Poster Abstract: Burstiness and Scaling in the Structure of Low-Power Wireless Links

Tal Rusak[a,b]
rusakt@stanford.edu

Philip Levis[b]
pal@cs.stanford.edu

[a]Computer Science Department, Cornell University, Ithaca, New York, United States
[b]Computer Systems Laboratory, Stanford University, Stanford, California, United States

We observe that low-power wireless links have non-trivial time-scaling characteristics at both the physical- and link-layers. Packet reception rate (PRR) analysis shows that links are bursty rather than constant, i.e., their reception quality varies greatly from the overall packet reception rate at different times. Furthermore, this variation is seen at many time-scales. We provide a possible explanation for burstiness using wavelet analysis of RSSI traces from a variety of wireless links. We show that these traces can be considered as consistent with statistical self-similarity but not with long range dependence. Using the variance in RSSI, we suggest a way to easily characterize when scaling occurs. Finally, while current simulators do not capture scaling, we propose and validate a possible modeling technique for network links that conforms to scaling phenomena.

I. Introduction

Low-power wireless networks are becoming increasingly pervasive in a society that expects information anytime, anywhere. These networks include consumer- and research-grade wireless networks, such as IEEE 802.15.4 sensor networks, Zigbee networks, Bluetooth, and IEEE 802.11 Wi-Fi networks.

Various studies have attempted to model, explain, and build protocols optimized for the physical- or link-layers of the wireless channel. These efforts focus on modeling low-level details of physical- and link-layers, such as interference and scheduling [11] or adaptive power control [10]. Some analytical models [7] of wireless networks, especially the well-studied cellular phone network, make simplifying assumptions such as time independence of packet reception. While many studies are complementary [8, 10, 13, 14], some are also contradictory, coming to different conclusions even while examining the same data traces [2, 6]. The purpose of this paper is not to support or reject any individual conclusion of the aforementioned work, but to suggest a fundamentally different way of understanding and modeling the wireless channel. We believe that a parsimonious model may be possible for wireless links.

The main observation of this paper is that the signal power time series of many wireless links is consistent with statistical self-similarity. We consider the generality of this property and possible modeling techniques that take into account the observed temporal variations.

II. Background and Related Work

II.A. Burstiness in Wireless Networks

Previous studies have observed burstiness in wireless links following the IEEE 802.11 and 802.15.4 standards. Aguayo et al. [2] observed bursty losses in some links of an 802.11b mesh network deployed in an urban environment and analyzed the Allan deviation of burst lengths in order to discover characteristic lengths. Likewise, Srinivasan et al. [14] studied various 802.15.4 and 802.11 links and defined a metric, β, which quantifies the burstiness at the level of individual packets—the shortest possible time-scale. In focusing on one time-scale, or by attempting to discover a characteristic time-scale for bursts, such work does not capture the full extent of burstiness in many low-power wireless links. We show that in many such links, there is no single time-scale for bursts of packet receptions or losses. Previous work [13, 14] has shown that bursts occur at shorter time-scales, possibly due to external interference and noise sources. We observe bursts at longer time-scales as well, characterized by a coherent scaling property that we discuss in this work.

II.B. Methodology and Datasets

In order to characterize the physical-layer of low-power wireless networks, we begin with experimental data about the behavior of these links. The experiments are conducted with TelosB [12] or MICAz motes running TinyOS 2.0.2. Networks used for the experiments have a mote that sends packets at a constant rate of 100 Hz. From each experiment with n

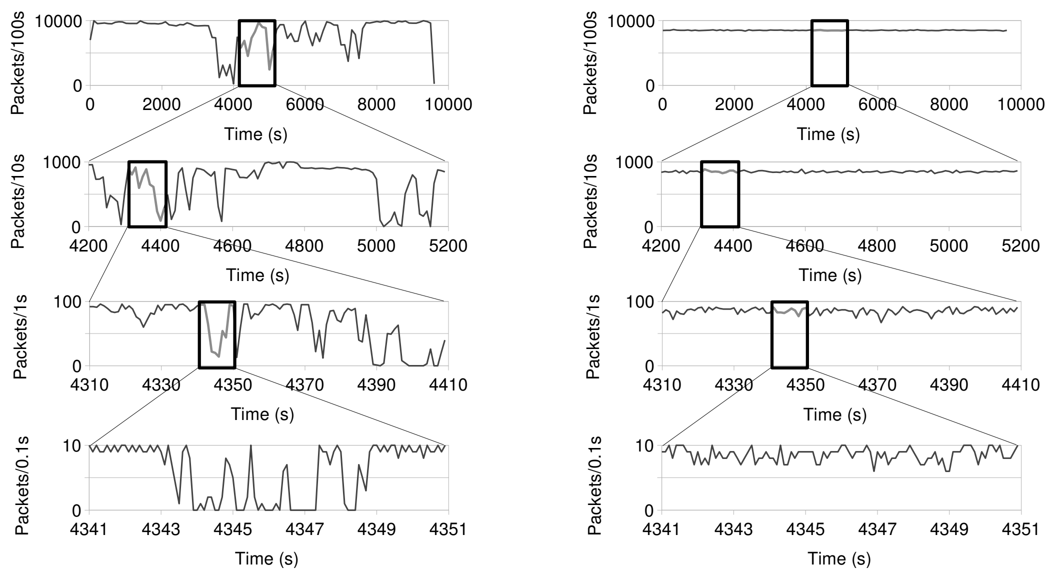

(a) The number of packets received over different time-scales in a real 802.15.4 experimental link. Periods of good and bad receptions are observed over a wide range of time-scales, suggesting pervasive burstiness.

(b) The number of packets received over different time-scales in a simulation using the TOSSIM simulator 2.0.2 with the same overall PRR. Even at short time-scales, packet reception exhibits limited burstiness.

Figure 1: Each diagram is a zoomed-in version of the boxed, highlighted section of the preceding packet reception trace. The overall packet reception rate of both the experiment and the simulated trace is about 84%.

successfully received and decoded packets, we derive a time series $\mathcal{R} = r_1, r_2, \cdots, r_n$, where r_i is the RSSI value in dBm for the i-th received packet ($1 \leq i \leq n$). We use $\mathcal{R}$ as a representation of the physical-layer of the network, since signal power is being measured. For each experiment with k transmitted packets, we derive a time series $\mathcal{P} = p_1, p_2, \cdots, p_k$ where p_j is 1 if the j-th packet ($1 \leq j \leq k$) transmitted is received and correctly decoded and 0 otherwise. Since $\mathcal{P}$ measures the reception characteristics of packets, we use it to characterize the link-layer of the network.

Such datasets were collected over periods of six hours to two weeks in IEEE 802.15.4 testbeds at Gates Hall at Stanford University, at an apartment, and at Intel Research Berkeley (Mirage) [4].

II.C. Scaling and Self-Similarity in Time Series

The presence of self-similarity or other coherent structure over many time-scales is called *scaling*, and can be detected using the wavelet transform and a tool called the *logscale diagram* [1]. The logscale diagram plots octaves (base-2 scales, i.e., octave 1 is 2^1 packets, octave 2 is 2^2 packets, etc.) on the horizontal axis and y_j on the vertical axis. y_j can be understood as the logarithm of an estimator for the variance of the discrete wavelet transform process $d_X(j, \cdot)$ defined by Abry et al. [1].

The slope of a linear portion of the logscale diagram, α, can be used to check for various types of scaling. If $\alpha > 1$, where the linear regime includes the largest octaves available, then the underlying time series can be considered as consistent with the self-similarity hypothesis, but not with long range dependence. The octave at which the linear regime starts is known as the *onset point* of scaling.

A rich body of previous work suggests ways to generate self-similar time series analytically. We evaluate the use of two such methods, fractional Brownian motion [5] and Infinitely Divisible Cascades [3], for their effectiveness in modeling the physical-layer of low-power wireless links.

III. Observation and Validation of Scaling Phenomena

The observation of scaling comes from considering Figure 1. Figure 1(a) plots the reception characteristics of a real network link over non-overlapping time windows of varying sizes, and the variations are heuristically "similar" over different time-scales, suggesting self-similarity and scaling. Figure 1(b) plots the equivalent link characteristic for the state-of-the-art TOSSIM simulator for low-power wireless sensor networks. This figure is simply a way of representing $\mathcal{P}$ for the Gates Hall experiment. Current simulation

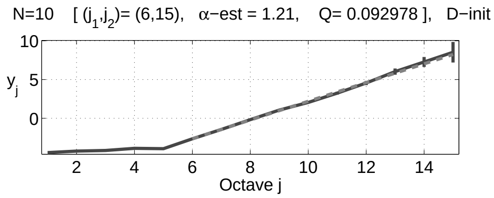

Figure 2: Logscale diagram for $\mathcal{P}$ for actual experimental data for the Gates Hall link. The solid line shows the actual value of y_j, while the dotted line is a linear fit over a region including the highest scales.

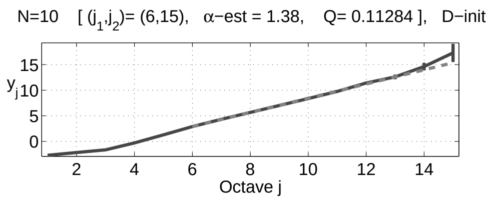

Figure 3: Logscale diagram for $\mathcal{R}$ for experimental data for the Gates Hall link. Same conventions as in Figure 2.

models, which assume signal power to be constant, do not capture burstiness and scaling correctly.

To quantify and make more rigorous the notion of scaling in the trace studied in Figure 1(a), we plot its logscale diagram for $\mathcal{P}$ (the link-layer reception trace) in Figure 2. As expected for a time series consistent with self-similarity, there is a long linear regime in the asymptotic domain of this plot, i.e., including the largest octaves. Here, we find $\alpha = 1.21 > 1$, which is consistent with the self-similarity hypothesis.

A natural question to ask is why the link-layer trace $\mathcal{P}$ is consistent with self-similarity. We can provide insight by considering the scaling behavior of the physical-layer trace $\mathcal{R}$, which approximates signal power. It might be expected that scaling and

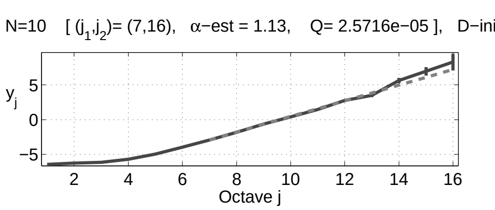

Figure 4: Logscale diagram for $\mathcal{P}$ from a physically-based simulation of the Gates Hall link using the real RSSI trace for signal power ($\mathcal{S}$) and the closest-fit pattern matching algorithm for noise. Same conventions as in Figure 2.

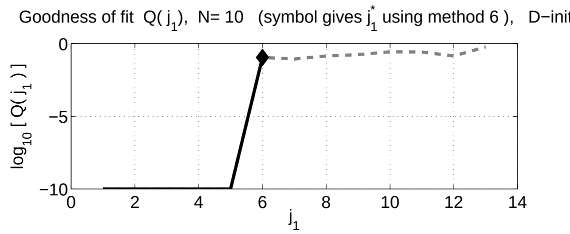

Figure 5: Estimation of the point of the onset of scaling in the logscale diagram in Figure 2. The solid line shows regions of non-decreasing Q and the dotted line shows the Q values following the non-decreasing regime. The estimate of the onset scale is given by the diamond in the figure.

self-similarity at the link-layer suggests similar behavior of the underlying physical-layer of the channel. The logscale diagram in Figure 3 shows that the physical-layer is also consistent with self-similarity, since $\alpha = 1.38 > 1$, computed including the largest time-scales available in the data.

In addition, we plot the logscale diagram of $\mathcal{P}$ for a physically-based simulation [13] of the Gates Hall link in Figure 4. For this simulation, we used the actual RSSI trace for a signal power time series, which we denote by $\mathcal{S}$, and modeled noise using the closest-fit pattern matching algorithm [9] based on an experimental trace. We see that $\alpha = 1.13 > 1$, indicating that the data is consistent with self-similar scaling. This observation provides validation to the simulator used below to evaluate models for signal power that take into account the scaling behavior.

IV. Onset Point and Implications for Protocols

We noted above that in each of the logscale diagrams, scaling starts at some onset point. The onset point is estimated using an algorithm of Veitch et al. [15]. Figure 5 shows how the onset point is estimated for the logscale diagram in Figure 2; the onset octave shows a large increase in Q, a goodness of fit metric defined by Veitch et al. [15]. Experimental work [14] has shown that waiting 500 ms before retransmitting lost packets greatly increases link reliability in terms of packet reception rate. The onset scale observed in this experiment, $j_1 = 6$, corresponds to $10 \text{ ms} \times 2^6 = 640 \text{ ms}$, a relatively close value to the previous experimental observation [14]. The onset point may correspond to the octave at which random variations stop impacting the wireless link and the self-similar scaling starts to dominate the variations of the link. Protocols that understand this parameter can use it to optimize times to

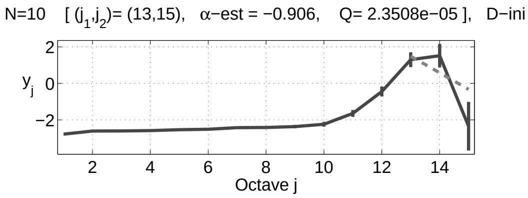

Figure 6: Logscale diagram for $\mathcal{P}$ for simulation with constant signal power (i.e., $\mathcal{S} = c$). This logscale diagram shows no apparent scaling. Same conventions as in Figure 2.

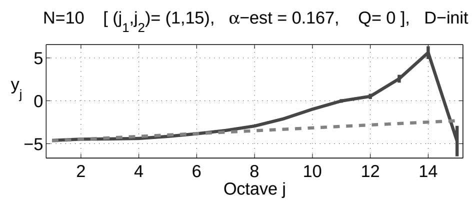

Figure 7: Logscale diagram for $\mathcal{P}$ for simulation with $\mathcal{S} = B_{H=0.192}(t)$ (fractional Brownian motion). Despite the self-similarity of $\mathcal{S}$, self-similarity is not observed at the link-layer of the simulation. Same conventions as in Figure 2.

transmit packets so that the likelihood of packet loss is lessened.

V. Modeling Self-Similar Links

State-of-the-art simulators of wireless networks assume signal power to be constant for a given link. To evaluate the effectiveness of this technique, we consider scaling at the link-layer of the simulation; a logscale diagram is given in Figure 6. We see that the link-layer simulated by the constant signal power assumption does not conform to self-similarity, since $\alpha < 1$; thus, the commonly applied model does not correctly capture link behavior.

Next, we use two self-similar time series synthesis techniques, fractional Brownian motion and Infinitely Divisible Cascading random walks, to generate self-similar traces of signal power synthetically. We set only one parameter, $H = \frac{\alpha - 1}{2}$, in each of these models (computed from α of the logscale diagram of the physical-layer trace $\mathcal{R}$), and modify the range and minimum and maximum points to correspond to the experimental range and extrema. We confirmed that each simulated signal power trace $\mathcal{S}$ is self-similar.

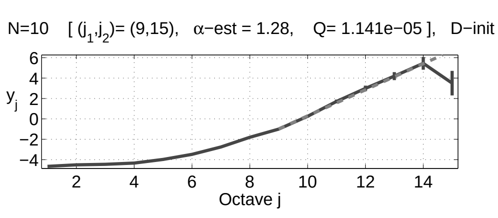

Figure 8: Logscale diagram for $\mathcal{P}$ for simulation with $\mathcal{S} = V_{H=0.192}(t)$ (Infinitely Divisible Cascading random walk). Same conventions as in Figure 2.

Figure 7 shows the result of modeling $\mathcal{S}$ with fractional Brownian motion $B_{H=0.192}(t)$. We conclude from this figure that despite its self-similarity, fractional Brownian motion is not an appropriate model for the signal power variations of the wireless links being studied, since scaling is not observed at the link-layer. Figure 8, in contrast, shows that modeling $\mathcal{S}$ with an Infinitely Divisible Cascading random walk $V_{H=0.192}(t)$ leads to a packet reception trace $\mathcal{P}$ for which $\alpha = 1.28 > 1$ using a fit including the largest time-scales. Because it captures scaling at the link-layer, we have found that among the models that we have investigated, Infinitely Divisible Cascading random walks are the best model for synthesizing signal power traces for the Gates Hall wireless link.

We are interested in pursuing future work accounting for the variations of y_j at each time-scale, instead of relying only on the statistical fit used to calculate α. It is especially important to understand whether the drop in y_j at the largest time-scale is due to the limited data at this scale or some more fundamental variation in the model. This can be done by generating longer Infinitely Divisible Cascading random walks and simulating the corresponding link.

VI. Generality of the Scaling Phenomenon

Although many links are consistent with self-similarity, we observe that not all links exhibit this property. Figure 9 plots the probability, over groups of ten links, that aL, the lower 95% confidence bound for α, is greater than 1 versus the logarithm of the variance in the RSSI traces $\mathcal{R}$ of these links. We observe a phase transition in the consistency of these links with self-similarity. Stable, low-variance links are never self-similar, while high-variance links are almost always self-similar at the physical-layer. Any realistic model of wireless networks has to take this critical point into account, and high-variance links should be simulated with a self-similar model.

VII. Conclusions

By considering the physical-layer structure of wireless links, we identified the self-similar nature at both

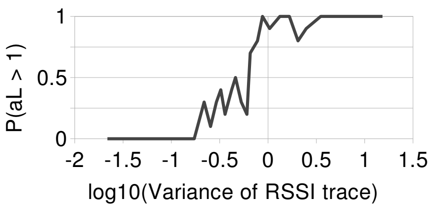

Figure 9: This plot considers groups of 10 Mirage links each, sorted by order of increasing variance, and plots the probability that $aL > 1$ versus the base-10 logarithm of the average variance.

the link- and physical-layers of such networks in many experiments. This concise explanation for the structure of IEEE 802.15.4 networks provides a possible reason for the difficulty in discovering common patterns in studies that have delved deeply into the details of particular networks. The simplicity of this explanation, as evidenced by the fact that we need to set only one parameter and the extrema in the model traces, is a strength of the observations of this paper.

We reviewed the phenomenon of self-similarity and showed how it applies to the physical- and link-layers of many IEEE 802.15.4 links. We suggested an approach to modeling links that exhibit this property, and we showed a phase transition in the presence of self-similarity at the physical-layer as compared to the elementary variance of the link's RSSI trace. We believe that this general, parsimonious property that applied to many links in our experiments offers a new perspective from which to study wireless networks and to design effective, reliable protocols.

Acknowledgments

This work was supported by generous gifts from the Bartels Family, Microsoft Research, Cisco Research, Intel Research, DoCoMo Capital, Foundation Capital, the National Science Foundation under grants #0615308 ("CSR-EHS") and #0627126 ("NeTS-NOSS"), and a Stanford Terman Fellowship. We appreciate insightful discussions with Jon Kleinberg, Lars Backstrom, and Kannan Srinivasan.

References

[1] P. Abry, P. Flandrin, M. Taqqu, and D. Veitch. Wavelets for the analysis, estimation and synthesis of scaling data. *Self-Similar Network Traffic and Performance Evaluation*, pages 39–88, 2000.

[2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Proceedings of SIGCOMM'04*, pages 121–132.

[3] P. Chainais, R. Riedi, and P. Abry. Scale invariant infinitely divisible cascades. In *Proceedings of PSIP'03*.

[4] B. Chun, P. Buonadonna, A. AuYoung, C. Ng, D. Parkes, J. Shneidman, A. Snoeren, and A. Vahdat. Mirage: a microeconomic resource allocation system for sensornet testbeds. In *Proceedings of EmNets'05*, pages 19–28.

[5] P. Flandrin. Wavelet analysis and synthesis of fractional Brownian motion. *IEEE Transactions on Information Theory*, 38(2):910–917, Mar. 1992.

[6] D. Gokhale, S. Sen, K. Chebrolu, and B. Raman. On the feasibility of the link abstraction in (rural) mesh networks. In *Proceedings of INFOCOM'08*, pages 61–65.

[7] A. Goldsmith. *Wireless Communications*. Cambridge University Press, 2005.

[8] J. Jeong, D. Culler, and J.-H. Oh. Empirical analysis of transmission power control algorithms for wireless sensor networks. In *Proceedings of INSS'07*.

[9] H. Lee, A. Cerpa, and P. Levis. Improving wireless simulation through noise modeling. In *Proceedings of IPSN'07*, pages 21–30.

[10] S. Lin, J. Zhang, G. Zhou, L. Gu, J. Stankovic, and T. He. ATPC: adaptive transmission power control for wireless sensor networks. In *Proceedings of SenSys'06*, pages 223–236.

[11] R. Maheshwari, S. Jain, and S. R. Das. A measurement study of interference modeling and scheduling in low-power wireless networks. In *Proceedings of SenSys'08*, pages 141–154.

[12] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. In *Proceedings of IPSN'05*, pages 364–369.

[13] T. Rusak and P. Levis. Investigating a physically-based signal power model for robust low power wireless link simulation. In *Proceedings of MSWiM'08*, pages 37–46.

[14] K. Srinivasan, M. Kazandjieva, S. Agarwal, and P. Levis. The beta-factor: Measuring wireless link burstiness. In *Proceedings of SenSys'08*, pages 29–42.

[15] D. Veitch, P. Abry, and M. Taqqu. On the Automatic Selection of the Onset of Scaling. *Fractals*, 11(4):377–390, 2003.