Participation Incentives on a Wireless Random Access Erasure Collision Channel

[Short talk]

Vasilis Gkatzelis
Drexel University
Department of Computer Science
Philadelphia, PA
gkatz@drexel.edu

Steven Weber
Drexel University
Dept. of Electrical and Computer Engineering
Philadelphia, PA
sweber@coe.drexel.edu

ABSTRACT

Random medium access on wireless channels has two key characteristics: heterogeneous channel quality across users (due to the variable wireless channel conditions with the access point across users) and packet collisions (due to random access) at the access point. The design of optimal channel contention probabilities is non-trivial on account of the need to balance between under-utilization (no transmission attempts) and over-utilization (channel collisions) of the wireless channel. The wireless random access erasure collision channel presented in this paper is a parsimonious abstraction of these phenomena. We consider a scenario wherein the base station provides a reward to users in proportion to the rate of successfully received packets, and users incur a cost in proportion to their contention probability. The objective is to select a reward such that the sum-user delivered rate in the equilibrium of the induced game is (approximately) optimal, i.e., such that only higher quality users are incentivized to contend the channel. We use the price of anarchy (PoA) measure and study the extent to which appropriate reward mechanisms can yield good PoA bounds.

CCS CONCEPTS

•Networks →Network economics; •Theory of computation →Quality of equilibria;

KEYWORDS

random access; erasure channel; collision channel; price of anarchy.

ACM Reference format:
Vasilis Gkatzelis and Steven Weber. 2017. Participation Incentives on a Wireless Random Access Erasure Collision Channel. In *Proceedings of NetEcon'17, Cambridge, MA, USA, June 27, 2017,* 7 pages.
DOI: 10.1145/3106723.3106726

1 INTRODUCTION

We consider a wireless scenario consisting of users (devices), each wishing to transmit to one of several available access points (AP).

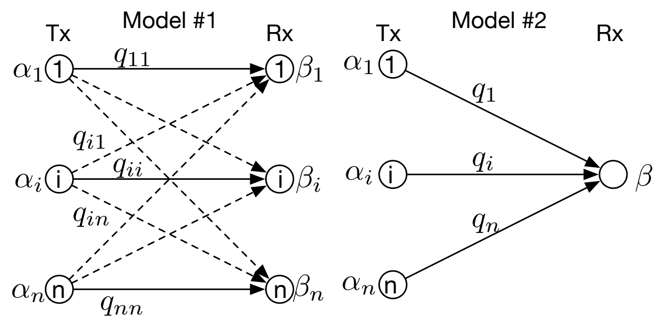

Figure 1: Two wireless random access collision erasure channel models. Model #1: n transmitter (Tx) receiver (Rx) pairs. Model #2: n Tx and 1 Rx. Each Tx selects a contention probability $\alpha_i \in [0,1]$, yielding a reception probability at each Rx $\beta_j \in [0,1]$. Channel nonerasure probabilities are denoted $Q = (q_{ij})$ (left) and $q = (q_i)$ (right). A message is successfully received if no other message arrives.

Random medium access on wireless channels has two key characteristics: heterogeneous channel quality across users (due to the inherent variable wireless channel conditions with the AP across users) and transmission collisions (due to random access) at the AP. While modern random medium access protocols (e.g., WiFi) employ collision avoidance mechanisms (e.g., carrier sensing), we ignore such mechanisms in this first paper in order to focus more directly on our core objective, namely, the design and performance analysis of mechanisms to incentivize users to select optimal contention probabilities with a given AP, minimizing the price of anarchy (PoA).

The selection of contention probabilities in a (time-slotted) random access protocol is non-trivial on account of the need to avoid both under- and over-utilization of the (shared) wireless channel. Contention probabilities that are too "low" result in under-utilization, i.e., too many idle time slots wherein no user attempts transmission. On the other hand, contention probabilities that are too "high" result in over-utilization, i.e., too many time slots wherein multiple users attempt transmission, all of which may be lost due to packet collision at the AP. The contention probability vectors define the rate of successfully received packets which, in the erasure collision channel described below, corresponds to the fraction of times that exactly one packet arrives at the AP.

NetEcon'17, Cambridge, MA, USA
© 2017 ACM. 978-1-4503-5089-1/17/06…$15.00
DOI: 10.1145/3106723.3106726

The contention probabilities that maximize the rate of successfully received packets at the AP will clearly depend upon the channel qualities: users with stronger channels may need to be given larger contention probabilities than users with weaker channels, since transmissions by stronger channels are more likely to be successfully received by the AP. Moreover, it is clear that the channel-optimal contention probabilities must be selected jointly. Finally, and most importantly, appropriate economic mechanisms may be required in order to incentivize the users to select (near)-optimal channel contention probabilities. Such mechanisms induce a game among the users, and we measure their quality over Nash equilibrium outcomes of this game. In particular, we use the price of anarchy (PoA) and price of stability (PoS) measures, i.e., the ratio of the packet delivery rate of the optimal contention probability vector over that of the worst-case (resp. best-case) Nash equilibrium contention probability vector. Characterizing the PoA and PoS in this setting, for simple reward-based mechanisms, is the primary objective of this paper.

1.1 Random access erasure collision channel

To achieve the above goals is it necessary to create a mathematical model of the wireless channel that retains the essential features mentioned above, i.e., heterogeneous channel quality and channel collisions. Our model combines two well-established (slotted time) wireless model abstractions: the random access channel and the random erasure channel, see Fig. 1. We briefly comment on the model here but defer a formal description to §2. The random access component is captured by *a*) the contention probabilities α_i for each user i, describing the probability that each user contends for access in each time slot, and *b*) the collision channel wherein any receiver i will successfully receive any single arriving packet, but will lose all packets when multiple packets arrive simultaneously. The random erasure component characterizes the channel between each transmitter (user) i and each receiver (AP) j via a nonerasure probability $q_{ij} \in [0, 1]$, such that each transmitted packet is independently delivered with probability q_{ij}, or erased with probability $1 - q_{ij}$.

While both the random access mechanism and the wireless channel are admittedly idealized, we will see that interesting relationships are found, even in this simple model, in the context of mechanism design and PoA analysis.

1.2 Related work

Our discussion of related work is brief due to space constraints, and we restrict our focus to the most pertinent papers that address PoA for wireless random access scenarios. [5] studies a random access game on a collision channel –without erasures – and moreover assumes users have a utility function that is the logarithm of the success probability, leading to proportional-fair optimal contention probabilities. [3] study the Nash equilibrium structure and the PoA for a congestion game for associations between stations and access points using the "airtime" as a congestion measure. [1] focuses on a class of collision-aware rate adaptation algorithms in the setting of wireless LANs. [4] studies a cognitive radio network with secondary users competing for transmissions on idle primary channels; their focus is on contention strategies across multiple channels. [2] study a spatial spectrum access game in terms of equilibrium structure and PoA. Key differences with our work include the channel model, utility function, and their inclusion of a backoff mechanism. Our more restrictive system model enables us to obtain somewhat more explicit results. Our own recent work on random access and random erasure channels has addressed broadcast delay on erasure channels [8], the Aloha stability region [7], and Aloha throughput–fairness tradeoffs [9].

1.3 Contributions and outline

We study two wireless random access erasure collision channel models which exhibit this interesting balance between under- and over-utilization (§2). In the first model we provide some results regarding the optimal contention probabilities (§3.1). In the second model which, as we show, can be viewed as a special case of the first model, we characterize the optimal contention probabilities in much more detail (§3.2). We then focus on a a game-theoretic analysis of the users' incentives. We study the performance of simple reward-based mechanisms and provide upper and lower bounds on the PoA and PoS of the induced games §4. Our main result is the existence of a simple mechanism that guarantees a PoS of 2 for all instances.

2 MODEL

Basic notation. Time is slotted, and the packet transmission time is one time slot. Natural numbers are denoted $\mathbb{N}$. $a \equiv b$ denotes a, b are equal by definition. A Bernoulli random variable (RV) with parameter $p \in [0, 1]$ is denoted $\mathsf{X} \sim \mathrm{Ber}(p)$, i.e., $\mathbb{P}(\mathsf{X} = 1) = p = 1 - \mathbb{P}(\mathsf{X} = 0)$; RVs are written in a sans-serif font, e.g., $\mathsf{X}, \mathsf{A}, \mathsf{E}$.

We consider two models (c.f. Fig. 1); the latter will be seen to be a special case of the former (c.f. Fact 1).

2.1 Model #1: n Tx-Rx pairs

The first model is referred to as "n Tx-Rx pairs", and is illustrated in Fig. 1 (left). There are $n \in \mathbb{N}$ transmitter-receiver (Tx-Rx) pairs, indexed $i \in [n]$, for $[n] \equiv \{1, \ldots, n\}$. Each Rx j wishes to receive information from its paired Tx j, and messages received at Rx j from other Txs $i \neq j$ are of no value to j. This model is natural in wireless scenarios where the association of Tx to Rx is fixed, as in certain cellular and *ad hoc* scenarios (e.g., the dumbell model).

Random access channel. The Txs employ slotted random access (i.e., slotted Aloha) to contend for the (shared) wireless channel: in each time slot each of the n Txs makes an independent random decision $\mathsf{A}_i \sim \mathrm{Ber}(\alpha_i)$ of whether or not to transmit (contend), with $\alpha_i \in [0, 1]$ the contention probability. Observe the contention decision RVs $\mathsf{A} = (\mathsf{A}_i, i \in [n])$ are independent in time and space, but not identically distributed. The contention probability vector $\alpha \equiv (\alpha_i, i \in [n])$ is the system control / design parameter.

Broadcast erasure channel. The wireless channel is a time slotted broadcast erasure channel characterized by the $n \times n$ matrix $Q = (q_{ij})$, where $q_{ij} \in [0, 1]$ is the nonerasure probability from Tx i to Rx j, i.e., the probability that a packet sent by Tx i arrives at Rx j. That is, the n Rxs are in a common interference domain. Erasures are independent across time slots and across the n^2 wireless channels, i.e., in each time slot we generate an independent $n \times n$ $\{0, 1\}$-valued matrix $\mathsf{E} = (\mathsf{E}_{ij})$, with $\mathsf{E}_{ij} \sim \mathrm{Ber}(q_{ij})$, and $\mathsf{E}_{ij} = 0$ (1) means channel (i, j) erases (does not erase) any message from Tx i en route to receiver j.

Collision channel. We assume each Rx $j \in [n]$ lies within a collision domain among Txs $\mathcal{T}_j = \{i \in [n] : q_{ij} > 0\}$, meaning that Rx j is able to successfully recover a message from its paired Tx j in a time slot provided *i)* Tx j contends (with probability α_j), *ii)* Tx j's message is not erased (with probability q_{jj}), and *iii)* there is no arriving message at Rx j from any other Tx $i \in \mathcal{T}_j$ (with probability $1 - \alpha_i q_{ij}$ for Tx i). Observe this last condition holds for Tx i if either *i)* Tx i does not contend in that time slot, or *ii)* Tx i contends in that time slot but its message is erased on the channel to Rx j. By the assumed temporal and channel independence property of the broadcast erasure channel, the probability of a successful delivery at Rx j is

$$\beta_j = \beta_j(\alpha, Q) \equiv \alpha_j q_{jj} \prod_{i \neq j} (1 - \alpha_i q_{ij}), \; j \in [n]. \tag{1}$$

Message arrival probabilities. Given the control α and the channel Q define the $n \times n$ message arrival probability matrix $P = (p_{ij})$, where $p_{ij} \equiv \alpha_i q_{ij}$ is the probability Tx i contends and the message is not erased en route to Rx j. If $A = (\alpha, \cdots, \alpha)$ is the $n \times n$ matrix comprised of n identical rows α, then $P = A \circ Q$, for $\circ$ the Hadamard product. In this notation (1) becomes, with $\bar{p}_{ij} \equiv 1 - p_{ij}$:

$$\beta_j(P) = p_{jj} \prod_{i \neq j} \bar{p}_{ij}, \; j \in [n]. \tag{2}$$

Optimal sum reception probability. Let $\beta = (\beta_j, j \in [n])$ be the reception probabilities at each Rx. A natural performance objective for evaluating candidate contention probability vectors is the resulting rate of successfully received packets.[1] Let $B \equiv \beta_1 + \cdots + \beta_n$ be the sum reception probability. The key optimization problem with the first model is to maximize this sum over all possible contention probabilities α, equivalently, over all $P \in \mathcal{P}(Q) \equiv \{A \circ Q : A = (\alpha, \cdots, \alpha), \alpha \in [0, 1]^n\}$:

$$\max_{P \in \mathcal{P}(Q)} B(P) = \sum_{j \in [n]} p_{jj} \prod_{i \neq j} \bar{p}_{ij} \tag{3}$$

Let $B^*(Q)$ denote the maximum value in (3), and let $\mathcal{A}^*_B(Q), \mathcal{P}^*_B(Q)$ denote the values of α and P, respectively, that achieve that maximum.

2.2 Model #2: n Tx and 1 Rx

We term the second model "n Txs and 1 Rx", illustrated in Fig. 1 (right). In this model we consider there to be n Txs, indexed $i \in [n]$, and a single Rx. In contrast to the previous model, in this scenario the Rx has value packets received from all n Txs. The nonerasure probabilities from the Txs to the Rx are denoted $q = (q_i, i \in [n])$, where without loss of generality we assume $1 > q_1 > \cdots > q_n > 0$. Under the collision channel, a message is successfully delivered at the Rx when exactly one message arrives at the Rx:

$$\tilde{B}(\alpha, q) = \sum_{i \in [n]} \alpha_i q_i \prod_{j \neq i} (1 - \alpha_j q_j). \tag{4}$$

Define $p \equiv (p_i, i \in [n])$, with $p_i \equiv \alpha_i q_i$ the message arrival probability from Tx i and $\bar{p}_i \equiv 1 - p_i$. The optimization problem is to maximize the reception probability (from any Tx) over all possible contention probabilities α, equivalently, over all $p \in \mathcal{P}(q) \equiv \{\alpha \circ q, \alpha \in [0, 1]^n\}$:

$$\max_{p \in \mathcal{P}(q)} \tilde{B}(p) = \sum_{j \in [n]} p_j \prod_{i \neq j} \bar{p}_i. \tag{5}$$

Let $\tilde{B}^*(q)$ denote the maximum value in (5), and let $\mathcal{A}^*_{\tilde{B}}(q), \mathcal{P}^*_{\tilde{B}}(q)$ denote the values of α and p, respectively, that achieve that maximum.

FACT 1. *Model #2 (with n-vector q) and its optimization problem (5) is a special case of model #1 (with $n \times n$ matrix Q) and its optimization problem (3). Namely, construct Q from q by $q_{jj} = q_j$ for all j and $q_{ij} = q_i$ for all $i \neq j$.*

3 OPTIMALITY RESULTS

3.1 Model #1

The following result shows it suffices to restrict the contention probabilities α from $[0, 1]^n$ to $\{0, 1\}^n$ when maximizing the sum reception probability B (due to space constraints, the proof is omitted).

PROPOSITION 3.1. *In Model #1, the maximum sum reception probability $B^*(Q)$ in (3) is attainable by extremal contention probabilities $\alpha \in \{0, 1\}^n$.*

By Prop. 3.1, the sum utility maximization (3) is a combinatorial maximization problem. If $\bar{q}_{ij} \equiv 1 - q_{ij}$, then for a given set S of Txs let

$$B(S) \equiv \sum_{j \in S} q_{jj} \prod_{i \in S - j} \bar{q}_{ij}, \tag{6}$$

be the sum reception probability if every Tx $i \in S$ contents with probability $\alpha_i = 1$ and every other Tx does not content. For "small" n (say fewer than 20), the most plausible case under the assumption that the receivers lie in a common collision domain, Prop. 3.1 ensures that we can solve (3) by simply enumerating $B(S)$ over all 2^n possible $S \subseteq [n]$. For unrestricted values of n, we show that this optimization problem is NP-hard.

LEMMA 3.2. *The problem of computing a set S of Txs that maximizes $B(S)$ in (3) is NP-hard.*

The function $B(S)$ is neither monotone nor submodular, which precludes the use of some standard algorithms for approximating this objective. To verify this, consider a simple instance with $n = 3$ identical Txs such that $q_{ij} = 3/4$ for all $i, j \in \{1, 2, 3\}$. In this case $B(\emptyset) = 0$, $B(\{1\}) = 3/4$, $B(\{1, 2\}) = 3/8$, and $B(\{1, 2, 3\}) = 9/64$. An open problem we plan to consider in future work is the approximability of this objective for general values of n. However, in the rest of this paper we focus on the interesting special case defined by Model #2, for which we can not only compute but also characterize the structure of the optimal solution.

3.2 Model #2

The following corollary is an immediate consequence of Fact 1 and Prop. 3.1.

COROLLARY 3.3. *In Model #2, the maximum reception probability in (5) is attainable via extremal contention probabilities, i.e., there exists $\alpha \in \{0, 1\}^n$ with $\tilde{B}(\alpha, q) = B^*(q)$.*

[1] Other objectives are also natural, e.g., incorporating a notion of fairness of the delivery rates across contending users.

Our following theorem (the proof of which is omitted due to space constraints) shows that the set of active Txs in the set that maximizes the reception probability in (5) comprises the Txs with the best nonerasure probabilities. The number of these Txs depends on q as follows:

$$k^*(q) = \begin{cases} n, & \text{if } \sum_{i=1}^{n} \frac{q_i}{\bar{q}_i} \leq 1 \\ \min\{k : \sum_{i=1}^{k} \frac{q_i}{\bar{q}_i} > 1\}, & \text{o/w.} \end{cases} \quad (7)$$

Theorem 3.4. *An optimal solution in Model #2 is* $\alpha^* = (1, \ldots, 1, 0, \ldots, 0)$, *comprising* $k^*(q)$ *ones (for the users with best nonerasure probabilities), followed by* $n - k^*(q)$ *zeros.*

Therefore the optimal solution in model #2 can be computed via "greedy packing": ordering the Txs by nonincreasing nonerasure probabilities q and then greedily add Txs to the active set in that order, until the sum of their $\tilde{s}_i(q_i)$ values exceeds 1, or until all of them are active.

4 INCENTIVES IN MODEL #2

Rather than assuming the existence of a centralized controller who can decide which Txs should transmit and to enforce that outcome, we now consider a more realistic decentralized model, where the Rx needs to provide appropriate incentives for the Txs.

We assume that each time a Tx transmits, it suffers some cost c, so the expected cost per slot of Tx i under contention probabilities α is $c \cdot \alpha_i$. To incentivize transmission, the Rx rewards each Tx i for every message transmitted by i that is successfully received. If π is the payment, or reward,[2] that a Tx receives for each such message, then the expected payment of Tx i in α is $\pi p_i \prod_{j \neq i} \bar{p}_j$. Facing the expected costs and rewards described above, the preferences of each Tx i are then defined by the quasilinear utility function

$$u_i(\alpha) \;\equiv\; \pi p_i \prod_{j \neq i} \bar{p}_j - c\, \alpha_i \;=\; \alpha_i \left(\pi q_i \prod_{j \neq i} (1 - \alpha_j q_j) - c \right).$$

We henceforth assume that a Tx is willing to transmit only if this leads to positive utility. That is, if every Tx $j \neq i$ is transmitting with probability α_j, then Tx i chooses to transmit with $\alpha_i > 0$ only if this yields $u_i(\alpha) > 0$. This holds only if

$$q_i \prod_{j \neq i} (1 - \alpha_j q_j) \;>\; \frac{c}{\pi}. \quad (8)$$

More generally, each Tx i, given every other Tx's transmission probabilities $\alpha_{-i} = (\alpha_1, \ldots, \alpha_{i-1}, \alpha_{i+1}, \ldots, \alpha_n)$, chooses α_i aiming to maximize $u_i(\alpha)$. We therefore focus on the Nash equilibria of the induced game.

Definition 4.1. An outcome with contention probabilities α is a Nash equilibrium of the induced game if for every Tx i and every $\alpha_i' \in [0, 1]$, $u_i(\alpha) \geq u_i(\alpha_i', \alpha_{-i})$.

In the rest of the paper we approach this problem from the perspective of the Rx who, given nonerasure probabilities q, can choose the reward value π, aiming to maximize $\tilde{B}(\alpha, q) = \sum_{i \in [n]} \alpha_i q_i \prod_{j \neq i} (1 - \alpha_j q_j)$. We study how the Nash equilibria of the induced game depend on the choice of π and we conclude by analyzing the performance of mechanisms that choose π as a function of q.

[2] Depending on the application at hand, this payment can either be monetary or take some other form, such as that of an artificial currency.

4.1 Equilibrium Structure

As we observed above, in every equilibrium α Tx i chooses $\alpha_i > 0$ only if Inequality (8) holds. In fact, if this inequality holds, the utility of Tx i is strictly increasing in α_i, and hence i will choose $\alpha_i = 1$ in equilibrium. Therefore, in every equilibrium of this game, if $\alpha_i > 0$ then $\alpha_i = 1$.

Proposition 4.2. *Every Nash equilibrium of this game corresponds to extremal contention probabilities* $\alpha \in \{0, 1\}^n$. *Hence, the only active Txs belong to a set* $S \subseteq [n]$, *and they all transmit with probability* 1.

Note that a change in the value of π may change the incentives of the Txs because it affects Inequality (8). As a result, such a change may also affect the set of equilibria, i.e, the set S of agents that may end up transmitting in equilibrium. Let $\gamma \equiv \frac{c}{\pi}$, and α be an equilibrium for the game induced by π and q. Given the cost parameter c, the choice of the reward parameter π by the Rx also defines γ, so we henceforth approach the problem as if the Rx can choose the γ as a function of q. If S is the set of Txs transmitting in some equilibrium α, then we know that every Tx $i \in S$ must have a positive utility, i.e.,

$$q_i \prod_{j \in S - i} \bar{q}_j \;>\; \gamma \qquad \forall i \in S. \quad (9)$$

Also, since α is an equilibrium, every other agent $i \notin S$ should have no incentive to start transmitting. Every such Tx has $\alpha_i = 0$, so it experiences no cost or reward in α, and $u_i(\alpha) = 0$. Hence, α is an equilibrium if and only if $u_i(\alpha_i', \alpha_{-i}) \leq 0$ for all $\alpha_i' \in [0, 1]$. According to the definition of $u_i(\cdot)$, this is equivalent to $\pi q_i \prod_{j \in S} \bar{q}_j - c \leq 0$, i.e.,

$$q_i \prod_{j \in S} \bar{q}_j \;\leq\; \gamma \qquad \forall i \notin S. \quad (10)$$

Inequalities (9) and (10) are both necessary and sufficient conditions for set S to be a set of active Txs in equilibrium. Combining these two inequalities we get the following observation.

Observation 1. *For a given* q *and a set* $S \subseteq [n]$, *there exists a choice of* γ *such that* S *is the set of active Txs in equilibrium if and only if for all* $i \in S$ *and* $j \notin S$:

$$q_j \prod_{h \in S} \bar{q}_h < q_i \prod_{h \in S - i} \bar{q}_h, \quad (11)$$

equivalently, $q_j < q_i / \bar{q}_i$.

Given some q and a set S, let

$$\ell_S \;\equiv\; \left(\max_{i \notin S} q_i \right) \prod_{j \in S} \bar{q}_j$$

$$r_S \;\equiv\; \min_{i \in S} \left(q_i \prod_{j \in S - i} \bar{q}_j \right) = \left(\min_{i \in S} \frac{q_i}{\bar{q}_i} \right) \prod_{j \in S} \bar{q}_j. \quad (12)$$

If $S = [n]$, i.e., there are no Txs other than those in S, we let $\ell_S = 0$. Note that for a given set S, $q_j \prod_{h \in S} \bar{q}_h$ takes its maximum value, ℓ_S, when $j \notin S$ is the Tx with the largest nonerasure probability q_j among the ones not in S. Also $q_i \prod_{h \in S - i} \bar{q}_h$ takes its minimum value, r_S, when $i \in S$ is the Tx with the smallest nonerasure probability q_i in S. With this notation, the observation above implies the following.

Observation 2. *A set $S \subseteq [n]$, is a potential set of active Txs in equilibrium if and only if*[3] $\ell_S < r_S$ *and* $\gamma \in [\ell_S, r_S)$.

Using these observations we now show the following lemma.

Lemma 4.3. *For any problem instance, if A is the optimal set of Txs transmitting then* $\ell_A < r_A$*. Hence, any choice of* $\gamma \in [\ell_A, r_A)$ *ensures that the optimal solution is an equilibrium, i.e., that the price of stability is* 1.

Proof. As shown in Thm. 3.4, the optimal A is the set of the top k^* Txs in terms of their nonerasure probabilities. Therefore, $\ell_A = q_{k+1}\prod_{h\in S}\bar{q}_h$, and $r_A = \frac{q_k}{\bar{q}_k}\prod_{h\in S}\bar{q}_h$. But, since $q_{k+1} \le q_k < \frac{q_k}{\bar{q}_k}$, we conclude that $\ell_A < r_A$. □

4.2 Price of Anarchy

We conclude this section by studying the extent to which γ can be chosen in a way that optimizes the performance in the worst equilibrium outcome, i.e., the price of anarchy. A mechanism with a price of anarchy of 1 would need to choose a value of γ as a function of q so that $\gamma \in [\ell_A, r_A)$ but $\gamma \notin [\ell_B, r_B)$ for any other (suboptimal) set B. We show that this is possible when $n \le 3$, but not possible for $n \ge 4$. However, we show that for any value of n, there exist mechanisms such that the price of anarchy is bounded by 2.

To gain some intuition, we first focus on instances where the number of Txs is $n = 2$, for which we show that there exists a choice of γ ensuring that the only equilibrium is the optimal solution.

Lemma 4.4. *For any instance with* $n = 2$ *with an optimal set of transmitters A, choosing* $\gamma = \ell_A + \epsilon$ *or* $\gamma = r_A - \epsilon$ *for some arbitrarily small constant* $\epsilon > 0$ *yields PoA* = 1.

Proof. If $q_1 \le 0.5$ then Theorem 3.4 implies that the optimal solution has both Txs transmitting. In this case, the range of γ values that make $\{1, 2\}$ an equilibrium is the following: Inequality (9) implies that $q_2\bar{q}_1 > \gamma$ and Inequality (10) does not apply since there are no Txs $i \notin \{1, 2\}$. But, for any $\gamma < q_2\bar{q}_1$ it is easy to verify that any suboptimal set $S \in \{\{1\}, \{2\}, \emptyset\}$ would not be an equilibrium since Inequality (10) would be violated, i.e., an agent $i \notin S$ would have incentive to transmit.

On the other hand, if $q_1 > 0.5$, the optimal solution comprises the first Tx alone, i.e., the optimal set of transmitters is $A = \{1\}$. For this set we have $\ell_A = q_2\bar{q}_1$ and $r_A = q_1$. It is easy to verify that any value of $\gamma \in [\ell_A, r_A)$ ensures that $\{1, 2\}$ would not be an equilibrium. In particular, this would imply that Tx 2 gets a positive utility, and it would contradict the fact that any $\gamma \in [\ell_A, r_A)$ makes the set A an equilibrium (Tx 2 would want to join). The interesting case is to ensure that $B = \{2\}$ will not be an equilibrium. Note that $\ell_B = q_1\bar{q}_2$ and $r_B = q_2$, so B can be an equilibrium only if $q_1\bar{q}_2 < q_2$. Note there exist q values such that this inequality is satisfied, e.g., $q_1 = 0.75$ and $q_2 = 0.5$. Also, note that these values satisfy the following inequalities

$$\ell_A < \ell_B < r_B < r_A. \tag{13}$$

This means that choosing a value of γ near the endpoints of the $[\ell_A, r_A)$ interval ensures that A is an equilibrium but B is not, hence ensuring that the PoA is 1. □

[3]For the special case when $\ell_S = 0$, there is no payment π that yields $c/\pi = 0$, so $\gamma \in (\ell_S, r_S)$ instead.

Next, consider instances with $n = 3$ Txs. In contrast to instances with 2 Txs, where the values of γ that yield a PoA of 1 lie near both endpoints of the $[\ell_A, r_A)$ interval, instances with 3 Txs require γ near the ℓ_A endpoint.

Lemma 4.5. *For any instance with* $n = 3$ *with an optimal set of transmitters A, choosing* $\gamma = \ell_A + \epsilon$ *for some arbitrarily small constant* $\epsilon > 0$ *yields price of anarchy 1.*

Proof. There are three possible optimal solutions in instances involving 3 Txs: $\{1\}, \{1, 2\}, \{1, 2, 3\}$. When the optimal solution is $A = \{1\}$, choosing $\gamma = \ell_A + \epsilon$ ensures that this is the only equilibrium, and the proof of that is analogous to the proof of Lemma 4.4. Also, when the optimal solution is $A = \{1, 2, 3\}$, then it is easy to verify that no subset would be an equilibrium when $\gamma = \ell_A + \epsilon = \epsilon$. In particular, for any such subset B we have $\ell_B > \epsilon$.

When the optimal solution is $A = \{1, 2\}$, then $\ell_A = q_3\bar{q}_1\bar{q}_2$. In this case, we know that for any $\gamma \in [\ell_A, r_A)$ the set $\{1, 2, 3\}$ could not be an equilibrium, as this would contradict the fact that these values of γ make A an equilibrium (Tx 3 would prefer to transmit). Similarly, no subset of A could be an equilibrium, so it suffices to show that the no $S \in \{\{1, 3\}, \{2, 3\}, \{3\}\}$ can be an equilibrium either. To verify that this is the case for $S = \{1, 3\}$, $S = \{2, 3\}$, and $S = \{3\}$, note that the first one leads to

$$\ell_S = q_2\bar{q}_1\bar{q}_3 > q_3\bar{q}_1\bar{q}_3 > q_3\bar{q}_1\bar{q}_2 = \ell_A, \tag{14}$$

the second one leads to

$$\ell_S = q_1\bar{q}_2\bar{q}_3 > q_3\bar{q}_2\bar{q}_3 > q_3\bar{q}_1\bar{q}_2 = \ell_A, \tag{15}$$

and the third one leads to $\ell_S = q_2\bar{q}_3 > q_2\bar{q}_1\bar{q}_3 > \ell_A$ (by Inequality (14)). Hence, setting $\gamma = \ell_A + \epsilon$ for an appropriately small ϵ ensures that these sets are not equilibria. □

These initial results provide some hope that an appropriate choice of γ always exists such that the only equilibrium in the induced game corresponds to the optimal solution. However, as the following lemma shows, such a choice of γ is not guaranteed to exist when $n \ge 4$. In particular, although a price of anarchy of 1 can be guaranteed when the optimal set is $\{1\}$, $\{1, 2, 3\}$, or $\{1, 2, 3, 4\}$, this is not true when the optimal set is $\{1, 2\}$.

Lemma 4.6. *There exists no mechanism that guarantees a price of anarchy of* 1 *for* $n \ge 4$.

Proof. Consider an instance involving $n = 4$ Txs with nonerasure probabilities $q_1 = 0.499$, $q_2 = 0.455$, $q_3 = 0.444$, and $q_4 = 0.433$. In this case, the optimal set is $A = \{1, 2\}$ and the interval $[\ell_A, r_A)$ is $[0.1212, 0.2279]$. For the suboptimal sets $S_1 = \{2, 3, 4\}$, $S_2 = \{1, 3\}$, and $S_3 = \{3, 4\}$, the corresponding intervals for S_1, S_2, and S_3 are approximately $[0.0857, 0.1312]$, $[0.1267, 0.2224]$, and $[0.1573, 0.2407]$. Note there is no value of $\gamma \in [\ell_A, r_A)$ that is not also contained in one of the intervals that stabilizes a suboptimal set. □

Since there is no mechanism that can guarantee a PoA of 1, the natural next question is whether there exists a mechanism with PoA bounded by some constant for all n. Our main result shows that choosing $\gamma = r_A - \epsilon$ for some arbitrarily small $\epsilon > 0$ for any instance where the optimal set is A guarantees a PoA less than 2. We first show that this choice of γ guarantees that the number of agents transmitting in equilibrium is not fewer than those in A.

Lemma 4.7. *For any given instance with A as the optimal set of Txs, if we let $\gamma = r_A - \epsilon$ for some arbitrarily small $\epsilon > 0$, any equilibrium set B satisfies $|B| \geq |A|$.*

Proof. Let $|A| = k$ and assume that $|B| = k' \leq k-1$, and $|A \cap B| = \tau$. Also, let $q_{\max}$ correspond to the highest q value among Txs in $A \setminus B$. Since B is assumed to be an equilibrium, the Tx in $A \setminus B$ whose nonerasure probability is $q_{\max}$ would not have positive utility if it unilaterally deviated and started transmitting, i.e.,

$$q_{\max} \prod_{j \in B} \bar{q}_j \;\leq\; q_k \bar{q}_1 \ldots \bar{q}_{k-1} - \epsilon \qquad \Longrightarrow$$
$$q_{\max} \prod_{j \in B \setminus A} \bar{q}_j \;<\; \frac{q_k}{\bar{q}_k} \prod_{j \in A \setminus B} \bar{q}_j.$$

Note that for all $j \in B \setminus A$ we have $q_j < q_k$, for all $j \in A \setminus B$ we have $q_j \geq q_k$, and for at least one $j \in A \setminus B$ we have $q_j = q_{\max}$. Hence,

$$(\bar{q}_k)^{k'-\tau} \;<\; \frac{q_k \bar{q}_{\max}}{q_{\max} \bar{q}_k} (\bar{q}_k)^{k-\tau-1} \qquad \Longrightarrow$$
$$1 \;<\; \frac{q_k \bar{q}_{\max}}{q_{\max} \bar{q}_k} \qquad \Longrightarrow$$
$$\frac{q_{\max}}{\bar{q}_{\max}} \;<\; \frac{q_k}{\bar{q}_k}.$$

As $q/\bar{q}$ is increasing in q, this is a contradiction. □

Theorem 4.8. *The price of anarchy of the game induced by choosing $\gamma = r_A - \epsilon$ for some arbitrarily small $\epsilon > 0$ (where A is the optimal set of Txs) is less than 2.*[4]

Proof. For every set $S \neq A$, if it is an equilibrium when $|A| = k$ and $\gamma = r_A - \epsilon$, then Inequality (9) yields

$$q_i \prod_{j \in S-i} \bar{q}_j \;>\; q_k \bar{q}_1 \ldots \bar{q}_{k-1} - \epsilon \qquad \forall i \in S. \tag{16}$$

This inequality implies that the contribution of every agent $i \in S$ to the $\tilde{B}(S)$ objective is at least as high as the minimum contribution among all agents in A the optimal objective value $\tilde{B}(A)$. If $q_{\min} \equiv \min_{i \in S} q_i$, note that $q_{\min} < q_k$ since $S \neq A$, and $|S| \geq |A|$ by Lemma 4.7. Hence, Inequality (16), when applied to the Tx $i \in S$ with nonerasure probability $q_{\min}$, implies that, for appropriately small values of ϵ and every $i \in A \cap S$

$$q_i \prod_{j \in S-i} \bar{q}_j \;\geq\; q_j \prod_{j \in A-i} \bar{q}_i \qquad \forall i \in A \cap S, \tag{17}$$

i.e., the contribution of any $i \in A \cap S$ to $\tilde{B}(S)$ is weakly higher than its contribution to $\tilde{B}(A)$.

Let $q_{\max} \equiv \max_{i \in A \setminus S} q_i$ be the highest nonerasure probability among the Txs in $A \setminus S$. Inequality (11) implies that

$$q_{\max} \;<\; \frac{q_{\min}}{\bar{q}_{\min}} \;<\; \frac{q_k}{\bar{q}_k}. \tag{18}$$

If $q_{\max} > 0.5$, then Theorem 3.4 implies that $|A| = 1$ and Inequality (16) implies that $S = A$, i.e., the PoA is 1. On the other hand, if $q_{\max} \leq 0.5$, i.e., $\bar{q}_{\max} \geq 0.5$, then Inequality (18) implies that

$$q_{\max} \bar{q}_k \;<\; q_k \;<\; 2 q_k \bar{q}_{\max}. \tag{19}$$

[4]Note that a Tx's utility in the induced game can be less than its marginal contribution to the objective, so there is no obvious connection with valid utility games [6].

If we let $\mathcal{S} \equiv \sum_{i \in A \cap S} q_i \prod_{j \in A-i} \bar{q}_j \leq \sum_{i \in A \cap S} q_i \prod_{j \in S-i} \bar{q}_j$ and use Inequality (17), then we get the first inequality below:

$$\begin{aligned}
\frac{\tilde{B}(A)}{\tilde{B}(S)} &= \frac{\sum_{i \in A} q_i \prod_{j \in A-i} \bar{q}_j}{\sum_{i \in S} q_i \prod_{j \in S-i} \bar{q}_j} \\
&= \frac{\sum_{i \in A \cap S} q_i \prod_{j \in A-i} \bar{q}_j + \sum_{i \in A \setminus S} q_i \prod_{j \in A-i} \bar{q}_j}{\sum_{i \in A \cap S} q_i \prod_{j \in S-i} \bar{q}_j + \sum_{i \in S \setminus A} q_i \prod_{j \in S-i} \bar{q}_j} \\
&\leq \frac{\mathcal{S} + \sum_{i \in A \setminus S} q_i \prod_{j \in A-i} \bar{q}_j}{\mathcal{S} + \sum_{i \in S \setminus A} q_i \prod_{j \in S-i} \bar{q}_j} \\
&\leq \frac{\sum_{i \in A \setminus S} q_i \prod_{j \in A-i} \bar{q}_j}{\sum_{i \in S \setminus A} q_i \prod_{j \in S-i} \bar{q}_j} \\
&\leq \frac{|A \setminus S| (q_{\max}/\bar{q}_{\max}) \prod_{j=1}^{k} \bar{q}_j}{|S \setminus A| (q_k/\bar{q}_k) \prod_{j=1}^{k} \bar{q}_j} \\
&\leq \frac{q_{\max}/\bar{q}_{\max}}{q_k/\bar{q}_k} < 2
\end{aligned}$$

The second inequality uses the fact that $\mathcal{S} > 0$: since the ratio is at least one, removing $\mathcal{S}$ from both the denominator and numerator increases it. The third inequality uses Inequality (16), Lemma 4.7, and the fact that the Tx with nonerasure probability $q_{\max}$ has the highest contribution to $\tilde{B}(A)$ among all Tx in $A \setminus S$. This holds since $|S| \geq |A|$ implies $|A \setminus S| \leq |S \setminus A|$. The last inequality uses (19). □

5 CONCLUSION

Our main results are for Model #2, where n devices seek to transmit over heterogeneous erasure channels to an access point. We provided an explicit characterization of the optimal subset of transmitters, and a bound of 2 on the PoA when the receiver uses a simple reward-based mechanism to incentivize the transmitters. Our ongoing work is towards more precise PoA characterizations.

ACKNOWLEDGEMENT

The authors wish to acknowledge assistance from Dr. Nan Xie on earlier versions of the proof of Thm. 3.4.

REFERENCES

[1] Prasanna Chaporkar, Alexandre Proutiere, and Bozidar Radunoviae. 2010. Rate Adaptation Games in Wireless LANs: Nash Equilibrium and Price of Anarchy. In *Proceedings of IEEE INFOCOM*. San Diego, CA.

[2] Xu Chen and Jianwei Huang. 2015. Spatial Spectrum Access Game. *IEEE Transactions on Mobile Computing* 14, 3 (March 2015), 646–659.

[3] Ozgur Ercetin. 2008. Association games in IEEE 802.11 wireless local area networks. *IEEE Transactions on Wireless Communications* 7, 12 (December 2008), 5136–5143.

[4] Lok Man Law, Jianwei Huang, and Mingyan Liu. 2012. Price of Anarchy for Congestion Games in Cognitive Radio Networks. *IEEE Transactions on Wireless Communications* 11, 10 (October 2012), 3778–3787.

[5] Pavan Nuggehalli, Jennifer Price, and Tara Javidi. 2008. Pricing and QoS in Wireless Random Access Networks. In *Proceedings of IEEE Global Communications Conference (GLOBECOM)*. New Orleans, LA.

[6] Adrian Vetta. 2002. Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19*. 416.

[7] Nan Xie, John MacLaren Walsh, and Steven Weber. 2017. Properties of an Aloha-like stability region. *IEEE Transactions on Information Theory* 63, 5 (May 2017), 3181–3208.

[8] Nan Xie and Steven Weber. 2017. Delay on broadcast erasure channels under random linear combinations. *IEEE Transactions on Information Theory* 63, 3 (March 2017), 1631–1661.

[9] Nan Xie and Steven Weber. accepted April, 2017. On the Aloha throughput-fairness tradeoff. *IEEE Transactions on Information Theory* (accepted April, 2017).