

Beyond Geometry: Towards Fully Realistic Wireless Models*

Marijke H.L. Bodlaender
ICE-TCS, School of Computer Science
Reykjavik University, Iceland.
mbodlaender@gmail.com

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

ABSTRACT

Signal-strength models of wireless communications capture the gradual fading of signals and the additivity of interference. As such, they are closer to reality than other models. However, nearly all theoretic work in the SINR model depends on the assumption of smooth geometric decay, one that is true in free space but is far off in actual environments. The challenge is to model realistic environments, including walls, obstacles, reflections and anisotropic antennas, without making the models algorithmically impractical or analytically intractable.

We present a simple solution that allows the modeling of arbitrary static situations by moving from geometry to arbitrary *decay spaces*. The complexity of a setting is captured by a *metricity* parameter ζ that indicates how far the decay space is from satisfying the triangular inequality. All results that hold in the SINR model in general metrics carry over to decay spaces, with the resulting time complexity and approximation depending on ζ in the same way that the original results depends on the path loss term α . For distributed algorithms, that to date have appeared to necessarily depend on the planarity, we indicate how they can be adapted to arbitrary decay spaces at a cost in time complexity that depends on a *fading* parameter of the decay space. In particular, for decay spaces that are *doubling*, the parameter is constant-bounded.

Finally, we explore the dependence on ζ in the approximability of core problems. In particular, we observe that the capacity maximization problem has exponential upper and lower bounds in terms of ζ in general decay spaces. In Euclidean metrics and related growth-bounded decay spaces, the performance depends on the exact metricity definition, with a polynomial upper bound in terms of ζ , but an exponential lower bound in terms of a variant parameter ϕ . The upper bound result is the first approximation of a capacity-type SINR problem that is subexponential in α .

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

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 the author(s) 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.

PODC'14, July 15–18, 2014, Paris, France.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2944-6/14/07 ...\$15.00.

<http://dx.doi.org/10.1145/2611462.2611476>.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*; F.1.2 [Computation by Abstract Devices]: Modes of Computation—*Probabilistic computation*

General Terms

Algorithms, Design, Theory

Keywords

Wireless Networks, SINR, Capacity, Distributed Algorithms

1. INTRODUCTION

Signal-strength models of wireless communications capture the gradual fading of signals and the additivity of interference that are consistent with electro-magnetic theory. In spite of the apparent complexity of such models, various fundamental problems have been resolved analytically in recent years. These models also seem essential for studying certain properties of wireless networks, such as capacity [23], or connectivity and aggregation, which can be achieved in logarithmic rounds in worst case [48, 31].

Nearly all theoretic work in signal-strength models have been done in the GEO-SINR model that assumes that signals decay as a smooth polynomial function of distance. This assumption about decay (or *path-loss*) is true in free space, but turns out to be far off in actual environments, as shown by a long history of experimental studies (e.g., [45]). Quoting a recent meta-study, [3], “link quality is not correlated with distance.” This questions the wisdom of studying “SINR models” analytically, given the added effort and complexity.

One hope might be that results in the “basic SINR model” could eventually carry some insights that would be of use in more detailed models that capture more of reality. However, real environments consist of assortments of walls, ceilings and obstacles, as well as complex interactions involving reflections, shadowing, multi-path signals, and anisotropic (or even directional) antennas. It might seem near impossible to capture all of that without making the resulting models hopelessly impractical for algorithm design and/or analytically intractable.

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 instance (e.g., for simulations) and for average-case analysis. For handling actual instances, their utility necessarily depends on the suitability of the assumptions. The alternative view, with deep roots in computer science theory, is to allow for worst-case behavior and obtain guarantees that hold for *all* instances. The aim of this work is to make such “any-case” analysis feasible, while avoiding assumptions that may or may not be reflected in actual instances.

Our contributions. We present a simple solution that allows the modeling of arbitrary static situations by moving from geometry to arbitrary *decay spaces*. The decay between two ordered nodes is the relative reduction in the strength of a signal sent from the first node to the second. These decays can be obtained by signal-strength measurements – which almost any cheap widget can perform today – capturing the *truth on the ground*. The complexity of a setting is captured by a *metricity* parameter ζ that indicates how far the decay space is from satisfying the triangular inequality.

The simple metricity definition has wide-ranging implications: *All* results that hold in the SINR model in general metrics carry over to decay spaces; the resulting time complexity and approximation depending on ζ in the same way that the original results depends on the path-loss term α .

For distributed algorithms, that to date have appeared to necessarily depend on the planarity, we introduce a *fading* parameter of the decay space and indicate they can be adapted to arbitrary decay spaces at a cost in time complexity that depend on a *fading* parameter of the decay space. In particular, for decay spaces that are *doubling*, the parameter is constant-bounded.

Finally, we explore the dependence on ζ in the approximability of core problems. In particular, we observe that the CAPACITY problem has exponential upper and lower bounds in terms of ζ in general decay spaces. In Euclidean metrics and related growth-bounded decay spaces, the performance depends on the exact metricity definition, with a polynomial upper bound in terms of ζ , but an exponential lower bound in terms of a variant parameter ϕ .

One may ask if we are being led to yet another model that will later been shown unrealistic. Fortunately, numerous experimental studies have verified the remaining key assumptions in wide range of situations and technology [52, 47, 7, 51, 22]: additivity of interference, SINR capture effectiveness (the near-thresholding relationship between SINR level and packet reception rate), and invariability of wireless conditions in static environments. Thus, we may finally be reaching a wireless model that is a close approximation of the reality at the PHY/MAC layer, yet usable algorithmically and analytically. That said, one should not discount the value of abstractions or the potentially value of simple models. Also, modeling dynamic and mobile situations, which is outside the scope of our work, remains a highly important (and largely open) issue.

Related work. The “abstract SINR” model captures, like decay spaces, arbitrary pairwise path-loss. Some positive results hold in that model, e.g., distributed power assignment of feasible sets [46], reductions involving Rayleigh fading [8], and special cases of capacity maximization [26]. However, for most problems of interest, extremely strong inapproximability results hold [19, 44]. Thus, it is essential to use near-metric properties of the decay space.

The introduction of general metrics (apparently first in [15, 14]) was a significant step in extending SINR theory beyond geometric assumptions. Fading metrics [24] were identified to capture the main property required from the planar setting. The concept of *inductive independence* [43, 36] has heralded a more systematic approach to SINR analysis, and can by itself be seen as parameter of the decay space. Same holds for *C-independence* [1, 10] in the case of uniform power.

In a sibling paper [22], we introduced decay spaces and metricity with a focus on validation by measurements in two testbeds. We found that the metricity parameter did appear to reflect the “complexity” of the environment. We also examined the impact of the availability of multiple channels and found that this can significantly reduce the measured metricity. Additionally, we verified that the key properties of the SINR formula in decay spaces closely approximate reality, even in highly complex environments, in alignment with previous experimental results (e.g., [52, 47, 7, 51]).

Outline of the rest of the paper. In the next section, we introduce decay spaces (formal definitions and the metricity parameter), and indicate how previous results in metric spaces carry over. In Sec. 3, we address the core requirement of *fading* for distributed algorithms, introduce a parameter that extends their reach to arbitrary spaces, and prove constant upper bounds in spaces with bounded doubling dimension. The impact of metricity parameters on approximability is treated in Sec. 4, concluding with some observations and open issues in Sec. 5.

2. DECAY SPACES

2.1 Signal-strength models

The *abstract SINR* model has two key properties: (i) signal decays as it travels from a sender to a receiver, and (ii) interference – signals from other than the intended transmitter – accumulates. Transmission succeeds if and only if the interference is below a given threshold.

Formally, a *link* $l_v = (s_v, r_v)$ is given by a pair of nodes, sender s_v and a receiver r_v . The *channel gain* G_{uv} denotes the multiplicative decay in the signal of l_u as received at r_v . The *interference* I_{uv} of sender s_u (of link l_u) on the receiver r_v (of link l_v) is $P_u G_{uv}$, where P_u is the power used by s_u . When $u = v$, we refer to I_{vv} as the *signal strength* of link l_v . If a set S of links transmits simultaneously, then the *signal to noise and interference ratio* (SINR) at l_v is

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

where N is the ambient noise.

We refer to the standard signal-strength model as the GEO-SINR model, which adds to the SINR formula the assumption of *geometric path-loss*: that signal decays proportional to a fixed polynomial of the distance, *i.e.*, $G_{uv} = d(s_u, r_v)^{-\alpha}$, where the *path-loss term* α is assumed to be an arbitrary but fixed constant between 1 and 6. This assumption is valid with $\alpha = 2$ in free space and perfect vacuum [18, Sec. 3.1].

The last assumption made in theoretical models is *thresholding*: the transmission of l_v is *successful* iff $\text{SINR}_v \geq \beta$, where $\beta \geq 1$ is a hardware-dependent constant. We shall also make this assumption. It’s been shown by Dams,

Kesselheim and Hoefer [8] that certain models that include a randomized filter in this decision can be efficiently simulated by thresholding algorithms.

2.2 Metrics and Decay Spaces

We seek to model arbitrary path-loss that is independent of distance. We capture this by a *decay* function f of pairs of points (or nodes) so that $G_{uv} = 1/f(s_u, r_v)$.

We shall formulate signal decay as *decay spaces*. Decays between distinct points are always positive. Exactly what happens at a given point (i.e., the value of $f(p, p)$) is immaterial to our consideration, since we may assume that all nodes are distinct.

Definition 1. A *decay space* is a pair $\mathcal{D} = (V, f)$, where V is a discrete set of nodes (or points) and f is a mapping (or matrix) $f : V \times V \rightarrow \mathbb{R}_{\geq 0}$ that associates values (*decays*) with ordered pairs of nodes. The decays satisfy: i) $f(p, q) \geq 0$ (non-negativity), and ii) $f(p, q) = 0$ if and only if $p = q$ (the identity of indiscernibles).

Decay spaces need not be symmetric ($f(p, q) = f(q, p)$) nor obey the triangular inequality ($f(p, q) \leq f(p, r) + f(r, q)$). Such spaces are known as *pre-metrics*. As shorthand, we write $f_{pq} = f(p, q)$.

Decay space can either represent the truth-on-the-ground, or its representation/approximation as data. Besides signal-strength measurements (typically available with precision of 1dB, or 26%), they can often be inferred by packet reception rates or predicted by heuristic or environmental models [18].

Metricity

We introduced in [22] a parameter that represents how close the decay matrix is to a distance metric.

Definition 2. The *metricity* $\zeta(\mathcal{D})$ of a decay space $\mathcal{D} = (V, f)$ is the smallest number such that, for every triplet $x, y, z \in V$,

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

We note that 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 $f_{\min} = \min_{x \neq y} f(x, y)$. 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, $\zeta \leq \alpha$.

We define *quasi-distances* between nodes in a decay space by $d(p, q) = f_{pq}^{1/\zeta}$. Let $d_{pq} = d(p, q)$ for short. These quasi-distances induce a *quasi-metric* $\mathcal{D}' = (V, d)$, i.e., a metric except for the possible lack of symmetry. In the Euclidean setting, quasi-distances are simply the Euclidean distances.

2.3 Theory transfer

The lion share of the theoretic literature on signal-strength models can be converted to decay spaces with minimal effort. We aim here to clarify and substantiate that observation. Our objective is for the non-specialist to be able to determine with limited effort which results do hold in the decay model and which don't, and additionally, when the question arises, which properties of metric and/or decay spaces are necessary for correct functioning.

In this section, we focus on what is needed for results to hold in arbitrary decay spaces. In the following section, we deal with results that require special space properties,

particularly in the context of distributed algorithms. By a *result*, we mean a combination of an algorithm or a protocol and its analysis.

The complexity of a result can be a function of the metric/space. Here, complexity refers to measures like time and message count, but also performance measures like approximability. In particular, these measures have nearly always been functions of the metric parameters, such as the path-loss term α , but this dependence is often hidden in big-oh notation.

We make the following sweeping observation:

PROPOSITION 2.1. *If a GEO-SINR result only requires metric properties (triangular inequality), then it holds equally well in arbitrary decay spaces. Symmetry is required of the decay space only if it was required in the original setting. The relevant complexity measure (time, approximation) grows with ζ in the same manner as for the original result in terms of α .*

PROOF. The quasi-distances d of a decay space $\mathcal{D} = (V, f)$ form a quasi-metric $\mathcal{D}' = (V, d)$, which becomes a metric iff $\mathcal{D}$ satisfies symmetry. Applying the original result to the metric $\mathcal{D}'$ with path loss constant $\zeta(\mathcal{D})$ gives an equivalent solution to the problem on the decay space $\mathcal{D}$. $\square$

For example, the following results on the following problems carry over without change: capacity maximization [27, 41], scheduling [14, 15], weighted capacity [24, 30], spectrum auctions [36, 35], relationship between power control regimes [53, 25], dynamic packet scheduling [42, 25], distributed scheduling [43, 32], and distributed capacity maximization with regret-minimization [1].

We remark that the resulting algorithms actually become *simpler* in decay spaces than in the standard setting. Namely, there is no longer any need to use or know the path-loss constant α . No known algorithms actually need ζ ; it is used purely for analysis purposes.

We can also make an immediate observation regarding methods that hold for restricted metrics.

Observation 1. If a result holds in GEO-SINR for a given class $\mathcal{M}$ of (quasi-)metrics, then it holds equally in those decay spaces whose induced (quasi-)metric is contained in $\mathcal{M}$.

Results that do not carry over to decay spaces. There remains a significant body of work in GEO-SINR that depends on *positions* (or distributions thereof). Such results are necessarily tied to geometry, although with some work it may be possible to extend them to other decay spaces.

A common use of positional information is by partitioning the plane, so as to make simultaneous communication non-conflicting. This is particularly an issue for deterministic distributed algorithms. Examples of this include deterministic distributed broadcast [38, 39] and local broadcast [37, 16]. Also, some centralized approximation algorithms and heuristics for CAPACITY and SCHEDULING of [21, 58]. Occasionally, angles are used, e.g. [19], which does not carry over (but see Sec. 4.1).

There is also a large literature on average case analysis, typically assuming a uniform distribution of points in the plane, starting with an influential paper of Gupta and Kumar [23] that first introduced GEO-SINR.

Finally, SINR diagrams [2] (and follow-up work of subsets of the authors) use intrinsically topological properties of Euclidean metrics.

2.4 Additional definitions: Power, affectance, separability

We will work with a total order $\prec$ on the links, where $l_v \prec l_w$ implies that $f_{vv} \leq f_{ww}$. A power assignment $\mathcal{P}$ is *monotone* if both $P_v \leq P_w$ and $\frac{P_w}{f_{ww}} \leq \frac{P_v}{f_{vv}}$ hold whenever $l_v \prec l_w$.¹ This captures the main power strategies, including uniform and linear power.

We modify the notion of *affectance* [19, 33, 43]: The affectance $a_w^{\mathcal{P}}(v)$ 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

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

where $c_v = \frac{\beta}{1 - \beta N / (P_v G_{vv})} > \beta$ is a factor 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 $a_v(S) = \sum_{l_w \in S} a_w(v)$ be the *out-affectance* of v on S and $a_S(v) = \sum_{l_w \in S} a_w(v)$ be the *in-affectance*. Assuming S contains at least two links we can rewrite Eqn. 1 as $a_S(v) \leq 1$ and this is the form we will use. A set S of links is *feasible* if $a_S(v) \leq 1$ and more generally *K-feasible* if $a_v(S) \leq 1/K$.

Define $d_{vv} = \min(d(s_v, r_w), d(s_w, r_v), d(s_v, s_w), d(r_v, r_w)) = d(l_v, l_w)$ as the (quasi-)distance between two links l_v and l_w . Let $d_{vv} = d(s_v, r_v)$. A link l_v is said to be η -separated from a set L of links, for parameter η , if $d(l_v, l_w) \geq \eta d_{vv}$ for every $l_w \in L$. A set L is η -separated if each link in L is η -separated from the rest of the set.

Let e refer to the base of the natural logarithm and let $\lg$ denote the base-2 logarithm. Recall that $1 + x \leq e^x$, for any value x .

3. FADING PROPERTIES AND DISTRIBUTED ALGORITHMS

In the study of distributed algorithms in GEO-SINR in the plane, the standard assumption is that the path-loss constant α is strictly larger than 2. The reason is that this ensures that spatially well-separated nodes cannot affect each other too much, a property that does not hold when $\alpha \leq 2$. This property is generalized to doubling metrics whose doubling dimension is strictly smaller than the path-loss constant α , dubbed *fading metrics* [24]. We call this property, that the sum of affectances from spatially separated transmitting nodes converges, the *fading* property. For the most common type of distributed algorithm to work, this has to be bounded.

We define a parameter γ that captures the fading effect. Let $\mathcal{X}(r)$ be the space of all r -separated subsets in V .

Definition 3. The *fading value* $\gamma_z(r)$ of a node z relative to a separation term r is

$$\gamma_z(r) = r \max_{X \in \mathcal{X}(r)} \sum_{x \in X} 1/f_{xz}.$$

¹This corresponds to *length monotone* and *sub-linear* power assignments in GEO-SINR.

The *fading parameter* γ of a decay space is the maximum fading value of a node in the space, $\gamma = \gamma(r) = \max_{z \in V} \gamma_z(r)$, relative to a given separation term r .

That is, the total interference $I_S(z)$ experienced by a node z from an r -separated set S (of senders) using uniform power P is at most $\gamma(r) \cdot P/r$. Thus, if the intended signal has decay f_{zz} at most r , then the resulting affectance is bounded by $a_S(z) \leq \frac{\gamma(r)P/r}{P/r} = \gamma(r)$.

Many distributed problems (e.g., broadcast, local broadcast) are primarily concerned with neighborhoods of particular size/radius; that radius then defines the separation term r and the fading value $\gamma(r)$ in consideration. Other problems (e.g., aggregation, capacity) are concerned with all possible distance (in a sense, “scale free”); in that case, we shall be interested in the fading term $\gamma = \max_r \gamma(r)$.

Until now, γ has been expected to be an absolute constant. However, we can now simply treat it as a parameter and thus handle arbitrary decay spaces by distributed algorithms, thus achieving significant increase in generality. The value of γ must naturally factor into the time complexity of the algorithms.

3.1 Fading spaces

We identify a large class of decay spaces for which the fading parameter is small. These are generalizations of fading metrics.

First, some additional notation. The t -ball $B(y, t) = \{x \in V \mid f(x, y) < t\}$ centered at y with radius t contains all points x for which decay to y is less than t . A set $Y \subseteq V$ is a t -packing if $f(x, y) > 2t$, for any $x, y \in Y$. Thus, Y is a t -packing iff the set $\{B(y, t)\}_{y \in Y}$ of balls are disjoint. The t -packing number $\mathcal{P}(\mathcal{B}, t)$ is the size of the largest t -packing into the body $\mathcal{B}$.

Intuitively, a space is *doubling* if the number of mutually unit-separated points within a given distance from a center increases by at most a polynomial of the distance.

Definition 4. Let $\mathcal{D} = (V, f)$ be a decay space. Define $g_{\mathcal{D}}(q) = \max_{x \in V} \max_{r \in \mathbb{R}^+} \mathcal{P}(B(x, r), r/q)$, as the size of the densest q -packing in $\mathcal{D}$. The *Assouad (doubling) dimension* A of $\mathcal{D}$ with parameter C is given by

$$A(\mathcal{D}) = \max_q \log_q \left(\frac{g(q)}{C} \right).$$

$A(\mathcal{D})$ is, in effect, the minimum k for which sizes of t -packings can be bounded by $O(t^k)$, for all t . Note that that $A(\mathbb{R}^k) = k$ [34].

Definition 5. A *fading space* is a decay space $\mathcal{D}$ with Assouad dimension strictly smaller than 1, $A(\mathcal{D}) < 1$, w.r.t. some absolute constant C .

Note that the (quasi-)distance metric of a decay space of doubling dimension A has doubling dimension $A \cdot \zeta$.

3.2 Annulus argument

Most randomized algorithms (e.g. in [4] and [62]) ensure that in any given neighborhood (defined as the set of nodes to which a given node can communicate directly), the expected number of transmissions in a slot is bounded above by a certain constant. This ensures that the total expected affectance from other nodes transmitting is also bounded by

a (different) constant. By adjusting the constants appropriately, one can focus only on the local behavior. Some deterministic algorithms similarly ensure a spatial separation of sending (and thus possibly interfering) nodes and use this property to bound the total affectance from these nodes.

All proofs of the discussed sort use a common approach. They define some type of separation between interfering nodes which can be a (probabilistic) constant density, a hard minimum distance between nodes or links or similar. Then the interference at a node v is bounded, either directly or (if the node is the receiver of a predefined link) in terms of the (possibly probabilistic) affectance on the node. To do this we draw concentric circles around v , cutting the space around v up into annuli. Using the separation of the interferers, we argue that the number of interferers that can be packed in the annulus at distance i is bounded by a polynomial depending on i and the Assouad dimension of the space.

We argue that a general version of this ‘annulus argument’ still holds when directly used in fading decay spaces, after which we indicate how other different variations carry over.

Recall the Riemann ζ -function, $\zeta(x) = \sum_{n \geq 1} n^{-x}$, which is known to converge for $x > 1$. We modify an earlier similar result for metric spaces [24].

THEOREM 3.1. *The fading parameter γ of a decay space $\mathcal{D} = (V, f)$ is bounded by $\gamma \leq C2^{A+1}(\zeta(2-A) - 1)$, where $A < 1$ is the Assouad dimension and C is the related constant.*

PROOF. Let r be a value and $R = r/2$. Let S be an R -packing. Since $\mathcal{D}$ is doubling, there is a constant C such that for any $t > 0$, the maximal size of an R -packing in a ball of radius tR centered around a point x is,

$$\mathcal{P}(B(x, tR), R) \leq Ct^A. \quad (3)$$

We bound the received signal $I_S(x)$ at a listening node $x \in S$. Let g be a natural number. Let $S_g = \{y \in S : f(y, x) < gR\}$ and let $T_g = S_g \setminus S_{g-1}$. Then $S_2 = \emptyset$ since S is r -spaced.

We first note that since $S_{g-1} \subseteq S_g$ and $S_2 = \emptyset$,

$$\begin{aligned} \sum_{g \geq 3} \frac{|S_g \setminus S_{g-1}|}{g-1} &= \sum_{g \geq 3} \frac{|S_g|}{g-1} - \sum_{g \geq 2} \frac{|S_g|}{g} \\ &= \sum_{g \geq 3} |S_g| \left(\frac{1}{g-1} - \frac{1}{g} \right) = \sum_{g \geq 3} \frac{|S_g|}{g(g-1)}. \end{aligned}$$

Since each sender $y \in T_g$ is of distance at least $(g-1)R$ from x the received signal from y on x is bounded by

$$I_y(x) = P/f(y, x) \leq \frac{P}{(g-1)R} \quad \forall y \in T_g.$$

Then,

$$\begin{aligned} I_S(x) &= \sum_{g \geq 3} I_{T_g}(x) \leq \sum_{g \geq 3} |S_g \setminus S_{g-1}| \frac{P}{(g-1)R} \\ &\leq \frac{P}{R} \sum_{g \geq 3} |S_g| \frac{1}{(g-1)^2}. \end{aligned}$$

By the doubling property of $\mathcal{D}$, the size of S_g is

$$|S_g| \leq \mathcal{P}(B(x, (g+1)R), R) \leq C(g+1)^A.$$

Thus, using that $g+1 \leq 2(g-1)$, since $g \geq 3$,

$$\frac{|S_g|}{(g-1)^2} \leq \frac{C(g+1)^A}{(g-1)^2} = \frac{C2^A}{(g-1)^{2-A}}.$$

Continuing,

$$\begin{aligned} I_S(x) &\leq \frac{P}{R} \sum_{g \geq 3} |S_g| \frac{1}{(g-1)^2} \leq \frac{2P}{r} \sum_{g \geq 3} \frac{C2^A}{(g-1)^{2-A}} \\ &\leq \frac{2P}{r} C2^A (\zeta(2-A) - 1) = \frac{\gamma(r) \cdot P}{r}, \end{aligned}$$

using the definitions of R and $\gamma(r)$. $\square$

3.3 Common usage of the annulus argument

We list some common types of lemmas in which the annulus argument is used and show how to use Theorem 3.1 in the proofs for these lemmas.

A common usage of the annulus argument is to prove the following: if L is a set of links, using a uniform power assignment P , with senders of a minimal mutual distance r and with the longest link of length at most a given constant times r , then L forms a q -feasible set. For sets as described in Theorem 3.1, where all nodes are r -separated and a maximum link decay f_{vv} at most constant r , the transition is straightforward. By the definition of affectance and Theorem 3.1, the affectance of L on link l_v with maximum decay f_{vv} is at most

$$a_L(v) \leq \frac{I_L(v)}{PG_{vv}} \leq \frac{f_{vv} \cdot \gamma(r)}{r},$$

where $I_L(v) = \sum_{l_w \in L} P/f_{wv}$. To obtain a q -feasible set, we simply set $r = f_{vv}\gamma(r)/q$.

However, if only a separation on senders is defined (e.g. [24]), we use the triangular inequality to bound the interference at r_v in terms of interference at s_v . Requiring $f_{vv} < R$, we obtain $I_L(r_v) \leq 2^\zeta I_L(s_v)$, since for any sender $s_x \in L$ by the triangle inequality

$$f(s_x, r_v)^{1/\zeta} \geq f(s_x, s_v)^{1/\zeta} - f(s_v, r_v)^{1/\zeta} \geq f(s_x, s_v)^{1/\zeta}/2,$$

using that $f_{vv} < R \leq f(s_x, s_v)^{1/\zeta}/2$. And thus $f_{xv} \geq f(s_x, s_v)/2^\zeta$, hence the argument holds as before by adjusting r with an extra 2^ζ factor. When $R \gg f_v$, the overhead factor is correspondingly smaller.

Examples of problems with centralized algorithms that use this form of annulus argument include: connectivity [48, 31], scheduling [5, 54], flow-based throughput [6], online capacity maximization [13], and bounds on the utility of conflict graphs [55].

For randomized algorithms, the annulus argument is used in a similar way to bound expected interference. The expected interference in a disk is bounded by arguments specific to the analyzed algorithm. These arguments may or may not translate to the decay space as discussed in Sec. 2.3. Instead of adjusting the separation term r , they typically adjust the transmission probabilities. Once the expected interference in a disk is bounded, however, the argumentation for bounding the total expected interference at a node x , $E(I_S(x))$, follows Theorem 3.1.

The probabilistic version of the annulus argument forms the core of the analysis for many randomized distributed algorithms which often carry over without any significant further adjustments. Example include (distributed) coloring

[61], local broadcast [20, 62, 29], broadcast [11] and multiple-message broadcast [59, 60], capacity [49], dominating set [50], and (multihop) connectivity [28, 4].

3.4 Beyond fading spaces

Fading spaces do not completely characterize spaces with a bounded fading parameter. One reason is that the definition of doubling metrics is scale-invariant in that the packing constraint holds for balls of any size, whereas we are often only interested in balls of a fixed size (or in a limited range of sizes).

Consider, for instance, the metric space formed by a star centered at node x_0 with k leaves $x_1, x_2, \dots, x_k$ at distance k^2 and one leaf x_{-1} at distance r , where $r \ll k$. Suppose the decay f_{xy} equals the distance (so $\zeta = 1$). The doubling dimension of this space is k , so unbounded. Suppose also that we are interested how well we can transmit from x_0 to x_{-1} in the presence of transmissions from the other nodes; that is, what matters is fading w.r.t. separation term r . If $r = o(k)$, we find that the total interference at node x_{-1} is $\sum_{i=1}^k 1/k^2 = 1/k$, which is asymptotically smaller than the signal received from x_0 . Thus, the fading property holds even if the metric is not doubling.

4. DEPENDENCE ON THE METRICITY IN APPROXIMATIONS

Having pinpointed the metricity parameter ζ as a key indicator of a decay space, the question arises how it affects the complexity of fundamental problems. This differs from GEO-SINR where the path-loss term α has traditionally been viewed as a constant.

We explore here the approximability of the CAPACITY problem as a function of innate properties of the decay space in question. Given a set L of links, the CAPACITY problem asks for maximum cardinality subset of L that is feasible. The CAPACITY problem is fundamental, not only because it addresses the basic question of how much wireless communication can coexist, but also because it has been the underlying core routine in other problems, including scheduling [19], throughput maximization (via flow) [56], spectrum auctions [36], spectrum sharing [30], and connectivity and aggregation [31, 28].

Our generic statement, Prop. 2.1, along with known approximation results [27, 41] in general metrics, implies that CAPACITY in decay spaces can be approximated within a function of ζ . Specifically, the approximation of [27] (for monotone power) is exponential in ζ , which was refined to 3^ζ in [22].

We can also observe that the known hardness construction for “abstract SINR” [19] (see also [27]) implies that $2^{\zeta(1-o(1))}$ -approximation for CAPACITY is computationally hard.

THEOREM 4.1 ([19, 27]). *CAPACITY of equi-decay links is hard to approximate within $2^{\zeta(1-o(1))}$ factor. This holds even if the algorithm is allowed arbitrary power control against an adversary that uses uniform power.*

This leaves the question whether better results are possible in the Euclidean metric and comparable decay spaces. Surprisingly, the answer depends on the exact definition of the metricity parameter. Specifically, CAPACITY with uniform power is then approximable within a polynomial of ζ ,

while for a natural variant of the ζ -parameter, exponential dependence is still necessary.

4.1 Improved Approximations in Bounded Growth Decay Spaces

We show here that CAPACITY with uniform power can be approximated within polynomial factors of ζ in Euclidean metrics. More generally, this holds for decay spaces of bounded growth, as we shall define shortly. Interestingly, it does not rely on the fading behavior of the plane (i.e., that $\alpha > 2$). This appears to be the first instance in the signal-strength literature where better results are shown to be obtainable in the plane independent of α than for general metrics.

The intuitive reason why uniform power in the plane proves to be easier is two-fold. The main cause for exponential dependence on ζ comes from the use of the triangular inequality. If one can ensure that one angle is highly acute, the overhead of the inequality goes down accordingly. In particular, the overhead in switching the reference from a receiver to a sender of a link goes down if the length of the link relative to the other distances is small.

We shall allow any positive ζ (or α); for this purpose, we define $\underline{\zeta} = \max(1, \zeta)$ and $\underline{\alpha} = \max(1, \alpha)$.

Bounded Growth Decay Spaces. We shall consider decay spaces that have upper bounds on two measures that restrict growth: the doubling dimension (from Sec. 3), and the independence dimension, defined in [19] for metrics and adapted as follows to decay spaces.

Definition 6. ([19]) A set I of points in a decay space $\mathcal{D} = (V, f)$ is *independent* w.r.t. a point $x \in V$ if $B(z, f_{zx}) \cap I = \{x\}$ for each $z \in I$. The *independence dimension* of $\mathcal{D}$ is the size of the largest independent point set.

Spaces of bounded independence dimension D have the following useful property: for any point $x \in V$, there is a set $J_x \subset V$ of at most D points that *guard* x in the following sense: $\min_{y \in J_x} d(z, y) \leq d(z, x)$, for any point $z \in V \setminus \{x\}$. A node y *guards* node x from node z if $d(z, y) \leq d(z, x)$.

Welzl [57] has made a number of useful observations of metrics of bounded independence dimension. He showed that the number of guards needed in a metric is indeed exactly its independence dimension. In a Euclidean space $\mathbb{R}^n$, it equals the maximum number of unit vectors that form pairwise angles of more than 60° . Therefore, the independence is at most the so-called kissing number, the maximum number of disjoint open balls of radius 1 that can touch the unit ball. This number grows exponentially in the dimensions but its exact value is not known for most dimensions.

As a simple example, let us see how six guards suffice in the plane. Given a point x , divide the plane into six 60° sectors around x and partition V accordingly into sets $S_1, S_2, \dots, S_6$. Let J_x consist of the nearest point to x in each of the six sectors. The guarding property follows from the fact that the angle $\angle g_i x y_i$ is at least 60° , for each point $y_i \in S_i$ and guard $g_i \in J_x$.

We define a decay space to be *bounded-growth* if it has bounded independence dimension and bounded doubling dimension.

The doubling and independence dimensions are actually incomparable. The uniform metric, with identical unit decays, is of independence dimension 1 but unbounded doubling dimension. The following curious construction of Welzl

[57] gives a metric of doubling dimension 1 whose independence dimension is unbounded: Let $V = \{v_{-1}, v_0, v_1, \dots, v_n\}$ with $d(v_{-1}, v_i) = 2^i - \epsilon$, for $0 < \epsilon \leq 1/4$, and $d(v_j, v_i) = 2^i$, for $i, j \neq -1, j < i$. We leave it to the curious reader to verify that any ball (only those of radius 2^i or $2^i - \epsilon$ matter) can be covered with two balls of half the radius and that $V \setminus \{v_{-1}\}$ are independent with respect to v_{-1} .

Amicability. The following definition originates in [10] (with roots in [1]) as *C-independent conflict graphs*.

Definition 7. A set L of links is *h-amicable* if, for any feasible subset $S \subseteq L$, there is a subset $S' \subseteq S$ with $|S'| \geq 3|S|/4$ such that for any vertex $v \in S$, $a_v(S') \leq h$ (using uniform power).

It is known that instances of GEO-SINR in metric spaces are $2^{O(\alpha)}$ -amicable [1]. Various decentralized capacity-type problems with uniform power have been treated with no-regret minimization techniques, relying only on the amicability property of the instances [12, 1, 9, 17, 10].

To bound amicability, we first show how to turn feasible sets in doubling spaces into well-separated sets at limited cost.

We shall make use of the following technique.

LEMMA 4.2 (SIGNAL-STRENGTHENING [33]). *There is a polynomial-time algorithm that, for any given p, q , partitions any p -feasible set into $\lceil 2q/p \rceil^2$ sets, all q -feasible.*

We first argue that feasible sets under uniform power must be at least $1/\zeta$ -separated, independent of metric.

LEMMA 4.3. *Let S be an e^2/β -feasible set of links under uniform power. Then, S is $1/\zeta$ -separated.*

PROOF. Suppose otherwise. Then, there are two links l_v, l_w in S that are not $1/\zeta$ -separated. There are three cases, depending on which pairwise distance bound is violated.

Consider first the case when $d(s_v, r_w) < (1/\zeta) \max(d_{vv}, d_{ww})$. Since the two links co-exist in a $1/\beta$ -feasible set, $a_v(w) = c_w f_{ww}/f_{vw} \leq \beta$, and since $c_w \geq \beta$, it follows that $P/f_{ww} \geq P/f_{vw}$, i.e., the signal received by r_w from s_w is at least as strong as that from the other sender s_v . So, $d_{ww} \leq d(s_v, r_w)$, implying that $d(s_v, r_w) < (1/\zeta)d_{vv}$. Then, by the triangular inequality and these bounds,

$$\begin{aligned} d(s_w, r_v) &\leq d_{ww} + d(r_w, s_v) + d_{vv} \leq 2d(s_v, r_w) + d_{vv} \\ &< (1 + 2/\zeta)d_{vv}. \end{aligned}$$

Thus, $f_{vv} < (1 + 2/\zeta)^\zeta f_{vv} \leq e^2 f_{vv}$. It follows that

$$a_w(v) = c_v \frac{f_{vv}}{f_{vw}} \geq \frac{c_v}{e^2} > \frac{\beta}{e^2}.$$

This contradicts the assumption that l_v and l_w coexist in the same e^2/β -feasible set.

Consider next the case when $d(r_v, r_w) < (1/\zeta) \max(d_{vv}, d_{ww})$. Without loss of generality, assume $d(r_v, r_w) < d_{vv}/\zeta$. By the triangular inequality, $d_{vv} \leq d_{vv} + d(r_v, r_w) < d_{vv}(1 + 1/\zeta)$, implying that $f_{vv} < (1 + 1/\zeta)^\zeta f_{vv} \leq e \cdot f_{vv}$, leading to a contradiction as before. Finally, the case when $d(s_v, s_w) < \max(d_{vv}, d_{ww})$ is symmetric to the previous one when swapping senders and receivers. Hence, the claim. $\square$

We next show that in doubling metrics, the separation factor can be expanded at the cost of a polynomial factor.

LEMMA 4.4. *Let τ and η be positive parameters, $\tau < \eta$. Let S be a τ -separated set of links in a decay space whose quasi-distance metric has doubling dimension A' . Then, S can be partitioned into $O((\eta/\tau)^{A'})$ sets, each of which is η -separated.*

PROOF. Consider a link l_v in S . Let $R_v = \{r_u : \exists l_u = (s_u, r_u) \in S, d(l_u, l_v) \leq \eta d_{vv}, d_{uu} \geq d_{vv}\}$ be the set of links in S that are at least as long as l_v and whose receivers are within distance ηd_{vv} from r_v . We first bound the cardinality of R_v .

Draw a ball of radius $\tau d_{vv}/2$ around each receiver in R_v . Since S is τ -separated, and the links lengths are at least d_{vv} , the $|R_v|$ balls are disjoint; also, they are properly contained in a ball of radius of $(\eta + \tau/2)d_{vv}$ (around r_v). By the definition of the Assouad dimension,

$$|R_v| \leq C \left(\frac{\eta + \tau/2}{\tau/2} \right)^{A'} = C ((2\eta + 1)/\tau)^{A'}. \quad (4)$$

We now form the graph $G_S = (V, E)$, where $V = S$ and $(l_v, l_w) \in E$ iff $l_v \in R_w$ or $l_w \in R_v$. Let $\rho = \max_{l_v \in S} |R_v| \leq C((2\eta + 1)/\tau)^{A'} = O((\eta/\tau)^{A'})$. Form a total order $\prec$ on the nodes by non-increasing link length. By (4), each node has at most ρ neighbors that follow it in the ordering (because if $d_{vv} < d_{ww}$ then $l_v \notin R_w$). That is, $\prec$ is a ρ -inductive (or, ρ -degenerate) ordering of G . Coloring the graph first-fit according to $\prec$ uses at most $\rho + 1$ colors. To complete the proof, we observe that a set of links is η -separated if and only if the corresponding set of vertices in the graph is independent (graph-theoretically). $\square$

Put together, we obtain a sparsity-strengthening lemma in doubling spaces. Recall that $\zeta = \max(1, \zeta)$.

LEMMA 4.5. *Let S be a feasible set of links in a decay space whose quasi-distance metric has doubling dimension A' . Then, S can be partitioned into $O(\zeta^{2A'})$ sets, all of which are ζ -separated.*

PROOF. Recall that by the signal strengthening Lemma 4.2, S can be partitioned into at most $(2e^2/\beta + 1)^2$ sets each of which is e^2/β -feasible. Let S' be such a set. By Lemma 4.3, S' is $1/\zeta$ -separated. If $\zeta \leq 1$, we are done. Otherwise, $\zeta = \zeta$ and by Lemma 4.4, S' can be partitioned into $O(\zeta^{2A'})$ sets, each of which is ζ -separated. $\square$

We are now ready to prove the structural result of this section.

THEOREM 4.6. *Let L be a set of links in a decay space of independence dimension D and whose quasi-distance metric has doubling dimension A' . Then, L is $O(D\zeta^{2A'})$ -amicable.*

PROOF. Let $S \subseteq L$ be a feasible subset of L . By Lemma 4.5, there is a partition of S into ζ -separated sets $X_1, X_2, \dots, X_s$, with $s = O(\zeta^{2A'})$. Let i be an index, $1 \leq i \leq s$. Let $Y_i = \{l_v : a_v(X_i) \leq 4\}$ be the subset of links in X_i with low out-affectance. Note that $\sum_{l_v \in X_i} a_v(X_i) = \sum_{l_v \in X_i} a_{X_i}(v) \leq |X_i|$, by feasibility, so the average out-affectance of links in X_i is at most 1. By Markov's inequality, at least three fourths of the links in X_i will have out-affectance at most four times the average. Thus, $|Y_i| \geq 3|X_i|/4$.

Consider any link $l_v \in L$. Let $J_v = \{g_1, g_2, \dots, g_t\}$ be the indices of senders in $|Y_i|$ that guard the sender s_v of l_v ,

where $t \leq D$. Partition Y_i into sets $S_1, S_2, \dots, S_t$, where s_{g_j} is contained in S_j and guards s_v from the senders of other links in S_j . Consider any set S_j and let l_x be a link in S_j . Since s_{g_j} guards s_v from s_x , $d(s_{g_j}, s_x) \leq d(s_v, s_x)$. Then, additionally using the triangular inequality and that S_j is ζ -separated,

$$d(s_{g_j}, s_x) \leq d(s_v, s_x) \leq d_{vx} + d_{xx} \leq (1 + 1/\zeta)d_{vx}.$$

So, $f(s_{g_j}, s_x) = d(s_{g_j}, s_x)^\zeta \leq (1 + 1/\zeta)^\zeta f_{vx} \leq e \cdot f_{vx}$. In a similar way, we obtain that $d_{g_j x} \leq d(s_{g_j}, s_x) + d_{xx} \leq (1 + 1/\zeta)d(s_{g_j}, s_x)$, so $f_{g_j x} \leq e \cdot f(s_{g_j}, s_x)$. Combining, we get that $f_{g_j x} \leq e \cdot f(s_{g_j}, s_x) \leq e^2 f_{vx}$. We can then bound the out-affectance of l_v on S_j by

$$\begin{aligned} a_v(S_j) &= \sum_{l_x \in S_j} a_v(x) = \sum_{l_x \in S_j} c_x \cdot \frac{f_{xx}}{f_{vx}} \\ &\leq a_v(g_j) + \sum_{l_x \in S_j \setminus \{l_{g_j}\}} c_x \cdot \frac{e^2 \cdot f_{xx}}{f_{g_j x}} \\ &\leq 1 + e^2 \cdot a_{g_j}(S_j) \leq 1 + 4e^2, \end{aligned}$$

using the definition of Y_i in the last inequality. Then, $a_v(Y_i) = \sum_j a_v(S_j) \leq (1 + 4e^2)D$. Let $S' = \cup_i Y_i$ and recall that $|S'| = \sum_i |Y_i| \geq \frac{3}{4} \sum_i |X_i| = \frac{3}{4}|S|$. Finally, $a_v(S') = \sum_i a_v(Y_i) \leq (1 + 4e^2)D \cdot O(\zeta^{2A'})$. Thus, L satisfies the definition of amicability with $h = O(D\zeta^{2A'})$. $\square$

This bound on amicability directly implies improved approximation for numerous distributed problems (from the previous, usually implicit, $\exp(O(\alpha))$ bounds). Note that for the plane $A' = 2$.

COROLLARY 4.7. *There is a distributed algorithm based on no-regret learning for CAPACITY with $O(\alpha^4)$ -approximation in the plane [12, 1]. This holds also in the presence of jamming [9], changing spectrum availability [10] and for online requests against stochastic adversaries [17].*

4.2 Inapproximability results for a variant of metricity

Metricity variant φ . Various alternative measures of the metric-like behavior of a space $\mathcal{D} = (V, f)$ can be concocted. A particularly natural one is the parameter φ that bounds the *multiplicative* factor within which f satisfies a relaxed triangular inequality:

$$\varphi = \max_{x, y, z \in V} \frac{f_{xy} + f_{yz}}{f_{xz}}.$$

So, φ is the smallest value such that $f_{xz} \leq \varphi(f_{xy} + f_{yz})$, for every $x, y, z \in V$. For comparison with ζ , we define $\phi = \lg \varphi$.

Examining the proofs of the various results for CAPACITY and *inductive independence* [36], we find that the triangular inequality is applied to compare lengths that are within constant factor of each other, in which case the overhead is comparable to the case of ζ . Thus, the results hold also in terms of ϕ .

Observation 2. CAPACITY, both with monotone power [27, 22] and arbitrary power control [40], is approximable within $2^{O(\phi)}$. Other results with effective (exponential) approximations in terms of similar bounds hold for inductive independence [36, 25] and relationships between power control and monotone power [25].

Bounds on inductive independence also have numerous implications, including connectivity and aggregation [31, 28], spectrum auctions [36, 35], dynamic packet scheduling [42], and distributed scheduling [43, 32].

We can observe that $\zeta \leq \phi$. Namely, for any nodes x, y, z , $f_{xz}^{1/\zeta} \leq f_{xy}^{1/\zeta} + f_{yz}^{1/\zeta} \leq 2 \max(f_{xy}^{1/\zeta}, f_{yz}^{1/\zeta}) = 2(\max(f_{xy}, f_{yz}))^{1/\zeta} \leq 2(f_{xy} + f_{yz})^{1/\zeta}$, using the definition of ζ . Thus, $f_{uv} \leq 2^\zeta(f_{uw} + f_{vw})$. Hence, lower bounds in terms of ζ carry over to lower bounds in terms of $\phi = \log \varphi$, so exponential approximations in terms of ϕ are best possible in general metrics.

A converse relation between ζ and ϕ does not exist, however. Consider the instance on three points $V = \{a, b, c\}$ with $f_{ab} = 1$, $f_{bc} = q$ and $f_{ac} = 2q$, with $q \gg 1$. Then, one can verify that $\phi \leq 2$, while $\zeta = \theta(\log q / \log \log q)$, which is unbounded.

Hardness in terms of φ . We find that CAPACITY in bounded-growth spaces is still exponentially hard in terms of ϕ . We give a construction that is embedded on a pair of lines, that holds for arbitrary values of a parameter α . For decays within the lines, it uses the usual distance function raised to power α , while between the lines, it uses two fixed decays: $n^{\alpha-1}$ and n^α . It then also shows that strong hardness holds even when none of the decay functions are particularly fast growing.

THEOREM 4.8 ([19]). *CAPACITY of equi-decay links in bounded-growth decay spaces is hard to approximate within $2^{\phi(1-o(1))}$ -factor. This holds even if the algorithm is allowed arbitrary power control against an adversary that uses uniform power.*

PROOF. By reduction from the maximum independent set problem in graphs. Let α be arbitrary value satisfying $\alpha > 1$, denoting the maximum path-loss term and let $\alpha' = \alpha - 1$. Assume for simplicity that $N = 0$ and $\beta = 1$. Let $d_2(\cdot)$ refer to the standard Euclidean distance.

Given graph $G = (V, E)$, form a set L of links with link $l_i = (s_i, r_i)$ for each vertex $v_i \in V$ located in the plane. The senders are located on the vertical line segment $[(0, 0), (0, n)]$ and the receivers on the segment $[(n, 0), (n, n)]$: s_i at point $(0, i)$ and r_i at point (n, i) .

Decays between points on the same line (both senders or both receivers) are set to their distance to the power of α' . For decays between points on different lines, we use two fixed decays: $n^{\alpha'}$ and $n^{\alpha'+1}$. Formally, for links l_i and l_j , let

$$f_{ij} = f(s_i, r_j) = \begin{cases} d_2(s_i, r_j)^{\alpha'} = n^{\alpha'} & \text{if } i = j \\ n^{\alpha'} - \delta & \text{if } v_i v_j \in E \\ n^{\alpha'+1} & \text{if } v_i v_j \notin E, \end{cases}$$

where $0 < \delta < 1/2$. Also, let $f(s_i, s_j) = f(r_i, r_j) = |i - j|^{\alpha'}$.

With uniform power P , we have that for each $i \neq j$,

$$a_i(j) = \frac{P/f_{ij}}{P/f_{jj}} = \frac{n^{\alpha'}}{f_{ij}} \begin{cases} > 1 & \text{if } v_i v_j \in E \\ \leq 1/n & \text{if } v_i v_j \notin E. \end{cases}$$

Hence, a set $S \subset L$ of links is feasible iff $V_S = \{v_i \in V : l_i \in S\}$ is an independent set.

For the case of power control, consider a pair of links l_v, l_w and let $\mathcal{P}$ be any power assignment on the links. If $(v, w) \in E$, then $f_{vw} \cdot f_{vw} = (n^{\alpha'} - \delta)^2$, which implies that

$$a_v^{\mathcal{P}}(w) \cdot a_w^{\mathcal{P}}(v) \geq \beta^2 \frac{f_{vv} \cdot f_{ww}}{f_{vw} \cdot f_{vw}} = \beta^2 \frac{n^{2\alpha'}}{(n^{\alpha'} - \delta)^2} > \beta^2 = 1.$$

So, at least one of $a_v^{\mathcal{P}}(w)$ and $a_w^{\mathcal{P}}(v)$ must be greater than one, implying that no power assignment allows l_v and l_w to be simultaneously feasible. Hence, any feasible set S must correspond to an independent set in G , and we know that any independent set in G can be made feasible in L using uniform power. Solutions to CAPACITY on L are therefore in one-one correspondence with solutions to MAX INDEPENDENT SET on G , preserving solution size.

Regarding φ , observe that $f(s_i, s_j) = f(r_i, r_j) \geq 1$. Then, we can verify that for any triplet x, y, z of points used in L ,

$$f_{xz} \leq 2n \max(f_{xy}, f_{yz}) .$$

Thus, $\varphi = O(n)$. Hence, if CAPACITY is approximable within $f(\varphi)$ factor, then MAX INDEPENDENT SET is approximable within $O(f(n))$ factor. In particular, the $\Omega(n^{1-o(1)})$ -computational hardness of MAX INDEPENDENT SET [44] implies equivalent $\Omega(\varphi^{1-o(1)})$ -hardness for CAPACITY.

Finally, we examine the bounded-growth properties of the space. Any packing consists of packings in a smaller target ball in each of the two lines. A t -packing on the line within a target ball of size R has cardinality at most $(R/t)^{\alpha'}$. Thus, we find that the decay space is doubling with Assouad dimension $A = 1/\alpha' = 1/(\alpha - 1)$ and constant $C = 2$. As for independence, all nodes on a line are closer to each other than they are to any node on the other line. Thus, an independent set with respect to a point x contains at most two points from the same line as x and at most one point from the other line, for an independence dimension of 3. $\square$

We note that the decays used in the construction were all in the range $d^{\alpha'}$ and $d^{\alpha'+1}$ between pairs of distance d . This result thus shows that huge decays (or, path-loss) are not needed *per se* to get large approximation hardness. Rather, it is the differences in decay among spatially related points that is the cause.

5. CONCLUSIONS

We have proposed studying wireless algorithms in abstract decay spaces to capture complex environments combined with parameters of closeness to a metric space. The wholesale theory transfer that follows appears to validate recent efforts to elucidate theoretical properties of the SINR model. The lesson for algorithm designers seems to be that GEO-SINR is fine but strictly Euclidean properties are to be avoided. Finally, our results may also indirectly strengthen the case for carrier sense or measurement mechanisms.

Several directions remain to be addressed. Different parameterizations, possibly tolerating local perturbations, is one. Capturing probabilistic behavior is another. Dynamic settings are extremely important yet lightly explored. Finally, recent advances at the physical layer (alignment, beam-forming) open up completely new dimensions worth exploring analytically.

6. REFERENCES

- [1] E. Ásgeirsson and P. Mitra. On a game theoretic approach to capacity maximization in wireless networks. In *INFOCOM*, 2011.
- [2] C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty. SINR diagrams: Convexity and its applications in wireless networks. *J. ACM*, 59(4), 2012.
- [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 Netw.*, 8(4):34, 2012.
- [4] M. H. Bodlaender, M. M. Halldórsson, and P. Mitra. Connectivity and Aggregation in Multihop Wireless Networks. In *PODC*, 2013.
- [5] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Cross-layer latency minimization for wireless networks using SINR constraints. In *Mobihoc*, 2007.
- [6] D. Chafekar, V. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. Approximation algorithms for computing capacity of wireless networks with SINR constraints. In *INFOCOM*, 2008.
- [7] Y. Chen and A. Terzis. On the mechanisms and effects of calibrating RSSI measurements for 802.15.4 radios. In *EWSN*, pages 256–271, 2010.
- [8] J. Dams, M. Hoefer, and T. Kesselheim. Scheduling in wireless networks with Rayleigh-fading interference. In *SPAA*, pages 327–335, 2012.
- [9] J. Dams, M. Hoefer, and T. Kesselheim. Jamming-resistant learning in wireless networks. *arXiv:1307.5290*, 2013. To appear in *ICALP* 2014.
- [10] J. Dams, M. Hoefer, and T. Kesselheim. Sleeping experts in wireless networks. In *DISC*, pages 344–357. Springer, 2013.
- [11] S. Daum, S. Gilbert, F. Kuhn, and C. C. Newport. Broadcast in the ad hoc SINR model. In *DISC*, pages 358–372, 2013.
- [12] M. Dinitz. Distributed algorithms for approximating wireless network capacity. In *INFOCOM*, 2010.
- [13] A. Fanghänel, S. Geulen, M. Hoefer, and B. Vöcking. Online capacity maximization in wireless networks. *Journal of Scheduling*, 16(1):81–91, 2013.
- [14] A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *PODC*, pages 220–229, August 2009.
- [15] A. Fanghänel, T. Kesselheim, and B. Vöcking. Improved algorithms for latency minimization in wireless networks. In *ICALP*, pages 447–458, 2009.
- [16] F. Fuchs and D. Wagner. On local broadcasting schedules and CONGEST algorithms in the SINR model. In *ALGOSENSORS*, pages 170–184, 2013.
- [17] O. Göbel, M. Hoefer, T. Kesselheim, T. Schleiden, and B. Vöcking. Online independent set beyond the worst-case: Secretaries, prophets, and periods. *CoRR*, abs/1307.3192, 2013. To appear in *ICALP* 2014.
- [18] A. Goldsmith. *Wireless Communications*. Cambridge Univ. Press, 2005.
- [19] O. Goussevskaia, M. M. Halldórsson, R. Wattenhofer, and E. Welzl. Capacity of Arbitrary Wireless Networks. In *INFOCOM*, pages 1872–1880, April 2009.
- [20] O. Goussevskaia, T. Moscibroda, and R. Wattenhofer. Local Broadcasting in the Physical Interference Model. In *DialM-POMC*, August 2008.
- [21] O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. Complexity in geometric SINR. In *Mobihoc*, 2007.
- [22] H. Gudmundsdottir, E. I. Ásgeirsson, M. Bodlaender, J. T. Foley, M. M. Halldórsson, G. M. Jarvelä, H. Ulfarsson, and Y. Vigfusson. Measurement based

- interference models for wireless scheduling algorithms. *arXiv:1401.1723*, 2014.
- [23] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Trans. Information Theory*, 46(2):388–404, 2000.
 - [24] M. M. Halldórsson. Wireless scheduling with power control. *ACM Trans. Algorithms*, 9(1):7, 2012.
 - [25] M. M. Halldórsson, S. Holzer, P. Mitra, and R. Wattenhofer. The power of non-uniform wireless power. In *SODA*, 2013.
 - [26] M. M. Halldórsson and P. Mitra. Wireless capacity with arbitrary gain matrix. In *ALGOSENSORS*, pages 215–224, 2011.
 - [27] M. M. Halldórsson and P. Mitra. Wireless Capacity with Oblivious Power in General Metrics. In *SODA*, 2011.
 - [28] M. M. Halldórsson and P. Mitra. Distributed connectivity of wireless networks. In *PODC*, 2012.
 - [29] M. M. Halldórsson and P. Mitra. Towards tight bounds for local broadcasting. In *FOMC*, page 2, 2012.
 - [30] M. M. Halldórsson and P. Mitra. Wireless capacity and admission control in cognitive radio. In *INFOCOM*, pages 855–863, 2012.
 - [31] M. M. Halldórsson and P. Mitra. Wireless Connectivity and Capacity. In *SODA*, 2012.
 - [32] M. M. Halldórsson and P. Mitra. Nearly optimal bounds for distributed wireless scheduling in the SINR model. *Dist. Comput.*, 2014. To appear. CoRR abs/1107.4981. Earlier version in *ICALP* 2011.
 - [33] M. M. Halldórsson and R. Wattenhofer. Wireless Communication is in APX. In *ICALP*, 2009.
 - [34] J. Heinonen. *Lectures on Analysis in Metric Spaces*. Springer, 1999.
 - [35] M. Hoefer and T. Kesselheim. Secondary spectrum auctions for symmetric and submodular bidders. In *EC*, pages 657–671, 2012.
 - [36] M. Hoefer, T. Kesselheim, and B. Vöcking. Approximation algorithms for secondary spectrum auctions. In *SPAA*, 2011.
 - [37] T. Jurdzinski and D. R. Kowalski. Distributed backbone structure for algorithms in the SINR model of wireless networks. In *DISC*, pages 106–120, 2012.
 - [38] T. Jurdzinski, D. R. Kowalski, M. Rozanski, and G. Stachowiak. Distributed randomized broadcasting in wireless networks under the SINR model. In *DISC*, pages 373–387, 2013.
 - [39] T. Jurdzinski, D. R. Kowalski, and G. Stachowiak. Distributed deterministic broadcasting in uniform-power ad hoc wireless networks. In *FCT*, 2013.
 - [40] T. Kesselheim. A Constant-Factor Approximation for Wireless Capacity Maximization with Power Control in the SINR Model. In *SODA*, 2011.
 - [41] T. Kesselheim. Approximation algorithms for wireless link scheduling with flexible data rates. In *ESA*, 2012.
 - [42] T. Kesselheim. Dynamic packet scheduling in wireless networks. In *PODC*, pages 281–290, 2012.
 - [43] T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *DISC*, August 2010.
 - [44] S. Khot and A. K. Ponnuswami. Better inapproximability results for maxclique, chromatic number and min-3lin-deletion. In *ICALP*, 2006.
 - [45] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *MSWiM*, 2004.
 - [46] Z. Lotker, M. Parter, D. Peleg, and Y. A. Pignolet. Distributed power control in the SINR model. In *INFOCOM*, pages 2525–2533, 2011.
 - [47] 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.
 - [48] T. Moscibroda and R. Wattenhofer. The complexity of connectivity in wireless networks. In *INFOCOM*, 2006.
 - [49] G. Pei and A. K. S. Vullikanti. Distributed approximation algorithms for maximum link scheduling and local broadcasting in the physical interference model. In *INFOCOM*, 2013.
 - [50] C. Scheideler, A. W. Richa, and P. Santi. An $O(\log n)$ dominating set protocol for wireless ad-hoc networks under the physical interference model. In *MobiHoc*, pages 91–100, 2008.
 - [51] V. Sevani and B. Raman. SIR based interference modeling for wireless mesh networks: A detailed measurement study. In *COMSNETS*. IEEE, 2012.
 - [52] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of concurrent transmission in wireless sensor networks. In *SenSys*. ACM, 2006.
 - [53] T. Tonoyan. Algorithms for scheduling with power control in wireless networks. In *TAPAS*, 2011.
 - [54] T. Tonoyan. On some bounds on the optimum schedule length in the SINR model. In *ALGOSENSORS*, pages 120–131, 2012.
 - [55] T. Tonoyan. Comparing schedules in the SINR and conflict-graph models with different power schemes. In *ADHOC-NOW*, pages 317–328, 2013.
 - [56] P. Wan, O. Frieder, X. Jia, F. Yao, X. Xu, and S. Tang. Wireless link scheduling under physical interference model. In *INFOCOM*, 2011.
 - [57] E. Welzl. Blue-dominant centers. Unpublished memorandum, June 2008.
 - [58] X. Xu, J. Cao, and X.-Y. Li. MLLS: Minimum length link scheduling under physical interference model. *CoRR*, abs/1208.0627, 2012.
 - [59] D. Yu, Q.-S. Hua, Y. Wang, H. Tan, and F. C. M. Lau. Distributed multiple-message broadcast in wireless ad-hoc networks under the SINR model. In *SIROCCO*, pages 111–122, 2012.
 - [60] D. Yu, Q.-S. Hua, Y. Wang, J. Yu, and F. C. M. Lau. Efficient distributed multiple-message broadcasting in unstructured wireless networks. In *INFOCOM*, 2013.
 - [61] D. Yu, Y. Wang, Q.-S. Hua, and F. C. M. Lau. Distributed $(\Delta + 1)$ -coloring in the physical model. In *ALGOSENSORS*, pages 145–160, 2011.
 - [62] D. Yu, Y. Wang, Q.-S. Hua, and F. C. M. Lau. An $O(\log n)$ distributed approximation algorithm for local broadcasting in unstructured wireless networks. In *DCOSS*, pages 132–139, 2012.