

Extending Wireless Algorithm Design to Arbitrary Environments via Metricity*

Helga Gudmundsdottir
School of Computer Science,
CRESS
Reykjavik University, Iceland
helgag10@ru.is

Joseph T. Foley
School of Science and
Engineering
Reykjavik University, Iceland
foley@ru.is

Eyjólfur I. Ásgeirsson
School of Science and
Engineering, ICE-TCS
Reykjavik University, Iceland
eyjo@ru.is

Magnús M. Halldórsson
School of Computer Science,
ICE-TCS
Reykjavik University, Iceland
mmh@ru.is

Marijke H. L. Bodlaender
School of Computer Science,
ICE-TCS
Reykjavik University, Iceland
marijke12@ru.is

Ymir Vigfusson
School of Computer Science,
ICE-TCS, CRESS
Reykjavik University, Iceland
ymir@ru.is

ABSTRACT

Efficient spectrum use in wireless sensor networks through spatial reuse requires effective models of packet reception at the physical layer in the presence of interference. Despite recent progress in analytic and simulations research into worst-case behavior from interference effects, these efforts generally assume geometric path loss and isotropic transmission, assumptions which have not been borne out in experiments.

Our paper aims to provide a methodology for grounding theoretical results into wireless interference in experimental reality. We develop a new framework for wireless algorithms in which distance-based path loss is replaced by an arbitrary gain matrix, typically obtained by measurements of received signal strength (RSS). We experimentally evaluate the framework in two indoors testbeds with 20 and 60 motes, and confirm superior predictive performance in packet reception rate for a gain matrix model over a geometric distance-based model.

At the heart of our approach is a new parameter ζ called *metricity* which indicates how close the gain matrix is to a distance metric, effectively measuring the complexity of the environment. A powerful theoretical feature of this parameter is that *all* known SINR scheduling algorithms that work in general metric spaces carry over to arbitrary gain matrices and achieve equivalent performance guarantees in terms of ζ as previously obtained in terms of the path loss constant. Our experiments confirm the sensitivity of ζ to the nature of the environment. Finally, we show analytically and empirically how multiple channels can be leveraged to improve metricity and thereby performance. We believe our contributions will facilitate experimental validation for recent advances in algorithms for physical wireless interference models.

*Supported by grant-of-excellence no. 120032011 from the Icelandic Research Fund.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MSWiM'14, September 21–26, 2014, Montreal, QC, Canada.
Copyright 2014 ACM 978-1-4503-3030-5/14/09 ...\$15.00.
<http://dx.doi.org/10.1145/2641798.2641811>.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Communication

Keywords

Wireless Ad-hoc Networks; Wireless Interference Models; SINR model; Wireless Algorithm Design; Experiments

1. INTRODUCTION

There is mounting demand for tomorrow's wireless networks to provide higher performance while lowering costs. A central challenge in meeting this demand is to improve the utilization of the wireless spectrum to enable simultaneous communications at the same radio frequency. To accommodate research into efficient use of wireless channels, for instance through spatial reuse, we require practical models of signal propagation behavior and reception at the physical layer in the presence of wireless interference.

Early models of worst-case wireless communication under interference were graph-based, most commonly based on distances. In comparison, physical models, or *SINR* (signal to interference and noise ratio) models, capture two important features of reality: signal strength decays as it travels (rather than being a binary property) and interference accumulates (rather than being a pairwise relation).

Analytic work on SINR – introduced by Gupta and Kumar [14] in an average-case setting and Moscibroda and Wattenhofer [26] in worst-case – has generally assumed *geometric path loss*, referred to here as the *GEO-SINR* model: signals decay as a fixed polynomial of the distance traveled.

While free space exhibits geometric decay, the reality for *real-world* wireless environments is more complex. When located above an empty plane, a signal bounces off the ground, resulting in complicated patterns of superpositions known as *multi-path fading*. Most real scenarios are more complex, with walls and obstructions. In particular, cityscape and indoor environments are notoriously hard to model. Moreover, the simple range-based models often make further assumptions into geometric path loss that do not concord with experiments, such as smooth and isotropic polynomial decrease in the signal strength. In fact, quoting recent meta-analysis [3], “link quality is not correlated with distance.”

Various stochastic extensions of geometric path loss have been proposed to address the observed variability in signal propagation. The most common are *log-normal shadowing* and *Rayleigh fading* for addressing long- and short-distance variability, respectively. Both modify the signal strength multiplicatively by an exponentially distributed random variable. These models are highly useful both for generating input for signal propagation simulations and for average-case analysis of wireless interference algorithms.

A complementary view to stochastic studies, with deep roots in computer science theory, is to allow for worst-case behavior and obtain guarantees that hold for *all* instances to the problem at hand. To avoid such results becoming too pessimistic, proper characterizations or parameterizations are often essential. Our goal is to contribute to such “any-case” analysis that avoids making assumptions about the environment that may not be reflected in actual real-world scenarios.

Our contributions. We propose moving theoretical algorithm design away from assuming geometric path loss models to an abstract SINR formulation with a matrix representing the fading (or signal decay) between pairs of nodes in an arbitrary environment. The matrix would typically be generated from direct measurements of *received signal strength* (RSS) provided by nodes, as proposed by experimentalists [34, 32, 24]. The RSS matrix could also be generated by other means, such as by inference, history, stochastic models or by accurate environmental models.

Following this approach, worst-case algorithmic analysis is heavily contingent on the contents of the RSS matrix, with unconstrained settings causing computational intractability. We introduce a new measure that reflects the attenuation complexity of the environment described by the RSS matrix. Dubbed *metricity* and denoted ζ , this parameter intuitively represents how close the RSS matrix is to a distance metric. From a theoretical standpoint, the definition of metricity has extensive implications: **All** SINR algorithms that work in arbitrary metric spaces work seamlessly in the abstract model, with performance ratio in terms of metricity that is equivalent to the original dependence on the path loss constant.

In an experimental evaluation on two testbeds of 20 and 60 nodes, our measurements indicate that the metricity parameter corresponds to the complexity of the environment. The experiments also suggest that the SINR model – without the geometric assumption – is of high fidelity, capturing signal propagation and reception well, even in environments with obstacles and lack of line-of-sight.

We further address the effect of multi-path fading by giving transmitters the choice of several channels/frequencies. Empirically, we find that in an environment with extensive multi-path propagation (but otherwise simple), the choice improves the metricity parameter significantly. Analytically, we show that a known algorithm for capacity maximization can be extended to handle multiple channels without loss in performance.

Roadmap. In the following section, we formally define our concepts, describe and calibrate our experimental setup in and validate the basic premises of our framework. We analyze the metricity parameter ζ in Section 3 and present experimental results. By leveraging the metricity concept in our framework, we introduce an approach for tackling multi-path fading using multiple frequencies in Section 4 and present experimental and theoretical results. We survey related work in Section 5 and conclude in Section 6.

2. MODEL VALIDATION

Our first order of action is to verify that the SINR model, without the geometric assumption, is faithful to reality. We assess the predictability of packet reception rate (PRR) under interference, and the assumption of the additivity of interference, by comparing our

abstract model to the original GEO-SINR through experiments. The experiments are conducted in two testbeds (Fig 1): one in the middle of a large open classroom (TB-20) and another in a challenging basement corridor (TB-60).

2.1 The Physical Model

The SINR model is based on two key principles: (i) a signal decays as it travels from a sender to a receiver, and (ii) interference – signals from other sources than the intended transmitter – accumulates. A transmission is successfully received if and only if the strength of the received signal relative to interference is above a given threshold.

Formally, a *link* $\ell_v = (s_v, r_v)$ is given by a pair of nodes, sender s_v and a receiver r_v . The channel *gain* G_{uv} denotes the reciprocal of the signal decay of ℓ_u as received at r_v . If a set S of links transmits simultaneously, then the SINR at ℓ_v is

$$\text{SINR}_v := \frac{P_v G_{vv}}{N + \sum_{u \in S} P_u G_{uv}}, \quad (1)$$

where P_v is the power used by the sender s_v of ℓ_v , and N is the ambient noise. In the *thresholded* SINR model, the transmission of ℓ_v is *successful* iff $\text{SINR}_v \geq \beta$, where β is a hardware-dependent constant.

The common assumption of *geometric path loss* in SINR models states that the gain is inversely proportional to a fixed polynomial of the distance traveled, *i.e.*, $G_{uv} = d(s_u, r_v)^{-\alpha}$, where the range of the *path loss constant* α is normally between 1 and 6. The geometric path loss assumption is valid in free space; we have $\alpha = 2$ in perfect vacuum.

The MB-SINR model refers to the SINR formula (1) applied to a general *gain matrix* G obtained through pairwise RSS measurements.

2.2 Experimental Setup




Figure 1: Topologies of our 20-node testbed (TB-20) (left) and 60-node curved corridor testbed (TB-60) (right).

Wireless hardware. Since a motivating goal of our study is to understand raw interference between wireless transceivers, we elected to operate at the physical-layer of a wireless device. We needed a mote with granular control over MAC-level capabilities, such as power and frequency control, over one tailored to specific protocol stacks, such as the 802.11 suite. For example, we require the ability to disable low-level features such as clear channel assessment (CCA).

We chose the Pololu Wixel, a development board for the TI CC2511F32 [35], as our mote hardware platform. The CC2511F32 is an 8051 micro controller SystemOnChip with integrated 2.4 GHz FM-transceiver stage (CC2500). In addition to meeting our functional requirements, the Wixels are inexpensive (14–20 USD), enabling larger scale deployments.

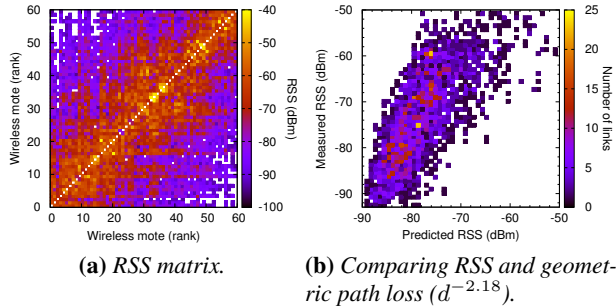


Figure 2: (a) *RSS matrix.* Gain between directed pairs of nodes in TB-60, measured by RSS and averaged over 1000 packets. (b) *Comparing measured and predicted RSS.* Correlation between RSS as predicted by distance with geometric path loss ($\alpha = 2.18$) and measured in TB-60 testbed.

Configuration. In our experiments, every sender node in each trial transmits a burst of 1000 packets with 4 ms delay between consecutive packets to facilitate successful delivery to the receiver. The length of each packet is 22 bytes, including a 16-bit CRC. Only packets that pass a CRC check are considered successful transmissions, with all error correction capabilities on the mote disabled. The radio is configured to use data whitening and Minimum Shift Keying (MSK) modulation format. During experiments, the wireless motes report details about packets sent or received to an auxiliary log via USB which also provides control signals and power for the experiments. Packet details include the received signal strength (RSS) as an integer in dBm.

TB-20 testbed. In the first testbed, we arranged 20 wireless motes on an 4×5 grid with 1 m spacing in an empty classroom. The motes in the grid were mounted on wooden poles 1 m from the ground in order to minimize reflection and attenuation from the ground; see Fig. 1 for the topology. The testbed was deployed temporarily for a focused set of tests.

TB-60 testbed. In the second testbed, we suspended 60 wireless motes about 0.3 m from metal wire trays and 2.5 m from the concrete floor in a curved basement corridor; see Fig. 1 for the topology. The corridor provides a challenging environment: limited line of sight between motes, obstacles such as water pipes and thick electric cables, and reinforced concrete walls. Approximately 94% of the directed links are in range for communication. The length of the corridor is 40.1 m, the longest distance (direct line) between any two motes is 21.8 m while the shortest distance is 0.4 m. The TB-60 testbed is a more permanent setup, with the experiments conducted over the span of several weeks.

2.3 Model calibration

We ran several experiments on the testbeds to gather calibration data for the MB-SINR and GEO-SINR models. The figures with error bars show the median, and upper and lower quartiles of the distribution of experimental trials.

Ambient noise parameter N . We evaluated N by sampling the noise level registered by each mote, over several hours in both early morning and during nighttime. All of our experiments use the 2.44 GHz frequency unless otherwise stated. We found the average ambient noise in TB-60 to be around -99.1 dBm, but considerably higher in TB-20 at around -94.4 dBm, in part due to external interference from 802.11 infrastructure.

Power setting P . We configured the wireless mote’s transmission power to 1 mW in all of our experiments.




Figure 3: *Packet reception rate for (a) MB-SINR (b) GEO-SINR.* Fraction of packets correctly decoded by receivers in TB-60 as the SINR is varied by evaluating different pairs, possibly invoking multiple senders. The plots show a transition from 0 to about 100% PRR as the SINR grows.

RSS matrix. We measured the RSS for all directed node pairs (s_v, r_v) in both testbeds. In each time slot, a chosen node transmits 1000 packets in a sequence, while other nodes act as receivers. The procedure was repeated for each pair of directed nodes. For temporal robustness, including day and night variations, the experiments were repeated at different times of the day.

Fig. 2a illustrates the RSS_{vv} for all node pairs (s_v, r_v) in testbed TB-60. The motes on both axes are ordered by the angle of their polar coordinates due to the arced positioning of the corridor, thus making neighboring motes in the testbed likely to be adjacent in Fig. 2a. The figure further demonstrates that every mote can hear some other mote in the testbed, that some mote pairs cannot communicate, and that transmissions are not fully symmetric.

Path loss constant α . Assuming geometric signal decay, the best linear least-squares fit for the path loss constant α given link lengths and the RSS values from Fig. 2a was $\alpha = 2.18 \pm 0.07$. By using geometric path loss $d(s_v, r_v)^{-2.18}$ to predict RSS values, we plot the correlation between the predicted RSS and measured RSS values in Fig. 2b. If the prediction were perfect, the points would fall on the $y = x$ diagonal line. Instead, the results in Fig. 2b confirm that geometric path loss is not a reliable predictor for RSS.

2.4 Comparison of GEO-SINR and MB-SINR

We next evaluate the predictive power of the two models in experiments with varying interference.

Controlling wireless interference in practice. One of the challenges with hardware experiments is the synchronization of the wireless motes. Our focus on measuring interference requires us to ensure that interfering motes are transmitting at the same time as the sender. Although we investigate sets of links that are transmitting at the same time, our analysis is focused on the performance of individual links. We therefore circumvent the problem of synchronizing the motes by running our experiments for each individual link in the link set. To analyze how a single link ℓ_v would perform in the presence of the other sender-receiver pairs in a set S , the links in $S \setminus \{\ell_v\}$ transmit continuously while we measure the transmission of ℓ_v . This continuous transmission ensures that the receiver in link ℓ_v experiences interference from other links.

Experimental design. With the synchronization issues in mind, we devised an experiment to compare the predictive power of the two models. We repeatedly select a random pair of nodes to act as sender and receiver, and a subset of 1–10 other nodes to cause interference. During the trial, the interfering nodes continuously transmit packets on the same frequency. We deploy low-level packet




Figure 4: Additivity in (a) MB-SINR and (b) GEO-SINR. Correlation between predicted and measured (RSS) in TB-60.

filtering at the receiver to minimize processing overhead due to interfering packets.

Packet reception rate by model. For every possible link in each testbed, we generated over 15,000 packet transmission trials over the link and measured the packet reception rate as a function of the SINR as calculated by the two models. Fig. 3a shows the PRR as a function of the SINR in MB-SINR, calculated using the RSS matrix shown in Fig. 2a. Corresponding results for GEO-SINR where the SINR is based on distances between nodes are shown in Fig. 3b.

The MB-SINR behaves as expected: generally the PRR and SINR values are either both small or both large. There is a swift transition from low to high PRR as the SINR value increases. We note that occasional trials produce a small PRR value despite SINR being large, as indicated by the two large error bars where $\text{SINR} \approx 10^3$. These outliers stem from occasionally no packets being received even for a large SINR, likely due to details of the testbed topology, such as destructive interference caused by signal reflection. Contrary to the GEO-SINR model, the MB-SINR model has a discernible threshold for successful transmissions.

Additivity of interference. Among the assumptions made by the SINR model is that interference is additive. In other words, if multiple senders transmit simultaneously, the RSS at the receiver can be estimated as the sum of the individual signals. Fig. 4a and 4b show the actual RSS as a function of the predicted RSS as given by MB-SINR (Fig. 4a) and GEO-SINR (Fig. 4b). We note that the variability evident in the measured RSS arises due to sparsity of data in those regions. If the additivity assumption is true, we would expect the values in the figures to fall on the diagonal line $y = x$.

The GEO-SINR appears more closely described by a pair of line segments with different slopes than a linear fit. Using linear regression, the coefficient of variation between the axes is low ($r^2 \simeq 0.031$), implying a low goodness-of-fit. Conversely, the MB-SINR model has a strong linear trend, with linear regression to the diagonal line incurring only 3.2% error and producing a large coefficient of variation ($r^2 \simeq 0.968$) between the predicted and measured RSS. The MB-SINR therefore more closely captures the additivity of interference than the canonical GEO-SINR model.

Sensitivity and specificity analysis. The GEO-SINR and MB-SINR models can be viewed as binary classifiers that compare the SINR to a threshold (β) to determine whether or not a transmission will be successful. We say that a transmission is experimentally successful if $\text{PRR} \geq T_{high}$ and declare it to be a failure if $\text{PRR} < T_{low}$. We focus on those links that were clearly either feasible or infeasible in our experiments, and set $T_{high} = 0.8$ and $T_{low} = 0.2$. Roughly 6% of the tested links fall within the 0.2 – 0.8 range and are thus not considered. A single instance in an SINR binary classifier can have four outcomes:

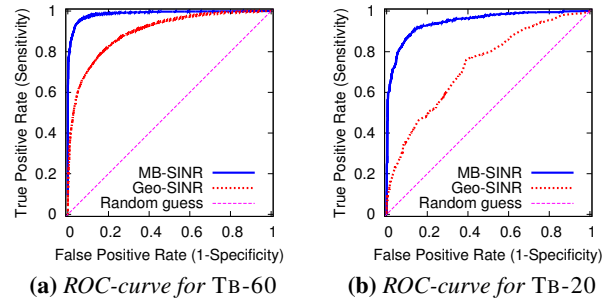


Figure 5: ROC-curves for (a) TB-60 and (b) TB-20. Comparison of MB-SINR and GEO-SINR as estimators for successful transmission of packets as the acceptance threshold β is varied. Each trial consists of 1000 packets exchanged in TB-60. A positive trial outcome has $\text{PRR} \geq 80\%$ whereas a negative one has $\text{PRR} \leq 20\%$.

- True positive (TP): $\text{SINR} \geq \beta, \text{PRR} \geq T_{high}$
- False positive (FP): $\text{SINR} < \beta, \text{PRR} \geq T_{high}$
- True negative (TN): $\text{SINR} < \beta, \text{PRR} < T_{low}$
- False negative (FN): $\text{SINR} \geq \beta, \text{PRR} < T_{low}$

A binary classifier incurs an inherent trade-off between true positive rate (sensitivity), defined as $\frac{\text{TP}}{\text{TP}+\text{FN}}$, and false positive rate (1-specificity), defined as $1 - \frac{\text{TN}}{\text{FP}+\text{TN}} = \frac{\text{FP}}{\text{FP}+\text{TN}}$. The trade-off balance can normally be tuned by a threshold parameter of the classifier, in this case β . By varying β , the trade-off can be graphically depicted on a ROC-curve (Receiver Operating Characteristic) that shows true and false positive rates on two axes. If $\beta = 0$, the classifier predicts that all transmissions will be successful, and if β is large, the classifier predicts that all transmissions will fail. A naive classifier making uniformly random guesses would fall on the diagonal line from (0, 0) to (1, 1), whereas the (0, 1) point denotes perfect classification.

Fig. 5a and 5b show the ROC-curves for the TB-60 and TB-20 testbeds, respectively. MB-SINR provides significantly better classification than GEO-SINR. In TB-60, the best trade-off between true and false positive rates occurs when $\beta = 2.15$, with a true positive rate of 94.8% and false positive rate of 5.2%.

In contrast, the predictions made by the canonical GEO-SINR on the same testbed plateaus at true positive rates of 81.4% and false positive rates of 18.6%. Both models give less accurate predictions for TB-20 compared to TB-60. The topology for TB-20 is more compact than TB-60. However, the larger and more variable ambient noise in TB-20 makes it more difficult to accurately predict the outcome of a transmission than in the TB-60 testbed.

As expected, MB-SINR provides significantly superior predictive power for PRR than GEO-SINR on both testbeds. Our results are robust against modifying the thresholds to $T_{high}, T_{low} = 0.5$. Note that RSS measurements used to compute SINR in MB-SINR were performed weeks in advance of these experiments. This suggests that the RSS matrix is resilient to temporal factors, with MB-SINR correctly predicting nearly 95% of all instances.

3. METRICITY

A plethora of important wireless interference algorithms rely on the GEO-SINR model, many of which have no obvious generalization to arbitrary metrics. Moreover, several problems in the domain, such as finding the maximum set of links that can simultaneously transmit (the LINK CAPACITY problem), have been proved to be computationally hard in an unconstrained SINR model [11].

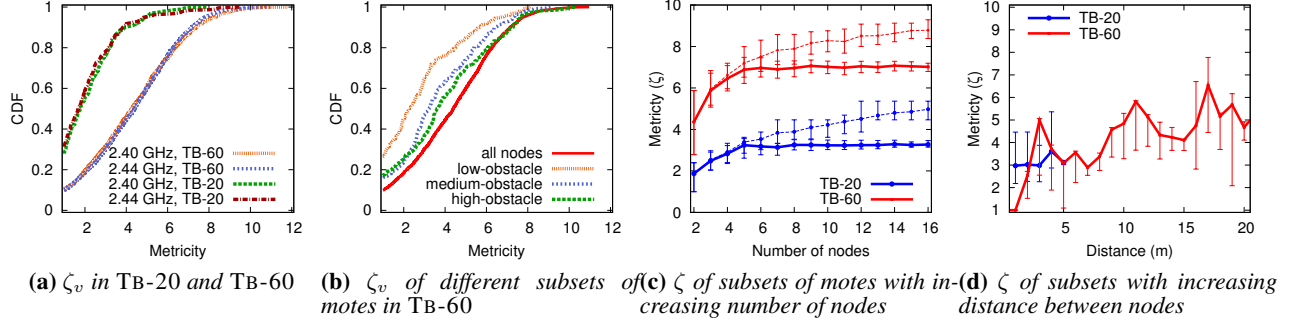


Figure 6: (a) Metricity (ζ_v). CDF comparison of computed ζ_v values for RSS matrices on both testbeds. The lines show ζ_v values for two different frequencies, which are representative of other frequencies in the respective testbeds. **(b) Environmental factors.** CDF comparison of computed ζ_v values on three different subsets of nodes. low-obstacle has the fewest environmental obstacles; high-obstacle has the most. **(c) Impact of scale on ζ .** Metricity compared to the number of nodes calculated as the average metricity of 200 randomly generated subsets of nodes of size at most 16, calculated for both testbeds. The thick lines represent the average of the 95th percentile ζ of subsets of the same size, whereas the thin lines represent the average maximum ζ of the subsets of the same size. The error bars represent the inter-quartile range. **(d) Influence of pairwise distances on ζ .** Metricity calculated on subsets of nodes with similar distances on both testbeds. The lines represent the 95th percentile of the ζ value of each set. The error bars represent the interval between the highest value of ζ and the median.

To facilitate algorithmic analysis under the more realistic MB-SINR model, we introduce a *metricity* parameter ζ that reflects how well signal decay resembles a metric space.

In what follows, we assume arbitrary path loss with gain $G_{uv} = 1/f(s_u, r_v)$ for some function f of pairs of points. Note that the RSS between sender s_u to receiver r_v is $P_u G_{uv}$.

Definition. The *metricity* $\zeta(x, y)$ of a given node pair (x, y) in gain matrix G is defined to be the smallest number satisfying for any mote z with links also in G ,

$$f(x, y)^{1/\zeta(x, y)} \leq f(x, z)^{1/\zeta(x, y)} + f(z, y)^{1/\zeta(x, y)}. \quad (2)$$

For notational simplicity we use $\zeta_v = \zeta(s_v, r_v)$.

Definition. We define ζ as the maximum value over all ζ_v in G unless specified otherwise. ζ is well defined, namely, consider $\zeta = \zeta_0 := \log_2(f_{max}/f_{min})$, where $f_{max} = \max_{x, y} f(x, y)$ and $f_{min} = \min_{x, y} f(x, y)$. Then, we can see that the LHS of (2) is at most $f_{max}^{1/\zeta_0} = 2f_{min}^{1/\zeta_0}$, while the RHS is at least that value. In the case of geometric path loss, $\zeta \leq \alpha$, since $f(x, y) = d(x, y)^\alpha$ and the distance function $d(x, y)$ satisfies the ordinary triangle inequality.

Theoretical implications. The ζ parameter has the advantage that theoretical results in the GEO-SINR can be imported *without significant changes* to the MB-SINR. Specifically, the following is true.

All results that hold for general metrics in the GEO-SINR carry over to the MB-SINR with trivial modifications, giving practically identical performance ratios in terms of ζ as the original result had in terms of α .

Results that were proven for general metric spaces therefore do not depend on the particular value of the path loss constant α or that the value holds homogeneously. The results rely upon the triangle inequality, for which Eqn. 2 applies equally well in arbitrary gain matrices.

Taking the LINK CAPACITY problem as an example, approximation results carry over for numerous cases: fixed power [16]; arbitrary power control [20, 21]; distributed setting based on regret minimization [2] and under jamming [7]; and the weighted version with linear power [17].

In a sibling paper [5], we have examined in detail how exactly the current body of analytic results carries over to the MB-SINR. With only few exceptions, results that have been derived specifically for the Euclidean plane also hold in the MB-SINR model, making only an elementary assumption about the convergence of interference: that the collective interference of uniformly distributed nodes does not tend to infinity.

A consequence of these translations between models is that previous theoretical work in the GEO-SINR model is in fact highly robust to spatial signal variability.

3.1 Experimental evaluation

Method and evaluation. Using values obtained for the RSS matrices in both testbeds, shown for TB-60 in Fig. 2a, we evaluate the minimal value of ζ_v , (as defined in Eqn. 2) for every directed node pair within communication range. The cumulative distribution function (CDF) of the resulting ζ_v values is shown for two frequencies (2.40 GHz and 2.44 GHz) in Fig. 6a. In both testbeds the values for ζ_v range up to 12. However, note that the metricity values in TB-60 are generally a bit larger than in TB-20. The discrepancy is to be expected, since the challenging TB-60 environment both contains longer links with more variable signal strength as well as more variable signal attenuation due to obstacles.

In both testbeds, we find that a small fraction of the links have comparatively higher values of ζ_v , as seen by the long tapering at the top in Fig. 6a, which in turn drives the value for ζ to a relatively high number. To quantify, we find that the 95th and 99th percentile of the ζ_v value of links are significantly smaller than the global maximum ζ . In TB-60 some of the weakest links are only able to communicate at particular frequencies. These links correspond to pairs that have weak signals, and limited or no line-of-sight. In TB-20, while all pairs have line-of-sight and can communicate with one another, the large ζ value is an artifact of relatively high ζ_v values for only a handful of links. We investigate the effect further in Section 4.

Factors affecting ζ . One of the drivers behind the definition of ζ is to measure the “complexity” of the environment. To be more precise, we investigated the impact of obstacles in the environment, the number of nodes of a network and the distances between nodes in order to better pinpoint what features most influence the ζ value.




Figure 7: (a) Metricity (ζ_v) with log-normal shadowing in TB-60. CDF of comparison on the metricity ζ_v using log-normal shadowing on distances in TB-60 averaged over 10 instances compared to actual measurements. (b) Metricity (ζ) with log-normal shadowing vs. standard deviation. Metricity ζ (as the 95th percentile on ζ_v) using log-normal shadowing on distances in both testbeds with growing standard deviation σ .

To examine the impact of different environmental characteristics, we divide the nodes in TB-60 into three subsets of 20 nodes in Fig. 6b. The set of nodes with the lowest values of ζ_v has the fewest number of obstacles (*low-obstacle*), whereas the set with the higher ζ_v values (*high-obstacle*) has a variety of barriers, such as electric cables suspended in the ceiling and greater distances between nodes. The *medium-obstacle* group has an average number of barriers while also having the most condensed topology. We calculated the values for ζ_v on the different links in the induced gain matrix for each set to obtain Fig. 6b. The figure suggests that the complexity of the environment, in the form of physical obstacles, might have a significant impact on the value of ζ .

We further examine the impact of the size of the set of nodes on ζ . Using the RSS matrices for TB-60 and TB-20, we calculate ζ for randomly generated subsets of different sizes. Fig. 6c shows metricity as a function of set size. The thick lines represent ζ as the average 95th percentile of the values for ζ_v , which is relatively stable with increased sizes for sets of five or more nodes. The corresponding thin lines, which represent the average ζ as the highest value for ζ_v in each testbed, demonstrate that the global maximum ζ continues to grow with the set size. The increase is to be expected, as larger sets are more likely to include the links with the highest ζ_v values.

We also looked into the relationship between node distances and the ζ value of a set. We took different subsets of nodes that have roughly the same distances to one another, as shown in Fig. 6d. The figure suggests that the distances between nodes are poorly correlated with the value of ζ .

3.2 Log-normal shadowing and metricity

One of the most commonly used stochastic extensions of geometric path loss to address observed variability in signal propagation is log-normal shadowing. According to this model, signal decay follows the geometric model d^α , but with a multiplicative exponentially distributed factor:

$$f(s_u, r_v) = d(s_u, r_v)^\alpha \cdot e^X,$$

where X is a normally distributed random variable with zero mean.

We note that log-normal shadowing is an approach to introduce non-geometric properties into gain matrices. Since metricity is a




Figure 8: Metricity (ζ_v) in TB-20 and TB-60. CDF comparison of computed ζ_v values for RSS matrices on both testbeds. The thick lines represent values computed for the \widehat{RSS} matrices. The dashed lines show ζ values for the 2.44 GHz frequency, which is representative of other frequencies in the respective testbeds.

measure of such discrepancy, it might be instructive to calculate metricity on instances generated with log-normal shadowing.

We used the topology of the TB-60 testbed and generated log-normal distributions with standard deviations $\sigma = 1.0$ and 2.0 , and computed the CDF of the metricity of the resulting gain matrices. Averages over 10 instances are shown in Fig. 7a, interposed with the actual ζ_v measurements. We note a similarity between the measured values and the log-normal values with $\sigma = 1.0$, although the tail is measurably heavier in the latter.

We also examined how global metricity ζ grows with the standard deviation σ in the generated log-normal distributed instances in both testbeds. We show the results on Fig. 7b, where we display the 95th percentile of the distribution of ζ_v values. The plot exhibits a linear relationship, as shown in Fig. 7b, or roughly $\zeta = 2.5 + 4.9\sigma$. The agreement accords with the expected signal variations produced by a multiple of an exponentially distributed random variable.

4. FINESSING MULTI-PATH FADING WITH MULTIPLE CHANNELS

In our experiments, high metricity values were primarily caused by a handful of links. In particular, experiments in the simple environment of the TB-20 testbed exhibited a higher value of ζ than we suspected. One potential source of the complexity may be due to adverse signal reflection.

When signals travel along different paths, the superposition of the different signals produces patterns of signal cancellation and amplification known as multi-path fading ([10, Sec. 2]). This effect is particularly pronounced and systematic in simpler settings. The interference pattern will necessarily shift with frequency. Hence, the influence of multi-path fading on signal reception between a pair of points is likely to vary greatly with the chosen channel.

We propose to tackle the problem of signal cancellation and destructive interference by supplying algorithms with multiple channels (frequencies) from which to choose. First, we propose a variation of the metricity parameter ζ and evaluate the difference from the experimental data from our testbeds. As a case study, we then formulate a multi-channel version of the link LINK CAPACITY problem and obtain worst-case approximation results.

4.1 Experimental Evaluation

We performed additional experiments to obtain RSS matrices RSS^f for 8 different frequencies f ranging from 2.40 GHz to

2.48 GHz or wavelengths between 7 and 7.15 cm. These frequencies mean that multi-path alignment can shift from fully destructive to fully constructive interference when the difference in the path lengths is at least 1.4 m.

We calculated ζ_v values separately for each RSS^f . Although they can vary significantly on a per link basis, we found the differences in the overall distributions to be insignificant. Fig. 8 shows $f = 2.40$ GHz — other frequencies had similar distributions.

To factor out frequency-dependent fading, we computed for each node pair (s_v, r_v) the median $\widehat{\text{RSS}}_v$ of the eight RSS_v^f values ranging over the different frequencies. We observe that ζ_v values of the matrix RSS are significantly lower, as the thick lines in Fig. 8 indicate.

The frequency dependency is more apparent in TB-20, which suggests that signal reflection plays a relatively large role in that environment. The increased reliance on a particular frequency can be explained by the regular grid structure and condensed setting of TB-20 (Fig. 1), which makes the testbed a good candidate to observe (frequency dependent) multi-path fading phenomena. However, the links in TB-60 are on average longer and thus reflection plays a smaller role in signal attenuation. Furthermore, the greater number of obstacles in TB-60 may also explain the decreased dependency on frequency.

The observation that the channels have different fading properties brought us to introduce a new version of the LINK CAPACITY problem to incorporate different frequencies.

4.2 Link capacity with multiple frequencies

The empirical indications – that having a choice of channels use results in smaller values of ζ_v , and thus better approximation factors – motivate us to generalize the LINK CAPACITY problem. Namely, in MULTI-CHANNEL LINK CAPACITY, we assign links to a set of frequencies, but each link is only *eligible* to use a subset of the frequencies. As before, we want to assign as many links as possible with the constraint that those assigned to a given frequency form a feasible set.

This formulation considers links that experience significant frequency-dependent fading as not usable in that frequency. It does not take into consideration the possible decrease in *interference* due to such fading. One reason is that such fading is too unpredictable to expect any algorithm to utilize that to obtain better solutions than otherwise, and thus it is also not fair to compare with such a strong adversary. The other reason is that with arbitrary fading patterns, we are back in the *abstract SINR* model, for which very strong in-approximability results hold [11].

The classic LINK CAPACITY problem is to find a maximum subset $S \subseteq L$ of a given set L of links that can successfully transmit simultaneously. We modify the LINK CAPACITY problem to fit our observations on the use of multiple frequencies:

MULTI-CHANNEL LINK CAPACITY

Given: A set L of n links, and k subsets $L_1, L_2, \dots, L_k \subseteq L$.

Find: Sets S_1, S_2, \dots, S_k with $S_i \subseteq L_i$ and S_i feasible, for $i = 1, 2, \dots, k$.

Maximize: $|S_1 \cup S_2, \dots \cup S_k|$.

Here L_i represents the links that are eligible for frequency i and S_i those scheduled for that frequency.

Additional definitions. To simplify notation we write $f_{uv} = f(s_u, r_v)$ and $f_v = f_{vv}$. We assume a total order \prec on the links, where $\ell_v \prec \ell_w$ implies that $f_v \leq f_w$. We use the shorthand notation $\ell_v \prec L$ to denote that $\ell_v \prec \ell_u$ for all links ℓ_u in L . A power assignment \mathcal{P} is *decay monotone* if $P_v \leq P_w$ whenever $\ell_v \prec \ell_w$, *reception monotone* if $\frac{P_w}{f_w} \leq \frac{P_v}{f_v}$ whenever $\ell_v \prec \ell_w$, and simply

monotone if both properties hold.¹ This captures the main power strategies, including uniform and linear power.

We modify the notion of *affactance* [11, 22]: The affactance $a_w^{\mathcal{P}}(v)$ of link ℓ_w on link ℓ_v under power assignment \mathcal{P} is the interference of ℓ_w on ℓ_v normalized to the signal strength (power received) of ℓ_v , or

$$a_w(v) = \min \left(1, c_v \frac{P_w G_{wv}}{P_v G_{vv}} \right) = \min \left(1, \frac{P_w f_v}{P_v f_{wv}} \right), \quad (3)$$

where $c_v = \frac{\beta}{1 - \beta N / (P_v G_{vv})} > \beta$ is a constant depending only on universal constants and the signal strength G_{vv} of ℓ_v , indicating the extent to which the ambient noise affects the transmission. We drop \mathcal{P} when clear from context. Furthermore let $a_v(v) = 0$. For a set S of links and link ℓ_v , let $a_v(S) = \sum_{\ell_w \in S} a_w(v)$ and $a_S(v) = \sum_{\ell_w \in S} a_w(v)$. Assuming S contains more than two links we can rewrite Eqn. 1 as $a_S(v) \leq 1$ and this is the form we will use. Observe that affactance is additive and thus $a_S(v) = a_{S_1}(v) + a_{S_2}(v)$ for any partition (S_1, S_2) of S .

We define a weight function $W_+(v, w) = a_v(w) + a_w(v)$, when $\ell_v \prec \ell_w$ and $W_+(v, w) = 0$, otherwise. The plus sign is to remind us that weights are from smaller to larger decay links. Also, $W_+(X, v) = \sum_{\ell_w \in X} W_+(w, v)$, representing the sum of the in- and out-affactances (as in Eqn. 3) of a link v to and from those links in set X that have smaller decay.

A set S of links is *anti-feasible* if $a_v(S) \leq 2$ for every link $\ell_v \in S$ and *bi-feasible* if both feasible and anti-feasible [15]. More generally, for $K \geq 1$, S is *K-feasible* (*K-anti-feasible*) if $a_v(S) \leq 1/K$ ($a_S(v) \leq 2/K$), and *K-bi-feasible* if both.

Approximation of MULTI-CHANNEL LINK CAPACITY. We extend a greedy algorithm for LINK CAPACITY [16] and show that it gives equally good approximation algorithm for MULTI-CHANNEL LINK CAPACITY, even in MB-SINR. We assume that the links are assigned monotone power.

Algorithm 1 MULTI-CHANNEL LINK CAPACITY in MB-SINR

```

Let  $L$  be a set of links using monotone power  $\mathcal{P}$  and  $L_1, L_2, \dots, L_k \subseteq L$ 
be subsets.
Set  $X_1, X_2, \dots, X_k \leftarrow \emptyset$ 
for  $\ell_v \in L$  in order of increasing  $f_v$  values do
  for  $i \leftarrow 1 \dots k$  do
    if  $W_+(X_i, v) \leq 1/2$  then
       $X_i \leftarrow X_i \cup \{\ell_v\}$ 
    break
  for each  $X_i$  do
     $S_i \leftarrow \{\ell_v \in X_i \mid a_{X_i}(v) \leq 1\}$ 
return  $(S_1, S_2, \dots, S_k)$ 

```

Note that the sets returned by Algorithm 1 are feasible by construction.

We turn to proving a performance guarantee for the algorithm. The following key result bounds the affactance of a feasible set to a (shorter) link outside the set to a constant. A similar but weaker bound was first introduced by Kesselheim and Vöcking [22].

LEMMA 1. *Let L be a $3^\zeta/\beta$ -bi-feasible set with monotone power assignment \mathcal{P} and let ℓ_v be a link (not necessarily in L) with $\ell_v \prec L$. Then, $W_+(v, L) = O(1)$.*

We prove Lemma 1 by splitting it into two lemmas bounding in-affactance for links in a feasible set and similarly bounding out-affactance for links in an anti-feasible set.

¹This corresponds to *length monotone* and *sublinear* power assignments in GEO-SINR [22].




Figure 9: We show that for each link ℓ_x it holds that $Q \geq \frac{F}{2^\zeta}$.

LEMMA 2. Let L be a $3^\zeta/\beta$ -feasible set with monotone power assignment \mathcal{P} and let ℓ_v be a link with $\ell_v \prec L$. Then, $a_L^{\mathcal{P}}(v) = O(1)$.

PROOF. The basic idea is to identify a “proxy” for ℓ_v within the set L . Namely, we bound the affectance of L on ℓ_v in terms of the affectance on the “nearest” link ℓ_u in L , which is small since L is feasible and contains ℓ_u .

Formally, consider the link $\ell_u = (s_u, r_u) \in L$ such that $F := f(r_v, r_u)$ is minimum and link $\ell_w = (s_w, r_w) \in S$ such that $f_{vw} = f(s_w, r_v)$ is minimum (possibly $\ell_u = \ell_w$). Let ℓ_x be an arbitrary link in L and define $Q = f_{xv}$. See Fig. 9.

We first show that

$$F \leq 2^\zeta Q. \quad (4)$$

Let ℓ_y, ℓ_z be renamings of the links ℓ_x, ℓ_w such that $f_y = \max(f_x, f_w)$ and $f_z = \min(f_x, f_w)$.

By definition of ℓ_u and ℓ_w , it holds that $\max(f_{yv}, f_{zv}) \leq Q$. Thus, using the weak triangular inequality,

$$f(s_y, s_z)^{1/\zeta} \leq f_{yv}^{1/\zeta} + f_{zv}^{1/\zeta} \leq 2Q^{1/\zeta}. \quad (5)$$

Using Eqn. 5 and that $f_z \leq f_y$, it holds that

$$f_{yz}^{1/\zeta} \leq f_z^{1/\zeta} + f(s_y, s_z)^{1/\zeta} \leq f_y^{1/\zeta} + 2Q^{1/\zeta}. \quad (6)$$

By the feasibility condition on L , $a_y(z) \leq \beta/3^\zeta$, while by definition of affectance and reception monotonicity (i.e., $P_y/f_y \leq P_z/f_z$),

$$a_y(z) = c_z \frac{P_y}{P_z} \frac{f_z}{f_{yz}} \geq \beta \frac{f_y}{f_{yz}}.$$

Combining the two bounds on $a_y(z)$, we get that $3^\zeta \cdot f_y \leq f_{yz}$. That, combined with Eqn. 6 and canceling a f_y factor, gives that $3^\zeta \leq \left(1 + 2(Q/f_y)^{1/\zeta}\right)^\zeta$, which implies that $f_y \leq Q$ and further that $f_w \leq f_y \leq Q$. Then, by the definitions of F, ℓ_w and Q ,

$$F^{1/\zeta} = f(r_v, r_w)^{1/\zeta} \leq f_{vw}^{1/\zeta} + f_w^{1/\zeta} \leq 2 \cdot Q^{1/\zeta}, \quad (7)$$

implying Eqn. 4, as desired.

Now, using the weak triangular inequality, the definition of F , and Eqn. 7, we get that

$$f_{xv}^{1/\zeta} \leq f_{xw}^{1/\zeta} + f(r_v, r_u)^{1/\zeta} \leq Q^{1/\zeta} + F^{1/\zeta} \leq Q^{1/\zeta} + 2Q^{1/\zeta}.$$

Thus,

$$f_{xu} = 3^\zeta Q = 3^\zeta f_{xv}. \quad (8)$$

Observe that since $f_v \leq f_u$ and power is monotone, it holds that $c_v \leq c_u$. Then, using Eqn. 8 and the definition of affectance,

$$a_x(v) = c_v \frac{P_x}{f_{xv}} \frac{f_v}{P_v} \leq c_u \frac{3^\zeta P_x}{f_{xu}} \frac{f_u}{P_u} = 3^\zeta a_x(u).$$

Finally, letting $L_w = L \setminus \{\ell_w\}$, we sum over all links in L ,

$$a_L(v) = a_w(v) + a_{L_w}(v) \leq 1 + 3^\zeta a_{L_w}(u) \leq 1 + 3^\zeta \cdot \frac{\beta}{3^\zeta} = O(1),$$

using the feasibility assumption for the last inequality. \square

For anti-feasible sets a similar result holds with a nearly identical proof, swapping the roles of senders and receivers of the links.

LEMMA 3. Let L be a $3^\zeta/\beta$ -anti-feasible set with monotone power assignment \mathcal{P} and let ℓ_v be a link with $\ell_v \prec L$. Then, $a_v^{\mathcal{P}}(L) = O(1)$.

PROOF. Form the dual links L^* , which has a link $\ell_v^* = (s_v^*, r_v^*)$ for each link $\ell_v = (s_v, r_v) \in L$ such that $s_v^* = r_v$ and $r_v^* = s_v$. Clearly the lengths of ℓ_v^* and ℓ_v are the same so $f_{v^*v^*} = f_{vv}$. Also, $f_{v^*u^*} = f_{uv}$. Observe that the anti-feasibility assumption on L implies that L^* is $3^\zeta/\beta$ -feasible. Then, we can follow the proof of Lemma 2, applied to the set L^* , to get that

$$f_{ux} = f_{x^*u^*} \leq 3^\zeta f_{x^*v^*} = 3^\zeta f_{vx}.$$

This implies, using the monotonicity of power, that

$$a_v(x) = c_x \frac{P_v}{f_{vx}} \frac{f_x}{P_x} \leq c_x \frac{3^\zeta P_u}{f_{ux}} \frac{f_x}{P_x} = 3^\zeta a_u(x).$$

The rest of the proof is identical. \square

Combining Lemma 2 and 3 implies Lemma 1.

Finally, to analyze the performance ratio of Algorithm 1, we will use an adaptation of the following signal-strengthening lemma from [8, Prop. 8] and a lemma generalizing a popular argument used to show that the size of a subset of links of another set of links is large.

LEMMA 4 ([8]). Let L be a feasible set and $K \geq 1$ be a value. Then, there exists a K -bi-feasible subset of L of size $\Omega(|L|/K)$.

The approximation that we can prove for MULTI-CHANNEL LINK CAPACITY has actually better dependence on ζ than what follows for LINK CAPACITY from [16]. A lower bound of $\Omega(2^{\zeta - o(1)})$ on the approximability of LINK CAPACITY [16, 5] implies that the bound is close to best possible.

THEOREM 5. Algorithm 1 yields a $O(3^\zeta)$ -approximation for MULTI-CHANNEL LINK CAPACITY.

PROOF. Let L be a set of links and let $L_1, L_2, \dots, L_k \subseteq L$ be subsets of L where L_i contains the links that are eligible in frequency i . Let $OPT = OPT_1 \cup OPT_2 \cup \dots \cup OPT_k$ be an optimum solution to MULTI-CHANNEL LINK CAPACITY on L . Let $K = 3^\zeta/\beta$. By Lemma 4, there is a K -bi-feasible subset OPT'_i in OPT_i of size $\Omega(|OPT_i|/K)$, for each $i \in \{1, \dots, k\}$. Let $OPT' = OPT'_1 \cup OPT'_2 \cup \dots \cup OPT'_k$.

Let $S = S_1 \cup S_2 \cup \dots \cup S_k$ and $X = X_1 \cup X_2 \cup \dots \cup X_k$ be the sets computed by Algorithm 1 on input L . We first bound $|S|$ in terms of $|X|$ and then $|X|$ in terms of $|OPT'|$. To bound $|S|$ to $|X|$ we bound $|S_i|$ to $|X_i|$ for every $i \in \{1 \dots k\}$. Note that by the construction of X_i , $a_{X_i}(X_i) = W_+(X_i, X_i) \leq |X_i|/2$, and thus the average in-affectance of links in X_i is at most $1/2$. Since each S_i consists of the links in X_i of affectance at most $1/2$, by Markov's inequality, $|S_i| \geq |X_i|/2$.

By the definition of the algorithm, $W_+(X_i, \ell_w) > 1/2, \forall \ell_w \in OPT' \setminus X, X_i \in X$. Summing over all links ℓ_w in $OPT' \setminus X$, we get that $W_+(X_i, OPT' \setminus X) > |OPT' \setminus X|/2$. Furthermore,

since OPT' contains k K -feasible sets, it follows by Lemma 1 that $W_+(\ell_v, OPT') = O(k)$, for each $\ell_v \in X$. Summing over all links in X_i , we get that $W_+(X_i, OPT) = O(k|X_i|)$. Combining yields that for any set X_i we have $|OPT' \setminus X|/2 < W_+(X_i, OPT' \setminus X_i) \in O(k|X_i|)$, giving that $|X_i| = \Omega(|OPT' \setminus X|/k)$.

Summing over i then gives $|X| = \sum_i |X_i| = \Omega(|OPT' \setminus X|)$. Thus, the solution output by the algorithm satisfies $|S| \geq |X|/2 = \Omega(|OPT'|) = \Omega(|OPT|/K) = \Omega(|OPT|/3^\zeta)$. \square

In summary, the metricity definition implies that a large range of algorithmic results from GEO-SINR carries over without change. Thus, MB-SINR has both the desired generality and amenability to algorithmic analysis. We also extend known results on LINK CAPACITY to handle frequency-sensitive links, and improve the dependence of the approximation on ζ along the way.

5. RELATED WORK

Numerous experimental results have indicated that simplistic range-based models of wireless reception are insufficient [9, 37, 23, 1, 29, 38, 36]. Besides directionality and asymmetry, signal strength is not well predicted by distance. Interference patterns are also insufficiently explained by pairwise relationships, suggesting the need for additive interference models, both experimentally [23, 27, 24] and analytically [27, 26, 19].

The weakness of the known prescriptive models for interference and packet reception has led experimentalists to form models based on measurements. Son, Krishnamachari and Heidemann [34] showed that the SINR formula, using separately measured RSS values, is the main factor in predicting PRR. They found PRR to be dependent on the number of interferers, which was not supported in later studies [24, 6] and attributed to hardware variability or the quality of the CC1000 radios used. Reis et al. [32] independently proposed a similar approach on a 802.11 platform. They found substantial variability across nodes, and that similarity across time was sufficient over moderate time scales of minutes to hours, but that prediction accuracy degrades over longer periods.

Maheshwari, Jain and Das [24] compared different models of interference using two testbeds with variations in hardware, power level, and indoor/outdoor. They concluded that the physical model gives best accuracy, albeit less than perfect. In their followup workshop paper [25], they focus on the relationship of *joint interference* (SINR with multiple interferers) to PRR. They gave strong evidence that the basic formula works, and verify the *additivity* of the SINR model.

Chen and Terzis [6] proposed a method for calibrating RSSI readings to combine interference measurements from different motes. They found that Tmote Sky motes consistently report RSSI values inaccurately, even reporting non-injective relationships. By aligning measurements from different motes, they obtained much better SINR vs. PRR relationship, reducing the width of the intermediate range significantly. They suggested that this may explain much of the imperfect relationship observed in [24].

Measurement-based approaches have also been proposed in the context of 802.11 [13, 31, 33], where carrier sense and control packets complicate the picture. Recent efforts have focused on reducing the required measurements by deducing interference using, e.g., linear algebra [30] or regression [18]. Boano et al. [4] also studied the impact of external interference on sensor-network MAC protocols, and identified mechanisms to improve their robustness.

The engineering literature has introduced various extensions in order to capture reality more faithfully. In the *two-ray* model [10], which captures reflection off the ground, signal decays in the near-term as a certain polynomial, but as a higher degree polynomial

further away. More generally, the *multi-ray* model has the signal (in dBm, log scale) decaying via a piecewise linear model with segments of increasing slopes. The function is typically empirically determined. We note that these models can be captured by MB-SINR, with ζ as the steepest slope.

There are also *empirical models* [10], such as the *Okumura* and *Hata* models, that take the environment into account. These could also be used to generate a gain matrix. Also, accurate estimates can be obtained via general *ray tracing* when highly detailed information is available.

Some studies have allowed signal strength to fluctuate from the geometric path loss by up to a constant factor [28, 11]. This is of limited help in general, however, since even minor fluctuations of the value of α can cause arbitrarily large changes in signal strength [12].

Various probabilistic models also exist. On one hand, they are means to prescribe non-geometric components to signal reception, which is useful for simulation studies (which MB-SINR cannot provide), but which could be captured more accurately by actual measurements. On the other hand, these can also model aspects that are necessarily random, in which case they could complement the deterministic MB-SINR.

6. CONCLUSION

Effective use of the wireless spectrum requires an understanding of interference from theoretical and experimental vantage points. A growing body of algorithmic work on worst-case wireless interference under the SINR threshold model assumes that signals decay geometrically with distance, the GEO-SINR model.

We outlined an approach for incorporating realism into the interference model while seamlessly generalizing previous theoretical results. By leveraging a matrix of pairwise RSS between wireless motes instead of geometric path loss, the MB-SINR model predicted PRR performance significantly better in our experiments on two indoors testbeds. The RSS matrix also appears resilient to temporal factors, with prediction accuracy of 95% when using an RSS matrix created weeks in advance.

We defined the notion of metricity, which quantifies the proximity of the RSS matrix to a distance metric. Through experiments, we showed how the concept effectively measures the complexity of the underlying environment. With metricity as a harness, worst-case theoretical results that hold under general metrics in the GEO-SINR model can now be translated with trivial modifications to the more realistic MB-SINR model. Moreover, the translation retains almost identical performance ratios for all such algorithms, with the metricity parameter ζ replacing the path loss constant α in the GEO-SINR model.

As a case study of the metricity concept, we addressed multi-path fading by allowing transmitters to choose between several frequencies. We found that environments with extensive multi-path effects exhibited better metricity values. By fusing empirical measurements into an analytical model commonly used for worst-case theoretical analysis, our approach suggests a methodology for harmonizing algorithmic theory of wireless interference with real-world observations.

Acknowledgments

This research is supported by grant-of excellence no 120032011 from the Icelandic Research Fund. We thank Henning Ulfarsson and members of RU Syslab for their valuable contributions and assistance with experiments, and the anonymous reviewers for their constructive feedback.

7. REFERENCES

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. *ACM SIGCOMM Computer Communication Review*, 34(4):121–132, 2004.
- [2] E. I. Ásgeirsson and P. Mitra. On a game theoretic approach to capacity maximization in wireless networks. In *INFOCOM*, 2011.
- [3] N. Baccour, A. Koubaa, L. Mottola, M. A. Zuniga, H. Youssef, C. A. Boano, and M. Alves. Radio link quality estimation in wireless sensor networks: a survey. *ACM Trans. Sensor Networks (TOSN)*, 8(4):34, 2012.
- [4] C. A. Boano, T. Voigt, N. Tsiftes, L. Mottola, K. Römer, and M. A. Zúniga. Making sensor MAC protocols robust against interference. In *Wireless Sensor Networks*, pages 272–288. Springer, 2010.
- [5] M. H. Bodlaender and M. M. Halldórsson. Beyond geometry: Towards fully realistic wireless models. In *PODC*, 2014. arXiv:1402.5003.
- [6] Y. Chen and A. Terzis. On the mechanisms and effects of calibrating RSSI measurements for 802.15.4 radios. In *Wireless Sensor Networks*, pages 256–271. Springer, 2010.
- [7] J. Dams, M. Hoefer, and T. Kesselheim. Jamming-resistant learning in wireless networks. In *ICALP*, 2014. arXiv:1307.5290.
- [8] A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *PODC*, pages 220–229, August 2009.
- [9] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical report, UCLA/CSD-TR 02, 2002.
- [10] A. Goldsmith. *Wireless Communications*. Cambridge Univ. Press, 2005.
- [11] O. Goussevskaia, M. M. Halldórsson, and R. Wattenhofer. Algorithms for wireless capacity. *IEEE/ACM Trans. Netw.*, 22(3):745–755, 2014.
- [12] Z. Gu, G. Wang, and Y. Wang. Path-loss fluctuations towards robust scheduling algorithms in the SINR model. In *MASS*, pages 416–424. IEEE, 2012.
- [13] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan. Understanding and mitigating the impact of RF interference on 802.11 networks. In *SIGCOMM*, pages 385–396. ACM, 2007.
- [14] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Trans. Information Theory*, 46(2):388–404, 2000.
- [15] M. M. Halldórsson and P. Mitra. Nearly optimal bounds for distributed wireless scheduling in the SINR model. In *ICALP*, 2011.
- [16] M. M. Halldórsson and P. Mitra. Wireless Capacity with Oblivious Power in General Metrics. In *SODA*, 2011.
- [17] M. M. Halldórsson and P. Mitra. Wireless capacity and admission control in cognitive radio. In *INFOCOM*, pages 855–863, 2012.
- [18] J. Huang, S. Liu, G. Xing, H. Zhang, J. Wang, and L. Huang. Accuracy-aware interference modeling and measurement in wireless sensor networks. In *ICDCS*, pages 172–181, 2011.
- [19] A. Iyer, C. Rosenberg, and A. Karnik. What is the right model for wireless channel interference? *IEEE Transactions on Wireless Communications*, 8(5):2662–2671, 2009.
- [20] T. Kesselheim. A Constant-Factor Approximation for Wireless Capacity Maximization with Power Control in the SINR Model. In *SODA*, 2011.
- [21] T. Kesselheim. Approximation algorithms for wireless link scheduling with flexible data rates. In *ESA*, pages 659–670, 2012.
- [22] T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *DISC*, pages 163–178, August 2010.
- [23] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *MSWiM*, pages 78–82. ACM, 2004.
- [24] R. Maheshwari, S. Jain, and S. R. Das. A measurement study of interference modeling and scheduling in low-power wireless networks. In *SenSys*, pages 141–154, 2008.
- [25] R. Maheshwari, S. Jain, and S. R. Das. On estimating joint interference for concurrent packet transmissions in low power wireless networks. In *WinTech*, pages 89–94. ACM, 2008.
- [26] T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *INFOCOM*, 2006.
- [27] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol Design Beyond Graph-Based Models. In *Hotnets*, 2006.
- [28] T. Moscibroda, R. Wattenhofer, and A. Zollinger. Topology control meets SINR: The scheduling complexity of arbitrary topologies. In *MOBIHOC*, pages 310–321, 2006.
- [29] J. Padhye, S. Agarwal, V. N. Padmanabhan, L. Qiu, A. Rao, and B. Zill. Estimation of link interference in static multi-hop wireless networks. In *IMC*, 2005.
- [30] X. Qi, Y. Wang, Y. Wang, L. Xu, and C. Hu. EIM: Efficient online interference measurement in wireless sensor networks. In *WiOpt*, pages 548–555. IEEE, 2013.
- [31] L. Qiu, Y. Zhang, F. Wang, M. K. Han, and R. Mahajan. A general model of wireless interference. In *MobiCom*, pages 171–182, 2007.
- [32] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Measurement-based models of delivery and interference in static wireless networks. In *SIGCOMM*, pages 51–62, 2006.
- [33] V. Sevani and B. Raman. SIR based interference modeling for wireless mesh networks: A detailed measurement study. In *COMSNETS*, pages 1–10. IEEE, 2012.
- [34] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of concurrent transmission in wireless sensor networks. In *SenSys*, pages 237–250. ACM, 2006.
- [35] TI Corporation. CC2510Fx/CC2511Fx: Low-power SoC (system-on-chip) with MCU, memory, 2.4 GHz RF transceiver, and USB controller. <http://www.ti.com/product/cc2511f32>. [Online; accessed 27-Nov-2013].
- [36] M. Z. Zamalloa and B. Krishnamachari. An analysis of unreliability and asymmetry in low-power wireless links. *ACM Trans. Sensor Networks (TOSN)*, 3(2):7, 2007.
- [37] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *SenSys*, pages 1–13. ACM, 2003.
- [38] G. Zhou, T. He, S. Krishnamurthy, and J. A. Stankovic. Models and solutions for radio irregularity in wireless sensor networks. *ACM Transactions on Sensor Networks (TOSN)*, 2(2):221–262, 2006.