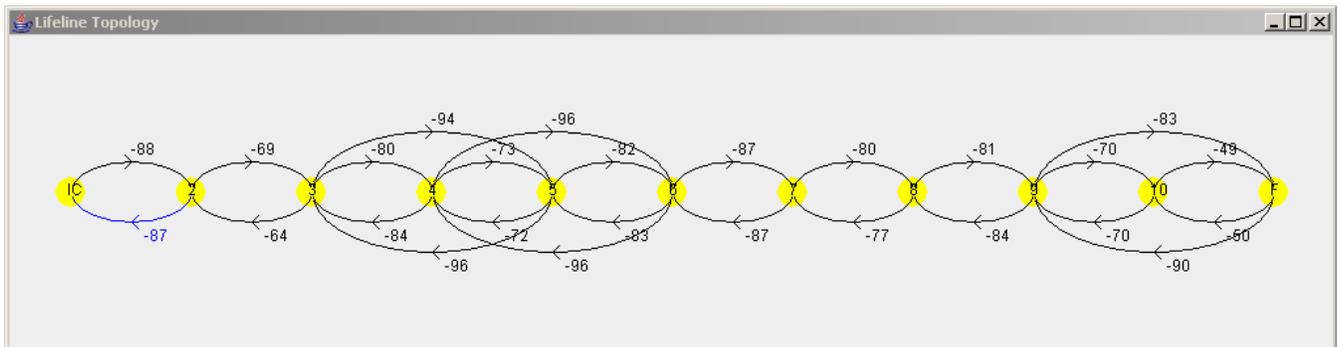
The scenario is that of a base node (e.g., incident command) located outside a building and a mobile node (e.g., first responder radio) that is carried into and about a building. The mobile node continuously executes the link assessment and deployment algorithm and gives a visual indication when to deploy a new relay. When such an indication is given, a relay is turned on, placed on the floor, and begins to forward traffic originated by either the mobile node or base node. This process continues as the mobile node moves further away from the base, creating a multihop network for monitoring of and communication with the mobile user.

The base node is connected to a laptop displaying a graphical user interface (see Fig. 10). Likewise, the mobile node is connected to a PDA for user input/output (see Fig. 11). Supported applications include monitoring sensor data reported by the mobile node, two-way text messaging between the base and mobile nodes, and RFID-assisted localization of the mobile node.

Sensor measurements collected every 5 seconds by the mobile node are displayed graphically at the base node (see Fig. 10(a)). Currently, the sensor values are temperature, light and movement (based on accelerometer readings); in practice, the node would report the users’s vital signs, such



(a) "Incident command" window. Left: mobile node sensor readings; above right: RFID localization map; below right: text message chat window.



(b) Logical network topology

Figure 10: Graphical user interfaces at base node



Figure 11: Graphical user interface at mobile node

as heart rate, breathing rate, body temperature, etc. A chat window displays text messages exchanged between the base node and the mobile node. Using neighbor information sent by each node, the base node's GUI also shows the logical topology of the network along with the average RSS (in dBm) measured on each link (see Fig. 10(b)).

The approximate location of the mobile node is provided by passive RFID tags pre-installed in the building (e.g., at each doorway). The mobile node is integrated with an RFID reader that reads a tag's ID and sends it to the base node. At the base node, the tag ID is used to display the location of the tag on a building floor plan.

6.2 Prototype Hardware

The hardware components of the system are as follows.

1. **Base Node.** The base node consists of a MICA2 mote connected via a USB gateway to a laptop computer (see Fig. 12).
2. **Mobile Node.** The mobile node consists of a MICA2 mote and sensor board (Crossbow MTS310) connected via a serial gateway to a handheld PocketPC PDA (HP iPAQ H3800). A 13.56 MHz RFID reader (ACGPass Plug-In Reader Module) is plugged into the CompactFlash slot of the PDA's expansion pack (see Fig. 13).
3. **Relay.** A deployed relay is a MICA2 mote with a smaller antenna (see Fig. 14).

On most of the motes we soldered an external antenna directly to the circuit board, replacing the quarter-wavelength monopole antenna plugged into the MMCX connector, which we found to be a loose connection in some cases. For the relays, we used a compact 15 mm length antenna (Linx Technologies ANT-916-JJB-RA), and for the base node, we used an 8 cm monopole antenna (Linx Technologies ANT-916-PW-QW). The mobile node uses the MMCX-connected monopole antenna that comes with the MICA2 mote.

Development of the prototype has focused on communications and networking aspects of the solution. Components

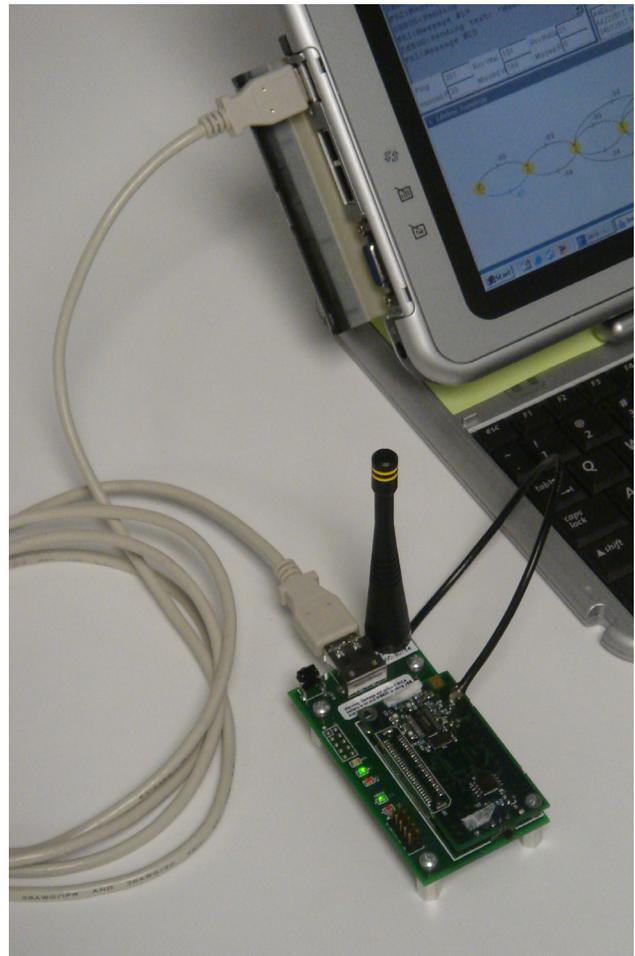


Figure 12: Base node

of a final system would, of course, be enclosed in ruggedized packaging appropriate for the target environment (e.g., fire-proof packaging for firefighters). Furthermore, rather than requiring the mobile user to actively deploy a relay upon indication by the deployment algorithm, the final system would likely activate and deploy relays automatically from a pack or canister of relays attached to the user's belt. The MICA2DOT mote (Crossbow MPR5x0) shown in Fig. 15 shows the potential size of a relay in the final system.

6.3 Routing Protocol

Existing TinyOS routing solutions which we are aware of are too slow to react to the rapid changes in topology of a network deployed as mission-critical data is being sent. In evaluating alternatives, we chose to base our implementation on Destination-Sequenced Distance Vector (DSDV) routing [9] because it isolates next-hop changes in routing tables to nodes in the recently-deployed part of the network. We adapted DSDV to incorporate link quality in the route metric in order to improve performance in the presence of links with varying quality.

As in standard DSDV, each node maintains a routing table with routes to destinations. (In our case, the base and mobile are the only valid destinations.) A routing entry for a given destination consists of the ID of the next-hop node



Figure 13: Mobile node with RFID reader

of that route, the number of hops to the destination, the sequence number of that route (originated by the destination), and the time of the latest update of that route. In our adaptation of DSDV, a routing entry also includes the number of “weak links” in the route and the average RSS of the weakest link, S_{\min} . A weak link is defined as a link with a measured average RSS below a threshold signal strength, S_{weak} .

Each node periodically broadcasts a route advertisement, with period T_{adv} , listing its best route information for each destination. Advertised route information consists of the destination of the route, the number of hops to the destination, the number of weak links in the path, the average RSS of the weakest link, the advertising node’s next-hop node for that route, and the sequence number. Each time a destination advertises a route to itself, it increments the sequence number of the route by one. Upon receiving an advertised route, a node creates or updates the corresponding entry in its table by incrementing the advertised number of hops by one. The advertised number of weak links is incremented if the measured average RSS of the link on which this route information was received is less than S_{weak} . In addition, the node updates S_{\min} if the RSS on this link is less than the advertised S_{\min} . Finally, the sequence number of this route information is recorded, as well as the time it was received. An advertised route with next-hop node ID equal to the receiving node’s ID is ignored (an obvious loop).

A node stores up to three routes to each destination. Among multiple routes to a given destination, the “best



Figure 14: Relay

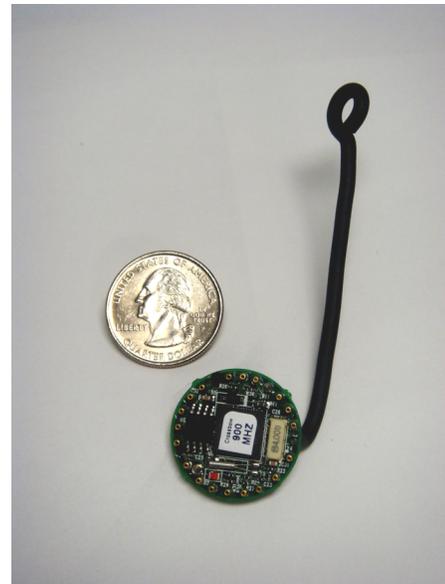


Figure 15: MICA2DOT relay

route” is the minimum hop route with no weak links. Ties in the number of hops are broken by the route with the larger sequence number. If only routes with weak links are available, then the best route is the one with the largest value of S_{\min} . Ties in S_{\min} are broken by the number of weak links. Finally, ties in the number of weak links are broken by the route with the larger sequence number.

Two link layer protocols are supported: acknowledged and unacknowledged. Unacknowledged messages are simply transmitted once to the next-hop node of the destination. Acknowledged messages are retransmitted, if necessary, using Stop-and-Wait Automatic Repeat Request (ARQ) with retransmission time-out T_{out} and maximum number of retransmissions N_{retx} .

A routing table entry is deleted if it is more than $3T_{\text{adv}}$ old. An entry is also deleted if a message sent to the next hop is not acknowledged after the maximum number of retransmissions.

Table 2: Routing Protocol Parameters

S_{weak}	Weak link threshold	-90 dBm
T_{adv}	Route advertisement period	2 s
T_{out}	Retransmission time-out	125 ms
N_{retx}	Maximum num. of retransmissions	10

Table 3: Message Delivery Rates for Hi-Rise Trials

Trial No.	Deploy	Stop			Return
	Ping	Ping	Text Δ	Text ∇	Ping
With Local Placement Assistance					
1	99 %	100 %	100 %	80 %	95 %
2	98 %	100 %	100 %	99 %	100 %
3	96 %	99 %	100 %	100 %	98 %
Without Local Placement Assistance					
4	90 %	99 %	100 %	99 %	100 %
5	98 %	98 %	99 %	90 %	100 %
6	55 %	35 %	28 %	36 %	71 %

Values for the routing parameters used in our implementation are summarized in Table 2.

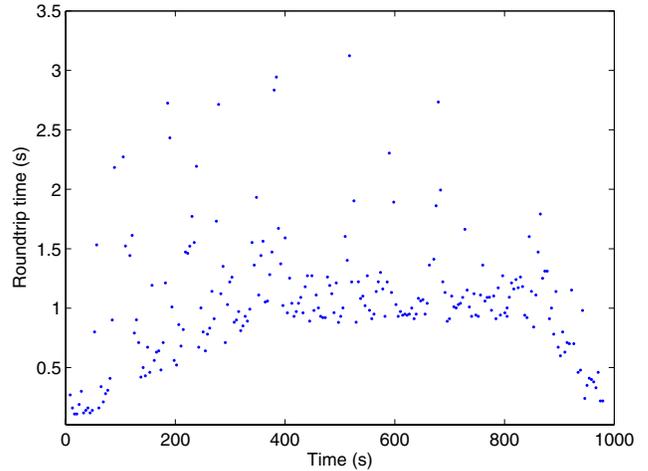
6.4 Experimental Trials

The prototype for real-time network deployment was tested in the eleven-story Administration building on the main campus of the National Institute of Standards and Technology. In each trial, the base node was located in the ground floor lobby. The mobile node was started next to the base, was walked to a stairwell and then up to the top floor, with relays being placed on the floor when indicated by the deployment algorithm. After stopping at the top for data collection, the mobile node was then walked down the same path to the base node on the ground floor, passing the relays that were deployed on the way up. While the path going up tested the deployment aspect of the system, the path coming down tested the ability of the routing protocol to handle the relatively rapid changes in routes as the mobile node made its way back to the base. Deployment was tested with and without the local placement assistance feature described in Section 5.3.

In most cases, we were able to reach the 10th or 11th floor with 9 deployed relays. Typically, 2 relays were deployed between the base node and the stairwell door, and the remainder were deployed inside the stairwell, roughly one relay per $1\frac{1}{2}$ floors.

During the deployment phase, the stop phase, and the return phase, message traffic was automatically generated by the base node application to measure delivery rates and round-trip delays. Specifically, a ping-like message was sent every 4 s to the mobile node’s mote, which sent a reply to the base. While at the top of the building, the base node application also sent automatically generated text messages every 4 s to the peer application on the mobile node’s PDA. The PDA logged each message that was received and replied to it with a text message. Round-trip delay and delivery rate were measured from the ping messages, and one-way delivery rates were measured from the auto-text messages.

Table 3 lists the delivery rates for six representative trials, three with and three without local placement assistance. The results shown are for relaying with link layer ARQ as described in Section 6.3, but without higher layer retransmis-

**Figure 16: Ping roundtrip delay vs. time of trial 3**

sion which would otherwise hide the performance delivered by the multihop network. We observe that with local placement assistance, most delivery rates are at 95% or greater, while without local placement assistance, the network is at greater risk of low quality links impacting end-to-end performance, as evidenced by trial 6. Traffic that is intolerant of loss, such as text messages and RFID localization tags, would employ retransmission at the application layer, in practice.

Fig. 16 plots ping roundtrip delay over the course of trial 3. During the deployment phase (between 0 and 400 s), roundtrip delay increases from about 100 ms to around 1 s, or about 100 ms per hop. Some roundtrip delays reach 2 or 3 s, due at least in part to the latency of routing updates when new relays are deployed. During the return phase (between 800 and 1000 s), roundtrip delay decreases as the mobile node approaches the base node and hop lengths shrink. Such data can be used to adjust the retransmission timeout value of an application-layer retransmission scheme.

The logical topology of the deployed network for trial 1 with local placement assistance is shown in Fig. 17(a). The leftmost node represents the base node, or incident command (IC) of an emergency response application, while the rightmost node represents the mobile node, or first responder (F). Numbered nodes represent relays in the order in which they were deployed. Directional links are labeled with the average RSS in dBm measured at the receiver. In some cases, relays have links to more than just their immediate predecessor and successor (i.e., relays 3, 5 and 9, in this example). As expected, with local placement assistance no relay is left without a link above -90 dBm, the threshold used by the placement assistance algorithm to prompt for a better local position. By contrast, Fig. 17(b) shows the logical topology of trial 4 deployed without local placement assistance. Here, links 5-6 and 8-9 exhibit an RSS of -90 dBm or below.

In both topologies shown, the adapted DSDV routing protocol avoids weak links when possible, opting for more reliable, if longer, routes to the destination. For example in the topology of Fig. 17(a), though an 8-hop path exists between the base (IC) and mobile (F) nodes, packets are routed over the 10 hops consisting of links above -90 dBm.

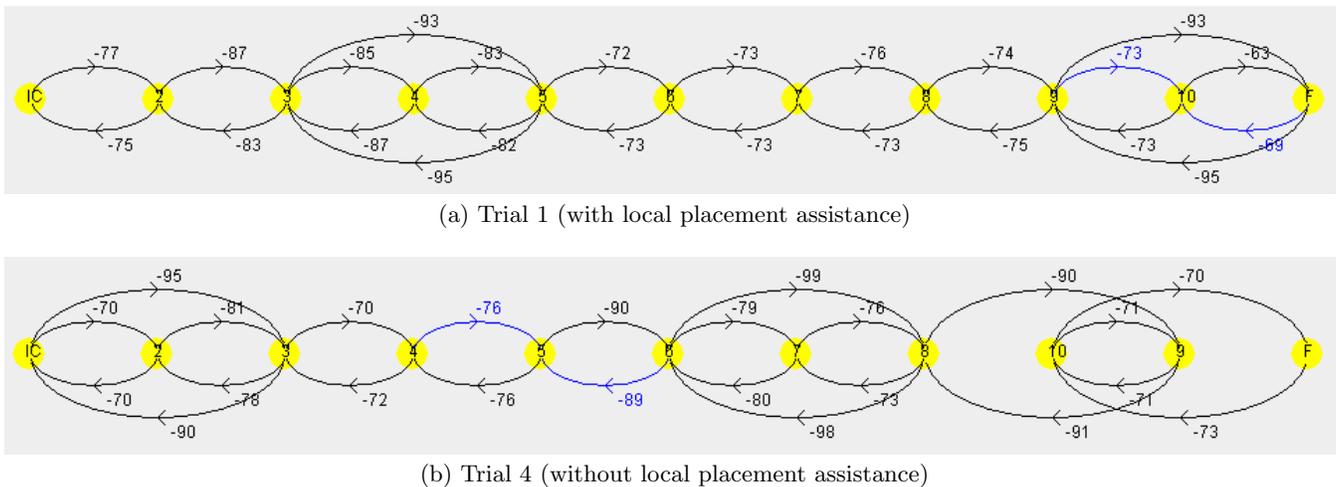


Figure 17: Logical topologies and average link RSSs (in dBm) for two hi-rise building trials

In general, through the course of these and other trials conducted in this 11-story building as well as a shorter, 6-story office building, the deployment algorithm resulted in networks with consistently reliable links when used with local placement assistance, and instances of disconnection when used without it. For example, in trial 6 above, a weak link (below -90 dBm) existed between the base node and the first deployed relay which led to periods of network disconnection during which no messages were received. However, as indicated earlier, alternative wideband physical layer technologies can mitigate to a great extent the effects of multipath fading. With such capabilities, the need for local placement assistance may be eliminated.

In work done subsequent to these experiments, we have also implemented a power control feature in the prototype which allows each relay to adjust its transmission power to strengthen weak links and, conversely, to reduce interference when the measured RSS is more than sufficient. In this implementation, each node monitors the RSS of its next-hop neighbors and issues power adjustment commands (up/down) to its neighbors in an additional field of its advertised routing table. If a given node receives power-up (power-down) commands from either (both) of its up-stream and downstream neighbors, it increases (decreases) its transmission power. Preliminary experience suggests that this power control feature reduces the incidence of unreliable links. It also enables the network to adapt to varying post-deployment conditions affecting signal strength.

7. CONCLUSION

Multihop wireless networks have been proposed as a means to extend radio coverage in hard to reach areas such as large buildings made of heavy construction, disaster sites and underground structures. They are particularly attractive in emergency situations for which rapid deployment and turn-key operation are desired. However, other than general heuristics requiring significant user involvement, little work has been done to formalize a strategy for real-time deployment of such networks.

This paper has investigated the feasibility of an automated deployment strategy that adapts to the environment at hand

and requires little if any user involvement. We proposed a relay deployment algorithm based on rapid measurement of the physical layer of the wireless link, and suggested ways for mitigating the impact of multipath fading. This algorithm, along with associated MAC and routing layer protocols, was implemented on a platform of 900 MHz TinyOS motes and was used to demonstrate and test its application to incident area first responder communications, supporting two-way messaging, sensor data and RFID-assisted localization. Results of experimental trials in a hi-rise office building validate the deployment strategy and point out areas for further work, including more robust automated deployment leveraging power control and diversity. Other areas of planned work include migrating to a higher bandwidth platform (e.g., at 2.4 GHz) to test the feasibility of supporting voice and video.

Acknowledgments

The authors would like to thank Julio Perez and Kamran Sayrafian for providing the data for Fig. 2, and Andreas Wapf and Huiqing You for their assistance in collecting the trial data presented in Section 6.4.

In memory of our colleague and friend, Leonard E. Miller (1942–2006).

8. REFERENCES

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Proc. ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 121–132, 2004.
- [2] A. Cerpa, J. L. Wong, M. Potkonjak, and D. Estrin. Temporal properties of low power wireless links: Modeling and implications on multi-hop routing. In *Proc. 6th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 414–425, May 2005.
- [3] L. Drahos. Chicago fire dept., Berkeley testing rescue communication system. *Responder Safety*, November 14, 2006.

- [4] D. Ganesan, D. Estrin, A. Woo, D. Culler, B. Krishnamachari, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. UCLA Computer Science Technical Report, UCLA/CSD-TR 02-0013, 2002.
- [5] C. L. Holloway, G. Koepke, D. Camell, K. A. Remley, and D. Williams. Radio propagation measurements during a building collapse: Applications for first responders. In *Proc. Intl. Symp. Advanced Radio Technology (ISART)*, pages 61–63, Mar. 2005.
- [6] S. Lee, B. Bhattacharjee, and S. Banerjee. Efficient geographic routing in multihop wireless networks. In *Proc. 6th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 230–241, May 2005.
- [7] H. Lundgren, E. Nordström, and C. Tschudin. Coping with communication gray zones in IEEE 802.11b based ad hoc networks. In *Proc. 5th ACM Int. Workshop on Wireless Mobile Multimedia*, pages 49–55, 2002.
- [8] S. F. Midkiff and C. W. Bostian. Rapidly-deployable broadband wireless networks for disaster and emergency response. In *Proc. First IEEE Workshop on Disaster Recover Networks (DIREN '02)*, June 2002.
- [9] C. E. Perkins and P. Bhagwat. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *Proc. ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 234–244, 1994.
- [10] BreadCrumb(R) wireless network user guide for the BreadCrumb(R) wireless network release 9.0. Rajant Corporation, Aug. 18, 2006.
- [11] T. S. Rappaport. *Wireless Communications: Principles & Practice*. Upper Saddle River, New Jersey: Prentice Hall PTR, 1996.
- [12] M. R. Souryal, L. Klein-Berndt, L. E. Miller, and N. Moayeri. Link assessment in an indoor 802.11 network. In *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, Apr. 2006.
- [13] UC Berkeley Mechanical Engineering Department. FIRE project. Available online at <http://fire.me.berkeley.edu/>.
- [14] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Proc. SenSys'03*, Nov. 2003.
- [15] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proc. SenSys'03*, Nov. 2003.
- [16] M. Zuniga and B. Krishnamachari. Analyzing the transitional region in low power wireless links. In *Proc. First Annual IEEE Commun. Society Conf. on Sensor and Ad Hoc Commun. and Networks (SECON)*, Oct. 2004.

[‡]Certain commercial equipment, instruments, or materials are identified in this paper in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.