A Converse Result on Convergence Time for Opportunistic Wireless Scheduling

1

Michael J. Neely

University of Southern California

https://viterbi-web.usc.edu/~mjneely/

Abstract

This paper proves an impossibility result for stochastic network utility maximization for multi-user wireless systems, including multiple access and broadcast systems. Every time slot an access point observes the current channel states for each user and opportunistically selects a vector of transmission rates. Channel state vectors are assumed to be independent and identically distributed with an unknown probability distribution. The goal is to learn to make decisions over time that maximize a concave utility function of the running time average transmission rate of each user. Recently it was shown that a stochastic Frank-Wolfe algorithm converges to utility-optimality with an error of $O(\log(T)/T)$, where T is the time the algorithm has been running. An existing $\Omega(1/T)$ converse is known. The current paper improves the converse to $\Omega(\log(T)/T)$, which matches the known achievability result. It does this by constructing a particular (simple) system for which no algorithm can achieve a better performance. The proof uses a novel reduction of the opportunistic scheduling problem to a problem of estimating a Bernoulli probability p from independent and identically distributed samples. Along the way we refine a regret bound for Bernoulli estimation to show that, for any sequence of estimators, the set of values $p \in [0, 1]$ under which the estimators perform poorly has measure at least 1/6.

I. INTRODUCTION

This paper establishes the fundamental learning rate for network utility maximization in wireless opportunistic scheduling systems, such as multiple access systems and broadcast systems. The recent work [2] shows that a stochastic Frank-Wofe algorithm with a vanishing stepsize achieves a utility optimality gap that decays like $O(\log(T)/T)$, where T is the time the algorithm is in operation. It does this without a-priori knowledge of the channel state probabilities. This paper establishes a matching converse. A simple example system is constructed for which all algorithms have an error gap of at least $\Omega(\log(T)/T)$. Specifically, we construct a system with channel states parameterized by an unknown probability $q \in [0, 1]$ such that for any algorithm, there is a set $Q \subseteq [0, 1]$ with measure at least 1/6 under which the algorithm performs poorly. This is done by a novel reduction of the opportunistic scheduling problem to a problem of estimating a Bernoulli probability p from independent and identically distributed (i.i.d.) Bernoulli samples. Along the way, a refined statement regarding the regret of Bernoulli estimation is developed.

A general structure for the class of opportunistic scheduling systems is as follows: The system is assumed to operate over slotted time $t \in \{0, 1, 2, ...\}$. There are *n* users. Every slot $t \in \{0, 1, 2, ...\}$ an access point allocates a vector $X[t] = (X_1[t], ..., X_n[t])$ for transmission of independent data belonging to each user. In the case of wireless multiple access systems, the *n* users transmit their data over uplink channels to the access point. It is assumed they use a coordinated scheme that allows successful decoding of all transmissions at the scheduled bit rates X[t]. In the case of wireless broadcast systems, the access point transmits data for each user over downlink channels at the scheduled bit rates X[t].

The set of all transmission rate vectors that are available on a particular slot t can change from one slot to the next. This can arise from time-varying connection properties such as channel states that vary due to device mobility. We model this time-variation by a random state vector $S[t] \in \mathbb{R}^m$ that is observed by the access point at the start of every slot t(where m is a positive integer that can be different from n). Assume that $\{S[t]\}_{t=0}^{\infty}$ is i.i.d. over slots with some distribution $F_S(s) = P[S[0] \leq s]$ for all $s \in \mathbb{R}^m$ (where inequality is taken entrywise). The distribution function $F_S(s)$ is unknown. Define $\mathcal{D}(S[t])$ as the decision set for slot t, being the set of all $(X_1[t], \ldots, X_n[t])$ vectors that can be chosen on slot t when the channel state vector is S[t].

The structure of $\mathcal{D}(S[t])$ depends on the network. For example, a multiple access network might allow only one user to transmit per slot. In this case we can define $S[t] = (S_1[t], \ldots, S_n[t])$ as a vector of *channel states*, where $S_i[t]$ represents the transmission rate available to user *i* on slot *t* if that user is selected for transmission. Then $\mathcal{D}(S[t])$ is a set that contains *n* vectors:

$$\mathcal{D}(S[t]) = \{(S_1[t], 0, 0, ..., 0), (0, S_2[t], 0, ..., 0), ..., (0, 0, ..., 0, S_n[t))\}$$

where the *i*th vector in this set corresponds to choosing user *i* for transmission. More sophisticated wireless signaling schemes can allow the set $\mathcal{D}(S[t])$ to contain vectors with multiple positive components. The set $\mathcal{D}(S[t])$ can be uncountably infinite

A partial version of this work was accepted for presentation at the IEEE INFOCOM conference, 2020 [1].

This work was supported in part by one or move of: NSF CCF-1718477, NSF SpecEES 1824418.

in the case when transmission rates depend on an uncountably infinite set of power allocation levels that are available for scheduling.

Every slot $t \in \{0, 1, 2, ...\}$ the access point observes S[t] and chooses $X[t] \in \mathcal{D}(S[t])$ in such a way that, over time, the following problem is solved:

Maximize:
$$\liminf_{T \to \infty} \phi\left(\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\left[X[t]\right]\right)$$
(1)

Subject to:
$$X[t] \in \mathcal{D}(S[t]) \quad \forall t \in \{0, 1, 2, \ldots\}$$
 (2)

where $\phi(x_1, \ldots, x_n)$ is a given real-valued utility function of the average user transmission rates. The function ϕ is assumed to be concave and entrywise nondecreasing. Let ϕ^* be the optimal utility value, which considers all possible algorithms that operate over an infinite time horizon, including algorithms that have perfect knowledge of the probability distribution F_S , and even including *non-causal* algorithms that have knowledge of future states $\{S[t]\}_{t=0}^{\infty}$.¹ It is challenging to design a (causal) scheduling algorithm that achieves a utility that is close to ϕ^* , particularly when the distribution F_S is unknown to the network controller. Algorithms that are causal (so that they have no knowledge of the future) and that have no a-priori knowledge of the distribution F_S shall be called *statistics-unaware* algorithms. This paper establishes the fundamental convergence delay required for any statistics-unaware algorithm to achieve utility that is close to the optimal value ϕ^* .

A general statistics-unaware algorithm may incorporate some type of learning or estimation of the distribution F_S or some functional of this distribution. Observations of past channel states can be exploited when making online decisions. Consider some statistics-unaware algorithm that makes decisions over time $t \in \{0, 1, 2, ...\}$. For each positive integer T, the expression

$$\phi\left(\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\left[X[t]\right]\right)$$

is the utility associated with running the algorithm over the first T slots $\{0, 1, 2, ..., T - 1\}$. This utility includes decisions X[t] made at each step of the way (including the decision X[0] made at time t = 0 based only on the observation S[0]). Decisions must be made intelligently at each step of the way and fast learning is crucial. How close can the achieved utility get to the optimal value ϕ^* ? What time T is required?

A. Example utility functions

Different concave utility functions can be used to provide different types of performance (with corresponding fairness properties). For example, consider the linear utility

$$\phi(x_1,\ldots,x_n) = \sum_{i=1}^n a_i x_i$$

where a_1, \ldots, a_n are given nonnegative weights. Under this utility function, the problem (1)-(2) seeks to maximize a weighted sum of average transmission rates of each user. This linear utility is a trivial special case: The statistics-unaware algorithm of observing S[t] at the start of every slot t and choosing $X[t] \in \mathcal{D}(S[t])$ to greedily maximize $\phi(X[t])$, called the *greedy algorithm*, can be shown to lead to immediate convergence.² This is because the time average expectation can be pushed inside the linear function ϕ and so maximization of immediate rewards translates into maximization of long term rewards.

This is not the case for concave but *nonlinear* utility functions. This is because the goal is to maximize a concave function of the time average, not to maximize the average of a concave function. This goal is crucial to network fairness. The greedy algorithm can be far from optimal for general concave but nonlinear utility functions.³

From a fairness perspective, linear utilities are undesirable. For example, suppose there are two users, at most one user can transmit per slot, and user 1 always has a strictly better channel condition than user 2 (perhaps because user 1 is closer to the access point). Maximizing the linear utility function $\phi(x_1, x_2) = x_1 + x_2$ results in the algorithm that *always* chooses user 1, so that user 2 receives a time average rate of 0. One way to be fair to user 2 is to change the utility function to

$$\phi(x_1, x_2) = \min[x_1, x_2]$$

¹The only assumption about the algorithms over which ϕ^* is optimized is that the algorithms are *probabilistically measurable* so that they produce random vectors X[t] with well defined expectations on all slots t, see [3] for details. This is a mild assumption. It precludes (impractical) algorithms that make decisions using the axiom of choice on every slot t to produce non-measurable X[t] vectors.

²The existence of a maximizer for $\phi(x)$ over all $x \in \mathcal{D}(S[t])$ holds under the mild additional assumption that ϕ is continuous and $\mathcal{D}(S[t])$ is a compact subset of \mathbb{R}^n for each slot t.

³Consider a 2-user example with utility $\phi(x_1, x_2) = \log(x_1) + \log(x_2)$, a function known to have *proportional fairness* properties [4][5]. Suppose there is no channel variation and $\mathcal{D}(S[t]) = \{(20, 0), (0, 19)\}$ for all t. Choosing a fraction of time p to transmit with user 1 and maximizing $\log(20p) + \log(19(1-p))$, which is the utility function applied to the long term transmission rate vector, leads to $p^* = 1/2$ and yields a long term transmission rate vector $(\overline{X}_1, \overline{X}_2) = (10, 9.5)$. In contrast, the greedy strategy chooses user 1 on every slot so that $(\overline{X}_1, \overline{X}_2) = (20, 0)$ and user 2 never gets a chance to transmit!

Under this concave (but non-smooth) utility function, the problem is to maximize the minimum average rate given to the users. Another common type of (smooth) utility function is

$$\phi(x_1, x_2) = \log(x_1) + \log(x_2)$$

This logarithmic utility function results in a type of fairness called *proportional fairness* [4][5]. The logarithmic utility function is often modified to remove the singularity at zero:

$$\phi(x_1, x_2) = \log(1 + cx_1) + \log(1 + cx_2)$$

where c > 0 is a constant. Large values of c can be used to approximate the $\log(x_1) + \log(x_2)$ function. Other types of concave and nonlinear utility functions can be used for other types of fairness, such as α -fair utility functions [6][7][8].

B. Prior work

The work [9][10] develops statistics-unaware Frank-Wolfe type algorithms (with various step size rules) for solving the problem (1)-(2) for smooth utility functions using a fluid limit analysis. An alternative statistics-unaware drift-plus-penalty algorithm of [11][3] can be used to solve (1)-(2) for smooth or nonsmooth utility functions, and this achieves utility within ϵ of optimality with convergence time $O(1/\epsilon^2)$. Drift-plus-penalty can also be used for extended problems of multi-hop networks with power minimization and constraints [12], and related algorithms for these extended problems are in [13][14][15][16].

Recent work in [2] shows that, for smooth utility functions, a Frank-Wolfe algorithm with a constant stepsize also has convergence time $O(1/\epsilon^2)$, while a Frank-Wolfe algorithm with a vanishing stepsize yields an improved $O(\log(1/\epsilon)/\epsilon)$ convergence time. In particular, in the latter case we obtain

$$\phi\left(\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\left[X[t]\right]\right) \ge \phi^* - \frac{c\log(T)}{T} \quad \forall T \in \{1, 2, 3, \ldots\}$$

where c > 0 is a particular system constant. The work [2] also provides a near-matching converse of $\Omega(1/T)$. The problem of closing the logarithmic gap between the achievability bound and the converse bound was left as an open question. We resolve that open question in this paper by showing that the $O(\log(T)/T)$ gap is optimal.

A related logarithmic convergence time result is developed by Hazan and Kale in [17] for the context of online convex optimization with strongly convex objective functions. The prior work [17] provides an example online convex optimization problem that immediately reduces to a problem of estimating a Bernoulli probability from i.i.d. Bernoulli samples. They then provide a deep analysis of the Bernoulli estimation problem to show, via a nested interval argument, that for any sequence of Bernoulli estimators there exists a probability $p \in [1/4, 3/4]$ under which the estimators have a sum mean square error that grows at least logarithmically in the number of samples. This prior work inspires the current paper. We show that certain opportunistic scheduling problems can also be reduced to Bernoulli estimation; then we can use the Bernoulli estimation result of [17]. However, this reduction is not obvious. Online convex optimization problems have a different structure than opportunistic scheduling problems and the same reduction techniques cannot be used. New techniques are used to establish the converse, including a novel reduction of the opportunistic scheduling problem to a Bernoulli estimation problem.

A conceptually related problem on the fundamental time required to minimize a deterministic convex function with noisy observations of the gradients/subgradients is treated in [17] (where an achievability result is developed concerning the error versus time) and from an information theoretic perspective in [18] (where converse results are developed); see also early work on computational complexity for this problem in [19]. At a high level, the convergence time and learning concepts for that problem are similar to the current paper. However, the structure and analysis of that problem is quite different. For example: (i) Unlike the problem of this paper, the fundamental asymptotic tradeoffs for that problem depend significantly on whether or not the convex function is strongly convex; (ii) The fundamental tradeoffs for that problem do not have the same logarithmic properties as the problem of the current paper.

Prior work on fundamental bounds on the time and accuracy for estimation problems is in, for example, [20] for classification problems, and work in [21][22] treats bounds for estimation of functionals of discrete distributions.

C. Our contributions

- 1) This paper proves an $\Omega(\log(T)/T)$ converse for opportunistic scheduling. This matches an existing achievability result and resolves the open question in [2] to show that this performance is optimal.
- 2) This paper shows that *strongly concave* utility functions cannot be used to improve the asymptotic convergence time for opportunistic scheduling problems in comparison to functions that are concave but not strongly concave. This is surprising because strong convexity/concavity provides convergence improvements in other contexts, including online convex optimization problems [23][24], deterministic minimization via gradient descent [25], and deterministic minimization via stochastic gradients [18][19][17]. This emphasizes the unique properties of opportunistic scheduling problems.

- 3) The technique for reducing opportunistic scheduling to Bernoulli estimation can more broadly impact future work on more complex networks (see open questions in this direction in the conclusion).
- 4) This paper refines the regret analysis for Bernoulli estimation theory in [17] to show that for any sequence of estimators, not only does *there exist* a probability $p \in [1/4, 3/4]$ for which the regret grows at least logarithmically, but *the set of all such values p has measure at least* 1/6. This is used to establish a 1/6 result for opportunistic scheduling: If any particular statistics-unaware algorithm is used, and if nature selects the channel according to a Bernoulli process with parameter p that is independently chosen over the unit interval, then with probability at least 1/6 the algorithm will be limited by the $\Omega(\log(T)/T)$ converse bound. Shouldn't algorithms *always* be limited by this bound? No. Imagine a scheduling algorithm that makes an *a-priori guess* $\hat{q} \in [0, 1]$ about the true network probability q, and then makes decisions that are optimal under the assumption that the guess is exact. In the "lucky" situation when $\hat{q} = q$, this algorithm would perform optimally and would not be limited by the $\Omega(\log(T)/T)$ converse. Nevertheless, our analysis shows that every algorithm (including algorithms that attempt to make lucky guesses) will fail to beat the $\Omega(\log(T)/T)$ converse with probability at least 1/6.

II. BERNOULLI ESTIMATION

This section gives preliminaries on estimating an unknown probability $p \in [0, 1]$ from i.i.d. Bernoulli samples $\{W_n\}_{n=1}^{\infty}$ with

$$P[W_n = 1] = p, \quad P[W_n = 0] = 1 - p$$

A. Estimation functions

Let $\{W_n^p\}_{n=1}^{\infty}$ be a sequence of i.i.d. Bernoulli random variables with $P[W_n^p = 1] = p$ (called a Bernoulli-*p* sequence). The value of $p \in [0, 1]$ is unknown. On each time step *n* we observe the value of W_n^p and then make an estimate of *p* based on all observations that have been seen so far.

A general estimation method can be characterized as follows: Let $\{\hat{A}_n\}_{n=1}^{\infty}$ be an infinite sequence of functions such that each function $\hat{A}_n(u, w_1, ..., w_n)$ maps a binary-valued sequence $(w_1, ..., w_n) \in \{0, 1\}^n$ and a random seed $u \in [0, 1)$ to a real number in the interval [0, 1]. That is, for all $n \in \{1, 2, 3, ...\}$ we have

$$\hat{A}_n : [0,1) \times \{0,1\}^n \to [0,1]$$
(3)

The \hat{A}_n functions shall be called *estimation functions*. Let U be a random variable that is uniformly distributed over [0,1) and that is independent of $\{W_n^p\}_{n=1}^{\infty}$. For $n \in \{1, 2, ...\}$, let A_n^p denote the estimate of p based on observations of W_1^p, \ldots, W_n^p :

$$A_n^p = \hat{A}_n(U, W_1^p, W_2^p, \dots, W_n^p) \quad \forall n \in \{1, 2, 3, \dots\}$$
(4)

The random variable U is used to facilitate possibly randomized decisions.⁴ The functions \hat{A}_n in (3) are only assumed to be *probabilistically measurable* so that A_n^p defined in (4) is a valid random variable for all $n \in \{1, 2, 3, ...\}$.

For a given sequence of estimator functions, let A_n^p be the estimate at time n, as defined by (4). For each $n \in \{0, 1, 2, ...\}$ define $\mathbb{E}_p[(A_n^p - p)^2]$ as the mean square estimation error at time n. The expectation is with respect to the random seed U and the random Bernoulli sequence. The random seed U is assumed to be independent of the sequence of Bernoulli variables $\{W_n^p\}_{n=1}^\infty$. Thus, if we condition on U = u for a particular $u \in [0, 1)$, the conditional expectation $\mathbb{E}_p[(A_n^p - p)^2 | U = u]$ is with respect to the probability measure associated only with the random vector (W_1^p, \ldots, W_n^p) . For a given random seed $u \in [0, 1)$ and for two different probabilities $p, q \in [0, 1]$, the values $\mathbb{E}_p[(A_n^p - p)^2 | U = u]$ and $E_q[(A_n^q - q)^2 | U = u]$ are the mean square errors at time n associated with the same deterministic estimation function \hat{A}_n but assuming a Bernoulli-p process and a Bernoulli-q process, respectively. The following theorem is due to Hazan and Kale in [17].

Theorem 1: (Bernoulli estimation from [17]) Fix any sequence of measurable estimation functions $\{\hat{A}_n\}_{n=1}^{\infty}$ of the form (3). There is a probability $p \in [1/4, 3/4]$ such that

$$\sum_{n=1}^{N} \mathbb{E}_{p}[(A_{n}^{p} - p)^{2}] \ge \Omega(\log(N)) \quad \forall N \in \{1, 2, 3, ...\}$$

where A_n^p is defined by (4).

It is important to distinguish the result of Theorem 1 from the Cramer-Rao estimation bound (see, for example, [26]). The Cramer-Rao bound is most conveniently applied to *unbiased estimators*. While biased versions of the Cramer-Rao bound exist, they require additional structural assumptions, such as knowledge of a (differentiable) bias function b'(p) with a derivative that is bounded away from -1 so that a term $(1 + b'(p))^2$ does not vanish. Moreover, Cramer-Rao bounds are typically applied

⁴A real number has infinite precision. Thus, in principle, a single random variable U that is uniform over [0, 1) is sufficient to generate all desired randomness for a sequence of randomized estimates. Indeed, from U we can create an i.i.d. sequence $\{U_i\}_{i=1}^{\infty}$ of uniformly distributed random variables such that each U_i is a deterministic function of U. To do this, first obtain an i.i.d. sequence of digits $\{D_i\}_{i=1}^{\infty}$ by writing U in its decimal expansion $U = \sum_{n=1}^{\infty} D_n 10^{-n}$. Then let $f : \mathbb{N}^2 \to \mathbb{N}$ be a bijection and define $U_i = \sum_{j=1}^{\infty} D_{f(i,j)} 10^{-j}$ for each $i \in \{1, 2, 3, \ldots\}$.

to a single estimator for time step n. In contrast, the Hazan and Kale theorem above treats the *sum* mean square error over a *sequence* of estimators, which is essential for establishing connections to the regret of online scheduling algorithms.

Using the nested interval techniques of [17], the asymptotic bound $\Omega(\log(N))$ in Theorem 1 can be written as an explicit function $b \log(N) - c$ where b and c are system constants that do not depend on N. Unfortunately, there is a minor constant factor error in Lemma 15 of [17]. That constant minor factor error does not affect correctness of the $\Omega(\log(N))$ result established in [17].⁵ For convenience, the minor error is identified and fixed in Appendix A.

B. Positive measure in the unit interval

Theorem 1 shows that for any sequence of Bernoulli estimators, there is a probability $p \in [1/4, 3/4]$ under which the estimators perform poorly, in the sense of having sum mean square error that grows at least logarithmically in N. The next theorem shows that, not only does such a probability p exist, the set of all such probabilities p is measurable and has measure at least 1/6 within the interval [1/4, 3/4]. It also generalizes to treat arbitrary powers of absolute error. For each $\alpha > 0$ define:

$$V_m(\alpha) = \sum_{n=1}^m (1/n)^{\alpha/2} \quad \forall m \in \{1, 2, 3, \ldots\}$$

Notice that

$$V_m(\alpha) \ge \int_1^{m+1} (1/t)^{\alpha/2} dt = \begin{cases} \log(m+1) & \text{if } \alpha = 2\\ \frac{(m+1)^{1-\alpha/2} - 1}{1-\alpha/2} & \text{if } \alpha \neq 2 \end{cases}$$

Theorem 2: Fix $\alpha \in (0,2]$. Fix any sequence of measurable estimation functions $\{\hat{A}_n\}_{n=1}^{\infty}$ of the form (3). Define $\mathcal{Q} \subseteq [1/4, 3/4]$ as the set of all $p \in [1/4, 3/4]$ such that:

$$\limsup_{m \to \infty} \left[\frac{1}{V_m(\alpha)} \sum_{n=1}^m \mathbb{E}_p[|A_n^p - p|^\alpha] \right] \ge \frac{1}{c^{\alpha} 2^{3+2\alpha}}$$
(5)

where $c = \sqrt{8/3}$. Then the set Q is Lebesgue measurable and has measure $\mu(Q) \ge 1/6$. Thus, a randomly and uniformly chosen $p \in [1/4, 3/4]$ will satisfy the lim sup inequality (5) with probability at least 1/3. In particular, the inequality (5) implies:

• If $\alpha = 2$ then

$$\limsup_{m \to \infty} \left[\frac{1}{\log(m+1)} \sum_{n=1}^m \mathbb{E}_p[|A_n^p - p|^2] \right] \ge \frac{3}{2^{10}}$$

• If $\alpha = 1$ then

$$\limsup_{m \to \infty} \left[\frac{1}{[(m+1)^{1/2} - 1]} \sum_{n=1}^{m} \mathbb{E}_p[|A_n^p - p|] \right] \ge \frac{1}{2^4 \sqrt{8/3}}$$

Proof: See Appendix B.⁶

Appendix C has details on the tightness of these lower bounds by comparing to a particular estimation method that achieves regret to within a constant factor of these bounds.

III. THE CONVERSE BOUND

This section constructs a simple 2-user opportunistic scheduling system with state vectors S[t] described by a single probability parameter $q \in [1/4, 3/4]$. It produces a converse bound on the utility optimality gap by mapping the problem to a Bernoulli estimation problem and then using Theorem 1 and Theorem 2.

A. The example 2-user system

Consider a 2-user wireless system that operates in slotted time $t \in \{0, 1, 2, ...\}$. Suppose the system state is described by a sequence of i.i.d. Bernoulli variables $\{S[t]\}_{t=0}^{\infty}$ with

$$P[S[t] = 1] = q; \quad P[S[t] = 0] = 1 - q$$

⁵It should be emphasized that the techniques in [17] are novel and deep and are in no way diminished by this constant factor error.

⁶The conference version of this paper in [1] stated a weaker version of Theorem 2 that showed $\mu(Q) \ge 1/8$ (based on an earlier version of this technical report). The current technical report proves the stronger result $\mu(Q) \ge 1/6$. An unfortunate error in the conference version [1] stated Theorem 2 as holding for $\alpha > 0$, whereas it should be stated to hold for $0 < \alpha \le 2$, which ensures $\lim_{m\to\infty} V_m(\alpha) = \infty$ and $\lim_{m\to\infty} V_m(\alpha+1)/V_m(\alpha) = 0$ (used in the proof of Theorem 2). The case $\alpha \in (0, 2]$ is the main case of interest for asymptotic tradeoff analysis: For $\alpha > 2$ the sum regret does not necessarily go to infinity with the number of samples because, for the simple estimator $A_n^p = \frac{1}{n} \sum_{j=1}^n W_j^p$, we have $\sum_{n=1}^{\infty} \mathbb{E}_p[|A_n^p - p|^{\alpha}] < \infty$.

where $q \in [0,1]$ is an unknown probability. Every slot $t \in \{0,1,2,\ldots\}$ the system controller observes S[t] and chooses a *transmission rate vector* $X[t] = (X_1[t], X_2[t])$ from a decision set $\mathcal{D}(S[t])$ given by

$$\mathcal{D}(S[t]) = \begin{cases} \{(1,0)\} & \text{if } S[t] = 0\\ \{(r,1-r^2) : r \in [0,1]\} & \text{if } S[t] = 1 \end{cases}$$
(6)

In particular, if S[t] = 0 then the controller has no choice but to allocate X[t] = (1, 0), which gives no transmission rate to user 2. On the other hand, if S[t] = 1 then the controller is free to allocate $X[t] = (r, 1 - r^2)$ for some $r \in [0, 1]$, which allows giving a nonzero transmission rate to user 2. Observe that under any system state and any decision, it holds that $0 \le X_1[t] \le 1$, $0 \le X_2[t] \le 1$, and $X_2[t] = 1 - X_1[t]^2$ for all t. The set of all points $(X_1[t], X_2[t]) \in \mathcal{D}(1)$ available when S[t] = 1 is shown as the solid curve in Fig. 1.



Fig. 1. An illustration of the decision set $\mathcal{D}(1)$ as compared to an alternative logarithmic decision set. The solid curve shows the $(X_1[t], X_2[t])$ points in $\mathcal{D}(1)$ when S[t] = 1; the dashed curve shows the points $(C_1[t], C_2[t])$ defined by (7)-(8).

While this example decision set $\mathcal{D}(S[t])$ is very specific, it is representative of the following physical scenario: Imagine that user 2 goes offline independently every slot t with probability 1 - q (possibly due to a time-varying channel condition, or because it allocates its resources to other tasks according to a randomized schedule). Hence, user 1 can allocate a full rate of 1 on those slots (corresponding to slots t such that S[t] = 0). On the other hand, during the slots in which users 1 and 2 are both online (corresponding to S[t] = 1), the users can simultaneously transmit but, due to interference, they cannot both transmit at the full rate of 1. During such slots t for which S[t] = 1, there is a tradeoff between the rates $X_1[t]$ and $X_2[t]$ that can be allocated, so that $X_2[t]$ is a nonincreasing function of $X_1[t]$. The particular nondecreasing function $X_2[t] = 1 - X_1[t]^2$ that is used is shown in Fig. 1. This function is chosen for mathematical convenience (it simplifies the proof to be given). Similar proofs can be given for curves that are qualitatively similar but that have more physical meaning: For example, for slots t such that S[t] = 1, suppose the total bandwidth available is B and the rates of users 1 and 2 are chosen by allocating fractions of the bandwidth $\theta_1[t]$ and $\theta_2[t]$ to users 1 and 2, so that user 1 is allocated a total bandwidth of $B\theta_1[t]$, $\theta_2[t]$ are chosen as nonnegative values that sum to 1. The users thus transmit over frequency-separated channels. Assuming each channel is an additive white Gaussian noise channel (with noise density uniform over the given frequency spectrum) and given these particular frequency division allocations on slot t, the point-to-point Shannon capacity of each channel is [26]:

$$C_1[t] = \theta_1[t]B\log\left(1 + \frac{P}{\theta_1[t]N}\right) \tag{7}$$

$$C_2[t] = \theta_2[t]B\log\left(1 + \frac{P}{\theta_2[t]N}\right) \tag{8}$$

where P and N are fixed positive parameters. The expression $\frac{P}{\theta_i[t]N}$ represents the signal-to-noise ratio for channel $i \in \{1, 2\}$ and the noise $\theta_i N$ is proportional to the bandwidth used on channel *i*.

The $(C_1[t], C_2[t])$ values traced out by all possible $(\theta_1[t], \theta_2[t])$ allocations are given in Fig. 1 for a particular choice of parameters B = 0.7, P/N = 3. The mathematical curve is different from the curve $(r, 1 - r^2)$, but it is qualitatively similar. In particular, like the curve $(r, 1 - r^2)$, it can be shown to have a *strongly concave structure*. The proof of our converse can

be extended to apply to this particular $(C_1[t], C_2[t])$ curve, and to similar curves that are strongly concave. We use the curve $(r, 1 - r^2)$ because it is simple and yields the most direct proof of the desired converse result.

B. The example network utility maximization problem

For positive integers T define:

$$\overline{X}_i(T) = \frac{1}{T} \sum_{t=0}^{T-1} X_i[t] \quad \forall i \in \{1, 2\}$$

and define $\overline{X}(T) = (\overline{X}_1(T), \overline{X}_2(T))$. Let $\phi : [0,1]^2 \to \mathbb{R}$ be a continuous and concave utility function. The goal of the network controller is to allocate X[t] over time to solve

Maximize:
$$\liminf_{T \to \infty} \phi\left(\mathbb{E}\left[\overline{X}(T)\right]\right)$$
(9)

Subject to:
$$X[t] \in \mathcal{D}(S[t]) \quad \forall t \in \{0, 1, 2, \ldots\}$$
 (10)

This problem indeed has the structure of (1)-(2). Define ϕ^* as the optimal utility for the above problem. That is, ϕ^* is the supremum value of (9) over all algorithms that satisfy (10). The specific utility function that we first consider is:

$$\phi(x_1, x_2) = \log(1 + x_1) + \log(1 + x_2)$$

This function is entrywise nondecreasing and *c*-strongly concave over the domain $(x_1, x_2) \in [0, 1]^2$ with parameter c = 1/4, meaning that the function $s: [0,1]^2 \to \mathbb{R}$ given by

$$s(x_1, x_2) = \phi(x_1, x_2) + \frac{c}{2}(x_1^2 + x_2^2)$$

is concave. The main result of the paper is below:

Theorem 3: Consider the 2-user example problem (9)-(10) with utility function $\phi(x_1, x_2) = \log(1 + x_1) + \log(1 + x_2)$ and with an unknown parameter $q \in [0, 1]$. Under any (possibly randomized) statistics-unaware control algorithm, there is a probability $q \in [1/4, 3/4]$ such that

$$\phi\left(\mathbb{E}\left[\overline{X}_{1}[T]\right], \mathbb{E}\left[\overline{X}_{2}[T]\right]\right) \le \phi^{*} - \Omega\left(\frac{\log(T)}{T}\right) \quad \forall T \in \{1, 2, 3, \ldots\}$$

$$(11)$$

Furthermore, there is a measurable subset $Q \subseteq [1/4, 3/4]$ with measure at least 1/6 such that if $q \in Q$ then

$$\limsup_{T \to \infty} \frac{T\phi^* - T\phi\left(\mathbb{E}\left[\overline{X}_1[T]\right], \mathbb{E}\left[\overline{X}_2[T]\right]\right)}{\log(T)} \ge \frac{3\beta^2}{2^{13}}$$
(12)

where $\beta = \frac{2}{3} - \frac{\sqrt{7}}{6} \approx 0.2257$. The proof of Theorem 3 is developed in the following subsections. Observe that since the utility function $\phi(x_1, x_2) =$ $\log(1+x_1) + \log(1+x_2)$ is smooth, the result in [2] ensures that the statistics-unaware Frank-Wolfe algorithm with vanishing stepsize can be used, without knowledge of the parameter $q \in [0, 1]$, to ensure that for all values $q \in [0, 1]$ we have

$$\phi\left(\mathbb{E}\left[\overline{X}_{1}[T]\right], \mathbb{E}\left[\overline{X}_{2}[T]\right]\right) \ge \phi^{*} - O\left(\frac{\log(T)}{T}\right) \quad \forall T \in \{1, 2, 3, \ldots\}$$

In particular, the $\log(T)/T$ asymptotic converse bound of Theorem 3 can be achieved. Hence, the asymptotic $\log(T)/T$ bound is tight and the corresponding Frank-Wolfe algorithm (with vanishing stepsize) is asymptotically optimal.

C. Optimality over stationary policies

Results in [3] show that optimality for the problem (9)-(10) can be characterized by the set C of all one-shot expectations $\mathbb{E}[(X_1[0], X_2[0])]$ that can be achieved on slot t = 0. Consider the set of all $(x_1, x_2) \in \mathbb{R}^2$ that satisfy

$$1 - q \le x_1 \le 1 \tag{13}$$

$$x_1 + x_2 \ge 1 \tag{14}$$

$$x_2 \le 2(1-x_1) - \frac{1}{q}(1-x_1)^2 \tag{15}$$

The set of all such points is shown in Fig. 2 for the case q = 0.5. The next result shows that this set is equal to C.

Lemma 1: (The set C) Fix $q \in (0,1]$ and assume $\{S[t]\}_{t=0}^{\infty}$ is an i.i.d. Bernoulli process with P[S[t] = 1] = q. The set C of all one-shot expectations $\mathbb{E}[(X_1[0], X_2[0])]$ achievable in the 2-user system is

$$\mathcal{C} = \text{Conv}(\{(1 - q + qr, q(1 - r^2)) \in \mathbb{R}^2 : r \in [0, 1]\})$$
(16)



0.8

0.9

1

Fig. 2. The set C of all one-shot expectations for the case q = 0.5, which includes all points on or between the upper and lower boundary curves. The lower boundary is determined by (14) and the upper boundary is determined by (15).

0.7

×₁

where Conv(·) denotes the convex hull. This set C is closed, bounded, convex, and is equivalently described as the set of all $(x_1, x_2) \in \mathbb{R}^2$ that satisfy the inequalities (13), (14), (15).

Proof: Define the set C_1 by

$$\mathcal{C}_1 = \{ (1 - q + qr, q(1 - r^2)) \in \mathbb{R}^2 : r \in [0, 1] \}$$

The set $\text{Conv}(\mathcal{C}_1)$ corresponds to the set defined in the right-hand-side of (16). The set \mathcal{C}_1 is closed and bounded and so $\text{Conv}(\mathcal{C}_1)$ is convex, closed, and bounded. Define \mathcal{C}_2 as the set of all points (x_1, x_2) that satisfy the three inequality constraints (13), (14), (15). Define \mathcal{C} as the set of all expectations $\mathbb{E}[(X_1[0], X_2[0])]$ achievable on slot t = 0. We want to show that $\text{Conv}(\mathcal{C}_1) = \mathcal{C} = \mathcal{C}_2$.

We first show $\text{Conv}(\mathcal{C}_1) \subseteq \mathcal{C}$. Fix $r \in [0, 1]$. Consider the following decision policy for slot t = 0:

$$(X_1[0], X_2[0]) = \begin{cases} (1,0) & \text{if } S[0] = 0\\ (r, 1 - r^2) & \text{if } S[0] = 1 \end{cases}$$

Then

$$\mathbb{E}\left[(X_1[0], X_2[0])\right] = (1, 0)P[S[0] = 0] + (r, 1 - r^2)P[S[0] = 1]$$

= (1, 0)(1 - q) + (r, 1 - r^2)q
= ((1 - q) + qr, q(1 - r^2))

Hence, any point in the set C_1 can be achieved as an expectation on slot t = 0. Any point in $Conv(C_1)$ can be achieved by randomizing over policies that achieve particular points in C_1 . Thus, $Conv(C_1) \subseteq C$.

We now show that $C \subseteq C_2$. Consider any point $(x_1, x_2) \in C$. By definition of C, the point (x_1, x_2) can be achieved as an expectation on slot t = 0. Let $(X_1[0], X_2[0]) \in \mathcal{D}(S[0])$ denote the (possibly randomized) decision vector on slot 0 that satisfies $\mathbb{E}[(X_1[0], X_2[0])] = (x_1, x_2)$. For simplicity of notation define S = S[0] and $(X_1, X_2) = (X_1[0], X_2[0])$ so that:

$$(X_1, X_2) \in \mathcal{D}(S)$$
$$\mathbb{E}\left[(X_1, X_2)\right] = (x_1, x_2)$$

By the structure of the decision set $\mathcal{D}(S)$ we have $X_1 \in [0,1]$ always, and

0.6

0.5

0.4

0.2

0.1

0└ 0.4

0.5

0.6

×[∼] 0.3

$$(X_1, X_2) = \begin{cases} (1,0) & \text{if } S = 0\\ (X_1, 1 - X_1^2) & \text{if } S = 1 \end{cases}$$
(17)

Since P[S = 1] = q we have

$$x_{1} = \mathbb{E} [X_{1}]$$

= $\mathbb{E} [X_{1}|S = 0] (1 - q) + \mathbb{E} [X_{1}|S = 1] q$
= $(1 - q) + \mathbb{E} [X_{1}|S = 1] q$

Thus

$$x_1 = (1 - q) + \mathbb{E}[X_1 | S = 1]q$$
(18)

Since $X_1 \in [0,1]$ we have $0 \leq \mathbb{E}[X_1|S=1] \leq 1$ and so (18) implies that (13) holds.

To show that (14) also holds, observe from (17) that regardless of whether S = 0 or S = 1 we have:

$$X_1 + X_2 \ge 1$$

which holds because $X_1 \in [0,1]$ and $X_1^2 \leq X_1$, and so

$$\mathbb{E}\left[X_1 + X_2\right] \ge 1$$

which implies that inequality (14) holds.

To show that (15) holds, observe that

$$x_{2} = \mathbb{E}[X_{2}]$$

$$\stackrel{(a)}{=} 0P[S = 0] + \mathbb{E}[1 - X_{1}^{2}|S = 1] P[S = 1]$$

$$= q - q\mathbb{E}[X_{1}^{2}|S = 1]$$

$$\stackrel{(b)}{\leq} q - q\mathbb{E}[X_{1}|S = 1]^{2}$$

$$\stackrel{(c)}{=} q - q\left(\frac{x_{1} - 1 + q}{q}\right)^{2}$$

$$= 2(1 - x_{1}) - \frac{(1 - x_{1})^{2}}{q}$$
(19)
(20)

where (a) holds by (17); (b) holds by Jensen's inequality; (c) holds by (18). Thus, inequality (15) holds.

It follows that $\text{Conv}(\mathcal{C}_1) \subseteq \mathcal{C} \subseteq \mathcal{C}_2$. Finally, it is not difficult to show that the upper boundary of the set \mathcal{C}_2 , defined by all points (x_1, x_2) that satisfy $1 - q \leq x_1 \leq 1$ and that satisfy inequality (15) with equality, is equal to the set \mathcal{C}_1 (see Fig. 2 for the case q = 0.5). Further, the convex hull of this upper boundary is the entire set \mathcal{C}_2 , so that $\text{Conv}(\mathcal{C}_1) = \mathcal{C}_2$. Results in [3] imply that the optimal utility ϕ^* for problem (9)-(10) is equal to the maximum of the continuous and concave

function $\phi(x_1, x_2)$ over all (x_1, x_2) in the closed, bounded, and convex set C. That is,

$$\phi^* = \max_{(x_1, x_2) \in \mathcal{C}} \phi(x_1, x_2)$$

The utility function of interest can be shown to be strongly concave and so the maximizer (x_1^*, x_2^*) is unique. The maximizer is given in terms of q in the next lemma.

Lemma 2: (Optimal operating point) Fix $q \in (0,1]$. Define $\phi : [0,1]^2 \to \mathbb{R}$ by $\phi(x_1, x_2) = \log(1+x_1) + \log(1+x_2)$. The unique maximizer of $\phi(x_1, x_2)$ over $(x_1, x_2) \in \mathcal{C}$ is a vector (x_1^*, x_2^*) that satisfies

$$(x_1^*, x_2^*) = \left((1-q) + qr, q(1-r^2)\right)$$
(21)

for some particular $r \in [0,1]$, so that $\phi(x_1^*, x_2^*) = \phi^*$. Furthermore the optimal value $r \in [0,1]$ in (21) satisfies

$$\frac{1}{1+x_1^*} = \frac{2r}{1+x_2^*} \tag{22}$$

and is exactly equal to

$$r = \frac{-(2-q) + \sqrt{4q^2 - q + 4}}{3q} \tag{23}$$

The expression on the right-hand-side of (23) increases from 1/4 to $\frac{-1+\sqrt{7}}{3}$ as q slides between 0 and 1, where $\frac{-1+\sqrt{7}}{3} \approx 0.54858$.

Proof: Considering the set C defined by (16), it is clear that (because $\phi(x_1, x_2)$ is entrywise nondecreasing) a maximizer occurs on the upper boundary curve $(1 - q + qr, q(1 - r^2))$ for $r \in [0, 1]$. The utility function associated with $r \in [0, 1]$ is

$$\log(2 - q + qr) + \log(1 + q(1 - r^2))$$

This is concave over $r \in [0, 1]$. If a point of zero derivative can be found over $r \in [0, 1]$ then that point must be optimal. Taking a derivative with respect to r and setting the result to 0 yields

$$\frac{q}{2-q+qr} + \frac{-2qr}{1+q(1-r^2)} = 0$$

10

Since q > 0, dividing by q > 0 and rearranging terms gives

$$\frac{1}{2-q+qr} = \frac{2r}{1+q(1-r^2)}$$

which yields (22) by the substitution $x_1^* = 1 - q + qr$, $x_2^* = q(1 - r^2)$. Rearranging the above equality again yields a quadratic equation in r that is solved by taking the only nonnegative solution, which is given in (23). It can be checked that the expression in (23) increases from 1/4 to $\frac{-1+\sqrt{7}}{3}$ as q slides between 0 and 1. In particular, a zero-derivative point $r \in [0, 1]$ can indeed be found for all values of q that are considered.

Lemma 3: (The bijection h) Define the function $h: [0,1] \rightarrow [1/4, \frac{-1+\sqrt{7}}{3}]$ by

$$h(q) = \begin{cases} 1/4 & \text{if } q = 0\\ \frac{-(2-q) + \sqrt{4q^2 - q + 4}}{3q} & \text{otherwise} \end{cases}$$
(24)

Then $q \in (0, 1]$ implies r = h(q) (where r is defined in (23)). Further, function h is strictly increasing, so that it has an inverse function:

$$h^{-1}: \left[1/4, \frac{-1+\sqrt{7}}{3}\right] \to [0, 1]$$

Finally, h is continuously differentiable and satisfies $h'(q) \ge h'(1)$ for all $q \in [0,1]$, and so defining $\beta > 0$ by

$$\beta = h'(1) = \frac{2}{3} - \frac{\sqrt{7}}{6} \approx 0.2257 \tag{25}$$

we have the following "expansion" property of h:

$$|h(a) - h(b)| \ge \beta |a - b| \quad \forall a, b \in [0, 1]$$

$$(26)$$

Proof: Taking a derivative of h at $q \in (0, 1]$ gives

$$h'(q) = \frac{q - 8 + 4\sqrt{4q^2 - q} + 4}{6q^2\sqrt{4q^2 - q} + 4}$$

This is continuous at all $q \in (0, 1]$. Further,

$$\lim_{q \to 0^+} h'(q) = \frac{21}{64}$$

This is consistent with the right-derivative of h at 0:

$$\lim_{q \to 0^+} \frac{h(q) - 1/4}{q} = \frac{21}{64}$$

Thus, h is continuously differentiable for all $q \in [0, 1]$. From the expression for h'(q) it follows that $h'(1) = \beta = \frac{2}{3} - \frac{\sqrt{7}}{6}$. As q slides over the interval [0, 1], the function h'(q) at first increases but eventually decreases to reach a minimum value at q = 1. Thus, $h'(q) \ge \beta$ for all $q \in [0, 1]$. The expansion property (26) follows by fixing $a, b \in [0, 1]$, without loss of generality assuming a < b, and using the fundamental theorem of calculus:

$$h(b) - h(a) = \int_{a}^{b} h'(q) dq$$
$$\geq \int_{a}^{b} \beta dq$$
$$= \beta(b - a)$$

where the inequality holds because $h'(q) \ge \beta$ for all $q \in [0, 1]$.

D. Statistics-unaware algorithms for utility maximization

This subsection completes the proof of Theorem 3 for the 2-user problem (9)-(10). Fix $q \in [1/4, 3/4]$ and recall that q = P[S[t] = 1]. Consider any *statistics-unaware* algorithm for choosing $X[t] = (X_1[t], X_2[t]) \in \mathcal{D}(S[t])$ over $t \in \{0, 1, 2, ...\}$, where $\mathcal{D}(S[t])$ is given in (6). In particular, the algorithm has no a-priori knowledge of q. For each $t \in \{0, 1, 2, ...\}$ define

$$x[t] = (x_1[t], x_2[t]) = \mathbb{E}[(X_1[t], X_2[t])]$$

Fix T as a positive integer. Define $\overline{X}[T]$ as the time average over the first T slots

$$\overline{X}[T] = \frac{1}{T} \sum_{t=0}^{T-1} (X_1[t], X_2[t])$$

Taking expectations of both sides of the above equality and using the definition of x[t] gives

$$\mathbb{E}\left[\overline{X}[T]\right] = \frac{1}{T} \sum_{t=0}^{T-1} x[t]$$
(27)

Let $x^* = (x_1^*, x_2^*)$ be the optimal operating point defined in (21) of Lemma 2. Let $\phi'(x^*)^{\top}$ denote the gradient at x^* expressed as a row vector:

$$\phi(x_1, x_2) = \log(1 + x_1) + \log(1 + x_2)$$

$$\implies \phi'(x^*)^\top = \left[\frac{\partial \phi(x^*)}{\partial x_1}; \frac{\partial \phi(x^*)}{\partial x_2}\right] = \left[\frac{1}{1 + x_1^*}; \frac{1}{1 + x_2^*}\right]$$
(28)

By concavity of ϕ we have:

$$\phi(\mathbb{E}\left[\overline{X}(T)\right]) \stackrel{(a)}{\leq} \phi(x^*) + \phi'(x^*)^\top (\mathbb{E}\left[\overline{X}(T)\right] - x^*) \\ \stackrel{(b)}{=} \phi(x^*) + \phi'(x^*)^\top \cdot \frac{1}{T} \sum_{t=0}^{T-1} (x[t] - x^*) \\ \stackrel{(c)}{=} \phi^* + \frac{1}{T} \sum_{t=0}^{T-1} \left[\frac{x_1[t] - x_1^*}{1 + x_1^*} + \frac{x_2[t] - x_2^*}{1 + x_2^*} \right] \\ \stackrel{(d)}{=} \phi^* + \frac{1}{T} \sum_{t=0}^{T-1} \left[\frac{2r(x_1[t] - x_1^*)}{1 + x_2^*} + \frac{(x_2[t] - x_2^*)}{(1 + x_2^*)} \right]$$
(29)

where (a) holds by the gradient inequality for concave functions; (b) holds by (27); (c) holds by (28) and the fact that $\phi(x^*) = \phi^*$ for the vector $x^* = (x_1^*, x_2^*)$ defined in Lemma 2; (d) holds by (22).

Now consider a particular $t \in \{0, 1, ..., T-1\}$. Define F[t] as the history of channel states over the first t slots:

$$F[t] = (S[0], S[1], \dots, S[t-1])$$

with S[-1] defined as a nonrandom constant so that F[0] provides no information about the channel. Define $z_{F[t]}$ by

$$z_{F[t]} = \mathbb{E}\left[X_1[t]|S[t] = 1, F[t]\right]$$
(30)

By Jensen's inequality we have

$$z_{F[t]}^2 \le \mathbb{E}\left[X_1[t]^2 | S[t] = 1, F[t]\right]$$
(31)

We have

$$\mathbb{E}\left[(X_{1}[t], X_{2}[t])|F[t]\right]
\stackrel{(a)}{=} \mathbb{E}\left[(X_{1}[t], X_{2}[t])|S[t] = 0, F[t]\right](1-q)
+ \mathbb{E}\left[(X_{1}[t], X_{2}[t])|S[t] = 1, F[t]\right]q
\stackrel{(b)}{=} (1,0)(1-q) + \mathbb{E}\left[(X_{1}[t], 1-X_{1}[t]^{2})|S[t] = 1, F[t]\right]q
\stackrel{(c)}{\leq} (1,0)(1-q) + (z_{F[t]}, 1-z_{F[t]}^{2})q
= (1-q+qz_{F[t]}, q(1-z_{F[t]}^{2}))$$
(32)

where (a) holds by conditioning on the events $\{S[t] = 1\}$ and $\{S[t] = 0\}$ (which are independent of F[t]); (b) holds by the decision set structure $\mathcal{D}(S[t])$ in (6); (c) holds as an entrywise inequality by (30) and (31). Taking expectations of (32) with respect to the random F[t] and using the law of iterated expectations gives (using $\mathbb{E}[X_i[t]] = x_i[t]$):

$$(x_1[t], x_2[t]) \le \left(1 - q + q\mathbb{E}\left[z_{F[t]}\right], q - q\mathbb{E}\left[z_{F[t]}^2\right]\right)$$
(33)

On the other hand, recall from Lemma 2 that

$$(x_1^*, x_2^*) = (1 - q + qr, q(1 - r^2))$$
(34)

Using (33) and (34) together gives

$$2r(x_{1}[t] - x_{1}^{*}) + (x_{2}[t] - x_{2}^{*})$$

$$\leq 2rq(\mathbb{E} [z_{F[t]}] - r) + q(r^{2} - \mathbb{E} [z_{F[t]}^{2}])$$

$$= -q\mathbb{E} [(z_{F[t]} - r)^{2}]$$

$$\leq -q\mathbb{E} \left[\left([z_{F[t]}]_{h(0)}^{h(1)} - r \right)^{2} \right]$$
(35)

where $[z_{F[t]}]_{h(0)}^{h(1)}$ projects $z_{F[t]}$ onto the interval [h(0), h(1)] and the final inequality holds because we know $r \in [h(0), h(1)]$, and the distance between $z_{F[t]}$ and r must be greater than or equal to their distances when projected onto the interval [h(0), h(1)]. Now we know that $h : [0, 1] \to [h(0), h(1)]$ is bijective and so we can define

$$\theta[t] = h^{-1} \left([z_{F[t]}]_{h(0)}^{h(1)} \right)$$
(36)

so that $[z_{F[t]}]_{h(0)}^{h(1)} = h(\theta[t])$. Substituting this and r = h(q) in (35) gives

$$2r(x_1[t] - x_1^*) + (x_2[t] - x_2^*) \le -q\mathbb{E}\left[\left(h(\theta[t]) - h(q)\right)^2\right]$$

Substituting the above inequality into (29) yields:

$$\phi(\mathbb{E}\left[\overline{X}(T)\right]) \leq \phi^* - \frac{q}{T(1+x_2^*)} \sum_{t=0}^{T-1} \mathbb{E}\left[(h(\theta[t]) - h(q))^2 \right]$$

$$\stackrel{(a)}{\leq} \phi^* - \frac{1}{8T} \sum_{t=0}^{T-1} \mathbb{E}\left[(h(\theta[t]) - h(q))^2 \right]$$

$$\stackrel{(b)}{\leq} \phi^* - \frac{1}{8T} \sum_{t=1}^{T-1} \mathbb{E}\left[(h(\theta[t]) - h(q))^2 \right]$$
(37)

where (a) holds by the fact that $x_2^* \in [0, 1]$ and $q \in [1/4, 3/4]$ so that $q/(1+x_2^*) \ge 1/8$; (b) holds by neglecting the nonnegative term for t = 0. By the expansion property of h in (26) we obtain

$$\phi(\mathbb{E}\left[\overline{X}(T)\right]) \le \phi^* - \frac{\beta^2}{8T} \sum_{t=1}^{T-1} \mathbb{E}\left[(\theta[t] - q)^2\right]$$
(38)

where $\beta = h'(1) \approx 0.2257$ is defined in Lemma 3.

Here is the crucial observation: For each slot $t \in \{1, 2, 3, ...\}$ we can view $\theta[t]$ as defined in (36) as a deterministic estimator of q based on the t observations $\{S[0], S[1], ..., S[t-1]\}$.

Indeed, $z_{F[t]}$ defined in (30) is a deterministic function of $F[t] = (S[0], \ldots, S[t-1])$, and $\theta[t]$ as defined in (36) is determined by first projecting $z_{F[t]}$ to the interval [h(0), h(1)] and then mapping the result through the continuous deterministic function $h^{-1} : [h(0), h(1)] \to [0, 1]$. In particular, $\theta[t] \in [0, 1]$. With this observation, we get from Theorem 1 that there exists a value $q \in [1/4, 3/4]$ such that:

$$\sum_{t=1}^{T-1} \mathbb{E}\left[(\theta[t] - q)^2 \right] \ge \Omega(\log(T)) \quad \forall T \in \{1, 2, 3, \ldots\}$$

Substituting this into (38) gives

$$\phi(\mathbb{E}\left[\overline{X}[T]\right]) \le \phi^* - \Omega(\log(T)/T)$$

This establishes the inequality (11) of Theorem 3.

Finally, assuming $T \ge 2$ and rearranging (38) gives

$$\frac{T\phi^* - T\phi(\mathbb{E}\left[\overline{X}[T]\right])}{\log(T)} \ge \frac{\beta^2}{8\log(T)} \sum_{t=1}^{T-1} \mathbb{E}\left[(\theta[t] - q)^2\right]$$
(39)

Again observing that $\{\theta[t]\}_{t=1}^{\infty}$ is a sequence of deterministic estimators of q, from Theorem 2 we know there is a set $Q \subseteq [1/4, 3/4]$ with measure $\mu(Q) \ge 1/6$ such that for all $q \in Q$ we have

$$\limsup_{T \to \infty} \frac{1}{\log(T)} \sum_{t=1}^{T-1} \mathbb{E}\left[(\theta[t] - q)^2 \right] \ge \frac{3}{2^{10}}$$

Substituting this into (39) yields

$$\limsup_{T \to \infty} \frac{T\phi^* - T\phi(\mathbb{E}\left[\overline{X}[T]\right])}{\log(T)} \ge \frac{\beta^2}{8} \frac{3}{2^{10}}$$

This completes the proof of Theorem 3. \Box

E. Discussion

The $O(\log(T)/T)$ achievability result derived in [2] holds for smooth and concave utility functions and does not require strong concavity. The $\Omega(\log(T)/T)$ converse bound of Theorem 3 was carried out using a smooth and *strongly concave* utility function. This was intentional: *This shows that, for these opportunistic scheduling problems, strong concavity cannot improve the fundamental convergence time.* This is surprising because strong convexity/concavity is known to significantly improve convergence time in other optimization scenarios, including deterministic subgradient minimization [25] and online convex programming [24][23].

F. Extension to other utility functions

Replace the utility function $\log(1+x_1) + \log(1+x_2)$ with the more general function $\phi: [0,1]^2 \to \mathbb{R}$:

$$\phi(x_1, x_2) = \phi_1(x_1) + \phi_2(x_2)$$

where $\phi_1(x)$ and $\phi_2(x)$ are concave and strictly increasing over [0, 1]. The converse proof can be repeated with mild additional assumptions on ϕ_1 and ϕ_2 . The main idea is to use the *implicit function theorem* of real analysis to show existence of a strictly increasing and continuously differentiable function $h: (0, 1) \to (0, 1)$ (different from the h function given for the log utility function in (24)) such that for each $q \in (0, 1)$, the value h(q) is the r value needed to define the optimal operating point $(x_1^*, x_2^*) \in C$ associated with this new utility:

$$(x_1^*, x_2^*) = ((1-q) + qr, q(1-r^2))$$

It must also be shown that there is a $\beta > 0$ such that $h'(q) \ge \beta$ for all $q \in [1/4, 3/4]$ so that the proof can proceed from (37) to (38). These properties are established in the next lemma. They allow the $\Omega(\log(T)/T)$ converse proof to be repeated using the modified estimator:

$$\theta[t] = h^{-1} \left(\left[\mathbb{E} \left[X_1[t] | S[t] = 1 \right] \right]_{h(1/4)}^{h(3/4)} \right)$$

Lemma 4: (General utilities) Suppose $\phi_1(x)$ and $\phi_2(x)$ are twice continuously differentiable functions that satisfy:

- Assumption 1: $\phi'_1(x) > 0$ and $\phi'_2(x) > 0$ for all $x \in [0, 1]$.
- Assumption 2: $\phi_1''(x) < 0$ and $\phi_2''(x) < 0$ for all $x \in [0, 1]$.
- Assumption 3: $\phi'_1(1) < 2\phi'_2(0)$.⁷

Then for each $q \in (0, 1)$ the equation:

$$\phi_1'(1 - q + qr) + \phi_2'(q(1 - r^2))(-2r) = 0$$
(40)

has a unique solution $r \in (0, 1)$. Further, there is a continuously differentiable function $h: (0, 1) \to (0, 1)$ with this property: (q, r) satisfies (40) if and only if r = h(q), and there is a $\beta > 0$ such that $h'(q) \ge \beta$ for all $q \in [1/4, 3/4]$.

Proof: Define g(q, r) for $q \in (0, 1)$ and $r \in [0, 1]$ by

$$g(q,r) = \phi_1'(1-q+qr) + \phi_2'(q(1-r^2))(-2r)$$
(41)

For $r \in (0,1)$, we see that g(q,r) = 0 if and only if (40) holds. Fix $q \in (0,1)$. Since ϕ_1 and ϕ_2 are twice differentiable, g(q,r) is a continuous function of r. We have

$$g(q,0) \stackrel{(a)}{=} \phi'_1(1-q) > 0$$

$$g(q,1) \stackrel{(b)}{=} \phi'_1(1) - 2\phi'_2(0) < 0$$

where (a) holds by Assumption 1 and (b) by Assumption 3. By the intermediate value theorem, there must exist a value $r \in (0,1)$ such that g(q,r) = 0. To show *uniqueness* of this value $r \in (0,1)$, it suffices to show that g(q,r) is strictly decreasing in r:

$$\frac{\partial g(q,r)}{\partial r} = \underbrace{\phi_1''(1-q+qr)q}_{<0} + \underbrace{\phi_2''(q(1-r^2))q(2r)^2}_{<0} + \underbrace{(-2)\phi_2'(q(1-r^2))}_{<0} < 0 \tag{42}$$

where the underbrace inequalities hold by Assumptions 1, 2, and the inequalities 0 < q < 1, 0 < r < 1. Thus, uniqueness holds, and for each $q \in (0,1)$ we can define h(q) as the unique value in (0,1) such that g(q,h(q)) = 0. Since g(q,r) is continuously differentiable (with respect to both q and r) over the open set $(q,r) \in (0,1) \times (0,1)$, and since $\partial g/\partial r \neq 0$, the *implicit function theorem* of real analysis can be applied to conclude the function h(q) is continuously differentiable.

It remains to prove existence of $\beta > 0$ such that that $h'(q) \ge \beta$ for all $q \in [1/4, 3/4]$. We have g(q, h(q)) = 0 for all $q \in (0, 1)$ and so by (41):

$$\phi_1'(1 - q + qh(q)) - 2h(q)\phi_2'(q(1 - h(q)^2)) = 0 \quad \forall q \in (0, 1)$$

⁷Note that Assumption 3 is implied by Assumption 2 in the special case $\phi_1(x) = \phi_2(x)$ for all $x \in [0, 1]$.

Taking a derivative with respect to q gives

$$0 = \phi_1''(1 - q + qh(q))(-1 + h(q) + qh'(q)) - 2h(q)\phi_2''(q(1 - h(q)^2))(1 - h(q)^2 - 2qh(q)h'(q)) - 2h'(q)\phi_2'(q(1 - h(q)^2))$$
(43)

Suppose there is a $q \in (0,1)$ for which $h'(q) \leq 0$. By Assumption 1 we know

$$2h'(q)\phi_2'(q(1-h(q))^2) \le 0$$

and so by (43):

$$0 \ge \underbrace{\phi_1''(1-q+qh(q))}_{<0} \underbrace{(-1+h(q)+qh'(q))}_{<0} \underbrace{(-2h(q)\phi_2''(1(1-h(q)^2))}_{>0} \underbrace{(1-h(q)^2-2qh(q)h'(q))}_{>0}}_{>0}$$

So 0 > 0, a contradiction. Thus, h'(q) is a continuous and always positive function over (0, 1). It must have a strictly positive minimum value over the compact interval [1/4, 3/4], call this minimum β .

The discussion in Section I-A shows that the $\Omega(\log(T)/T)$ converse does not hold for *linear* utility functions. For intuition on how the proof of Lemma 4 would fail with linear utilities, it is easy to see that if $\phi_1(x) = a_1x$, $\phi_2(x) = a_2x$ for some real numbers $a_1 > 0$, $a_2 > 0$, then the r value that solves (40) does not depend on q and hence h'(q) = 0 for all q, so there is no $\beta > 0$ such that $h'(q) \ge \beta$. Assumption 2 of Lemma 4 enforces nonlinearity. Assumption 2 implies that the ϕ function is strongly concave over the domain $[0, 1]^2$.

IV. CONCLUSION

This paper establishes a converse bound of $\Omega(\log(T)/T)$ on the utility gap for opportunistic scheduling. This matches a recently established achievability bound of $O(\log(T)/T)$. This means that $\log(T)/T$ is the optimal asymptotic behavior. The bound in this paper was proven for an example 2-user system with a strongly concave utility function. This demonstrates the surprising the result that strong concavity of the utility function cannot improve the asymptotic convergence time for opportunistic scheduling systems. This is in contrast to other optimization scenarios, such as online convex optimization, where strong convexity/concavity is known to significantly improve asymptotic convergence. The converse proof constructed a nontrivial mapping of the opportunistic scheduling problem to a Bernoulli estimation problem and used a prior result on the regret associated with Bernoulli estimation. The paper also develops a refinement on Bernoulli estimation to show that for any sequence of Bernoulli estimators, not only do probabilities exist for which the estimators perform poorly, but such probabilities have measure at least 1/6 in the unit interval. This is used to show that for any opportunistic scheduling algorithm, if nature chooses a Bernoulli state distribution by selecting the Bernoulli probability uniformly over the unit interval, the algorithm is limited by the $\Omega(\log(T)/T)$ bound with probability at least 1/6.

The converse bound of this paper was established for a simple 2-user system. This means that *there exist* systems that are limited by the $\Omega(\log(T)/T)$ bound. The techniques in this paper link opportunistic scheduling to estimation problems and can likely be used in future work to investigate bounds on more general networks, including networks with state variables S[t] that are described by more complex distributions. This motivates the following open questions: Can refined bounds be established for non-Bernoulli S[t] processes? Can more detailed coefficients of the $\log(T)/T$ curve be obtained in terms of simple parameters of the distribution on S[t]? The Cramer-Rao bound of estimation theory allows bounds for non-Bernoulli variables that depend on *Fisher information* of the underlying probability distribution. However, it is currently unclear how to reduce a general opportunistic scheduling problem to a generalized (non-Bernoulli) estimation problem, and it is not clear how to incorporate Fisher information concepts to provide "regret" type bounds for networks.

APPENDIX A — A REFINED VARIATION INEQUALITY

This appendix refines a Lemma in [17] about the total variation distance associated with the measure of i.i.d. Bernoulli random variables.⁸ Let Ω be a finite and nonempty sample space and consider the sigma algebra of all subsets of Ω . Let P[A] and P'[A] be two probability measures defined for all $A \subseteq \Omega$. Define the *total variation distance* v(P, P') by

$$v(P, P') = \sup_{A \subseteq \Omega} |P[A] - P'[A]|$$

⁸Lemma 15 of [17] contains a minor constant factor error: It claims $p \log(p/q) + (1-p) \log((1-p)/(1-q)) \le \epsilon^2/2$ for $q = p + \epsilon$ (a counter-example is $\epsilon = 1/16$, p = 1/2). This affects the statement of Lemma 15 in [17] by a constant factor that does not change the methodology or asymptotic $\Omega(\log(N))$ results proven in that paper.

The Pinsker inequality states

$$v(P, P') \le \min\left[\sqrt{\frac{1}{2}D_{KL}(P||P')}, \sqrt{\frac{1}{2}D_{KL}(P'||P)}\right]$$
(44)

where $D_{KL}(P||P')$ is the Kullback-Leibler divergence (in nats) between P and P':

$$D_{KL}(P||P') = \sum_{x \in \Omega} P[\{x\}] \log\left(\frac{P[\{x\}]}{P'[\{x\}]}\right)$$

For a fixed positive integer n, define the sample space $\{0,1\}^n$. We view all 2^n outcomes $(x_1,...,x_n) \in \{0,1\}^n$ as possible realizations of i.i.d. Bernoulli random variables $(X_1,...,X_n)$. For each $p \in [0,1]$, define B_n^p as the probability measure on $\{0,1\}^n$ associated with i.i.d. Bernoulli random variables with $P[X_i = 1] = p, P[X_i = 0] = 1 - p$. That is,

$$B_n^p[A] = \sum_{(x_1,\dots,x_n)\in A} \left[\prod_{i=1}^n p^{x_i} (1-p)^{1-x_i}\right] \quad \forall A \subseteq \{0,1\}^r$$

The next lemma is a refinement of Lemma 15 in [17] that fixes a constant factor error and also removes a restriction on |p-q| that was assumed there.

Lemma 5: Fix $p, q \in [1/4, 3/4]$. Then

$$v(B_p^n, B_q^n) \le c|p - q|\sqrt{n}$$

where $c = \sqrt{8/3} \approx 1.6329932$.

Proof: Without loss of generality assume p > q and define $\epsilon = p - q$. The Pinsker inequality (44) gives

$$v(P,P') \le \sqrt{\frac{1}{2}D_{KL}(B_n^p||B_n^q)} \tag{45}$$

By basic properties of the measure of i.i.d. Bernoulli random variables it holds that

$$D_{KL}(B_n^p || B_n^q) = n \left[p \log\left(\frac{p}{q}\right) + (1-p) \log\left(\frac{1-p}{1-q}\right) \right]$$

To compute the right-hand-side of the above equality, we have

$$p \log(\frac{p}{q}) + (1-p) \log(\frac{1-p}{1-q}) = p \log(1+\frac{\epsilon}{q}) + (1-p) \log(1-\frac{\epsilon}{1-q})$$

$$\stackrel{(a)}{\leq} p\left(\frac{\epsilon}{q}\right) + (1-p)\left(\frac{-\epsilon}{1-q}\right)$$

$$\stackrel{(b)}{=} \epsilon^2 \left[\frac{1}{q} + \frac{1}{1-q}\right]$$

$$\leq \epsilon^2 \sup_{q \in [1/4, 3/4]} \left[\frac{1}{q} + \frac{1}{1-q}\right]$$

$$= \frac{16\epsilon^2}{3}$$

where (a) uses the inequality $\log(1+x) \le x$ for all x > -1; (b) uses $p = q + \epsilon$. Hence

$$D_{KL}(B_n^p || B_n^q) \le \frac{16n\epsilon^2}{3}$$

Substituting this inequality into (45) proves the result.

We now utilize the above refined lemma. Let $\{\hat{A}_n\}_{n=1}^{\infty}$ be an infinite sequence of estimation functions, as defined in Section II-A, so that

$$\hat{A}_n: [0,1) \times \{0,1\}^n \to [0,1]$$
(46)

Fix $p, q \in [0, 1]$. Let $\{W_n^p\}_{n=1}^{\infty}$ and $\{W_n^q\}_{n=1}^{\infty}$ denote sequences of i.i.d. Bernoulli random variables with $P[W_n^p = 1] = p$ and $P[W_n^q = 1] = q$, respectively. Let U be a random variable that is uniformly distributed over [0, 1) and that is independent of both $\{W_n^p\}_{n=1}^{\infty}$ and $\{W_n^q\}_{n=1}^{\infty}$. For $n \in \{1, 2, 3, \ldots\}$, let A_n^p and A_n^q denote the estimates of p and q based on $(U, W_1^p, \ldots, W_n^p)$ and $(U, W_1^q, \ldots, W_n^q)$, respectively:

$$A_n^p = \hat{A}_n(U, W_1^p, W_2^p, \dots, W_n^p) \quad \forall n \in \{1, 2, 3, \dots\}$$
(47)

$$A_n^q = \hat{A}_n(U, W_1^q, W_2^q, \dots, W_n^q) \quad \forall n \in \{1, 2, 3, \dots\}$$
(48)

Fix $\alpha > 0$. We have

$$\mathbb{E}_p[|A_n^p - p|^{\alpha}] = \int_0^1 \mathbb{E}_p[|A_n^p - p|^{\alpha} \mid U = u]du$$
$$\mathbb{E}_p[|A_n^q - q|^{\alpha}] = \int_0^1 \mathbb{E}_q[|A_n^q - q|^{\alpha} \mid U = u]du$$

where $\mathbb{E}_p[\cdot]$ and $\mathbb{E}_q[\cdot]$ represent expectations with respect to the probability distributions that form the random vectors $(U, W_1^p, \ldots, W_n^p)$ and $(U, W_1^q, \ldots, W_n^q)$, respectively. Since U is independent of the samples, the conditional expectation $\mathbb{E}_p[|A_n^p - p|^{\alpha} | U = u]$ is with respect to the probability measure B_n^p associated only with the random vector of i.i.d. Bernoullip variables (W_1^p, \ldots, W_n^p) . Similarly, $\mathbb{E}_q[|A_n^q - q|^{\alpha} | U = u]$ considers the same estimation function $\hat{A}_n(\cdot)$ but is with respect to the probability measure B_n^q associated only with the random vector of i.i.d. Bernoulli-q variables (W_1^q, \ldots, W_n^q) . For $n \in \{1, 2, 3, \ldots\}$, these expectations can be evaluated in terms of the measures B_n^p and B_n^q for which the Pinkser inequality applies.

The following lemma generalizes Lemma 16 of [17], which treats mean square error, to treat general powers of the absolute error. The proof closely follows the structure developed in [17] but uses the refined lemma (Lemma 5 above) in a key place.

Lemma 6: Fix $\alpha > 0$. Fix any sequence of measurable estimation functions $\{\hat{A}_n\}_{n=1}^{\infty}$ of the form (46). Let p, q be probabilities that satisfy $p, q \in [1/4, 3/4]$. Then for all $n \in \{1, 2, 3, ...\}$ we have

$$E_p[|A_n^p - p|^{\alpha}] + E_q[|A_n^q - q|^{\alpha}] \ge \frac{|p - q|^{\alpha}}{2^{1 + \alpha}} \quad \text{whenever } |p - q| \le \frac{1}{2c\sqrt{n}}$$
(49)

where A_n^p and A_n^q are defined by (47) and $c = \sqrt{8/3}$.

Proof. Fix $n \in \{1, 2, 3, ...\}$. Define $\epsilon = |p - q|$ and assume that $\epsilon \leq \frac{1}{2c\sqrt{n}}$. It suffices to prove the following claim: For all $u \in [0, 1)$ we have

$$E_p[|A_n^p - p|^{\alpha} | U = u] + E_q[|A_n^q - q|^{\alpha} | U = u] \ge \frac{\epsilon^{\alpha}}{2^{1+\alpha}}$$
(50)

Without loss of generality assume $q \ge p$ so that $p = q + \epsilon$. If $\epsilon = 0$ then (50) trivially holds. Assume $\epsilon > 0$ and suppose (50) is false (we reach a contradiction). Then

$$E_p[|A_n^p - p|^{\alpha} | U = u] + E_q[|A_n^q - q|^{\alpha} | U = u] < \frac{\epsilon^{\alpha}}{2^{1+\alpha}}$$
(51)

Thus, there is a constant $\theta \in (0, 1)$ such that:⁹

$$E_p[|A_n^p - p|^{\alpha} \mid U = u] < \frac{\theta \epsilon^{\alpha}}{2^{1+\alpha}}$$
(52)

$$E_q[|A_n^q - q|^{\alpha} \mid U = u] < \frac{(1 - \theta)\epsilon^{\alpha}}{2^{1 + \alpha}}$$
(53)

Following the technique in [17], applying the Markov inequality to (52) and (53) gives

$$\begin{split} P_p[|A_n^p - p| \geq \epsilon/2 \mid U = u] &< \frac{\theta}{2} \\ P_q[|A_n^q - q| \geq \epsilon/2 \mid U = u] &< \frac{1 - \theta}{2} \end{split}$$

where $P_p[\cdot | U = u]$ and $P_q[\cdot | U = u]$ represent probabilities under the probability measures B_n^p and B_n^q , respectively. For simplicity of notation, for the remainder of this proof we suppress the explicit "U = u" conditioning, with the understanding that all probabilities are implicitly conditioned on U = u. With this simplified notation the above inequalities become

$$P_p[|A_n^p - p| \ge \epsilon/2] < \frac{\theta}{2}$$
(54)

$$P_q[|A_n^q - q| \ge \epsilon/2] < \frac{1-\theta}{2}$$
(55)

From (54) we obtain

$$P_p[A_n^p \ge p + \epsilon/2] < \frac{\theta}{2} \tag{56}$$

⁹Indeed from (51): If $E_p[|A_p^n - p|^{\alpha}] = 0$ then any $\theta \in (0, 1)$ satisfies (52) and we can choose $\theta \in (0, 1)$ sufficiently close to 0 to ensure (53). Else, if $E_q[|A_n^q - q|^{\alpha}] = 0$ then any $\theta \in (0, 1)$ satisfies (53) and we can choose $\theta \in (0, 1)$ sufficiently close to 1 to ensure (52). Else, define $\theta \in (0, 1)$ by $\theta = \frac{E_p[|A_p^n - p|^{\alpha}]}{E_p[|A_p^n - p|^{\alpha}] + E_q[|A_n^n - q|^{\alpha}]}$ and note that both (52)-(53) hold.

From (55) we obtain

$$P_{q}[|A_{n}^{q} - q| < \epsilon/2] > \frac{1}{2} + \frac{\theta}{2}$$

$$\implies P_{q}[A_{n}^{q} > q - \epsilon/2] > \frac{1}{2} + \frac{\theta}{2}$$

$$\stackrel{(a)}{\implies} P_{q}[A_{n}^{q} > p + \epsilon/2] > \frac{1}{2} + \frac{\theta}{2}$$
(57)

where (a) holds because $q = p + \epsilon$. Now define the set $C \subseteq \{0, 1\}^n$ as follows:

$$C = \{(x_1, \dots, x_n) \in \{0, 1\}^n : \hat{A}_n(x_1, \dots, x_n) > p + \epsilon/2\}$$

Then (56) implies $P_p[C] < \theta/2$ and (57) implies $P_q[C] > 1/2 + \theta/2$ and so

$$|P_p[C] - P_q[C]| > 1/2 \tag{58}$$

Then

$$1/2 \stackrel{(a)}{<} |P_p[C] - P_q[C]$$
$$\stackrel{(b)}{\leq} v(B_n^p, B_n^q)$$
$$\stackrel{(c)}{\leq} c|p - q|\sqrt{n}$$
$$= c\epsilon\sqrt{n}$$
$$\stackrel{(d)}{\leq} 1/2$$

where (a) holds by (58); (b) holds by definition of $v(\cdot)$ as the supremum absolute error over all possible events (including the event C); (c) follows by Lemma 5; (d) follows because we have assumed $c\epsilon\sqrt{n} \leq 1/2$. This gives the contradiction.

APPENDIX B — PROOF OF THEOREM 2

This appendix proves Theorem 2. For $\alpha > 0$ define

$$V_m(\alpha) = \sum_{n=1}^m (1/n)^{\alpha/2} \quad \forall m \in \{1, 2, 3, \ldots\}$$

Notice that

$$V_m(\alpha) \ge \int_1^{m+1} (1/t)^{\alpha/2} dt = \begin{cases} \log(m+1) & \text{if } \alpha = 2\\ \frac{(m+1)^{1-\alpha/2} - 1}{1-\alpha/2} & \text{if } \alpha \neq 2 \end{cases}$$

Fix $\alpha \in [0,2)$. Fix any sequence of measurable estimation functions $\{\hat{A}_n\}_{n=1}^{\infty}$ of the form (3). Define \mathcal{Q} as the set of all $p \in [1/4, 3/4]$ such that

$$\limsup_{m \to \infty} \left[\frac{1}{V_m(\alpha)} \sum_{n=1}^m \mathbb{E}_p[|A_n^p - p|^{\alpha}] \right] \ge \frac{1}{c^{\alpha} 2^{3+2\alpha}}$$
(59)

Let $\mu(\mathcal{Q})$ denote the total Lebesgue measure of the set \mathcal{Q} .

We first show that Q is Lebesgue measurable: Let \mathcal{X}_n be the set of all 2^n binary-valued vectors of size n. For each random seed $u \in [0,1)$, each function $\hat{A}_n(u, x_1, ..., x_n)$ assigns each vector $(x_1, ..., x_n) \in \mathcal{X}_n$ to a real number in the interval [0,1]. So

$$\mathbb{E}\left[|A_{n}^{p}-p|^{\alpha}|U=u\right] = \sum_{\vec{x}\in\mathcal{X}_{n}} |\hat{A}_{n}(u,\vec{x})-p|^{\alpha} \prod_{i=1}^{n} p^{x_{i}}(1-p)^{1-x_{i}}$$

and this is a continuous and bounded function of $p \in [0, 1]$. Then by the law of iterated expectations

$$\mathbb{E}\left[|A_{n}^{p}-p|^{\alpha}\right] = \sum_{\vec{x}\in\mathcal{X}_{n}} \prod_{i=1}^{n} p^{x_{i}} (1-p)^{1-x_{i}} \int_{0}^{1} |\hat{A}_{n}(u,\vec{x})-p|^{\alpha} du$$

where the integral is a well defined real number because $\hat{A}_n(u, \vec{x})$ is a bounded and measurable function of u and hence $|\hat{A}_n(u, \vec{x}) - p|$ can be integrated with respect to the uniform distribution over $u \in [0, 1]$. It follows that $\mathbb{E}[|A_n^p - p|^{\alpha}]$ is Lebesgue-measurable in p. The weighted sum of Lebesgue-measurable functions is again Lebesgue-measurable, and the set of all $p \in [1/4, 3/4]$ under which the limsup of these functions exceeds a threshold is again Lebesgue-measurable. Thus, Q is indeed Lebesgue-measurable.

We want to show $\mu(Q) \ge 1/6$.

Proof: (Theorem 2) Fix $\alpha \in (0, 2]$. For each $n \in \{1, 2, 3, \ldots\}$ define

$$\epsilon[n] = \frac{1}{2c\sqrt{n}}\tag{60}$$

with $c = \sqrt{8/3}$. It follows by Lemma 6 that if $p, q \in [1/4, 3/4]$ such that $|p - q| = \epsilon[n]$ then

$$\mathbb{E}_p[|A_n^p - p|^{\alpha}] + E_q[|A_n^q - q|^{\alpha}] \ge \frac{\epsilon[n]^{\alpha}}{2^{1+\alpha}}$$
(61)

For each $p \in [1/4, 3/4]$ define:

$$f_n(p) = \min\left[\mathbb{E}_p[|A_n^p - p|^{\alpha}], \frac{\epsilon[n]^{\alpha}}{2^{1+\alpha}}\right]$$
(62)

We first claim that if $p,q \in [1/4,3/4]$ and $|p-q| = \epsilon[n]$ then:

$$f_n(p) + f_n(q) \ge \frac{\epsilon[n]^{\alpha}}{2^{1+\alpha}} \tag{63}$$

Indeed, if either $f_n(p) \ge \epsilon[n]^{\alpha}/2^{1+\alpha}$ or $f_n(q) \ge \epsilon[n]^{\alpha}/2^{1+\alpha}$ then the result is trivially true, and otherwise the result holds by (61). Integrating the nonnegative function $f_n(p)$ gives

$$\begin{aligned} \frac{1}{1/2} \int_{1/4}^{3/4} f_n(p) dp &\stackrel{(a)}{=} \int_{1/4}^{3/4} f_n(p) dp + \int_{1/4}^{3/4} f_n(q) dq \\ &\stackrel{(b)}{\geq} \int_{1/4}^{3/4 - \epsilon[n]} f_n(p) dp + \int_{1/4 + \epsilon[n]}^{3/4} f_n(q) dq \\ &= \int_{1/4}^{3/4 - \epsilon[n]} [f_n(p) + f_n(p + \epsilon[n])] dp \\ &\stackrel{(c)}{\geq} \int_{1/4}^{3/4 - \epsilon[n]} \frac{\epsilon[n]^{\alpha}}{2^{1+\alpha}} dp \\ &= \frac{\epsilon[n]^{\alpha}}{2^{1+\alpha}} (1/2 - \epsilon[n]) \\ &\stackrel{(\underline{d})}{=} \frac{(1/n)^{\alpha/2} (1 - \frac{1}{cn^{1/2}})}{c^{\alpha} 2^{2+2\alpha}} \end{aligned}$$

where (a) holds by simply doubling the integral; for inequality (b) we note that $\epsilon[n] < 1/2$ for all $n \in \{1, 2, 3, ...\}$ and so $3/4 - \epsilon[n] > 1/4$ and $1/4 + \epsilon[n] < 3/4$; (c) holds by (63); (d) holds by definition $\epsilon[n]$ in (60). Summing over $n \in \{1, ..., m\}$ gives

$$\frac{1}{1/2} \int_{1/4}^{3/4} \left[\sum_{n=1}^{m} f_n(p) \right] dp \ge \frac{V_m(\alpha)}{c^{\alpha} 2^{2+2\alpha}} - \frac{V_m(\alpha+1)}{c^{\alpha+1} 2^{2+2\alpha}}$$
(64)

Now let Z be a random variable that is independent of all else and is uniform over [1/4, 3/4]. Define $H_m = \sum_{n=1}^m f_n(Z)$. Inequality (64) can be interpreted as

$$\mathbb{E}\left[H_m\right] \ge \frac{V_m(\alpha)}{c^{\alpha} 2^{2+2\alpha}} - \frac{V_m(\alpha+1)}{c^{\alpha+1} 2^{2+2\alpha}}$$
(65)

Inequality (62) implies a deterministic bound on H_m :

$$f_n(Z) \le \frac{\epsilon[n]^{\alpha}}{2^{1+\alpha}} = \frac{(1/n)^{\alpha/2}}{c^{\alpha}2^{1+2\alpha}}$$

Summing the above over $n \in \{1, \ldots, m\}$ gives

$$H_m \le \frac{V_m(\alpha)}{c^{\alpha} 2^{1+2\alpha}} \tag{66}$$

Define $A_m(\alpha)$ as the following event:

$$A_m(\alpha) = \left\{ H_m \ge \frac{V_m(\alpha)}{c^{\alpha} 2^{3+2\alpha}} \right\}$$

Define $A_m(\alpha)^c$ as the complement of this event. Then

$$\mathbb{E}[H_m] = \mathbb{E}[H_m \mid A_m(\alpha)^c] \left(1 - P[A_m(\alpha)]\right) + \mathbb{E}[H_m \mid A_m(\alpha)] P[A_m(\alpha)]$$
(67)

However

$$\mathbb{E}\left[H_m | A_m(\alpha)^c\right] \le \frac{V_m(\alpha)}{c^{\alpha} 2^{3+2\alpha}} \tag{68}$$

$$\mathbb{E}\left[H_m|A_m(\alpha)\right] \le \frac{V_m(\alpha)}{c^{\alpha}2^{1+2\alpha}} \tag{69}$$

where (68) holds by definition of the event $A_m(\alpha)$ and (69) uses the deterministic upper bound on H_m in (66). Substituting (68) and (69) into (67) gives

$$\mathbb{E}[H_m] \le \frac{V_m(\alpha)}{c^{\alpha} 2^{3+2\alpha}} + \frac{(3/4)V_m(\alpha)}{c^{\alpha} 2^{1+2\alpha}} P[A_m(\alpha)]$$

Substituting this inequality into (65) gives

$$\frac{V_m(\alpha)}{c^{\alpha}2^{2+2\alpha}} - \frac{V_m(\alpha+1)}{c^{\alpha+1}2^{2+2\alpha}} \le \frac{V_m(\alpha)}{c^{\alpha}2^{3+2\alpha}} + \frac{(3/4)V_m(\alpha)}{c^{\alpha}2^{1+2\alpha}}P[A_m(\alpha)]$$

Rearranging terms of the above inequality gives

$$P[A_m(\alpha)] \ge \frac{1}{3} - \frac{(2/3)V_m(\alpha+1)}{cV_m(\alpha)} \quad \forall m \in \{1, 2, 3, \ldots\}$$

Taking limits as $m \to \infty$ and using the fact that $\alpha \in (0,2]$ implies $\lim_{m\to\infty} \frac{V_m(\alpha+1)}{V_m(\alpha)} = 0$ yields:

$$\limsup_{m \to \infty} P[A_m(\alpha)] \ge 1/3$$

Thus

$$1/3 \leq \limsup_{m \to \infty} P[A_m(\alpha)]$$

$$\stackrel{(a)}{\leq} \limsup_{m \to \infty} P[\bigcup_{i=m}^{\infty} A_i(\alpha)]$$

$$\stackrel{(b)}{=} \lim_{m \to \infty} P[\bigcup_{i=m}^{\infty} A_i(\alpha)]$$

$$\stackrel{(c)}{=} P[A_m(\alpha) \quad i.o]$$

where (a) holds because $A_m(\alpha) \subseteq \bigcup_{i=m}^{\infty} A_i(\alpha)$; (b) and (c) hold by monotonicity of probability:

$$(\cup_{i=m}^{\infty} A_i(\alpha)) \searrow \{A_m(\alpha) \quad i.o\}$$
$$\implies \lim_{m \to \infty} P[\cup_{i=m}^{\infty} A_i(\alpha)] = P[A_m(\alpha) \quad i.o.]$$

where the notation "i.o." represents "infinitely often," that is, $P[A_m(\alpha) \quad i.o.]$ is the probability that the event $A_m(\alpha)$ occurs for an infinite number of indices m. Thus

$$1/3 \le P[A_m(\alpha) \quad i.o.]$$

However, by definition of the events $A_m(\alpha)$ we have

$$\{A_m(\alpha) \quad i.o.\} \subseteq \left\{ \limsup_{m \to \infty} \frac{H_m}{V_m(\alpha)} \ge \frac{1}{c^{\alpha} 2^{3+2\alpha}} \right\}$$

Thus

$$1/3 \le P\left[\limsup_{m \to \infty} \frac{H_m}{V_m(\alpha)} \ge \frac{1}{c^{\alpha} 2^{3+2\alpha}}\right]$$
(70)

Finally we note by definition of $f_n(p)$ in (62) that

$$H_m = \sum_{n=1}^m f_n(Z) \le \sum_{n=1}^m E_Z[|A_n^Z - Z|^{\alpha}]$$

Substituting this into (70) gives

$$P\left[\limsup_{m \to \infty} \frac{\sum_{n=1}^{m} E_Z[|A_n^Z - Z|^{\alpha}]}{V_m(\alpha)} \ge \frac{1}{c^{\alpha} 2^{3+2\alpha}}\right] \ge 1/3$$

Since Z is chosen uniformly over the size-(1/2) interval [1/4, 3/4] it follows that the measure of all values $p \in [1/4, 3/4]$ for which the above $\lim \sup$ inequality holds is at least 1/6, that is, $\mu(Q) \ge 1/6$.

Appendix C — Tightness of Bounds for Bernoulli Estimation

Let $\{W_n\}_{n=1}^{\infty}$ be i.i.d. Bernoulli random variables with $P[W_n = 1] = p$, where $p \in [0, 1]$ is an unknown parameter. Theorem 2 establishes a lower bound on regret for arbitrary (biased or unbiased) estimators of p. This appendix shows that the simple (and unbiased) estimators $\{A_n\}_{n=1}^{\infty}$ defined by

$$A_n = \frac{1}{n} \sum_{i=1}^n W_i \quad \forall n \in \{1, 2, 3, \ldots\}$$
(71)

achieve the regret bounds of Theorem 2 to within a constant factor. Fix $\alpha \in (0,2]$. Define

$$regret_n = \sum_{m=1}^n \mathbb{E}\left[|A_m - p|^{\alpha}\right]$$

• (Case $\alpha = 2$) For each $n \in \{1, 2, 3, ...\}$ we have from (71):

$$\mathbb{E}\left[\left(A_n - p\right)^2\right] = \mathbb{E}\left[\left(\frac{1}{n}\sum_{m=1}^n (W_n - p)\right)^2\right] = \frac{p(1-p)}{n}$$
(72)

and so

$$\begin{aligned} regret_{n} &= p(1-p) \sum_{m=1}^{n} \frac{1}{m} \\ \stackrel{(a)}{\leq} \frac{1}{4} \sum_{m=1}^{n} \frac{1}{m} \\ &\leq \frac{1}{4} \left[1 + \int_{1}^{n} \frac{1}{t} dt \right] \\ &= \frac{1}{4} + \frac{1}{4} \log(n) \quad \forall n \in \{1, 2, 3, \ldots\} \end{aligned}$$

where (a) holds because $p(1-p) \leq 1/4$ whenever $p \in [0,1]$. Thus

$$\limsup_{n \to \infty} \left[\frac{regret_n}{\log(n+1)} \right] \le \frac{1}{4}$$

This value 1/4 is within a factor of 85.4 from the lower bound $3/2^{10}$ of Theorem 2.

• (Case $0 < \alpha < 2$) Fix $m \in \{1, 2, 3, ...\}$. We have

$$\mathbb{E}\left[|A_m - p|^{\alpha}\right] = \mathbb{E}\left[\left(|A_m - p|^2\right)^{\alpha/2}\right]$$

$$\stackrel{(a)}{\leq} \mathbb{E}\left[|A_m - p|^2\right]^{\alpha/2}$$

$$\stackrel{(b)}{\equiv} \left(\frac{p(1-p)}{m}\right)^{\alpha/2}$$
(73)

where (a) holds by Jensen's inequality on the concave function $x^{\alpha/2}$; (b) holds by (72). Thus

$$\begin{aligned} \operatorname{regret}_{n} &= \sum_{m=1}^{n} \mathbb{E}\left[|A_{m} - p|^{\alpha}\right] \\ &= (p(1-p))^{\alpha/2} \sum_{m=1}^{n} \frac{1}{m^{\alpha/2}} \\ &\leq (1/4)^{\alpha/2} \left[1 + \int_{1}^{n} 1/t^{\alpha/2} dt\right] \\ &= \left[\frac{(1/2)^{\alpha}}{1 - \alpha/2}\right] \left[n^{1-\alpha/2} - \frac{\alpha}{2}\right] \quad \forall n \in \{1, 2, 3, \ldots\} \\ &\lim_{n \to \infty} \sup_{n \to \infty} \left[\frac{\operatorname{regret}_{n}}{(n+1)^{1-\alpha/2} - 1}\right] \leq \frac{(1/2)^{\alpha}}{1 - \alpha/2} \end{aligned}$$

Thus

This value is within a factor $\frac{2^{3+\frac{5\alpha}{2}}}{3^{\alpha/2}}$ of the lower bound $\frac{1}{c^{\alpha}2^{3+2\alpha}(1-\alpha/2)}$ of Theorem 2 (where $c = \sqrt{8/3}$). For $\alpha = 1$ the constant factor gap is $32\sqrt{2/3} \approx 26.12789$. For $\alpha \to 0^+$ the constant factor gap converges to 8.

REFERENCES

- [1] M. J. Neely. A converse result on convergence time for opportunistic wireless scheduling. Proc IEEE INFOCOM, 2020.
- [2] M. J. Neely. Convergence and adaptation for utility optimal opportunistic scheduling. *IEEE/ACM Transactions on Networking*, 27(3):904–917, June 2019.
- [3] M. J. Neely. Stochastic Network Optimization with Application to Communication and Queueing Systems. Morgan & Claypool, 2010.
- [4] F. Kelly. Charging and rate control for elastic traffic. European Transactions on Telecommunications, vol. 8, no. 1 pp. 33-37, Jan.-Feb. 1997.
- [5] F.P. Kelly, A. Maulloo, and D. Tan. Rate control for communication networks: Shadow prices, proportional fairness, and stability. *Journ. of the Operational Res. Society*, vol. 49, no. 3, pp. 237-252, March 1998.
- [6] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. IEEE/ACM Transactions on Networking, vol. 8, no. 5, Oct. 2000.
- [7] E. Altman, K. Avrachenkov, and A. Garnaev. Generalized α-fair resource allocation in wireless networks. In *Proc. Conference on Decision and Control*, Dec. 2008.
- [8] T. Lan, D. Kao, M. Chiang, and A. Sabharwal. An axiomatic theory of fairness in network resource allocation. In Proc. IEEE INFOCOM, March 2010.
- H. Kushner and P. Whiting. Asymptotic properties of proportional-fair sharing algorithms. Proc. 40th Annual Allerton Conf. on Communication, Control, and Computing, Monticello, IL, Oct. 2002.
- [10] R. Agrawal and V. Subramanian. Optimality of certain channel aware scheduling policies. Proc. 40th Annual Allerton Conf. on Communication, Control, and Computing, Monticello, IL, Oct. 2002.
- [11] M. J. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. *IEEE/ACM Transactions on Networking*, vol. 16, no. 2, pp. 396-409, April 2008.
- [12] M. J. Neely. Energy optimal control for time varying wireless networks. IEEE Transactions on Information Theory, vol. 52, no. 7, pp. 2915-2934, July 2006.
- [13] A. Eryilmaz and R. Srikant. Joint congestion control, routing, and MAC for stability and fairness in wireless networks. *IEEE Journal on Selected Areas in Communications, Special Issue on Nonlinear Optimization of Communication Systems*, vol. 14, pp. 1514-1524, Aug. 2006.
- [14] A. Eryilmaz and R. Srikant. Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control. IEEE/ACM Transactions on Networking, vol. 15, no. 6, pp. 1333-1344, Dec. 2007.
- [15] A. Stolyar. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. *Queueing Systems*, vol. 50, no. 4, pp. 401-457, 2005.
- [16] A. Stolyar. Greedy primal-dual algorithm for dynamic resource allocation in complex networks. Queueing Systems, vol. 54, no. 3, pp. 203-220, 2006.
- [17] E. Hazan and S. Kale. Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization. Journal of Machine Learning Research, vol. 15 (July):pp. 2489–2512, 2014.
- [18] A. Agarwal, P. L. Bartlett, P. Ravikumar, and M. J. Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Transactions on Information Theory*, vol. 58, no. 5:pp. 3235–3249, May 2012.
- [19] A. S. Nemirovski and D. B. Yudin. Problem Complexity and Method Efficiency in Optimization. Chichester-Wiley, 1983.
- [20] P. Massart and E. Nédélec. Risk bounds for statistical learning. Annals of Statistics, 34(5):2326–2366, 2006.
- [21] Y. Han, J. Jiao, and T. Weissman. Minimax estimation of discrete distributions under l_1 loss. *IEEE Transactions on Information Theory*, 61(11):6343–6354, Nov. 2015.
- [22] J. Jiao, K. Venkat, Y. Han, and Tsachy Weissman. Minimax estimation of functionals of discrete distributions. IEEE Transactions on Information Theory, 61(5):2835–2885, May 2015.
- [23] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, vol. 69, no. 2, pp. 169-192, Dec. 2007.
- [24] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. Proc. 20th International Conference on Machine Learning (ICML), 2003.
- [25] Y. Nesterov. Introductory Lectures on Convex Optimization: A Basic Course. Kluwer Academic Publishers, Boston, 2004.
- [26] T. M. Cover and J. A. Thomas. Elements of Information Theory. New York: John Wiley & Sons, Inc., 1991.