

Wireless Testbed Bonsai

Vinayak Naik, Emre Ertin, Hongwei Zhang, Anish Arora
The Ohio State University,
Columbus, OH-43210
Email: {naik,ertine,zhangho,anish}@cse.ohio-state.edu

Abstract—This paper focuses on the spatial scaling of wireless network behaviors, ranging from low-level signals to network protocol executions to high-level application behaviors. This task is complicated as wireless channels exhibit complex fading patterns across space and time. And there have been several instances of “phase transitions”, where wireless network protocols and applications that perform well at small scales fail to do so at higher scales. Presently, experimentation with wireless testbeds is preferred to simulations for obtaining high fidelity testing, but at the same time experimentation is also known to be more inconvenient and expensive, especially for large scale networks. We therefore investigate the ability of high fidelity spatial downscaling of such networks, for making more convenient the testing and predictability of designs for large scale networks.

Specifically, we present necessary and sufficient conditions for high fidelity scaling of wireless networks and evaluate the feasibility of scaling an indoor multihop IEEE 802.11b network by reducing the separation between the nodes by a constant factor. Our validation is in terms of experimental studies of down-scaling performed in the *Kansei* testbed¹ at Ohio State University, and compare the performance at the physical, link, and messaging/dissemination layers before and after scaling.

I. INTRODUCTION

The scale of ad hoc wireless network deployments is growing rapidly. Already there have been deployments with several hundreds of devices networked in a peer-to-peer, multihop manner, so to provide dense communication coverage over significantly sized regions. The complex nature of the wireless channels and, in several cases, the limited resources available to the devices introduce significant challenges in developing protocols and ap-

plications for these network deployments. Existing protocols and mechanisms (for example, those that have succeeded in wired internetworks or in small scale wireless deployments) are not always suited to large scale wireless networks. And simulation has been shown to be a poor option for validating new protocols at scale in a high fidelity manner because efficient modeling of the electromagnetic wave propagation in a multipath environment with high accuracy has yet to be achieved for large scale networks. Thus, experimentation with wireless network testbeds is preferred, yet it is often prohibitive to deploy testbeds at full scale (for reasons of space, providing/maintaining power supply, maintenance, and cost).

Most existing testbeds, even those with several hundreds of devices, accommodate nodes in a space that is compact relative to common case deployment spaces. To compensate for reduced internode distances, a number of them, including the *Kansei* [10], *MiNT* [5] and *Orbit* [8] testbeds, attempt to shrink a wireless network into a smaller space while maintaining link characteristics through power control. They reduce transmission power via software control and/or radio frequency (RF) attenuator hardware, and reception power via augmented environmental noise and/or attenuator hardware.

This paper focuses on the fidelity of spatial scaling of wireless network behaviors, ranging from low-level signals to network protocol executions to high-level application behaviors. The problem of high-fidelity spatial downscaling of behaviors is motivated by the need for more convenient testing and predictability of protocols for large scale networks. Specifically we study, via experiments in the *Kansei* testbed, *how the behaviors of a network at one spatial scale and power level can be related to*

¹This work was sponsored by DARPA NEST contract OSU-RF program F33615-01-C-1901 and NSF grants CNS-0520222 CCR-0341703.

the behaviors of the network at another spatial scale by suitable choice of its power level. We assume that the same power level is used by all network nodes. Note that, even at one spatial scale, behaviors are affected by other factors such as the variability in the node hardware, antenna orientation, node placement error, temporal variation of environments, etc.

A case study in spatial scaling. We begin with anecdotal evidence of both positive and negative results in reproducing behaviors of network protocols across spatial scales with power control. The evidence is from the *ExScal* project [1], [3], which designed and fielded a large scale wireless sensor system in December 2004 in an open field in a forested area in central Florida. The deployment, which spanned a 1.3km by 300m area, included a lower-tier about 1200 “mote” wireless sensor device network, and a about 200 Stargate nodes higher-tier multi-hop peer-to-peer 802.11b wireless network. Leading up to the full *ExScal* deployment, we conducted spatial scaling tests on our higher-tier protocols on a 7×7 grid of Stargates nodes 45m and with 90m separations respectively and at multiple power levels. We observed that for a certain protocol (*Sprinkler* [9], a bulk data dissemination service there were power levels for which the behaviors at the two separations (obtained from many runs of the protocols) were essentially identical. We were not able to replicate such a result for another protocol (*LOF* [13], a beacon-free routing service).

Contributions of the paper. In this paper, we reproduce the *Sprinkler* and *LOF* spatial scaling results on the indoor *Kansei* testbed.

Towards explaining these differing results, we show that the power control approach to scaling is valid only under a specific large scale fading model. And the two networks at various scales are equivalent in a probabilistic sense. It follows from this result that while there is a strong relationship between the behaviors of any protocol on *the set of all instances* of both networks (where the instances allows for variation of node hardware, antenna orientation, placement error, etc.), there need not be a strong similarity between the protocol’s behaviors on *specific instances* of the both networks.

We then study a conjecture of the similarity of

protocol behavior in different scales empirically and show that the particular links used by the protocol affects the similarity of the behaviors at the different scales. Specifically, “short” (aka, inner-band) links are matched deterministically between different scales while “long” links exhibit predictable statistical variation. We argue that this conjecture explains the positive and negative results respectively of spatial scaling on *Sprinkler* and *LOF*.

Related work. *MinT* researchers recently considered [5] the relationship of behaviors at different scales and discussed a criterion for choosing values for hardware attenuation at both the transmitter and the receiver ends (based on matching the mean RSSI values at various distances). With such attenuation selection, they presented examples of behavior preservation in protocols executed at different scales (e.g. the flapping of a specific route upon a node failure or the measured bandwidth in a two-node connection).

In contrast to their work, we do not find that behaviors in a specific instance of a network can always be identically reproduced in any other instance of that network at a different spatial scale. This finding is in part due to our consideration of a larger set of nodes (while they use 2 nodes to 4 nodes for their experiments, we use 105 nodes). We can thus, for example, observe the non-trivial variation in the RSSI experiments on different instances of linear networks at a particular spatial scale, and find statistical support for our result on the probabilistic relationship between the behaviors at the different spatial scales.

The *Orbit* testbed uses increased noise levels to reduce signal strength as opposed to transmission power attenuation. Testbed nodes are used to inject noise. Given the finite number of grid nodes and the fact that the nodes’ path losses can assume only discrete values, there is significant limitation on the shrinking factor and the fidelity of shrinking. For example, for a 20 by 20 network, *Orbit* can faithfully shrink a range of 57dB of transmission power with an average mapping error of 2dB, which could have significant impact on the fidelity of performance at the higher network layers.

Organization of the paper. In Section II, we

theoretically derive the necessary conditions for the scaling to work. In Section III, we describe the network topology and wireless configuration used in our experiments. We identify the attenuation factor for shrinking in Section IV. We observe the performance at the link layer in Section V. We introduce messaging layer services Sprinkler and LOF in Section VI, and also describe the performance properties, provide a metric for performance similarity, and compare the performance similarity for the two services. We summarize our observations in Section VII. Finally, we mention future work in Section VIII.

II. NETWORK SCALING

The performance of a wireless communication system is largely determined by the radio channel characteristics between transmitter receiver pairs in the network. Unlike wired channels, wireless channels exhibit a high degree of variation and unpredictability. Electromagnetic wave propagation is influenced by the obstructions in the path between the terminals, operating frequency, mobility patterns of the wireless terminals and other environmental factors. The complex nature of the wireless channel has led communication system designers to derive statistical models from *in situ* empirical measurements. The empirical approach to channel modeling fits an analytical curve to data points to predict average as well as the expected variation in received signal strength (RSS) at a particular location.

A common model of large scale fading is log-normal shadowing model, which is applicable to both indoor and outdoor channels, where the average received signal power decreases logarithmically with distance:

$$R(d) = R(d_0) - 10n \log(d/d_0) + N_\sigma \quad (1)$$

where n is the path loss exponent $R(d)$ is the average received signal strength at distance of d measured in dB and N_σ is a zero-mean normally distributed random variable with standard deviation σ . The random variable N_σ captures the variation of clutter between different transmitter receiver pairs with the same interdistance d . For free space the

path loss exponent $n = 2$. In practice, received signal power data collected at various locations is fitted a linear model as a function of log-distance to determine the path loss exponent n and standard deviation σ . The empirically derived values for n ranges between 1.5 and 3.5 and σ varies between 3 and 9 dBs in indoor environments such as grocery store, office and factory buildings [2]. The log-normal shadowing model implies the following simple theorem on scaling network behavior.

Theorem 1: Given a wireless network \mathcal{W} with the set of internode distances $\{d_i\}_{i=1}^m$ and its scaled version $\tilde{\mathcal{W}}$ with the set of internode distances $\{\tilde{d}_i = \alpha d_i\}_{i=1}^m$. Assuming log-normal shadowing model for large scale fading, there exists a constant attenuation factor for all transmitters, such that the set of link properties for \mathcal{W} and $\tilde{\mathcal{W}}$ are sampled from the same multivariate gaussian probability distribution whose mean vector is given by $\{R(d_i)\}$ calculated from Equation 1 and covariance matrix σI .

Proof: The received signal strength set for the set of links $\{\tilde{d}_i\}_{i=1}^m$ in $\tilde{\mathcal{W}}$ is given by:

$$\begin{aligned} R(\tilde{d}_i) &= R(d_0) - 10n \log(\alpha d_i/d_0) + N_\sigma \\ &= R(d_i) - 10n \log(\alpha) \end{aligned}$$

In other words, all transmitter powers have to be changed by the same factor $10n \log(\alpha)$, which is negative if $\tilde{\mathcal{W}}$ is a scaled down version of \mathcal{W} (i.e. $\alpha < 1$). ■

The log-normal shadowing model gives a sufficient condition for fidelity scaling with constant power attenuation. We make the following observations about the necessary conditions:

- 1) The linear relationship between mean power and log-distance is a necessary condition for a constant attenuation across all transmitter receiver pairs. For all other functions of distance the attenuation constant has to be varied based on the particular receiver transmitter pair. Although in principle one can attenuate each packet using a separate attenuation factor based on its intended receiver. This will lead to a different interference pattern since the same packet will interfere with the reception of multiple receivers.
- 2) The second necessary condition is that the

distribution of signal strength values around the mean value has to be independent of distance. For a Gaussian distribution this condition is equivalent that the standard deviation σ has to be a constant function of distance. If σ varies based on distance, then the scaled network will provide the same average strength but a different statistical variation around these averages.


- 3) The two networks have to be in the same environment with same coefficients of n and σ for the scaling to be correct in this probabilistic sense. Different environments will again necessitate variable attenuation across the different links.

The equivalence between \mathcal{W} and its scaled version $\tilde{\mathcal{W}}$ is a probabilistic one. Particular realizations of each network will differ in their link realizations. However, repeated experimentation with different spatial configurations will result statistically equivalent set of outcomes. Repeated experimentation is not always possible and the network protocol designers are usually restricted to use a particular realization of the scaled network. Therefore the effect of link variations in the different network realizations on protocol performance is a key issue in protocol design and validation. In general, variation in path-losses between the different network realizations leads to variations in signal to noise ratio (SNR) in the received signals. The specific relationship between SNR and packet error rate depends on the modulation scheme, but it follows in general a sigmoid shape saturating at high and low SNR regions to 100% and 0% rates respectively. Therefore for links with high and low received signal strength the variations in path losses between different network realizations have minimal impact on the realized packet error rate. This leads us the following conjecture on protocol behavior, which we will study empirically in the next sections.

Conjecture 2: Consider a wireless network \mathcal{W} and its scaled version $\tilde{\mathcal{W}}$ with power control. A network protocol designed to use exclusively the inner-band links with high signal-to-noise ratio will have identical performance on \mathcal{W} and its scaled version $\tilde{\mathcal{W}}$.

III. EXPERIMENTAL SETUP

Description of Kansei testbed. A stargate is a single board linux-based computer [7]. It uses a Intel's 400 MHz X-Scale[®] processor (PXA255). It has 64 MB SDRAM, 32 MB FLASH, and a type II PCMCIA slot. A stargate is equipped with a SMC2532W-B high power IEEE 802.11b card, which is connected to a 3dBi antenna of length 1.82m via a fixed attenuator of -20 dB as shown in figure 1(a). We have 210 stargates in a 15×14 grid, with an internode separation of 3 feet in X and Y axes, inside Kansei testbed [10].



(a) Anatomy of a wireless node (b) Arrangement of nodes in Kansei testbed

Fig. 1. Kansei testbed

The nodes are raised from the ground at 4 feet. Each stargate is connected to a PC via wired ethernet. All the control and data traffic is communicated via wired ethernet to avoid interference with experimental traffic as seen in Figure 1(b).

Network architecture. The network is configured in IEEE 802.11b ad hoc mode. The frequency is set to 2.462 GHz (channel number 11). The environment contains another IEEE 802.11 network in an infrastructure mode, which is operating at 2.437 GHz (channel number 6). According to IEEE 802.11 standard, frequencies 2.462 GHz and 2.437 GHz (channel separation of 5) are non-interfering, which was corroborated experimentally by Robinson et. al. [11] The sensitivity threshold of the card is set to maximum. In other words, the signal at the receiver is not attenuated via software control.

For our experiments, we use a topology of 7×3 nodes at two different inter-node separations, viz. $D_1=3$ feet and $D_2=6$ feet. For each distance, we chose attenuation factors that results in similar performances.

IV. PHYSICAL LAYER

With the aim of statistical characterization of the relationship between received signal strength and hop distance, we have collected received signal strength on 25 lines of different spatial orientations. In each row a transmitter sends 250 packets to each of the 14 receivers at hop lengths $k \times 3$ ft. The received signal strength measurements as reported by the prism2 chipset of the 802.11 card is collected for each transmitter receiver pair. We have determined the mapping of the RSS values reported by the NIC card to dBm experimentally with a precise variable attenuator directly connecting the two cards.

The resulting data set is given in Figure 2. Linear regression results in a log-normal model fit with $n = 1.59$ and $\sigma = 4.22$ dB. The Lilliefors test (Kolmogorov-Smirnov test with adjusted p-values) of normality supports Gaussian distributed variations hypothesis around the log-normal fit at 5% significance level. The log-normal model fit and $\pm\sigma$ variation is also given in Figure 2. With the possible exception of the first hop the log-normal provides a good fit. We should note that in *Kansei* testbed only the first hop provides a line-of-sight to the transmitter which can result in higher signal strength for the first hop. The variation around the mean is uniform and does not depend on the distance, consistent with the assumptions of the scaling result given in Theorem 1. In particular, path loss exponent $n = 1.59$ indicates that a constant 4.8 dB attenuation is required for all transmitters for a scaling factor of 2.






Fig. 2. The relationship between received signal strength and hop distance

V. LINK LAYER

Figures 3(a) and 3(b) show the boxplot of link



(a) At power level 0 dBm



(b) At power level 3.95 dBm

Fig. 3. Link reliability

reliability at different transmitter-receiver distances when the transmission power is 0 and 3.95 dBm respectively. The selections of the transmission power levels are motivated to ensure a multihop network and by the attenuation factor of about 4.98dB measured in section II. We see that, at each transmission power level, there is a threshold distance (i.e., 18 and 9 feet respectively) such that (a) link reliability is relatively high (e.g., greater than 85%) and stable when the transmitter-receiver distance is within the threshold value, and (b) link reliability is relatively low and unstable when the transmitter-receiver distance is greater than the threshold value. For convenience, we divide communication range into three regions: *inner band* where the transmitter-receiver distance is less than the threshold value, *middle band* where the distance is greater than the threshold value and link reliability is greater than 0, and *outer band* where link reliability is 0.

From Figures 3(a) and 3(b), we can see that, if we use transmission power 0 dBm for D1 and transmission power 3.95 dBm for D2, link reliability will be such that two nodes no more than 3 grid-hops apart are within inner band of each other, but nodes more than 3 grid-hops apart are in the middle or outer band of each other. More specifically, Table I shows the median link reliability at different grid-distances. The reason we use median instead of

grid-hops	1	2	3	4	5	6
D1	93.80	86.00	86.80	79.62	75.80	36.40
D2	85.60	83.00	86.50	36.40	51.80	76.50

TABLE I
MEDIAN LINK RELIABILITY (%) AT DIFFERENT
GRID DISTANCES

mean is because the distribution of link reliability at a transmitter-receiver distance is not symmetric. Note that while 1 grid-hop for D1 is 3 feet, it is 6 feet for D2. In Table I, we see that the median link reliability within inner band at D1 and D2 are more similar as compared to those in the middle band. Also, from Figures 3(a) and 3(b), we see that the variation in link reliability is lesser within inner band. This is in accordance with Conjecture 2. The median link reliability values are not monotonically decreasing function of hop count is due to the heterogeneity in radio hardware, node placement, environment, etc.

VI. MESSAGING LAYER

A. *Sprinkler*

In this subsection, we first describe *Sprinkler*, the reliable data dissemination protocol *Sprinkler*, then we discuss the experiment design and experimental results.

Sprinkler. *Sprinkler* [9] is a reliable data dissemination service for wireless embedded devices which are constrained in energy. *Sprinkler* uses a near optimal connected dominating set (CDS) of the devices to avoid redundant transmissions, and a transmission schedule to avoid collisions. Informally, a CDS is subset of a set of nodes, such that the nodes in the CDS form a single connected network and every node in the network has at least one neighbor in the CDS.

Experiment design. In the 7×3 grid, we configure the node at location [0,1] as the source for data dissemination. The table I, shows that inner bands for D1 and D2 end at 9ft and 18 ft respectively. Therefore, in the 7×3 grid, we configure the node at location [3,1] as the second CDS node. The node at [3,1] is 9 ft from the source for D1 and 18

ft for the source in D2. We broadcasts a payload composed 101 packets for D1 and D2. We repeat each broadcast session 20 times.

In both D1 and D2, we exam the following properties of *Sprinkler*:

- *Number of transmissions*: the total number of packet transmissions in the network for a broadcast session.
- *Latency*: the total time taken from the start of a session at the source until the entire payload is received at all the nodes for a broadcast session.

The number of packet transmissions measure the energy efficiency of *Sprinkler*. The latency captures the real time quality of *Sprinkler*. Both of the above mentioned properties reflect the link reliability.

Experimental results. Figures 4 and 5 show the boxplots of number of transmissions and latency respectively in D1 and D2. Table II show that the median values for the performance properties at D1 and D2. It also shows the ratios for each of the performance properties at D1 and D2. The ratios quantify the similarity in performance. In Subsection VI-C, we will analyze the data.




Fig. 4. Number of transmissions

Metrics	D1	D2	Ratio of D1 and D2
# transmissions	204	206.5	1.01
Latency (ms)	10579.09	11002.53	1.04

TABLE II
MEDIANS OF NUMBER OF PACKET
TRANSMISSIONS AND LATENCY




Fig. 5. Latency

B. LOF

In this subsection, we first describe the routing protocol LOF, then we discuss the experiment design and experimental results.

Learn on the Fly (LOF). LOF [13] addresses the challenge of high-fidelity link property (such as reliability) estimation and routing in wireless (sensor) networks. In stead of estimating link properties via broadcast beacon exchange between neighbors, LOF estimates link properties based on unicast data traffic itself. Therefore, LOF is able to precisely estimate link properties according to network traffic patterns, and LOF chooses routes which incur shorter latency and consume less energy than those chosen by beacon-based protocols such as ETX [4] and ETT [6].

Experiment design. To study the properties of LOF in D1 and D2, we use the traffic trace extracted from ExScal [13]. In each case of D1 and D2, we let the node at [0,0] act as the source node, and we let the node at [6,2] be the destination. For each case, we run the same experiment for 50 times (and about 1,000 packets are generated by the source).

In both D1 and D2, we examine the following properties of LOF:

- *Per-hop link length*: the geographic length of each link used in LOF. It illustrates the structural property of LOF, via which we can examine other properties (such as reliability) of the links used by LOF.
- *Per-hop MAC latency*: the time taken for the MAC to transmit a packet across each hop chosen by LOF. MAC latency reflects both link reliability and energy efficiency [13].

- *End-to-end MAC latency*: the amount of time that a packet spends in MAC along the route from source to destination. It affects the network throughput [6], [13].

Experimental results. Figure 6 shows the boxplot




Fig. 6. Per-hop link length

of per-hop link length in D1 and D2. Interestingly, the median in D1 and D2 are 6 feet and 12 feet respectively. Therefore, each link spans 2 grid hops on average in both D1 and D2, and the routing structure has similar link length in D1 and D2.

Figures 7 and 8 show the per-hop MAC latency




Fig. 7. Per-hop MAC latency




Fig. 8. End-to-end MAC latency

and end-to-end MAC latency respectively. Table

III shows the median values for the performance properties at D1 and D2. It also shows the ratios for each of the performance properties at D1 and D2. The ratios quantify the similarity in performance. In Subsection VI-C, we will analyze the data.

Metrics	D1	D2	Ratio of D1 and D2
Per-hop latency (ms)	3.45	2.96	0.85
End-to-end latency(ms)	21.37	19.38	0.91

TABLE III
MEDIANS OF PER-HOP AND END-TO-END MAC
LATENCY

C. Analysis of similarity in performance

As shown in tables II and III, the performance ratios for Sprinkler are closer to 1 as compared to that for LOF. In other words, performance of Sprinkler is more similar after shrinking as compared to that of LOF. The reason behind the difference in performance is attributed to the fact that while Sprinkler relies on inner band links, LOF relies on middle band link in addition to the inner band links. And, there is greater variability/instability in the middle band of D1. This leads further support to Conjecture 2.

VII. CONCLUSION

In this paper, we presented a necessary and sufficient condition for scaling up/down a set of links. Our experimental studies showed that this log-normal large scale fading condition essentially holds in the *Kansei* testbed. The relationship between a link set and its scaled version is, however, a probabilistic one. So, even with careful selection of transmission power attenuation for downscaling, the experimenter is not assured that for any given protocol its behaviors on the link set before and after scaling will be identical. (Indeed our experiments on the *LOF* protocol yielded a lack of strong similarity.) The experimenter can be assured however that repeated experimentation of the protocol on different instances of the network at each spatial scale will yield statistically equivalent sets of behaviors; this sort of experimentation is however not particularly convenient.

We also argued that for a link set with only high and low received signal strength links, the variations in path losses on different instances of the network have minimal impact. For such link sets, we claim that the experimenter is assured that for any given protocol its behaviors can be reproduced on any instance of a spatially scaled version of the network. (Our experiments on the *Sprinkler* protocol illustrated this reproducibility.)

There are several problems for future study: Given a particular link set at a large spatial scale, (i) how to design a smaller scale testbed that has a similar set of link properties; (ii) is it feasible to find a “replica” of that link set from among the nodes in a spatially compact wireless testbed?; and (iii) how to change the power level in reaction to temporal environmental variations that affect both spatial scales similarly.

VIII. FUTURE WORK

In Section II, we attributed variation in similarity in performance to the different network realizations. The variation is more when middle band links are used. In this work, we have not used multiple network realizations for the link and messaging layers. We plan to study the performance of link and messaging layers at multiple network realizations, especially for the middle band links.

In addition to different network realizations, interference from the radios outside of Kansei may also be a contributor in performance variation [12]. Although, we try to minimize such interference by selecting the frequency with the least possible interfering frequency, we have not ruled out the interference completely. The iWWT testbed [12] controls the radio environment by shielding the other radio sources, which is orthogonal to the problem of different network realizations. It would be interesting to study impact of zero interference.

In this paper, we don’t use a consistent procedure for verifying the similarity in performance, e.g. we use Lilliefors test at the physical layer and an informal test of ratio of median values at the messaging layer. In the future, we will use a standard test and definition for similarity in performance.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for the suggestions.

REFERENCES

- [1] <http://www.cse.ohio-state.edu/exscal>.
- [2] J. Anderson, T. Rappaport, and S. Yoshida. Propagation measurements and models for wireless communications channels. *IEEE Communications Magazine*, Nov 1994.
- [3] A. Arora, R. Ramnath, P. Sinha, E. Ertin, S. Bapat, V. Naik, V. Kulathumani, H. Zhang, M. Sidharan, S. Kumar, H. Cao, N. Seddon, C. Anderson, T. Herman, C. Zhang, N. Trivedi, M. Gouda, Y. Choi, M. Nesterenko, R. Shah, S. Kulkarni, M. Arumugam, L. Wang, D. Culler, P. Dutta, C. Sharp, G. Tolle, M. Grimmer, B. Ferriera, and K. Parker. Exscal: Elements of an extrem scale wireless sensor network. In *11th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*, 2005.
- [4] D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *The 9th Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2003.
- [5] P. De, A. Raniwala, S. Sharma, and T. Chiueh. Mint: A miniaturized network testbed for mobile wireless research. In *The 24th IEEE International Conference on Computer Communication (INFOCOM)*, 2005.
- [6] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *The 10th Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2004.
- [7] Intel, <http://www.xbow.com/Products/XScale.htm>. *Stargate Specifications*.
- [8] J. Lei, R. Yates, L. Greenstein, and H. Liu. Wireless link snr mapping onto an indoor testbed. In *2nd International IEEE/Create-Net Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (Tridentcom)*, 2005.
- [9] V. Naik, A. Arora, P. Sinha, and H. Zhang. Sprinkler: A reliable and energy efficient data dissemination service for wireless embedded devices. In *The 26th IEEE Real-Time Systems Symposium*, December 2005.
- [10] V. Naik, S. Bapat, H. Zhang, C. Anderson, G. Fox, J. Wieseman, A. Arora, E. Ertin, and R. Ramnath. Kansei: Sensor testbed for at-scale experiments. Technical report, The Ohio State University, Computer Science and Engineering Department, <http://www.cast.cse.ohio-state.edu/exscal/content/CheckpointPresentations/ExscalOpenHouse/TestbedPoster-2005-01-31.pdf>, February 2005.
- [11] J. Robinson, K. Papagiannaki, X. Guo, C. Diot, and L. Krishnamurthy. Experimenting with a multi-radio mesh networking testbed. In *1st Workshop on Wireless Network Measurements (WiNMee)*, 2005.
- [12] N. Vaidya, J. Bernhard, V. Veeravalli, P. Kumar, and R. Iyer. Illinois wireless wind tunnel: a testbed for experimental evaluation of wireless networks. In *E-WIND '05: Proceeding of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, pages 64–69, New York, NY, USA, 2005. ACM Press.
- [13] H. Zhang, A. Arora, and P. Sinha. Learn on the fly: Data-driven link estimation and routing in sensor network backbones. In *The 25th IEEE International Conference on Computer Communication (INFOCOM)*, 2006.