

Braess's Paradox in Wireless Networks: The Danger of Improved Technology

Michael Dinitz^{1,2} and Merav Parter¹ *

¹ Department of Computer Science and Applied Mathematics,
The Weizmann Institute, Rehovot, Israel

{michael.dinitz, merav.parter}@weizmann.ac.il **

² Department of Computer Science, Johns Hopkins University

Abstract. When comparing new wireless technologies, it is common to consider the effect that they have on the capacity of the network (defined as the maximum number of simultaneously satisfiable links). For example, it has been shown that giving receivers the ability to do interference cancellation, or allowing transmitters to use power control, never decreases the capacity and can in certain cases increase it by $\Omega(\log(\Delta \cdot P_{\max}))$, where Δ is the ratio of the longest link length to the smallest transmitter-receiver distance and P_{\max} is the maximum transmission power. But there is no reason to expect the optimal capacity to be realized in practice, particularly since maximizing the capacity is known to be NP-hard. In reality, we would expect links to behave as self-interested agents, and thus when introducing a new technology it makes more sense to compare the values reached at game-theoretic equilibria than the optimum values.

In this paper we initiate this line of work by comparing various notions of equilibria (particularly Nash equilibria and no-regret behavior) when using a supposedly “better” technology. We show a version of Braess's Paradox for all of them: in certain networks, upgrading technology can actually make the equilibria *worse*, despite an increase in the capacity. We construct instances where this decrease is a constant factor for power control, interference cancellation, and improvements in the SINR threshold (β), and is $\Omega(\log \Delta)$ when power control is combined with interference cancellation. However, we show that these examples are basically tight: the decrease is at most $O(1)$ for power control, interference cancellation, and improved β , and is at most $O(\log \Delta)$ when power control is combined with interference cancellation.

* Recipient of the Google European Fellowship in distributed computing; research is supported in part by this Fellowship.

** Supported in part by the Israel Science Foundation (grant 894/09), the I-CORE program of the Israel PBC and ISF (grant 4/11), the United States-Israel Binational Science Foundation (grant 2008348), the Israel Ministry of Science and Technology (infrastructures grant), and the Citi Foundation.

1 Introduction

Due to the increasing use of wireless technology in communication networks, there has been a significant amount of research on methods of improving wireless performance. While there are many ways of measuring wireless performance, a good first step (which has been extensively studied) is the notion of *capacity*. Given a collection of communication links, the capacity of a network is simply the maximum number of simultaneously satisfiable links. This can obviously depend on the exact model of wireless communication that we are using, but is clearly an upper bound on the “usefulness” of the network. There has been a large amount of research on analyzing the capacity of wireless networks (see e.g. [15,14,2,19]), and it has become a standard way of measuring the quality of a network. Because of this, when introducing a new technology it is interesting to analyze its affect on the capacity. For example, we know that in certain cases giving transmitters the ability to control their transmission power can increase the capacity by $\Omega(\log \Delta)$ or $\Omega(\log P_{\max})$ [6], where Δ is the ratio of the longest link length to the smallest transmitter-receiver distance, and can clearly never decrease the capacity.

However, while the capacity might improve, it is not nearly as clear that the *achieved* capacity will improve. After all, we do not expect our network to actually have performance that achieves the maximum possible capacity. We show that not only might these improved technologies not help, they might in fact *decrease* the achieved network capacity. Following Andrews and Dinitz [2] and Ásgeirsson and Mitra [3], we model each link as a self-interested agent and analyze various types of game-theoretic behavior (Nash equilibria and no-regret behavior in particular). We show that a version of *Braess’s Paradox* [8] holds: adding new technology to the networks (such as the ability to control powers) can actually decrease the average capacity at equilibrium.

1.1 Our Results

Our main results show that in the context of wireless networks, and particularly in the context of the SINR model, there is a version of *Braess’s Paradox* [8]. In his seminal paper, Braess studied congestion in road networks and showed that adding additional roads to an existing network can actually make congestion *worse*, since agents will behave selfishly and the additional options can result in worse equilibria. This is completely analogous to our setting, since in road networks adding extra roads cannot hurt the network in terms of the value of the optimum solution, but can hurt the network since the *achieved* congestion gets worse. In this work we consider the physical model (also called the SINR model), pioneered by Moscibroda and Wattenhofer [23] and described more formally in Section 2.1. Intuitively, this model works as follows: every sender chooses a transmission power (which may be pre-determined, e.g. due to hardware limitations), and the received power decreased polynomially with the distance from the sender. A transmission is successful if the received power from the sender

is large enough to overcome the interference caused by other senders plus the background noise.

With our baseline being the SINR model, we then consider four ways of “improving” a network: adding power control, adding interference cancellation, adding both power control and interference cancellation, and decreasing the SINR threshold. With all of these modifications it is easy to see that the optimal capacity can only increase, but we will show that the equilibria can become worse. Thus “improving” a network might actually result in worse performance.

The game-theoretic setup that we use is based on [2] and will be formally described in Section 2.2, but we will give an overview here. We start with a game in which the players are the links, and the strategies depend slightly on the model but are essentially possible power settings at which to transmit. The utilities depend on whether or not the link was successful, and whether or not it even attempted to transmit. In a pure Nash equilibrium every player has a strategy (i.e. power setting) and has no incentive to deviate: any other strategy would result in smaller utility. In a mixed Nash equilibrium every link has a probability distribution over the strategies, and no link has any incentive to deviate from their distribution. Finally, no-regret behavior is the empirical distribution of play when all players use *no-regret* algorithms, which are a widely used and studied class of learning algorithms (see Section 2.2 for a formal definition). It is reasonably easy to see that any pure Nash is a mixed Nash, and any mixed Nash is a no-regret behavior. For all of these, the quality of the solution is the achieved capacity, i.e. the average number of successful links.

Our first result is for interference cancellation (IC), which has been widely proposed as a practical method of increasing network performance [1]. The basic idea of interference cancellation is quite simple. First, the strongest interfering signal is detected and decoded. Once decoded, this signal can then be subtracted (“canceled”) from the original signal. Subsequently, the next strongest interfering signal can be detected and decoded from the now “cleaner” signal, and so on. As long as the strongest remaining signal can be decoded in the presence of the weaker signals, this process continues until we are left with the desired transmitted signal, which can now be decoded. This clearly can increase the capacity of the network, and even in the worst case cannot decrease it. And yet due to bad game-theoretic interactions it might make the achieved capacity worse:

Theorem 1. *There exists a set of links in which the best no-regret behavior under interference cancellation achieves capacity at most c times the worst no-regret behavior without interference cancellation, for some constant $c < 1$. However, for every set of links the worst no-regret behavior under interference cancellation achieves capacity that is at least a constant fraction of the best no-regret behavior without interference cancellation.*

Thus IC can make the achieved capacity worse, but only by a constant factor. Note that since every Nash equilibrium (mixed or pure) is also no-regret, this implies the equivalent statements for those type of equilibria as well. In this result (as in most of our examples) we only show a small network (4 links)

with no background noise, but these are both only for simplicity – it is easy to incorporate constant noise, and the small network can be repeated at sufficient distance to get examples with an arbitrarily large number of links.

We next consider power control (PC), where senders can choose not just whether to transmit, but at what power to transmit. It turns out that any equilibrium without power control is also an equilibrium with power control, and thus we cannot hope to find an example where the best equilibrium with power control is worse than the worst equilibrium without power control (as we did with IC). Instead, we show that adding power control can create worse equilibria:

Theorem 2. *There exists a set of links in which there is a pure Nash equilibrium with power control of value at most c times the value of the worst no-regret behavior without power control, for some constant $c < 1$. However, for every set of links the worst no-regret behavior with power control has value that is at least a constant fraction of the value of the best no-regret behavior without power control.*

Note that the first part of the theorem implies that not only is there a pure Nash with low-value (with power control), there are also mixed Nash and no-regret behaviors with low value (since any pure Nash is also mixed and no-regret). Similarly, the second part of the theorem gives a bound on the gap between the worst and the best mixed Nashes, and the worst and the best pure Nashes.

Our third set of results is on the combination of power control and interference cancellation. It turns out that the combination of the two can be quite harmful. When compared to either the vanilla setting (no interference cancellation or power control) or the presence of power control without interference cancellation, the combination of IC and PC acts essentially as in Theorem 2: pure Nash equilibria are created that are worse than the previous worst no-regret behavior, but this can only be by a constant factor. On the other hand, this factor can be super-constant when compared to equilibria that only use interference cancellation. Let Δ be the ratio of the length of the longest link to the minimum distance between any sender and any receiver.³

Theorem 3. *There exists a set of links in which the worst pure Nash with both PC and IC (and thus the worst mixed Nash or no-regret behavior) has value at most $O(1/\log \Delta)$ times the value of the worst no-regret behavior with just IC. However, for every set of links the worst no-regret behavior with both PC and IC has value at least $\Omega(1/\log \Delta)$ times the value of the best no-regret behavior with just IC.*

This theorem means that interference cancellation “changes the game”: if interference control were not an option then power control can only hurt the equilibria by a constant amount (from Theorem 2), but if we assume that interference

³ Note that this definition is slightly different than the one used by [17,3,16] and is a bit more similar to the definition used by [2,11]. The interested reader can see that this is in fact the appropriate definition in the IC setting, namely, in a setting where a receiver can decode multiple (interfering) stations.

cancellation is present then adding power control can hurt us by $\Omega(\log \Delta)$. Thus when deciding whether to use both power control and interference cancellation, one must be particularly careful to analyze how they act in combination.

Finally, we consider the effect of decreasing the SINR threshold β (this value will be formally described in Section 2.1). We show that, as with IC, there are networks in which a decrease in the SINR threshold can lead to *every* equilibrium being worse than even the worst equilibrium at the higher threshold, despite the capacity increasing or staying the same:

Theorem 4. *There exists a set of links and constants $1 < \beta' < \beta$ in which the best no-regret behavior under threshold β' has value at most c times the value of the worst no-regret behavior under threshold β , for some constant $c < 1$. However, for any set of links and any $1 < \beta' < \beta$ the value of the worst no-regret behavior under β' is at least a constant fraction of the value of the best no-regret behavior under β .*

Our main network constructions illustrating Braess's paradox in the studied settings are summarized in Fig. 1.

1.2 Related Work

The capacity of *random* networks was examined in the seminal paper of Gupta and Kumar [15], who proved tight bounds in a variety of models. But only recently has there been a significant amount of work on algorithms for determining the capacity of *arbitrary* networks, particularly in the SINR model. This line of work began with Goussevskaia, Oswald, and Wattenhofer [13], who gave an $O(\log \Delta)$ -approximation for the uniform power setting (i.e. the vanilla model we consider). Goussevskaia, Halldórson, Wattenhofer, and Welzl [14] then improved this to an $O(1)$ -approximation (still under uniform powers), while Andrews and Dinitz [2] gave a similar $O(\log \Delta)$ -approximation algorithm for the power control setting. This line of research was essentially completed by an $O(1)$ -approximation for the power control setting due to Kesselheim [19].

In parallel to the work on approximation algorithms, there has been some work on using game theory (and in particular the games used in this paper) to help design distributed approximation algorithms. This was begun by Andrews and Dinitz [2], who gave an upper bound of $O(\Delta^{2\alpha})$ on the price of anarchy for the basic game defined in Section 2.2. But since computing the Nash equilibrium of a game is PPAD-complete [10], we do not expect games to necessarily converge to a Nash equilibrium in polynomial time. Thus Dinitz [11] strengthened the result by showing the same upper bound of $O(\Delta^{2\alpha})$ for no-regret behavior. This gave the first distributed algorithm with a nontrivial approximation ratio, simply by having every player use a no-regret algorithm. The analysis of the same game was then improved to $O(\log \Delta)$ by Ásgeirsson and Mitra [3].

There is very little work on interference cancellation in arbitrary networks from an algorithmic point of view, although it has been studied quite well from an information-theoretic point of view (see e.g. [12,9]). Recently Avin et al. [4] studied the topology of *SINR diagrams* under interference cancellation, which

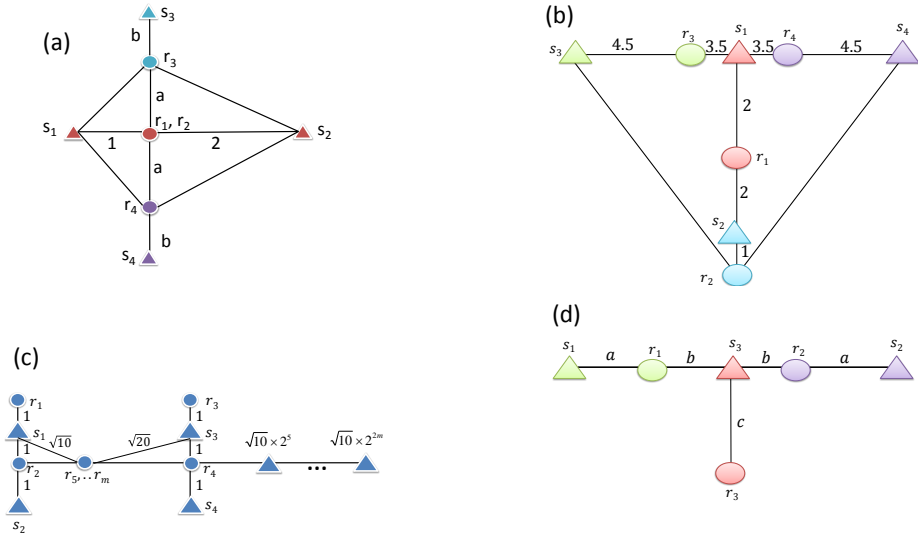


Fig. 1: Schematic illustration of the main lower bounds illustrating the Braess's paradox with (a) IC: a network in which every no-regret behavior without IC is better than any no-regret behavior solution with IC; (b) PC: a network in which there exists a pure Nash equilibrium with PC which is worse than any no-regret behavior with IC; (c) PIC: a network with a pure Nash equilibrium in the PIC setting which is $\Omega(\log \Delta)$ worse than any no-regret behavior in the IC setting but *without* power control; and (d) Decreased SINR threshold $\beta' < \beta$: a network in which every no-regret behavior with β' has a smaller value than any no-regret behavior with higher SINR threshold β . Edge weights represent distances.

are a generalization of the SINR diagrams introduced by Avin et al. [5] and further studied by Kantor et al. [18] for the SINR model without interference cancellation. These diagrams specify the reception zones of transmitters in the SINR model, which turn out to have several interesting topological and geometric properties but have not led to a better understanding of the fundamental capacity question.

2 Preliminaries

2.1 The Communication Model

We model a wireless network as a set of links $L = \{\ell_1, \ell_2, \dots, \ell_n\}$ in the plane, where each link $\ell_v = (s_v, r_v)$ represents a communication request from a sender s_v to a receiver r_v . The n senders and receivers are given as points in the Euclidean plane. The Euclidean distance between two points p and q is denoted

$d(p, q)$. The distance between sender s_i and receiver r_j is denoted by $d_{i,j}$. We adopt the physical model (sometime called the SINR model) where the received signal strength of transmitter s_i at the receiver r_j decays with the distance and it is given by $P_{i,j} = P_i/d_{i,j}^\alpha$, where $P_i \in [0, P_{\max}]$ is the transmission power of sender s_i and $\alpha > 0$ is the path-loss exponent. Receiver r_j successfully receives a message from sender s_j iff $\text{SINR}_j(L) = \frac{P_{j,j}}{\sum_{\ell_i \in L} P_{i,j} + \eta} \geq \beta$, where η denotes the amount of background noise and $\beta > 1$ denotes the minimum SINR required for a message to be successfully received. The *total interference* that receiver r_j suffers from the set of links L is given by $\sum_{i \neq j} P_{i,j}$. Throughout, we assume that all distances $d_{i,j}$ are normalized so that $\min_{s_i, r_j} d_{i,j} = 1$ hence the maximal link length is Δ , i.e., $\Delta = \max_i d_{i,i}$ and any received signal strength $P_{i,j}$ is bounded by $P_{i,j} \leq P_{\max}$.

In the vanilla SINR model we require that P_i is either 0 or P_{\max} for every transmitter. This is sometimes referred to as *uniform* powers. When we have power control, we allow P_i to be any integer in $[0, P_{\max}]$.

Interference cancellation allows receivers to cancel signals that they can decode. Consider link ℓ_j . If r_j can decode the signal with the largest received signal, then it can decode it and remove it. It can repeat this process until it decodes its desire message from s_j , unless at some point it gets stuck and cannot decode the strongest signal. Formally, r_j can decode s_i if $P_{i,j}/(\sum_{\ell_k: P_{k,j} < P_{i,j}} P_{k,j} + \eta) \geq \beta$ (i.e. it can decode s_i in the presence of weaker signals) and if it can decode s_k for all links ℓ_k with $P_{k,j} > P_{i,j}$. Link ℓ_j is successful if r_j can decode s_j .

The following key notion, which was introduced in [17] and extended to arbitrary powers by [20], plays an important role in our analysis. The *affectance* $a_w^P(v)$ of link ℓ_v caused by another link ℓ_w with a given power assignment vector P is defined to be $a_w^P(v) = \min\{1, c_v(\beta)P_{w,v}/P_{v,v}\}$, where $c_v(\beta) = \beta/(1 - \beta\eta d_{v,v}^\alpha/P_v)$. Informally, $a_w^P(v)$ indicates the amount of (normalized) interference that link ℓ_w causes at link ℓ_v . It is easy to verify that link ℓ_v is successful if and only if $\sum_{w \neq v} a_w^P(v) \leq 1$.

When the powers of P are the same $P_u = P_{\max}$ for every ℓ_u , (i.e. uniform powers), we may omit it and simply write $a_w(v)$. For a set of links L and a link ℓ_w , the total affectance caused by ℓ_w is $a_w(L) = \sum_{\ell_v \in L} a_w(v)$. In the same manner, the total affectance caused by L on the link ℓ_w is $a_L(w) = \sum_{\ell_v \in L} a_v(w)$. We say that a set of links L is β -feasible if $\text{SINR}_v(L) \geq \beta$ for all $\ell_v \in L$, i.e. every link achieves SINR above the threshold (and is thus successful even without interference cancellation). It is easy to verify that L is β -feasible if and only if $\sum_{\ell_w \in L} a_w(v) \leq 1$ for every $\ell_v \in L$.

Following [16], we say that a link set L is *amenable* if the total affectance caused by any single link is bounded by some constant, i.e., $a_u(L) = O(1)$ for every $u \in L$. The following basic properties of amenable sets play an important role in our analysis.

Fact 5 [3] (a) *Every feasible set L contains a subset $L' \subseteq L$, such that L' is amenable and $|L'| \geq |L|/2$.*

(b) *For every amenable set L' that is β -feasible with uniform powers, for every other link u , it holds that $\sum_{v \in L'} a_u(v) = O(1)$.*

2.2 Basic Game Theory

We will use a game that is essentially the same as the game of Andrews and Dinitz [2], modified only to account for the different models we consider. Each link ℓ_i is a player with $P_{\max} + 1$ possible strategies: broadcast at power 0, or at integer power $P \in \{1, \dots, P_{\max}\}$. A link has utility 1 if it is successful, has utility -1 if it uses nonzero power but is unsuccessful, and has utility 0 if it does not transmit (i.e. chooses power 0). Note that if power control is not available, this game only has two strategies: power 0 and power P_{\max} . Let S denote the set of possible strategies. A *strategy profile* is a vector in S^n , where the i 'th component is the strategy played by link ℓ_i . For each link ℓ_i , let $f_i : S^n \rightarrow \{-1, 0, 1\}$ be the function mapping strategy profiles to utility for link i as described. Given a strategy profile a , let a_{-i} denote the profile without the i 'th component, and given some strategy $s \in S$ let $f_i(s, a_{-i})$ denote the utility of ℓ_i if it uses strategy s and all other links use their strategies from a .

A pure Nash equilibrium is a strategy profile in which no player has any incentive to deviate from their strategy. Formally, $a \in S^n$ is a pure Nash equilibrium if $f_i(a) \geq f_i(s, a_{-i})$ for all $s \in S$ and players $\ell_i \in L$. In a mixed Nash equilibrium [24], every player ℓ_i has a probability distribution π_i over S , and the requirement is that no player has any incentive to change their distribution to some π' . So $\mathbb{E}[f_i(a)] \geq \mathbb{E}[f_i(\pi', a_{-i})]$ for all $i \in \{1, \dots, n\}$, where the expectation is over the random strategy profile a drawn from the product distribution defined by the π_i 's, and π' is any distribution over S .

To define no-regret behavior, suppose that the game has been played for T rounds and let a^t be the realized strategy profile in round $t \in \{1, \dots, T\}$. The *history* $\mathbf{H} = \{a^1, \dots, a^T\}$ of the game is the sequence of the T strategy profiles. The *regret* \mathcal{R}_i of player i in an history \mathbf{H} is defined to be

$$\mathcal{R}_i(\mathbf{H}) = \max_{s \in S} \frac{1}{T} \sum_{t=1}^T f_i(s, a_{-i}^t) - \frac{1}{T} \sum_{t=1}^T f_i(a^t).$$

The regret of a player is intuitively the amount that it lost by not playing some fixed strategy. An algorithm used by a player is known as a *no-regret* algorithm if it guarantees that the regret of the player tends to 0 as T tends to infinity. There is a large amount of work on no-regret algorithms, and it turns out that many such algorithms exist (see e.g. [25,22]). Thus we will analyze situations where every player has regret at most ϵ , and since this tends to 0 we will be free to assume that ϵ is arbitrarily small, say at most $1/n$. Clearly playing a pure or mixed Nash is a no-regret algorithm (since the fact that no one has incentive to switch to any other strategy guarantees that each player will have regret 0 in the long run), so analyzing the worst or best history with regret at most ϵ is more general than analyzing the worst or best mixed or pure Nash. We will call a history in which all players have regret at most ϵ an ϵ -regret history. Formally an history $\mathbf{H} = \{a^1, \dots, a^T\}$ is an ϵ -regret history if $\mathcal{R}_i(\mathbf{H}) \leq \epsilon$ for every player $i \in \{1, \dots, n\}$.

A simple but important lemma introduced in [11] and used again in [3] relates the average number of *attempted* transmissions to the average number of *successful* transmissions. Fix an ϵ -regret history, let s_u be the fraction of times

in which u successfully transmitted, and let p_u be the fraction of times in which u attempted to transmit. Note that the average number of successful transmissions in a time slot is exactly $\sum_u s_u$, so it is this quantity which we will typically attempt to bound. The following lemma lets us get away with bounding the number of attempts instead.

Lemma 1 ([11]). $\sum_u s_u \leq \sum_u p_u \leq 2\sum_u s_u + \epsilon n \leq O(\sum_u s_u)$.

Notation: Let L be a fixed set of n links embedded in \mathbb{R}^2 . Let $\mathcal{N}_{\min}(L)$ denote the minimum number of successful links (averaged over time) in any ϵ -regret history, and similarly let $\mathcal{N}_{\max}(L)$ denote the maximum number of successful links (averaged over time) in any ϵ -regret history. Define $\mathcal{N}_{\max}^{IC}(L)$, and $\mathcal{N}_{\min}^{IC}(L)$ similarly for the IC setting, $\mathcal{N}_{\max}^{PC}(L)$ and $\mathcal{N}_{\min}^{PC}(L)$ for the PC setting, and $\mathcal{N}_{\max}^{PIC}(L)$ and $\mathcal{N}_{\min}^{PIC}(L)$ for the setting with both PC and IC. Finally, let $\mathcal{N}_{\max}^{\beta}(L)$, $\mathcal{N}_{\min}^{\beta}(L)$ be for the corresponding values for the vanilla model when the SINR threshold is set to β (this is hidden in the previous models, but we pull it out in order to compare the effect of modifying β).

While we will focus on comparing the equilibria of games utilizing different wireless technologies, much of the previous work on these games instead focuses on a single game and analyzes its equilibria with respect to OPT, the maximum achievable capacity. The *price of anarchy* (PoA) is the ratio of OPT to the value of the worst mixed Nash [21], and the *price of total anarchy* (PoTA) is the ratio of OPT to the value of the worst ϵ -regret history [7]. Clearly $\text{PoA} \leq \text{PoTA}$. While it is not our focus, we will prove some bounds on these values as corollaries of our main results.

3 Interference Cancellation

We begin by analyzing the effect on the equilibria of adding interference cancellation. We would expect that using IC would result equilibria with larger values, since the capacity of the network might go up (and can certainly not go down). We show that this is not always the case: there are sets of links for which even the *best* ϵ -regret history using IC is a constant factor worse than the *worst* ϵ -regret history without using IC.

Theorem 6. *There exists a set of links L such that $\mathcal{N}_{\max}^{IC}(L) \leq \mathcal{N}_{\min}(L)/c$ for some constant $c > 1$.*

Proof. Let L' be the four link network depicted in Figure 1(a), with $b = 3/2$ and $a = \sqrt{8.8} - b$. We will assume that the threshold β is equal to 1.1, the path-loss exponent α is equal to 2, and the background noise $\eta = 0$ (none of these are crucial, but make the analysis simpler). Let us first consider what happens without using interference cancellation. Suppose each link has at most ϵ -regret, and for link ℓ_i let p_i denote the fraction of times at which s_i attempted to transmit. It is easy to see that link 1 will always be successful, since the received

signal strength at r_1 is 1 while the total interference is at most $(1/4) + 2(1/8.8) = 1/\beta$. Since ℓ_1 has at most ϵ -regret, this implies that $p_1 \geq 1 - \epsilon$.

On the other hand, whenever s_1 transmits it is clear that link ℓ_2 cannot be successful, as its SINR is at most $1/4$. So if s_2 transmitted every time it would have average utility at most $-(1 - \epsilon) + \epsilon = -1 + 2\epsilon < 0$ (since $\epsilon < 1/2$), while if it never transmitted it would have average utility 0. Thus its average utility is at least $-\epsilon$. Since it can succeed only an ϵ fraction of the time (when link 1 is not transmitting), we have that $\epsilon - (p_2 - \epsilon) \geq -\epsilon$ and thus $p_2 \leq 3\epsilon$. Since the utility of s_2 is at least $-\epsilon$, it holds that the fraction of times at which both s_1 and s_2 are transmitting is at most 2ϵ .

Now consider link ℓ_3 . If links ℓ_1 and ℓ_2 both transmit, then ℓ_3 will fail since the received SINR will be at most $\frac{1/b^2}{(1/(1+a^2))+(1/(4+a^2))} \approx 0.92 < 1.1$. On the other hand, as long as link ℓ_2 does not transmit then ℓ_3 will be successful, as it will have SINR at least $\frac{1/b^2}{(1/(1+a^2))+(1/(2a+b)^2)} \geq 1.2 > 1.1$. Thus by transmitting at every time step ℓ_3 would have average utility at least $(1 - 2\epsilon) - 2\epsilon = 1 - 4\epsilon > 0$ (since $\epsilon < 1/4$), and thus we know that ℓ_3 gets average utility of at least $1 - 5\epsilon$, and thus successfully transmits at least $1 - 5\epsilon$ fraction of the times. ℓ_4 is the same by symmetry. Thus the total value of any history in which all links have regret at most ϵ is at least $\mathcal{N}_{\min}(L) \geq 1 - \epsilon + 2(1 - 5\epsilon) = 3 - 11\epsilon$.

Let us now analyze what happens when using interference cancellation and bound $\mathcal{N}_{\max}^{IC}(L)$. Suppose each link has at most ϵ -regret, and for link ℓ_i let q_i denote the fraction of times at which s_i attempted to transmit. As before, ℓ_1 can always successfully transmit and thus does so in at least $1 - \epsilon$ fraction of times. But now, by using interference cancellation it turns out that ℓ_2 can also always succeed. This is because r_2 can first decode the transmission from s_1 and cancel it, leaving a remaining SINR of at least $\frac{1/4}{2/(a+b)^2} = \beta$. Thus ℓ_2 will also transmit in at least $1 - \epsilon$ fraction of times and hence so far $1 - \epsilon \leq q_1, q_2 \leq 1$. Note that since $a^2 + 1 < b^2$, it holds that r_3 cannot cancel s_1 or s_2 before decoding s_3 (i.e., $P_{1,3}, P_{2,3} < P_{3,3}$). Hence, cancellation is useless. But now at r_3 the strength of s_1 is $1/(1 + a^2) > 0.317$, the strength of s_2 is $1/(4 + a^2) > 0.162$, and the strength of s_3 is $1/b^2 = 4/9$. Thus r_3 cannot decode any messages when s_1, s_2 , and s_3 are all transmitting since its SINR is at most $0.92 < \beta$, which implies that ℓ_3 can only succeed on at most 2ϵ fraction of times. The link ℓ_4 is the same as the link ℓ_3 by symmetry. Thus the total value of any history in which all links have an ϵ -regret is at most $\mathcal{N}_{\max}^{IC}(L) \leq 2 + 4\epsilon$. Thus $\mathcal{N}_{\min}(L)/\mathcal{N}_{\max}^{IC}(L) \geq 3/2 - o(1)$ as required. ■

It turns out that no-regret behavior with interference cancellation cannot be much worse than no-regret behavior without interference cancellation – as in Braess's paradox, it can only be worse by a constant factor.

Theorem 7. $\mathcal{N}_{\min}^{IC}(L) \geq \mathcal{N}_{\max}(L)/c$ for any set of links L and some constant $c \geq 1$.

Proof. Consider an ϵ -regret history without IC that maximizes the average number of successful links, i.e. one that achieves $\mathcal{N}_{\max}(L)$ value. Let p_i de-

note the fraction of times at which s_i attempted to transmit in this history, so $\sum_{i \in L} p_i = \Theta(\mathcal{N}_{\max}(L))$ by Lemma 1. Similarly, let q_i denote that fraction of times at which s_i attempted to transmit in an ϵ -regret history *with* IC that achieves value of only $\mathcal{N}_{\min}^{IC}(L)$, and so $\sum_{i \in L} q_i = \Theta(\mathcal{N}_{\min}^{IC}(L))$.

Note that since the best average number of successful connections in the non-IC case is $\mathcal{N}_{\max}(L)$, there must exist some set of connections $A \subseteq L$ such that $|A| \geq \mathcal{N}_{\max}(L)$ and A is feasible without IC. Let $B = \{i : q_i \geq 1/2\}$ and let $A' = A \setminus B$. If $|B \cap A| \geq |A|/2$ then we are done, since then

$$\mathcal{N}_{\max}(L) \leq |A| \leq 2|B| \leq 4 \sum_{i \in L} q_i = 4 \cdot \mathcal{N}_{\min}^{IC}(L)$$

as required. So without loss of generality we will assume that $|B \cap A| < |A|/2$, and thus that $|A'| > |A|/2$. Note that A' is a subset of A , and so it is feasible in the non-IC setting.

Now let $\hat{A} = \{i \in A' : \sum_{j \in A'} a_i(j) \leq 2\}$ be an amenable subset of A' . By Fact 5(a), it holds that $|\hat{A}| \geq |A'|/2 \geq |A|/4$. Fact 5(b) then implies that for any link $i \in L$, its total affectance on A is small: $\sum_{j \in A} a_i(j) \leq c'$ for some constant $c' \geq 0$. Thus we have that

$$\sum_{i \in L} \sum_{j \in \hat{A}} q_i a_i(j) \leq c' \cdot (\sum_{i \in L} q_i). \quad (1)$$

On the other hand, we know that the q_i values correspond to the worst history in which every link has regret at most ϵ (in the IC setting). Let $j \in A'$. Then $q_j < 1/2$, which means the average utility of link ℓ_j is at most $1/2$. Let y_j be the fraction of time s_j would have succeeded had it transmitted in every round. Since the average utility of the best single action is at most $1/2 + \epsilon$ it holds that $y_j - (1 - y_j) \leq 1/2 + \epsilon$ or the that $y_j \leq \frac{3}{4} + \frac{\epsilon}{2}$. In other words, in at least $1 - y_j = \frac{1}{4} - \frac{\epsilon}{2}$ fraction of the rounds the affectance of the other links on the link ℓ_j must be at least 1 (or else j could succeed in those rounds even without using IC). Thus the expected affectance (taken over a random choice of time slot) on ℓ_j is at least $\sum_{i \in L} a_i(j) q_i \geq \frac{1}{4} - \frac{\epsilon}{2}$. Summing over all $j \in \hat{A}$, we get that

$$\sum_{j \in \hat{A}} \sum_{i \in L} a_i(j) q_i \geq \sum_{j \in A} \frac{1-2\epsilon}{4} \geq \Omega(|\hat{A}|). \quad (2)$$

Combining equations (1) and (2) (and switching the order of summations) implies that $|\hat{A}| \leq O(\sum_{i \in L} q_i)$. Since $|\hat{A}| \geq |A|/4 \geq \Omega(\mathcal{N}_{\max}(L)) \geq \Omega(\sum_{i \in L} p_i)$, we get that $\sum_i p_i \leq O(\sum_{i \in L} q_i)$ as desired. ■

As a simple corollary, we will show that this lets us bound the price of total anarchy in the interference cancellation model (which, as far as we know, has not previously been bounded). Let $OPT \subseteq L$ denote some optimal solution without IC, i.e., the set of transmitters forming a maximum β -feasible set, and let $OPT^{IC} \subseteq L$ denote some optimal solution with IC.

Corollary 1. *For every set of links L it holds that the price of total anarchy with IC is $O(\log \Delta)$, or $|OPT^{IC}|/\mathcal{N}_{\min}^{IC}(L) = O(\log \Delta)$.*

Proof. We begin by providing the following auxiliary lemma which demonstrates that any feasible set of links under the setting of IC and power control in the range $[1, P_{\max}]$ contains a non-negligible subset of links that are feasible with the same transmission powers *without* IC.

Lemma 2. *For every feasible set L with IC, there exists a subset $L' \subseteq L$ such that L is feasible without IC and satisfies $|L'| > |L|/O(\log_{\beta}(\Delta^{\alpha} \cdot P_{\max}))$*

Proof. Let L be a feasible set with IC and transmission power in $[1, P_{\max}]$. We first claim that every link $\ell_v \in L$ has canceled a subset of at most $x = \lceil \log_{\beta}(\Delta^{\alpha} \cdot P_{\max}) - 1 \rceil$ links in L (where the last cancellation corresponds to the signal of the designated transmitter s_v). Assume towards contradiction that there exists $\ell_v \in L$ that cancels $y \geq x + 1$ distinct signals transmitted by s_1, \dots, s_y , where $s_y = s_v$ is the last canceled signal.

It then holds that $P_{i,v} \geq \beta \cdot P_{i+1,v}$ for every $i \in \{1, \dots, y\}$, hence $P_{1,v} \geq \beta^y P_{v,v}$.

Since the maximum power is P_{\max} and the minimal transmitter-receiver distance $d(s_i, r_j)$ is at least 1, we have that $P_{\max} \geq P_{1,v} \geq \beta^y \cdot P_{v,v} \geq \beta^y / \Delta^{\alpha}$ where the last inequality follows by the fact that the maximal link length (after the normalization) is Δ and the minimum power level is 1. Hence $y \leq \log_{\beta}(\Delta^{\alpha} \cdot P_{\max}) - 1$, giving a contradiction. Thus every link has cancelled at most x links. It then holds that $P_{i,v} \geq \beta \cdot P_{i+1,v}$ for every $i \in \{1, \dots, y\}$, hence $P_{1,v} \geq \beta^y P_{v,v}$. Combining this with the fact (due to normalizing distances so the smallest is equal to 1) that every received power is at least $1/\Delta^{\alpha}$, we have that $P_{\max} \geq P_{1,v} \geq \beta^y \cdot P_{v,v} \geq \beta^y / \Delta^{\alpha}$. Hence $y \leq \log_{\beta}(\Delta^{\alpha} \cdot P_{\max}) - 1$, giving a contradiction. Thus every link has cancelled at most x links.

The subset L' of feasible links without IC is created as follows. Initially set $Q \leftarrow L$. Until Q is non-empty, consider some link $\ell_i \in Q$ and let M_i be the subset of links cancelled by receiver r_i (other than ℓ_i itself). Add ℓ_i to L' and set $Q = Q \setminus M_i$. Since for every link ℓ_i in L' there are at most x links in $L \setminus L'$, it holds that $|L'| \geq |L|/x$. Moreover, it follows by construction that L' is feasible without IC and using the same transmission powers as before. The lemma follows. ■

We proceed by showing that $|OPT^{IC}|/|\mathcal{N}_{\min}^{IC}(L)| = O(\log \Delta)$. According to Theorem 2 of [3], it holds that $\mathcal{N}_{\min}(L) \geq |OPT|/c'$ for some constant c' . Hence, by combining with Theorem 7 we get that

$$\mathcal{N}_{\min}^{IC}(L) \geq \mathcal{N}_{\min}(L)/c \geq |OPT|/(c \cdot c') \geq |OPT^{IC}|/O(\log \Delta),$$

where the last inequality follows by Lemma 2.

Finally, we show that this analysis is actually tight by exhibiting a network where there is a bad pure Nash equilibrium, and thus there are bad no-regret histories. Consider the $m = \lfloor (\log \Delta) \rfloor$ -linkset network illustrated in Fig. 2. The transmitters \tilde{s}_1 and \tilde{s}_2 are equidistant from the set of m -receivers $R' = \{r_1, \dots, r_m\}$. Since \tilde{s}_1 and \tilde{s}_2 are closer to R' than any other transmitter s_i , if both \tilde{s}_1 and \tilde{s}_2 transmit then none of the links $\ell_i = \langle s_i, r_i \rangle$ can be satisfied. By letting the links $\tilde{\ell}_j = \langle \tilde{s}_j, \tilde{r}_j \rangle$ for $j \in \{1, 2\}$ be sufficiently short, these 2 links

can always succeed no matter which other links transmit. Thus $\{\tilde{\ell}_1, \tilde{\ell}_2\}$ form a pure Nash equilibrium. On the other hand, clearly the set of links $\{\ell_i\}_{i \in [\log \Delta]}$ are feasible with interference cancellation. Thus $|OPT^{IC}|/\mathcal{N}_{\min}^{IC}(L) \geq \Omega(\log \Delta)$. ■




Fig. 2: Schematic illustration of a network in which the price of anarchy in the IC setting is $\Omega(\log \Delta)$.

4 Power Control

In the power control setting, each transmitter s_v either broadcasts at power 0 or broadcasts at some arbitrary integral power level $P_v \in [1, P_{\max}]$. Our main claim is that Braess's paradox is once again possible: there are networks in which adding power control can create worse equilibria. For illustration of such a network, see Fig. 1(b). We first observe the following relation between no-regret solutions with or without power control.

Observation 8 *Every no-regret solution in the uniform setting, is also a no-regret solution in the PC setting.*

Hence, we cannot expect the best no-regret solution in the PC setting to be smaller than the worst no-regret solution in the uniform setting. Yet, the paradox still holds.

Theorem 9. *There exists a configuration of links L satisfying $\mathcal{N}_{\min}^{PC}(L) \leq \mathcal{N}_{\min}(L)/c$ for some constant $c > 1$.*

Proof. Since any pure-Nash solution is also a no-regret solution, in this case it is sufficient to find a pure Nash equilibrium that uses power control that is a constant factor worse than $\mathcal{N}_{\min}(L)$, the worst no-regret behavior without power control. Let L be the set of links as illustrated in Fig. 1(b). Let $\eta = 0$, $\beta = 1.1$, $\alpha = 3$ and $P_{\max} = 2$. Let us first consider the case of power control, where each sender s_v transmits with power $P_v \in [1, P_{\max}]$. Let $N = \{\ell_1, \ell_2\}$ and s_1 transmits with $P_1 = P_{\max}$, s_2 transmits with $P_2 = 1$ and both other transmitters user power 0. It is easy to see that this is a pure Nash: the SINR of ℓ_1 is $\frac{P_{\max}/8}{1/8} > \beta$ and the SINR of ℓ_2 is $\frac{1/1}{P_{\max}/(5)^3} > \beta$, while even if they used power P_{\max} links ℓ_3 and ℓ_4 would not be able to overcome the interference caused by ℓ_1 since $P_{\max}/4.5^3 < 1/3.5^3$ (as $P_{\max} = 2$).

We will now analyze the worst no-regret behavior without power control. Given a history in which all players have regret at most ϵ , let p_i denote the fraction of times in which ℓ_i broadcasts. Note that ℓ_2 is always feasible, since under uniform powers the SINR of ℓ_2 is at least $1/(P_{1,2} + P_{3,2} + P_{4,2}) \geq \beta$ where $P_{1,2} = 1/125$ and $P_{3,2} = P_{4,2} = 1/(64 + 25)^{3/2}$. Thus $p_2 \geq 1 - \epsilon$, and link ℓ_2 succeeds at least $1 - \epsilon$ fraction of the time. Since $\beta > 1$ and the interfering sender s_2 is at the same distance to the receiver r_1 as its intended sender s_1 , it holds that ℓ_1 cannot succeed if s_2 transmits. Hence, if s_1 transmitted every time it would have average utility at most $-(1 - \epsilon) + \epsilon = -1 + 2\epsilon < 0$ (since $\epsilon < 1/2$), while if it never transmitted it would have average utility 0. Thus its average utility is at least $-\epsilon$. Since it can succeed only an ϵ fraction of the time (when link ℓ_2 is not transmitting), we have that $\epsilon - (p_1 - \epsilon) \geq -\epsilon$ and thus $p_1 \leq 3\epsilon$.

Now consider link ℓ_3 . Since s_1 is closer to r_3 than s_3 , the link ℓ_3 cannot succeed if s_1 transmits. However, it can succeed if s_1 is not transmitting and s_2 and s_4 are transmitting, since $P_{3,3} \geq \beta(P_{2,3} + P_{4,3})$ where $P_{3,3} = 1/4.5^3$, $P_{2,3} = 1/(16 + 3.5^2)^{3/2}$ and $P_{4,3} = 1/11.5^3$. Thus if s_3 had chosen to transmit at every time it would have average utility at least $1 - 3\epsilon - 3\epsilon = 1 - 6\epsilon > 0$ for $\epsilon \leq 1/6$. Thus ℓ_3 must have average utility of at least $1 - 7\epsilon$ and thus must succeed at least $1 - 7\epsilon$ fraction of the time. Link ℓ_4 is the same by symmetry. Thus the average number of successes in any ϵ -regret history is at least $1 - \epsilon + 2(1 - 7\epsilon) = 3 - 15\epsilon$, which proves the theorem. \blacksquare

We now prove that (as with IC) that the paradox cannot be too bad: adding power control cannot cost us more than a constant. The proof is very similar to that of Thm. 7 up to some minor yet crucial modifications.

Theorem 10. $\mathcal{N}_{\min}^{PC}(L) \geq \mathcal{N}_{\max}(L)/c$ for any set of links L and some constant $c \geq 1$.

Proof. Fix an arbitrary ϵ -regret history without power control, where all transmitters transmit with P_{\max} and let p_i be the fraction of rounds in which s_i attempts to transmit. Similarly, fix an arbitrary ϵ -regret history with power control, and let q_i be the fraction of rounds in which s_i attempts to transmit (at any nonzero power). By Lemma 1, it is sufficient to prove that $\sum_{i \in L} p_i \leq O(\sum_{i \in L} q_i)$.

Note that since the average number of successful connections in the history of the uniform case is $\sum p_i$, there must exist some set of connections $A \subseteq L$ that

transmitted successfully in some round $t \in [1, T]$ such that $|A| \geq \Omega(\sum_i p_i)$ and A is feasible when all senders transmit with power P_{\max} . Let $B = \{i : q_i \geq 1/2\}$ and let $A' = A \setminus B$. If $|B \cap A| \geq |A|/2$ then we are done, since then

$$\sum_{i \in L} p_i \leq O(|A|) \leq O(|B|) \leq O\left(\sum_{i \in L} q_i\right)$$

as required. So without loss of generality we will assume that $|B \cap A| < |A|/2$, and thus that $|A'| > |A|/2$. Note that A' is a subset of A , and so it is feasible in the uniform setting. Now let $\hat{A} = \{i \in A' : \sum_{j \in A'} a_i(j) \leq 2\}$ be an amenable subset of A' . By Fact 5(a), it holds that $|\hat{A}| \geq |A'|/2 \geq |A|/4$.

We have the following.

$$\sum_{i \in L} \sum_{j \in \hat{A}} q_i a_i(j) \leq c' \cdot \left(\sum_{i \in L} q_i\right). \quad (3)$$

where $\sum_{i \in L} \sum_{j \in \hat{A}} q_i a_i(j)$ is the average affectance of L on the set \hat{A} when every transmitter that transmitted in the PC history transmitted with power P_{\max} . The inequality follows by Fact 5(b).

On the other hand, we know that the q_i values correspond to an ϵ -regret history in the PC setting. Consider some $j \in \hat{A}$. Since $\hat{A} \subseteq A' = A \setminus B$, we know that $q_j < 1/2$ and thus the average utility of link ℓ_j is at most $1/2$. Let y_j be the fraction of time s_j would have succeeded has it transmitted in every round with full power P_{\max} . Since the average utility of the best single action is at most $1/2 + \epsilon$ it holds also that the utility of transmitting with full power is at most $1/2 + \epsilon$ as well, hence $y_j - (1 - y_j) \leq 1/2 + \epsilon$ and so $y_j \leq \frac{3}{4} + \frac{\epsilon}{2}$. In other words, in at least $1 - y_j = \frac{1}{4} - \frac{\epsilon}{2}$ fraction of the rounds the affectance of the other links on the link ℓ_j must be at least 1 (or else j could succeed in those rounds) when it attempted to transmit with P_{\max} . Thus the average affectance on ℓ_j in the PC history is at least $\frac{1}{4} - \frac{\epsilon}{2}$. Summing over all $j \in \hat{A}$, to get that

$$\sum_{j \in \hat{A}} \sum_{i \in L} q_i a_i(j) \geq \sum_{j \in \hat{A}} \frac{1 - 2\epsilon}{4} \geq \Omega(|\hat{A}|), \quad (4)$$

where the first inequality follows by the fact the in the true ϵ -regret history in the PC setting, the average affectance on ℓ_j is at least $\frac{1}{4} - \frac{\epsilon}{2}$ when s_j transmit with P_{\max} and all other transmitters $s_{j'}$ transmit with power at most P_{\max} .

Combining equations (3) and (4) (and switching the order of summations) implies that $|\hat{A}| \leq O(\sum_{i \in L} q_i)$. Since $|\hat{A}| \geq |A|/4 \geq \Omega(\sum_{i \in L} p_i)$, we get that $\sum_i p_i \leq O(\sum_{i \in L} q_i)$ as desired. \blacksquare

Corollary 2. *The price of total anarchy under the power control setting with maximum transmission energy P_{\max} is $\Theta(\log \Delta)$.*

Proof. The upper bound of $O(\log \Delta)$ is given by [3]. Let $m = \lceil \log \Delta \rceil$. The lower bound example is given by the nested link network described in Fig. 3. Let $OPT \subseteq L$ denote the solution of the optimal solution (i.e., maximum feasible set). According to [6] the link set $L \setminus \{\ell^*\}$ is feasible with exponentially increasing power level $P_i > 2P_{i-1}$. In particular, since $\beta > 1$, a necessary condition for feasibility is to maintain that $P_1 > P_2 > \dots > P_m$. Consider an ϵ -regret history in which $P^* = P_{\max}$. Since the link $\ell^* = (s^*, r^*)$ is sufficiently short, s^* always succeeds with any power level P^* and hence s^* transmits at least $1 - \epsilon$ fraction of the time. For every other link $\ell_i, i \in \{1, \dots, m\}$, if s^* transmits then s_i would fail to transmit for every transmission power $P_i \in [1, P_{\max}]$ as the interfering sender s^* is closer to r_i than the intended sender s_i and s^* transmits with full power. Consider link s_i and let p_i be the fraction of time it attempted to transmit (possibly with different power levels) in the ϵ -regret history. If s_i would have transmit in every round using any power $P_i^t \in [1, P_{\max}]$ in round $t \in \{1, \dots, T\}$, it would have average utility at most $-(1 - \epsilon) + \epsilon = -1 + 2\epsilon < 0$ (since $\epsilon < 1/2$), while if it never transmitted it would have average utility 0. Thus its average utility is at least $-\epsilon$. Since it can succeed only an ϵ fraction of the time (when link ℓ^* is not transmitting), we have that $\epsilon - (p_i - \epsilon) \geq -\epsilon$ and thus $p_i \leq 3\epsilon$. Overall, the average number of attempted transmission is at most $1 + 3 \cdot m/\epsilon \leq 2$ for $\epsilon = 1/(3m)$. Since $|OPT| = m$, the lemma follows. ■




Fig. 3: Price of anarchy in the PC setting: a network in which the gap between the optimal solution and a no-regret solution is $\Omega(\log \Delta)$.

5 Power Control with Interference Cancellation (PIC)

In this section we consider games in the power control with IC setting where *transmitters* can adopt their transmission energy in the range of $[1, P_{\max}]$ and in addition, *receivers* can employ interference cancelation. This setting is denote as *PIC* (power control+IC). We show that Braess's paradox can once again happen and begin by comparing the PIC setting to the setting of power control without IC and to the most basic setting of uniform powers.

Lemma 3. *There exists a set of links L and constant $c > 1$ such that*

- (a) $\mathcal{N}_{\min}^{PIC}(L) \leq \mathcal{N}_{\min}^{PC}(L)/c.$
- (b) $\mathcal{N}_{\min}^{PIC}(L) \leq \mathcal{N}_{\min}(L)/c.$

Proof. Consider the 6-transmitters network illustrated in Fig. 4. Let $\alpha = 8$, $\beta = 10$, $\eta = 0$, $b = 25^{1/8}$ and $P_{\max} = 2$. Set $\hat{\epsilon} \ll 1$ to be a sufficiently small constant. Intuitively, $\hat{\epsilon}$ should be sufficiently small so that the receivers r_4 , r_5 and r_6 consider the set of transmitters s_1, s_2 and s_3 are co-located at the same point, say at the point of transmitter s_1 . For simplicity, we therefore analyze the network example as if this is the case while keeping in mind that this effect can be achieved by setting $\hat{\epsilon}$ to be sufficiently small. In addition, if one insists on minimal transmitter-receiver distance $d(s_i, r_j) \geq 1$, then the entire set of distances can be multiplied by $1/\hat{\epsilon}$, without affecting the analysis (since there is no ambient noise, such normalization factor cancels out). We begin by considering the value of solutions in an ϵ -regret histories without IC. Let p_i be the fraction of time that s_i transmits in the worst PC setting for every $i \in \{1, \dots, 6\}$.

It is easy to see that link 1 will always be successful, since the length of the link $\hat{\epsilon}$ is set to be sufficiently small. Since ℓ_1 has at most ϵ -regret, this implies that $p_1 \geq 1 - \epsilon$.

On the other hand, whenever s_1 transmits it is clear that the link ℓ_2 cannot be successful even if s_2 transmits with full power and s_1 transmits with power 1 as $\beta/\hat{\epsilon}^8 > P_{\max}/(\beta^{16} \cdot \hat{\epsilon}^8)$.

So if s_2 transmitted every time it would have average utility at most $-(1 - \epsilon) + \epsilon = -1 + 2\epsilon < 0$ (since $\epsilon < 1/2$), while if it never transmitted it would have average utility 0. Thus its average utility is at least $-\epsilon$. Since it can succeed only an ϵ fraction of the time (when link 1 is not transmitting), we have that $\epsilon - (p_2 - \epsilon) \geq -\epsilon$ and thus $p_2 \leq 3\epsilon$. In the same manner, it also holds that $p_3 \leq 3\epsilon$.

Now consider link ℓ_4 . As long as links ℓ_2 and ℓ_3 do not transmit ℓ_4 always succeeds even if it transmits with power $P_4 = 1$ and s_1, s_5, s_6 transmit with full power P_{\max} . This holds since the amount of interference it suffers is at most $2/b^8 + 4/(16b^8) \leq 1/\beta$. Since the fraction of time that both ℓ_2 and ℓ_3 are transmitting is at most $p_2 + p_3 \leq 6\epsilon$, it holds that if ℓ_4 always transmits it succeeds at least $1 - 6\epsilon$ fraction of the time and hence its average utility is $1 - 6\epsilon - 6\epsilon = 1 - 12\epsilon$ which is strictly positive by taking a sufficiently small ϵ . Therefore it holds that in ϵ -regret history, the average utility of ℓ_4 is at least $1 - 13\epsilon$, concluding that $p_4 \geq 1 - 13\epsilon$. By symmetry, the same holds for ℓ_5

and ℓ_6 . Overall, the value of any no-regret solution in PC setting is at least $\mathcal{N}_{\min}^{PC}(L) = \sum_{i=1}^6 p_i \geq 1 - \epsilon + 6\epsilon + 3 - 39\epsilon \geq 4 + o(1)$.

Let us now analyze what happens when using interference cancellation with power control and bound $\mathcal{N}_{\min}^{PIC}(L)$. In this case, it is sufficient to consider a specific pure Nash solution. Let $N = \{s_1, s_2, s_3\}$ be transmitting with full power P_{\max} . It then holds that r_2 and r_3 can cancel the signal of s_1 and that both r_2 and r_3 can decode (resp., cancel) the signal of s_2 (this is achieved since s_1, s_2 and s_3 form an exponential chain with respect to r_1, r_2, r_3 though it looks the same with respect to the receivers of r_4, r_5, r_6). We now show that in this case, s_4 cannot succeed even if they transmit with full power P_{\max} . This holds since $\beta \cdot 6/b^8 > P_{\max}$. The same holds also for s_5 and s_6 , concluding that $\mathcal{N}_{\min}^{PIC}(L) \leq 3$ and that $\mathcal{N}_{\min}^{PC}(L)/\mathcal{N}_{\min}^{PIC}(L) = 4/3 + o(1)$ as required. Consider claim (b). This case is practically the same as the case of claim (a), in particular it holds that $\mathcal{N}_{\min}^{PC}(L) = \mathcal{N}_{\min}(L)$ (since $\ell_1, \ell_4, \ell_5, \ell_6$ always succeed when s_2 and s_3 are not transmitting and without IC, s_2 and s_3 cannot succeed if s_1 transmits). Therefore it also holds that $\mathcal{N}_{\min}(L)/\mathcal{N}_{\min}^{PIC}(L) = 4/3 + o(1)$. The lemma follows. \blacksquare




Fig. 4: Schematic illustration of a network in which playing IC with power control might generate no-regret solutions that are worse by a factor of $\Omega(1)$ than no-regret solution in a setting without interference cancellation, with or without power control.

Moreover, we proceed by showing that PIC can hurt the network by more than a constant when comparing PIC equilibria to IC equilibria. For an illustration of such a network, see Fig. 1(c).

Theorem 11. *There exists a set of links L and constant $c > 1$ such that the best pure Nash solution with PIC is worse by a factor of $\Omega(\log \Delta)$ than the worst no-regret solution with IC.*

Proof. Consider the following m -transmitters network for $m = \Omega(\log \Delta)$ described in Fig. 1(c). There is a set of $m - 4$ receivers $R' = \{r_5, \dots, r_m\}$ located at $(0, 0)$. Their corresponding transmitters s_i are located at $(\sqrt{10}(1 + 2^{i/2}), 0)$ for every $i \in \{5, \dots, m\}$. The remaining first links are located as follows: s_1 is located at $(-3, 1)$, s_2 is located at $(-3, -1)$, the receiver r_2 is located in the middle, between s_1 and s_2 and the receiver r_1 is located on top of s_1 at $(-3, 2)$. The receiver r_4 is located at $(\sqrt{19}, 0)$, its transmitter s_4 is located at $(\sqrt{19}, -1)$. Finally, the transmitter s_3 is located at $(\sqrt{19}, 1)$ and its receiver r_3 is located at $(\sqrt{19}, 2)$. Let $\beta = 1.5, \eta = 0, \alpha = 2, P_{\max} = 2$.

We begin by considering the PIC setting and showing that in this case there exists a unique Nash in which only s_1, \dots, s_4 are transmitting. Hence, the total value of any pure Nash is 4. First, observe that s_1 and s_3 always succeed even if they transmit with power 1 and all other links transmit with power P_{\max} . For example, for s_1 we get that the received signal strength is 1, the amount of interference from s_2 is at most $P_{2,1} = 2/9$ and $P_{3,1} + P_{4,1} \leq 4/20$ in addition the amount of interference from all other transmitters s_5, \dots, s_m is at most $1/4^4$. Hence $1 \geq \beta(2/9 + 2/20 + 1/4^4)$. Next, note that also s_2 and s_4 always succeed. Consider s_2 for example. There are two cases. If s_1 transmits with power $P_1 = 1$ then with interference cancellation, s_2 can succeed if it transmits with power 2. This holds since in this case $P_{2,2} = 2, P_{1,2} = 1, \sum_{i>3} P_i \leq 4/20 + 1/4^4$, hence r_2 can successfully decode s_2 as $P_{2,2} \geq \beta(1 + 21/80)$. Alternatively, if s_1 transmits with $P_1 = 2$ and $P_2 = 1$ then r_2 can successfully cancel s_1 (by the same computation as above) and subsequently it can decode s_2 since in this case the received signal strength of s_2 is $1 \geq \beta(4/20 + 1/4^4)$. By the same argumentation we also have that s_4 can always succeed as well. Note that any Nash solution has the following structure: exactly one of the transmitters s_1 and s_2 transmit with power 1 and the other with power 2, and exactly one of the transmitters s_3 and s_4 transmit with power 1 and the other with power 2. We now show that s_i always fails for every $i \geq 5$. Consider some pure Nash solution, then by the above, s_1, \dots, s_4 are active, without loss of generality, let s_1 be transmitting with power 1 and s_3 be transmitting with power 2. It then holds that the received signal strength of transmitter s_1 at r_i is $P_{1,i} = 1/10$ and the received signal strength of s_3 is $P_{3,i} = 2/20$. Hence, $P_{1,i} = P_{3,i}$ and in addition, these signals are stronger than any other signal, i.e., $P_{3,i} > P_{\max}/(10 \cdot 2^{2j})$ for every $j \geq 5$. Hence, r_i cannot cancel the strongest signal (as $P_{1,i} < \beta \cdot P_{3,i}$) and in particular it cannot decode its intended message. Concluding that any Nash solution in the IC setting consists of exactly 4 links.

We proceed by considering the worst no-regret solution in the IC setting (without power control). Let p_i be the fraction of time that s_i transmits in an IC history. Note that since the links of s_1 and s_3 are sufficiently short they always succeed if they attempt to transmit hence in any ϵ -regret history, they transmit at least $p_1, p_3 \geq 1 - \epsilon$ fraction of the time. Consider link ℓ_2 and note that it cannot succeed if s_1 transmits since $P_{1,2} = P_{2,2}$ hence if s_2 transmitted every time it would have average utility at most $-(1 - \epsilon) + \epsilon = -1 + 2\epsilon < 0$ (since $\epsilon < 1/2$), while if it never transmitted it would have average utility 0. Thus its average

utility is at least $-\epsilon$. Since it can succeed only an ϵ fraction of the time (when link 1 is not transmitting), we have that $\epsilon - (p_2 - \epsilon) \geq -\epsilon$ and thus $p_2 \leq 3\epsilon$. The same holds for p_4 . Finally, consider some link ℓ_j for $j \geq 5$. Note that ℓ_j can always succeed if s_2 and s_4 are not transmitting. This holds since $P_{1,j} \geq \beta(P_{3,j} + 1/32)$ (i.e., hence, s_j can cancel the strongest signal of s_1), $P_{3,j} \geq \beta \cdot 1/32$ (i.e., it can successfully cancel the second strongest signal) and in addition by the structure of the exponential transmitters chain, the remaining of cancelations of s_5, \dots, s_j are successful. Hence, if s_j transmitted every time it would have average utility at $1 - 6\epsilon - 6\epsilon$ and therefore it transmits at least $p_5 \geq 1 - 13\epsilon$ fraction of the time. Since it holds for every ℓ_j for $j \geq 5$ we have that the total value in any ϵ -regret history is at least $\mathcal{N}_{\min}^{IC} = \sum_{i \geq 5} p_i = \Omega(\log \Delta) - o(1)$. Since the total value of the unique pure Nash in the PIC setting is 4, the lemma follows. \blacksquare

Corollary 3. *There exists a set of links L satisfying that $\mathcal{N}_{\min}^{PIC}(L) \leq (c/\log \Delta) \cdot \mathcal{N}_{\min}^{IC}(L)$.*

As in the previous sections, we show that our examples are essentially tight.

Theorem 12. *For every set of links L it holds that there exists a constant $c \geq 1$ such that*

- (a) $\mathcal{N}_{\min}^{PIC}(L) \geq \mathcal{N}_{\max}(L)/c$.
- (b) $\mathcal{N}_{\min}^{PIC}(L) \geq \mathcal{N}_{\max}^{PC}(L)/c$.
- (c) $\mathcal{N}_{\min}^{PIC}(L) \geq \mathcal{N}_{\max}^{IC}(L)/(c \log \Delta)$.

Proof. Part (a) follows by the argumentation as in Lemma 2 and Part (b) follows trivially by Part (a) and by Theorem 10.

Consider part (c). Fix an arbitrary ϵ -regret history with IC and uniform powers and let p_i be the fraction of rounds in which s_i attempts to transmit. Similarly, fix an arbitrary ϵ -regret history with PIC, and let q_i be the fraction of rounds in which s_i attempts to transmit (at any nonzero power). By Lemma 1, it is sufficient to prove that $\sum_{i \in L} p_i \leq O(\sum_{i \in L} q_i)$. Note that since the average number of successful connections in the best history of the IC case is $\sum_{i \in L} p_i$, there must exist some set of connections $A \subseteq L$ that transmitted successfully in some round $t \in [1, T]$ such that $|A| \geq \sum_{i \in L} p_i$ and A is feasible when all transmitters transmit with power 1 and employing IC. Then by Cl. 2, there exists a subset $\tilde{A} \subseteq A$ of cardinality $|\tilde{A}| \geq |A|/O(\log \Delta)$ that is feasible without IC and with uniform power level P_{\max} . Let $B = \{i : q_i \geq 1/2\}$ and let $A' = \tilde{A} \setminus B$. If $|B \cap \tilde{A}| \geq |\tilde{A}|/2$ then we are done, since then

$$O(\log \Delta) \cdot \sum_{i \in L} p_i \leq |\tilde{A}| \leq 2|B| \leq 2 \sum_i q_i$$

as required. So without loss of generality we will assume that $|B \cap \tilde{A}| < |\tilde{A}|/2$, and thus that $|A'| > |\tilde{A}|/2$. Note that A' is a subset of \tilde{A} , and so it is feasible in the IC setting. Now let $\hat{A} = \{i : \sum_{j \in A'} a_i(j) \leq 2\}$ be an amenable subset of A' . By Fact 5(a), it holds that $|\hat{A}| \geq |A'|/2 \geq |A|/4$. By Fact 5(b) then implies that

for any link $\ell_i \in L$, its total affectance on \widehat{A} is small: $a_i(\widehat{A}) = \sum_{j \in \widehat{A}} a_i(j) \leq c'$ for some constant $c \geq 0$. Note that Fact 5(b) holds for the case where all interfering transmitters transmit with the same power P_{\max} . Since in the power control setting $P_j \leq P_{\max}$, the actual affectance can only be lower. We have that

$$\sum_{i \in L} \sum_{j \in \widehat{A}} q_i a_i(j) \leq c' \cdot \left(\sum_{i \in L} q_i \right), \quad (5)$$

where $\sum_{j \in \widehat{A}} q_i a_i(j)$ is the average affectance in the PIC history if all links that attempt to transmit, transmit with full power. The inequality follow by Fact 5(b). On the other hand, we know that the q_i values correspond to the worst ϵ -regret history in the PIC setting. Consider some $j \in \widehat{A}$, hence $q_j < 1/2$ (i.e., the average utility of link ℓ_j is at most $1/2$). Let y_j be the fraction of time s_j would have succeeded has it transmitted in every round with full power P_{\max} and in the IC setting. Since the average utility of the best single action is at most $1/2 + \epsilon$ in particular the average utility of transmitting with full power P_{\max} is at most $1/2 + \epsilon$ as well. It therefore holds that $y_j - (1 - y_j) \leq 1/2 + \epsilon$ or the that $y_j \leq \frac{3}{4} + \frac{\epsilon}{2}$. In other words, in at least $1 - y_j = \frac{1}{4} - \frac{\epsilon}{2}$ fraction of the rounds the affectance of the other links on the link ℓ_j must be at least 1 (or else j could succeed in those rounds). Thus the average affectance (over the length of the history) on ℓ_j (assuming it transmits with full power P_{\max}) is at least $\frac{1}{4} - \frac{\epsilon}{2}$. Summing over all $j \in \widehat{A}$, to get that

$$\sum_{j \in \widehat{A}} \sum_{i \in L} q_i a_i(j) \geq \frac{1 - 2\epsilon}{4} \geq \Omega(|\widehat{A}|). \quad (6)$$

Combining equations (5) and (6) (and switching the order of summations) implies that $|\widehat{A}| \leq O(\sum_{i \in L} q_i)$. Since $|\widehat{A}| \geq |A|/O(\log \Delta) \geq \sum_{i \in L} p_i/O(\log \Delta)$ as desired. \blacksquare

Finally, as a direct consequences of our result, we obtain a tight bound for the price of total anarchy in the PIC setting.

Corollary 4. *For every set of links L it holds that the price of total anarchy with PIC is $\Theta(\log(\Delta \cdot P_{\max}))$.*

Proof. Let $OPT \subseteq L$ denote some optimal solution without PIC, i.e., the set of transmitters forming a maximum β -feasible set, and let $OPT^{PIC} \subseteq L$ denote some optimal solution with PIC. We first show that $|OPT^{PIC}|/\mathcal{N}_{\min}^{PIC}(L) = O(\log(\Delta \cdot P_{\max}))$. According to Theorem 2 of [3], it holds that $\mathcal{N}_{\min}(L) \geq |OPT|/c'$ for some constant $c' > 1$. Hence, by combining with Theorem 12(a) we get that

$$\mathcal{N}_{\min}^{PIC}(L) \geq \mathcal{N}_{\min}(L)/c \geq |OPT|/(c \cdot c') \geq |OPT^{PIC}|/O(\log(\Delta \cdot P_{\max})),$$

where the last inequality follows by Lemma 2. Finally, we show that this is tight by noting that the example of Fig. 2 can be easily modified so that there

are $m = \Omega(\log \Delta \cdot P_{\max})$ transmitters (the location and powers should be carefully set) whose receivers are positioned at the origin. By placing the additional transmitters \tilde{s}_1 and \tilde{s}_2 at equidistance from the origin, we get that these two links block the cancellation sequence at the m receivers; in addition by setting these links to be sufficiently short, we get that in \tilde{s}_1 and \tilde{s}_2 transmit most of the time. Thus, any ϵ -regret history is of total value of $2 + o(1/n)$ but $OPT^{PIC} = \Omega(\log(\Delta \cdot P_{\max}))$. ■

6 Decreasing the SINR Threshold

We begin by showing that in certain cases the ability to successfully decode a message at a lower SINR threshold results in *every* no-regret solution having lower value than *any* no-regret solution at higher β . For an illustration of such a network, see Fig. 1(d).

Theorem 13. *There exists a set of links L and constants $1 < \beta' < \beta$ such that $\mathcal{N}_{\max}^{\beta'}(L) \leq \mathcal{N}_{\min}^{\beta}(L)/c$ for some constant $c > 1$.*

Proof. Let β' be slightly greater than 1 (say 101/100), and let $\beta = 4$. Let L be as in Fig. 1(d) with $\alpha = 2$, some small constant noise η , and $c = \sqrt{1/(\beta\eta)}$ and $a = b = \frac{2}{3}c$.

We first analyze $\mathcal{N}_{\min}^{\beta}(L)$. Fix an ϵ -regret history under threshold β . Note that the SINR of ℓ_3 is at most β , simply due to background noise. So by making β infinitesimally larger, ℓ_3 can never succeed, and thus it transmits in at most an ϵ fraction of the times. Whenever ℓ_3 does not transmit, ℓ_1 can succeed no matter what ℓ_2 does, since the SINR of ℓ_1 is at least

$$\frac{1/a^2}{\eta + \frac{1}{(3a)^2}} = \frac{(9/4)\beta\eta}{\eta + (1/4)\beta\eta} = \frac{9}{2} > \beta.$$

Hence choosing to transmit in every round would get ℓ_1 utility at least $1 - 2\epsilon$, and thus it must succeed in at least $1 - 3\epsilon$ fraction of the times. The same is true for ℓ_2 by symmetry. So we have that in any ϵ -regret history under threshold β , the average number of successes is at least $2 - 6\epsilon$.

We now analyze $\mathcal{N}_{\max}^{\beta'}(L)$. Fix an ϵ -regret history under threshold β' . We first claim that ℓ_3 can always succeed, no matter what ℓ_1 and ℓ_2 do. This is because its SINR is at least

$$\frac{1/c^2}{\eta + \frac{2}{(a+b)^2 + c^2}} = \frac{1/c^2}{\eta + \frac{2}{25c^2/9}} = \frac{1/c^2}{\eta + \frac{18}{25}\frac{1}{c^2}} = \frac{\beta\eta}{\eta + \frac{18}{25}\beta\eta} = \frac{4}{1 + (72/25)} = \frac{100}{97} > \beta'.$$

Thus ℓ_3 must transmit in at least a $1 - \epsilon$ fraction of the rounds. In any round where ℓ_3 transmits, neither ℓ_1 nor ℓ_2 can succeed since $a = b$ and $\beta' \geq 1$. Thus the average number of successes is at most $(1 - \epsilon) + 2\epsilon = 1 + \epsilon$. ■

We now show that the gap between the values of no-regret solution for different SINR threshold values is bounded by a constant.

Lemma 4. For every $1 \leq \beta' \leq \beta$ and every set of links L satisfying that $P_{vv} \geq 2\beta \cdot \eta$ for every $\ell_v \in L$, it holds that $\mathcal{N}_{\min}^{\beta'}(L) \geq \mathcal{N}_{\max}^{\beta}(L)/c$ for some constant $c \geq 1$.

For simplicity, for the rest of the proof we will assume that $P_{vv} \geq 2\beta \cdot \eta$ for every value of β we consider. This does not hold in the bad example of Theorem 13, but that example can easily be modified so that it holds (although it makes the example messier). This is a standard assumption, and limits strange phenomena due to thresholding. In addition, we isolate the affect of the SINR threshold value and restrict attention to uniform powers.

Given an SINR threshold β' and links ℓ_w and ℓ_v define

$$a_w^{\beta'}(v) = \min \{1, c_v(\beta')P_{w,v}/P_{v,v}\}.$$

The affectness of a subset of (resp. on) a subset of links L' is denoted by $a_{L'}^{\beta'}(w)$ ($a_w^{\beta'}(L')$). We now show a simple lemma which will help us prove that the above example is tight.

Observation 14 Let L be a 2β -feasible set without noise, $\eta = 0$ then L is a β -feasible set with noise $\eta > 0$ satisfying that $P_{vv} \geq 2\beta\eta$ for every $\ell_v \in L$.

Since L is a 2β -feasible set, it holds that $\sum_{u \in L} P_{uv}/P_{vv} \leq 1/(2\beta)$ for every $\ell_v \in L$. Since $P_{vv} \geq 2\beta\eta$, it holds that $1/(2\beta) \leq 1/\beta - \eta/P_{vv}$ and hence that $\sum_{u \in L} P_{uv}/P_{vv} \leq 1/\beta - \eta/P_{vv}$. Concluding that L is a β -feasible set with $\eta > 0$.

7 Conclusion

In this paper we have shown that Braess's paradox can strike in wireless networks in the SINR model: improving technology can result in worse performance, where we measured performance by the average number of successful connections. We considered adding power control, interference cancellation, both power control and interference cancellation, and decreasing the SINR threshold, and in all of them showed that game-theoretic equilibria can get worse with improved technology. However, in all cases we bounded the damage that could be done.

There are several remaining interesting open problems. First, what other examples of wireless technology exhibit the paradox? Second, even just considering the technologies in this paper, it would be interesting to get a better understanding of when exactly the paradox occurs. Can we characterize the network topologies that are susceptible? Is it most topologies, or is it rare? What about random wireless networks? Finally, while our results are tight up to constants, it would be interesting to actually find tight constants so we know precisely how bad the paradox can be.

References

1. J.G. Andrews. Interference cancellation for cellular systems: a contemporary overview. *Wireless Communications, IEEE*, 12:19 – 29, 2005.

2. M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proc. 28th Conf. of IEEE Computer and Communications Societies (INFOCOM)*, 2009.
3. E.I. Åsgeirsson and P. Mitra. On a game theoretic approach to capacity maximization in wireless networks. In *INFOCOM*, 2011.
4. C. Avin, A. Cohen, Y. Haddad, E. Kantor, Z. Lotker, M. Parter, and D. Peleg. Sinr diagram with interference cancellation. In *SODA*, pages 502–515, 2012.
5. C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty. SINR diagrams: Towards algorithmically usable sinr models of wireless networks. In *Proc. 28th Symp. on Principles of Distributed Computing (PODC)*, 2009.
6. C. Avin, Z. Lotker, and Y.-A. Pignolet. On the power of uniform power: Capacity of wireless networks with bounded resources. In *ESA*, pages 373–384, 2009.
7. Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, STOC '08*, pages 373–382, New York, NY, USA, 2008. ACM.
8. D. Braess. über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung*, 12(1):258–268, 1968.
9. Max H. M. Costa and Abbas A. El Gamal. The capacity region of the discrete memoryless interference channel with strong interference. *IEEE Trans. Inf. Theor.*, 33(5):710–711, 1987.
10. Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing, STOC '06*, pages 71–78, New York, NY, USA, 2006. ACM.
11. Michael Dinitz. Distributed algorithms for approximating wireless network capacity. In *INFOCOM*, pages 1397–1405, 2010.
12. R. H. Etkin, D. N.C. Tse, and Hua Wang. Gaussian interference channel capacity to within one bit. *IEEE Trans. Inf. Theor.*, 54(12):5534–5562, 2008.
13. O. Goussevskaia, Y.A. Oswald, and R. Wattenhofer. Complexity in geometric SINR. In *Proc. 8th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 100–109, 2007.
14. O. Goussevskaia, R. Wattenhofer, M.M. Halldórsson, and E. Welzl. Capacity of arbitrary wireless networks. In *INFOCOM*, pages 1872–1880, 2009.
15. P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Trans. Information Theory*, 46(2):388–404, 2000.
16. M.M. Halldórsson and P. Mitra. Wireless connectivity and capacity. In *SODA*, pages 516–526, 2012.
17. M.M. Halldórsson and R. Wattenhofer. Wireless communication is in apx. In *ICALP*, pages 525–536, 2009.
18. Erez Kantor, Zvi Lotker, Merav Parter, and David Peleg. The topology of wireless communication. In *Proceedings of the 43rd annual ACM symposium on Theory of computing, STOC '11*, pages 383–392. ACM, 2011.
19. T. Kesselheim. A constant-factor approximation for wireless capacity maximization with power control in the sinr model. In *Proc. ACM-SIAM Symp. on Discrete Algorithms (SODA 2011)*, 2011.
20. T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. *Distributed Computing*, 2010.
21. Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th annual conference on Theoretical aspects of computer science, STACS'99*, pages 404–413, 1999.

22. Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Inf. Comput.*, 108(2):212–261, February 1994.
23. T. Moscibroda and R. Wattenhofer. The complexity of connectivity in wireless networks. In *Proc. 25th Conf. of IEEE Computer and Communications Societies (INFOCOM)*, 2006.
24. John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
25. Y. Freund P. Auer, N. Cesa-Bianchi and R.E. Schapire. The nonstochastic multi-armed bandit problem. *SIAM J. Comput.*, 32(1):48–77, January 2003.