2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)
Allerton Park and Retreat Center
Monticello, IL, USA, September 24-27, 2019

Sum throughput on a random access erasure collision channel

Hariharan Narayanan
Dept. of ECE
Drexel University
Philadelphia, PA
hn384@drexel.edu

Steven Weber
Dept. of ECE
Drexel University
Philadelphia, PA
spw26@drexel.edu

***Abstract*—The random access erasure collision channel captures, in an abstracted manner, several important features of a wireless environment shared by uncoordinated radios. The radios employ random access and, when contending, transmit over independent heterogeneous erasure channels with the common access point. The access point is capable of only receiving a single message at a time, and so any colliding messages are lost. The combined effects of the channel heterogeneity and the collision rule give rise to a natural question: how does the expected sum throughput vary with the subset of radios that are active? The subset of radios achieving the optimal throughput is found by a simple *greedy packing* procedure — add the radios, sorted by nonerasure probability, until a target offered load is exceeded.**

***Index Terms*—wireless; random access; erasure channel; collision channel; throughput; majorization**

I. Introduction

The random access erasure collision channel, illustrated in Fig. 1, captures, in an abstracted manner, several important features of a wireless environment shared by uncoordinated radios. The classical collision channel for random access, in which a collection of radios contend in a distributed manner to communicate with a shared access point (AP), illustrates the fundamental tension in such a shared environment — the tradeoff between under-utilization (resulting in unused time slots) and overloading (resulting in collisions and packet loss).

This work generalizes the classical collision channel by incorporating channel heterogeneity through the use of independent erasure channels, characterized by nonerasure probabilities $q \equiv (q_1, \ldots, q_n)$. The same fundamental utilization tradeoff must still be navigated, with the complicating factor that the contention probabilities, $p \equiv (p_1, \ldots, p_n)$, must be chosen as a function of q so as to maximize the probability of precisely one message arriving at the AP.

Our previous work [1] analyzed this channel from the perspective of economic incentives for participation, and evaluated the proposed mechanism using the Price of Anarchy (PoA) framework. It was stated without proof that extremal contention probabilities, i.e., $p \in \{0,1\}^n$, maximized the sum expected throughput, thereby transforming the nonlinear optimization into a combinatorial optimization over all subsets of radios. It was also stated without proof that the optimal throughput is achieved by the *greedy packing* procedure of adding radios, sorted in decreasing nonerasure probability, until the load (defined later) exceeds a threshold (of one).

Support from the National Science Foundation, Award CNS-1816387, is gratefully acknowledged. An unpublished earlier version of this work was presented by S. Weber at the Information Theory and Applications (ITA) Workshop in San Diego, CA in February, 2019. S. Weber is the contact author.

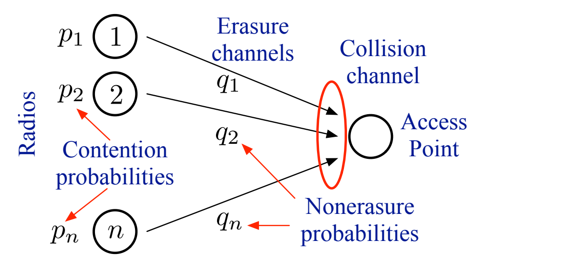

Fig. 1. The random access erasure collision channel consists of n radios employing random access (with heterogeneous contention probabilities $p_1, \ldots, p_n$) in a shared wireless environment, connected with a common access point (AP) by independent but heterogeneous erasure channels with nonerasure probabilities $q_1, \ldots, q_n$. The AP is subject to the constraint that multiple messages arriving in the same time slot collide and are lost.

The primary contributions of this paper are four theorems:

1) Thm. 2 bounds the throughput over all offered load vectors with specified max and sum load
2) Thm. 3 establishes that throughput maximization over offered load vectors is a combinatorial problem
3) Thm. 4 gives the throughput–optimal subset of radios for a given load vector
4) Thm. 5 bounds the throughput over all subsets with common cardinality.

This paper is part of the large literature on throughput maximization over a wireless medium using random access; due to space constraints we only highlight a few of the references most influential in our research. Ref. [2] maximizes throughput using distributed price signaling in a non-cooperative game. Ref. [3] generalizes slotted Aloha to allow a radio to continue to transmit until a collision occurs. Ref. [4] develops a distributed channel-aware random access mechanism. Ref. [5] describes coordinated and uncoordinated throughput maximization, focused on machine to machine (M2M) communications. Ref. [6] derives the maximum sum rate of slotted Aloha under a capture model. Ref. [7] maximizes throughput under the "semi-Poisson" model of the Aloha protocol with exponential back-off. Ref. [8] characterizes tradeoffs among throughput, reliability, latency, and bandwidth under random access. Finally, our prior work [9] explores the throughput–fairness tradeoff of slotted Aloha under a stability constraint.

The paper is organized as follows: §II introduces general notation, §III reviews key results from majorization theory used in the paper, §IV defines the random access erasure collision channel, §V defines sum expected throughput and its properties, §VI gives results that bound the throughput over varying loads and subsets, and §VII presents some numerical results. Some technical results are proved in the Appendix.

II. General notation

Let $a \equiv b$ denote a and b are equal by definition.

For $p \in [0,1]$, let $\bar{p} \equiv 1-p$ denote the *complement* of p.

Let $[m:n] \equiv \{m,\ldots,n\}$ for $m \leq n$, and $[n] \equiv [1:n]$.

Points $x \in \mathbb{R}^n$ are termed both *lists* and *vectors*; restrictions to $\mathbb{R}^k$ for $k<n$ are termed both *sublists* and *subvectors*.

For any $x \in \mathbb{R}^n$ let $x^{\downarrow} \equiv (x_{[1]},\ldots,x_{[n]})$ denote x as a list in descending order. Let $\mathscr{D}^n \equiv \{x \in \mathbb{R}^n : x_1 \geq \cdots \geq x_n\}$ and $\mathscr{D}^n_+ \equiv \{x \in \mathbb{R}^n : x_1 \geq \cdots \geq x_n \geq 0\}$, denote the set of all lists with elements arranged in a descending order and its nonnegative counterpart, respectively. Note $\mathscr{D}^n = \{x \in \mathbb{R}^n : x = x^{\downarrow}\}$ and $\mathscr{D}^n_+ = \{x \in \mathbb{R}^n_+ : x = x^{\downarrow}\}$.

Let A be a subset of $[n]$ ($A \subseteq [n]$) and i an element of $[n]$ ($i \in [n]$). The shorthand $A_{\cup i}$ for $A \cup \{i\}$ and $A_{\setminus i}$ for $A \setminus \{i\}$ is used. Any (unordered) subset $A \subseteq [n]$, say $A = \{a_1,\ldots,a_k\}$, is also an (ordered) list $A = (a_1,\ldots,a_k)$. When a set is denoted with a subscript, say A_k, then the subscript indicates the cardinality of the set, $|A|=k$.

Definition 1 (product order): Equal cardinality subsets of $[n]$, say A_k, B_k with $|A|=|B|=k$, each viewed as a list, say $A_k = (a_1,\ldots,a_k)$ and $B_k = (b_1,\ldots,b_k)$, obey *product order* (abbreviated *po*), denoted $A_k \leq_{\mathrm{po}} B_k$, if $a_i \leq b_i$ for $i \in [k]$.

Definition 2 (k-sets): For $k \in [n]$, a k-set, say A_k, is any subset of $[n]$ of cardinality k. Two special k-sets are given special notation: $i)$ $F_k = [k]$, the set of the first k indices, is termed the *forward-packed set*, and $ii)$ $R_k = [n-k+1:n]$, the set of the last k indices, is termed the *reverse-packed set*.

Fix $x \in \mathbb{R}^n$. For any $A \subset [n]$, the notation $x(A) = (x_i, i \in A)$ denotes elements of x indexed by A, termed the *sublist* or *subvector* of x indexed by A. Observe $x(F_k), x(R_k)$ are the first (last) k elements from x, and if $x \in \mathscr{D}^n$ then $x(F_k), x(R_k)$ are the k largest (smallest) elements from x.

Fact 1: $i)$ If A_k is a k-set then $F_k \leq_{\mathrm{po}} A_k \leq_{\mathrm{po}} R_k$. $ii)$ For $x \in \mathscr{D}^n$: if $A \leq_{\mathrm{po}} B$ then $x(B) \leq_{\mathrm{po}} x(A)$. $iii)$ For $x \in \mathscr{D}^n$ and a k-set A_k, $x(R_k) \leq_{\mathrm{po}} x(A_k) \leq_{\mathrm{po}} x(F_k)$.

III. Majorization and order-preserving functions

Pertinent concepts and results from majorization theory are reviewed, using (standard) notation and definitions from [10].

Definition 3 (majorization): For $x,y \in \mathbb{R}^n$, x is *majorized by* y (equivalently, y majorizes x), denoted $x \prec y$, if

$$\sum_{i=1}^k x_{[i]} \leq \sum_{i=1}^k y_{[i]},\ k \in [n-1],\ \sum_{i=1}^n x_{[i]} = \sum_{i=1}^n y_{[i]}. \quad (1)$$

Functions that preserve the ordering of majorization are defined below.

Definition 4: A function ϕ preserves majorization on $\mathscr{D}^n$ if

$$x \prec y \text{ on } \mathscr{D}^n \ \Rightarrow\ \phi(x) \leq \phi(y). \quad (2)$$

The class of functions that preserve majorization are termed *Schur-convex* functions. This class is characterized below.

Theorem 1 (3.A.3 in [10]): Let ϕ be a real-valued function, defined and continuous on $\mathscr{D}^n$, and continuously differentiable on the interior of $\mathscr{D}^n$. Then, ϕ is Schur-convex on $\mathscr{D}^n$ iff

$$\frac{\partial \phi(z)}{\partial z_1} \geq \cdots \geq \frac{\partial \phi(z)}{\partial z_n},\ z \in \mathscr{D}^n. \quad (3)$$

The above criterion may be restated as: ϕ is Schur-convex iff ϕ is *symmetric* (i.e., ϕ is permutation invariant) and

$$\nabla_{\phi}(z,i,j) \equiv (z_i - z_j)\left(\frac{\partial \phi}{\partial z_i} - \frac{\partial \phi}{\partial z_j}\right) \geq 0,\ z \in \mathbb{R}^n. \quad (4)$$

IV. The random access erasure collision channel

Erasure channels. Let $n \in \mathbb{N}$ denote the number of radios contending for transmission (uplink) to an access point (AP) or (equivalently, for this scenario) cellular base station (BS), on a shared wireless medium and let each radio be indexed by an $i \in [n]$. The channel between each radio and the AP is modeled as an erasure channel with nonerasure parameter $q_i \in (0,1)$, meaning the transmission is either received at the AP with probability q_i, or is erased with probability $\bar{q}_i \equiv 1-q_i$.

Channel access is slotted in time and time slots are synchronized across radios. Channels are assumed to be identically distributed in time, and independent across both space and time, i.e., if $\mathsf{x}_{i,t} \sim \mathrm{Ber}(q_i)$ is a Bernoulli ("Ber") RV indicating the erasure status of radio i in time slot t (with $\mathsf{x}_{i,t} = 1$ denoting nonerasure), then the sequence of indicators for radio i in time, denoted $(\mathsf{x}_{i,t})_{t \in \mathbb{N}}$, is independent and identically distributed (IID), and the sequence of indicators for time t across radios, denoted $(\mathsf{x}_{i,t})_{i \in [n]}$, is independent.

Collision. The receiver (AP or BS) operates under the classical collision model: it successfully receives any transmission that is the sole transmission in its time slot, but irreparably loses any and all transmissions in which there are multiple transmissions in the time slot.

Contention. Each radio is assumed to have an infinite backlog of packets for transmission, and to employ random access in contending for the shared channel. That is, the random contention decisions in time by each radio are IID across time, and the random contention decisions by the n radios in a given time slot are independent. Let $p_i \in [0,1]$ denote the contention probability for radio i, and $p = (p_i, i \in [n])$ denote the contention probability vector. Each packet transmission is of the same size, assumed to be matched to the time slot duration.

V. Throughput properties and bounds

Throughput is defined in §V-A, it is shown to be Schur-convex in §V-B, and it is bounded under given sum and max load in §V-C.

A. Definitions

Our figure of merit for the performance of the random access erasure collision channel is the (expected sum) throughput, denoted $T(r)$. The discussion in §IV makes clear that

$$T(r) \equiv \sum_{i \in [n]} r_i \prod_{j \in [n] \setminus i} \bar{r}_j, \tag{5}$$

where $r_i \equiv p_i q_i$ is the probability the message is *received* at the AP (in any given time slot), i.e., radio i contends and the message on channel i is not erased (it may still be lost due to collision). Without loss of generality, radios are labeled in terms of decreasing receive probability, i.e., $r_1 \geq \cdots \geq r_n$.

Define the parameter $\lambda_i \equiv r_i / \bar{r}_i$, observing that λ_i, termed the (offered) *load*, when viewed as the function $r_i/(1-r_i)$ of $r_i \in [0,1)]$, is convex increasing in r_i and onto $[0, \infty)$. As λ_i is increasing in r_i, the assumed ordering of $r_1 \geq \cdots \geq r_n$ ensures the ordering is preserved for $\lambda \equiv (\lambda_1, \ldots, \lambda_n)$, i.e., $\lambda_1 \geq \cdots \geq \lambda_n$. The throughput expression in the definition below follows from (5), the definition of λ_i, and algebra.

Definition 5 (Throughput and sum offered load): The (sum expected) throughput associated with offered loads λ is:

$$T(\lambda) = \frac{\Lambda(\lambda)}{\Pi(\lambda)}, \ \Lambda(\lambda) \equiv \sum_{i \in [n]} \lambda_i, \ \Pi(\lambda) \equiv \prod_{i \in [n]} (1 + \lambda_i). \tag{6}$$

The quantity $\Lambda(\lambda)$ is termed the *sum (offered) load*.

Definition 6 (Channel loading): A channel with offered loads λ is labeled either $i)$ *overloaded* when the sum (offered) load equals or exceeds one: $\Lambda(\lambda) \geq 1$, or $ii)$ *underloaded* when the sum (offered) load is strictly less than one: $\Lambda(\lambda) < 1$.

B. Throughput is Schur-convex

The following lemma establishes the throughput function to be Schur-convex (Def. 4), and will be used in the proofs of several subsequent results, including Thm. 2 and Thm. 4.

Lemma 1: The throughput $T(\lambda)$ in Def. 5 is Schur-convex.

Proof: Simple calculus yields:

$$\frac{\partial T(\lambda)}{\partial \lambda_i} = \frac{1 - \Lambda(\lambda) + \lambda_i}{\Pi(\lambda)}, \ i \in [n]. \tag{7}$$

It follows, recalling the definition in (4), that $\nabla_T(\lambda, i, j)$

$$\begin{aligned} &= (\lambda_i - \lambda_j) \left(\frac{1 - \Lambda(\lambda) + \lambda_i}{\Pi(\lambda)} - \frac{1 - \Lambda(\lambda) + \lambda_j}{\Pi(\lambda)} \right) \\ &= \frac{(\lambda_i - \lambda_j)^2}{\Pi(\lambda)} \geq 0. \end{aligned} \tag{8}$$

It is apparent that $T(\lambda)$ satisfies Schur's criterion (Thm. 1). ■

C. Throughput bounds for given sum and maximum load

Thm. 2 below gives lower and upper bounds on the throughput, $T(\lambda)$ over a given feasible set, defined as

$$\mathcal{L}^{(n)}(\Lambda, \bar{\lambda}) \equiv \{\lambda \in [0, \bar{\lambda}]^n : \Lambda(\lambda) = \Lambda\}. \tag{9}$$

Specifically, it shows how the dimension n, sum load Λ, and maximum load $\bar{\lambda}$ bound the throughput T.

Theorem 2: Given $(n, \Lambda, \bar{\lambda})$, the throughput has lower $(T_{\mathrm{l}}^{(n)}(\Lambda, \bar{\lambda}))$ and upper $(T_{\mathrm{u}}^{(n)}(\Lambda, \bar{\lambda}))$ bounds, for $\Lambda \in [0, n\bar{\lambda}]$:

$$\begin{aligned} \min_{\lambda \in \mathcal{L}^{(n)}(\Lambda, \bar{\lambda})} T(\lambda) &= \frac{\Lambda}{(1 + \Lambda/n)^n} \\ \max_{\lambda \in \mathcal{L}^{(n)}(\Lambda, \bar{\lambda})} T(\lambda) &= \frac{\Lambda}{(1+\bar{\lambda})^{\lfloor \frac{\Lambda}{\bar{\lambda}} \rfloor} \left(1 + \Lambda - \lfloor \frac{\Lambda}{\bar{\lambda}} \rfloor \bar{\lambda}\right)} \end{aligned} \tag{10}$$

with $\lfloor \cdot \rfloor$ denoting floor. These bounds are achievable: $i)$ the lower bound by $\lambda_{\min} \equiv (\Lambda/n, \ldots, \Lambda/n)$, and $ii)$ the upper bound by any permutation of λ of the form

$$\lambda_i = \begin{cases} \bar{\lambda}, & i \in [\lfloor \frac{\Lambda}{\bar{\lambda}} \rfloor] \\ \Lambda - \lfloor \frac{\Lambda}{\bar{\lambda}} \rfloor \bar{\lambda}, & i = \lfloor \frac{\Lambda}{\bar{\lambda}} \rfloor + 1 \\ 0, & i \in [\lfloor \frac{\Lambda}{\bar{\lambda}} \rfloor + 2 : n] \end{cases} \tag{11}$$

The λ in (11) has as few nonzero components as possible, subject to the sum and max constraints.

Proof: Define $\lambda_{\max}$ as any permutation of λ of the form in (11). Then, for any $\lambda \in \mathcal{L}(\Lambda, \bar{\lambda})$:

$$\lambda_{\min} \prec \lambda \prec \lambda_{\max}. \tag{12}$$

As $T(\lambda)$ is Schur-convex (Lem. 1), it follows (Def. 4) that $T(\lambda_{\min}) \leq T(\lambda) \leq T(\lambda_{\max})$. ■

Proof (alternate, lower bound): As Λ is fixed in $\mathcal{L}$, to prove the lower bound it suffices to show

$$\max_{\lambda \in \mathcal{L}^{(n)}(\Lambda, \bar{\lambda})} \Pi(\lambda) = (1 + \Lambda/n)^n. \tag{13}$$

This follows immediately from the AM-GM inequality:

$$\left(\prod_i (1 + \lambda_i) \right)^{\frac{1}{n}} \leq \frac{1}{n} \sum_i (1 + \lambda_i). \tag{14}$$

■

Remark 1: The following points merit comment.

$i)$ $T_{\mathrm{u}}^{(n)}(\Lambda, \bar{\lambda})$ is increasing in $\bar{\lambda}$, with limit

$$\lim_{\bar{\lambda} \uparrow \infty} T_{\mathrm{u}}^{(n)}(\Lambda, \bar{\lambda}) = \frac{\Lambda}{1 + \Lambda}. \tag{15}$$

$ii)$ $T_{\mathrm{l}}^{(n)}(\Lambda, \bar{\lambda})$ is decreasing in n with limit

$$\lim_{n \uparrow \infty} T_{\mathrm{l}}^{(n)}(\Lambda, \bar{\lambda}) = \Lambda \mathrm{e}^{-\Lambda}. \tag{16}$$

$iii)$ *Upper bound:* consider separately $n\bar{\lambda} \geq 1$ and $n\bar{\lambda} < 1$.

Suppose $n\bar{\lambda} \geq 1$. The upper bound is concave increasing over $\Lambda \in \left[0, \lceil \frac{1}{\bar{\lambda}} \rceil \bar{\lambda}\right]$ (with $\lceil \cdot \rceil$ denoting ceiling), noting $\lceil \frac{1}{\bar{\lambda}} \rceil \bar{\lambda} \geq 1$, achieving a maximum value at $\Lambda = \lceil \frac{1}{\bar{\lambda}} \rceil \bar{\lambda}$ of

$$T_{\mathrm{u}}^{(n)} \left(\left\lceil \frac{1}{\bar{\lambda}} \right\rceil \bar{\lambda}, \bar{\lambda} \right) = \frac{\lceil \frac{1}{\bar{\lambda}} \rceil \bar{\lambda}}{(1 + \bar{\lambda})^{\lceil \frac{1}{\bar{\lambda}} \rceil}} \tag{17}$$

and is thereafter decreasing in Λ. The maximum value of the upper bound corresponds to λ with $\lceil \frac{1}{\bar{\lambda}} \rceil$ components with value $\bar{\lambda}$, and the remaining components at zero. When $\bar{\lambda} \geq 1$, observe $\lceil \frac{1}{\bar{\lambda}} \rceil = 1$, and thus the upper bound is concave increasing over $\Lambda \in [0, \bar{\lambda}]$, achieving a maximum value of $\bar{\lambda}/(1+\bar{\lambda})$ (corresponding to the point $\lambda = (\bar{\lambda}, 0, \ldots 0)$) at $\Lambda = \bar{\lambda}$, and is thereafter decreasing in Λ.

Suppose $n\bar{\lambda} < 1$. The upper bound is increasing in Λ, achieving at $\Lambda = n\bar{\Lambda}$ (i.e., $\lambda = (\bar{\lambda}, \ldots, \bar{\lambda})$) the value

$$T_{\rm u}^{(n)}(n\bar{\lambda}, \bar{\lambda}) = \frac{n\bar{\lambda}}{(1+\bar{\lambda})^n}. \tag{18}$$

iv) Lower bound: Throughput is concave over $\Lambda \in [0, 2n/(n-1)]$, convex over $\Lambda \in (2n/(n-1), \infty)$, increasing over $\Lambda \in [0, n/(n-1)]$, decreasing over $\Lambda \in (n/(n-1), \infty)$, with maximum valule of $(1-1/n)^{n-1}$ at $\Lambda = n/(n-1)$.

VI. Throughput over subsets and Throughput sequences

The notation $T(\lambda), \Lambda(\lambda), \Pi(\lambda)$ is used in §V to emphasize the dependence of these quantities upon λ. By contrast, in this section the list of offered loads λ is viewed as fixed, and the focus is instead on which subset of indices (i.e., which radios) will be active on the channel. In particular, consider these same functions as the subset $A \subseteq [n]$ is varied, recalling that the subset A induces the sublist $\lambda(A) \equiv (\lambda_i, i \in A)$.

Definition 7: For fixed offered loads λ, the throughput associated with index set $A \subseteq [n]$ is

$$T(A) \equiv \frac{\Lambda(A)}{\Pi(A)},\ \Lambda(A) \equiv \sum_{i \in A} \lambda_i,\ \ \Pi(A) \equiv \prod_{i \in A}(1+\lambda_i), \tag{19}$$

adopting the shorthand $f(A) \equiv f(\lambda(A))$ for $f \in \{T, \Lambda, \Pi\}$.

Note Λ is *additive* and Π is *multiplicative*: for A, B disjoint

$$\Lambda(A \cup B) = \Lambda(A) + \Lambda(B),\ \ \Pi(A \cup B) = \Pi(A)\Pi(B). \tag{20}$$

Moreover, Λ and Π are both *set-increasing* in the sense that $A \subset B$ implies $\Lambda(A) < \Lambda(B)$, and $\Pi(A) < \Pi(B)$.

By letting A be each possible subset of $[n]$ we obtain the *throughput collection* (for the fixed λ), i.e., $\{T(A) : A \subseteq [n]\}$.

A. Throughput maximization is combinatorial

The first motivation for studying the throughput collection is the fact that maximization of throughput $T(\lambda)$ is achieved by extremizing λ. Define $\mathcal{L}(\underline{\lambda}, \bar{\lambda}) \equiv \{\lambda \in \mathbb{R}^n : \lambda \in [\underline{\lambda}, \bar{\lambda}]\}$ and $\mathcal{M}(\underline{\lambda}, \bar{\lambda}) \equiv \{\lambda \in \mathbb{R}^n : \lambda \in \{\underline{\lambda}, \bar{\lambda}\}\}$, where $(\underline{\lambda}, \bar{\lambda})$ are n-vectors obeying $\underline{\lambda} \leq \bar{\lambda}$. In words, $\mathcal{L}(\underline{\lambda}, \bar{\lambda})$ is a coordinate-wise box constraint on feasible values of λ, specifying lower- and upper-bounds on each coordinate in the form $\underline{\lambda}_i \leq \lambda_i \leq \bar{\lambda}_i$, and $\mathcal{M}(\underline{\lambda}, \bar{\lambda})$ is the corresponding discrete extremal set, consisting of points λ in which each coordinate takes either its minimum or maximum value, i.e., $\lambda_i \in \{\underline{\lambda}_i, \bar{\lambda}_i\}$.

The motivating example of the feasible set $\mathcal{L}(\underline{\lambda}, \bar{\lambda})$ is the maximization of expected sum throughput on the random access erasure collision channel, with nonerasure probabilities $q \in (0,1)^n$, over the set of feasible contention probabilities $p \in [0,1]^n$. Using $r_i = p_i q_i$ and $\lambda_i = r_i/\bar{r}_i$ yields the feasible set for λ with $\underline{\lambda} = 0$ and $\bar{\lambda}_i = q_i/\bar{q}_i$, for $i \in [n]$.

The following theorem, stated without proof in our prior work [1], asserts that throughput maximization is a combinatorial problem, i.e., when maximizing throughput over contention probablities it suffices to restrict attention to the case where each radio either never or always contends, i.e., $p^* \in \{0,1\}^n$, equivalently, it suffices to restrict λ to $\mathcal{M}(\underline{\lambda}, \bar{\lambda})$.

Theorem 3: Maximum throughput (over $\mathcal{L}(\underline{\lambda}, \bar{\lambda})$) is achieved by an extremal value of λ (in $\mathcal{M}(\underline{\lambda}, \bar{\lambda})$):

$$\max_{\lambda \in \mathcal{L}(\underline{\lambda}, \bar{\lambda})} T(\lambda) = \max_{\lambda \in \mathcal{M}(\underline{\lambda}, \bar{\lambda})} T(\lambda). \tag{21}$$

Proof: by contradiction. Suppose $\max_{\lambda \in \mathcal{L}} T(\lambda) > \max_{\lambda \in \mathcal{M}} T(\lambda)$. Then there exists $\lambda^* \in \mathrm{argmax}_{\lambda \in \mathcal{L}} T(\lambda)$ with an index, say k, with nonextremal value $\lambda^*_k \in (\underline{\lambda}_k, \bar{\lambda}_k)$. Define

$$\lambda_i^{*,+} = \left\{ \begin{array}{ll} \lambda^*_i, & i \neq k \\ \bar{\lambda}_k, & i = k \end{array} \right.,\ \ \lambda_i^{*,-} = \left\{ \begin{array}{ll} \lambda^*_i, & i \neq k \\ \underline{\lambda}_k & i = k \end{array} \right. \tag{22}$$

Let $\lambda^{*,+}, \lambda^{*,-}$ denote the corresponding vectors. We will show that either $T(\lambda^{*,+})$ or $T(\lambda^{*,-})$ exceeds $T(\lambda^*)$, contradicting the assumed optimality of λ^*. Define $\Delta^+ \equiv T(\lambda^{*,+}) - T(\lambda)$ and $\Delta^- \equiv T(\lambda^{*,-}) - T(\lambda^*)$. Then:

$$\Delta \equiv \max(T(\lambda^{*,+}), T(\lambda^{*,-})) - T(\lambda^*) \equiv \max(\Delta^+, \Delta^-) \tag{23}$$

Leveraging (19), (20), algebraic manipulation (omitted) yields

$$\begin{aligned} \Delta^+ &= \frac{(1-\Lambda(\lambda^*_{\setminus k}))(\bar{\lambda}_k - \lambda^*_k)}{\Pi(\lambda^*_{\setminus k})(1+\bar{\lambda}_k)(1+\lambda^*_k)} \\ \Delta^- &= \frac{(\Lambda(\lambda^*_{\setminus k})-1)(\lambda^*_k - \underline{\lambda}_k)}{\Pi(\lambda^*_{\setminus k})(1+\underline{\lambda}_k)(1+\lambda^*_k)} \end{aligned} \tag{24}$$

where $\lambda^*_{\setminus k} \equiv \lambda^*([n] \setminus k)$. If $\Lambda(\lambda^*_{\setminus k}) < 1$ then $\Delta^+ > 0$, while if $\Lambda(\lambda^*_{\setminus k}) > 1$ then $\Delta^- > 0$. In conclusion, $\Delta > 0$. ■

A second motivation also comes from [1], which considered the problem of designing incentives for (users of) radios to contend for access on the random access erasure collision channel. In this scenario, each transmission *attempt* (whether successful or not) incurs a cost $c > 0$ to the user, while each transmisison *success* earns the user a reward $\rho \geq c$. Users are assumed to possess a quasilinear utility function that captures the expected net utility (reward minus cost):

$$u_i(p_i) \equiv \rho r_i \prod_{j \neq i} \bar{r}_j - c p_i = p_i \rho \left(q_i \prod_{j \neq i} \bar{r}_j - \gamma \right), \tag{25}$$

where $\gamma \equiv \frac{c}{\rho}$ is the cost to reward ratio. The right side makes clear that the utility is in fact linear in p_i, and therefore, fixing the other users, user i will either elect $p_i \in \{0,1\}$, i.e., to "join" the channel and contend in each slot ($p_i = 1$) or to "leave" the channel and not contend at all ($p_i = 0$), depending upon the sign of $q_i \prod_{j \neq i} \bar{r}_j - \gamma$. A set A is then naturally defined as a *Nash equilibrium* if $u_i(1) > 0$ for $i \in A$ and $u_i(1) < 0$ for $i \notin A$, i.e., all users who have joined (left) receive positive (negative) net utility. The objective is to properly select γ so as to induce an equilibrium with near-optimal throughput. Using the price of anarchy (PoA) framework, the key result in [1] is that the PoA is upper bounded by two, i.e., the sum throughput induced by participation incentives is at worst half of the sum throughput achievable by selecting (by fiat) the optimal subset of users. This bound is conservative, however, motivating our interest in how $T(A)$ varies with $A \subseteq [n]$, the focus of this paper.

B. Definitions

Recall the channel loading definition (Def. 6). Denote the set of all underloaded (overloaded) sets as $\mathcal{A}_\mathrm{u}$ ($\mathcal{A}_\mathrm{o}$), respectively:

$$\begin{aligned} \mathcal{A}_\mathrm{u} &\equiv \{A \subseteq [n] : \Lambda(\lambda(A)) < 1\} \\ \mathcal{A}_\mathrm{o} &\equiv \{A \subseteq [n] : \Lambda(\lambda(A)) \geq 1\}. \end{aligned} \quad (26)$$

We further categorize the overloaded sets:

Definition 8 (Overloaded sets): Fix λ and let $A \in \mathcal{A}_\mathrm{o}$ be an overloaded set. Then A is in one of three categories:

- A is *critically-overloaded (oc)* if the channel *becomes underloaded* upon removing *any* radio from A:
$$\Lambda(\lambda(A \setminus i)) < 1 \leq \Lambda(\lambda(A)), \ \forall i \in A. \quad (27)$$
Let $\mathcal{A}_\mathrm{oc}$ denote the set of all critically over loaded sets.
- A is *pure-overloaded (op)* if the channel *stays overloaded* upon removing *any* radio from A:
$$\Lambda(\lambda(A \setminus i)) \geq 1, \ \forall i \in A. \quad (28)$$
Let $\mathcal{A}_\mathrm{op}$ denote the set of all pure overloaded sets.
- A is *impure-overloaded (oi)* if i) there exists one or more radios, say $i \in A$, such that the channel *becomes underloaded* upon removing i, and ii) there exists one or more radios, say $j \in A$, such that the channel *stays overloaded* upon removing j:
$$\exists i, j \in A : \Lambda(\lambda(A \setminus i)) < 1 \leq \Lambda(\lambda(A \setminus j)). \quad (29)$$
Let $\mathcal{A}_\mathrm{oi}$ denote the set of all impure overloaded sets.

Definition 9 (Critical cardinalities): The following are defined for overloaded λ, i.e., $\Lambda(\lambda) \geq 1$. Recall Def. 2.

- $\tilde{k}(\lambda)$ is the cardinality of the smallest overloaded forward-packed set: $\tilde{k} \equiv \min\{k : \Lambda(F_k) \geq 1\}$. A little thought shows this set is critically-overloaded: $F_{\tilde{k}} \in \mathcal{A}_\mathrm{oc}$.
- $\hat{k}(\lambda)$ is the cardinality of the smallest overloaded reverse-packed set: $\hat{k} \equiv \min\{k : \Lambda(R_k) \geq 1\}$. A little thought shows this set may be either critically- or impurely-overloaded: $R_{\hat{k}} \in \mathcal{A}_\mathrm{oc} \cup \mathcal{A}_\mathrm{oi}$. The set $R_{\hat{k}+1}$ is always pure overloaded: $R_{\hat{k}+1} \in \mathcal{A}_\mathrm{op}$.
- $k'(\lambda, \sigma)$, where $\lambda(\sigma)$ is the list of loads λ reordered under permutation σ, is the smallest value of k such that the first k components of the reordered loads are purely overloaded:
$$k'(\lambda, \sigma) \equiv \min\left\{k : \Lambda((\lambda_{\sigma_1}, \ldots, \lambda_{\sigma_k})) - \max_{i \in [k]} \lambda_{\sigma_i} \geq 1\right\}. \quad (30)$$
Using Fact 1, it can be shown that $\tilde{k}(\lambda) \leq k'(\lambda, \sigma) \leq \hat{k}(\lambda)$.

C. Throughput sequence properties

Let Σ be the set of all permutations of $[n]$ and $\sigma \equiv (\sigma_1, \sigma_2, \ldots, \sigma_n)$ be any permutation in the set of permutations, denoted Σ, with cardinality $|\Sigma| = n!$. Let $\lambda(\sigma) = (\lambda_{\sigma_1}, \lambda_{\sigma_2}, \ldots, \lambda_{\sigma_n})$ be the reordered offered loads associated with this permutation, and let $\lambda^i(\sigma) = (\lambda_{\sigma_1}, \lambda_{\sigma_2}, \ldots, \lambda_{\sigma_i})$ denote the first i elements of $\lambda(\sigma)$. Finally, let $T(\lambda^i(\sigma))$ denote the throughput associated with the loads $\lambda^i(\sigma)$, and define the *throughput sequence*

$$T(\lambda(\sigma)) \equiv (T(\lambda^i(\sigma)), \ i \in [n]). \quad (31)$$

In words, every throughput sequence gives n throughputs, where the throughput in position i is that associated with the first i offered loads of λ under permutation σ, i.e., $T(\lambda^i(\sigma))$.

Recall that a sequence $a = (a_1, \ldots, a_n)$ is *unimodal*[1] if a obeys the pattern below, for some $k \in [n]$:

$$a_1 \leq \cdots \leq a_k \geq \cdots \geq a_n \quad (32)$$

We call a_k as the mode of the sequence and k as the modal index. Note a is *monotone* (increasing or decreasing) if $k \in \{1, n\}$. Let $\Delta a(i) \equiv a_i - a_{i-1}$, for $i \in [2:n]$ be the signed change in a at i, and $\Delta a = (\Delta a(2), \ldots, \Delta a(n))$ be the sequence of signed changes. Note a is unimodal if Δa is any of the following:

- all positive (i.e., a is monotone increasing)
- all negative (i.e., a is monotone decreasing)
- positive then negative (i.e., increasing then decreasing)
- negative then positive (i.e., decreasing then increasing)

Proposition 1: The throughput sequence $T(\lambda(\sigma))$ defined in (31) is *unimodal*.

Proof: Fix $\sigma \in \Sigma$ and let λ^i denote $\lambda^i(\sigma)$. Then:

$$\begin{aligned} \Delta T(i) &\equiv T(\lambda^i) - T(\lambda^{i-1}) \\ &= \frac{\lambda_{\sigma_i}(1 - \Lambda(\lambda^{i-1}))}{\Pi(\lambda^i)} \end{aligned} \quad (33)$$

For a given (λ, σ), one of the following three cases holds true:

1) $\Lambda(\lambda^n) < 1$, i.e., the channel is *underloaded* when all radios are in the channel. Then $\Delta T(i) > 0$ and $T(\lambda(\sigma))$ is an *increasing* sequence.
2) $\Lambda(\lambda^1) \geq 1$, i.e., the channel is *overloaded* with the addition of the first radio of the permutation. Then $\Delta T(i) < 0$ and $T(\lambda(\sigma))$ is a *decreasing* sequence.
3) $\Lambda(\lambda^1) < 1 \leq \Lambda(\lambda^n)$. Then there exists an $i^* \in [2:n-1]$ such that $\Lambda(\lambda^{i^*-1}) < 1$ and $\Lambda(\lambda^{i*}) \geq 1$. This implies $\Delta T(i) \geq 0 \ \forall i \in [i^*]$ but $\Delta T(i) \leq 0 \ \forall i \in [i^* + 1 : n]$, which in turn implies $T(\lambda(\sigma))$ is an increasing then decreasing sequence.

■

Prop. 1 suggests that every sequence $T(\lambda(\sigma))$ has a modal set, which we shall denote as $M_{k(\sigma)}(\sigma)$ with cardinality $k(\sigma)$. (33) also suggests that the mode is achieved by a modal set that is unique for each permutation except for the trivial cases

[1]Although $a_1 \geq \cdots \geq a_k \leq \cdots \leq a_n$ is also unimodal, it is not pertinent to our concern here.

of transmitters with zero load. Denote the collection of all modal sets as

$$\mathcal{M} \equiv \{M_{k(\sigma)}(\sigma) : \sigma \in \Sigma\}. \quad (34)$$

The unimodality of the throughput sequence is a key proposition used to prove the main result in Thm. 4.

D. Results on throughput over subsets

The main results on throughput over subsets under given offered loads λ are: $i)$ characterize the subset of $[n]$ that maximizes throughput over all subsets of $[n]$, and $ii)$ establish bounds on the throughputs achieved by all the k-sets, i.e., bounds on $\{T(A_k) : A \subseteq [n], |A_k| = k\}$ as a function of k.

Theorem 4: The set $F_{\tilde{k}}$ (c.f. Def. 9), is the unique *forward-packed* critically-overloaded set. The cardinality of $F_{\tilde{k}}$, i.e., $\tilde{k}$, lower bounds that of any overloaded set:

$$\tilde{k} \leq \min_{A \in \mathcal{A}_o} |A|. \quad (35)$$

The set $F_{\tilde{k}}$ maximizes throughput over all sets ($A \subseteq [n]$):

$$T(F_{\tilde{k}}) = \max_{A \subseteq [n]} T(A) \equiv T^*. \quad (36)$$

Proof: we prove the claims regarding uniqueness, cardinality, and throughput in turn.

Uniqueness. Forward-packed sets F_k, for $k < \tilde{k}$, are underloaded, by definition of $\tilde{k}$, while forward-packed sets F_k, for $k > \tilde{k}$, are mixed- or pure-overloaded, since,

$$\min_{i \in [k]} \Lambda(F_k \setminus i) = \Lambda(F_{k-1}) \geq \Lambda(F_{\tilde{k}}) \geq 1. \quad (37)$$

Cardinality. Let $B \in \mathcal{A}_o$ be any overloaded set and let k denote $|B|$. We establish $\tilde{k} \leq k$ by contradiction: suppose $k < \tilde{k}$. Observe $i)$ the packed set F_k is underloaded, by definition of $\tilde{k}$, $ii)$ $F_k \leq_{\text{po}} B$ (c.f. Def. 1), and therefore $iii)$ $\Lambda(B) \leq \Lambda(F_k) < 1$ (Fact 1). Thus B is underloaded, in contradiction of the assumption it is overloaded.

Throughput. By the unimodality property established in Prop. 1, for every permutation σ, there exists a modal set that maximises the throughput for that permutation. Let the modal set for permutation σ have a cardinality $k(\sigma)$ and be denoted as $M_{k(\sigma)}(\sigma)$. In seeking to maximize throughput $T(A)$ over all subsets $A \subseteq [n]$, it suffices to restrict attention to the modal sets $\mathcal{M}$ defined in (34), i.e.,

$$\max_{A \subseteq [n]} T(A) = \max_{A \in \mathcal{M}} T(A), \quad (38)$$

as any non-modal set has, by construction, a throughput less than that of one or more modal sets.

Case i): $\Lambda([n]) \leq 1$. If $\Lambda([n]) \leq 1$, then all sets of transmitters are underloaded and by Prop. 1, for any permutation σ, the throughput sequences are strictly increasing and each attains its maximum values at $\tilde{k} = k(\sigma) = n$. Hence $\lambda(F_{\tilde{k}}) \equiv \lambda(M_{k(\sigma)}(\sigma)) \equiv \lambda([n])$ for each σ. Thereby,

$$T(F_{\tilde{k}}) = T(M_n(\sigma)), \forall \sigma \in \Sigma. \quad (39)$$

Case ii): $\Lambda([n]) > 1$. When the sum offered load exceeds 1, the following template is used to prove that the forward packed critically overloaded set $F_{\tilde{k}}$ has the maximum throughput.

- Create an intermediate list of transmitters F'_k by adding or reducing the load of $F_{\tilde{k}}$ to be equal to that of a modal set and then adding zeros to make the length equal to that of the modal set. We may now establish an ordering of majorization between $\lambda(F'_k)$ and the modal set.
- Show that $F_{\tilde{k}}$ has a higher throughput than F'_k by unimodality of throughput sequence and F'_k has a higher throughput than the modal set by Schur-convexity.

Fix some $\sigma \in \Sigma$ and simplify notation by suppressing the dependence upon σ, e.g., $M_k(\sigma) \equiv M_k$. Define x, y as:

$$\Lambda(M_k) \equiv 1 + y, \quad \Lambda(F_{\tilde{k}}) \equiv 1 + x. \quad (40)$$

Case ii-a) : $x = y$. Since $\tilde{k} \leq k$, construct a list $\lambda(F'_k) = (\lambda(F_{\tilde{k}}), 0, \ldots, 0)$ of length k. Since the throughput does not change by addition of zero-value components, $T(F'_k) = T(F_{\tilde{k}})$. Also, $\lambda(M_k) \prec \lambda(F'_k)$ and by the Schur convexity property of the throughput function, $T(M_k) \leq T(F'_k) = T(F_{\tilde{k}})$.

Case ii-b) : $x < y$. Construct $\lambda(F'_k) = (\lambda(F_{\tilde{k}}), y - x, 0, \ldots, 0)$, i.e., add one non-zero component with rate $y - x$ to reach the target sum load and then append the necessary number of zero-value components to reach the target vector length k. Then $T(F'_k) \leq T(F_{\tilde{k}})$, because $F_{\tilde{k}}$ is the modal set of the forward packed sequence and is overloaded. Addition of any other radio to an overloaded set will cause the throughput to decrease (Cor. 2). Moreover, $\Lambda(M_k) = \Lambda(F'_k)$ and $\lambda(M_k) \prec \lambda(F'_k)$ and by Schur convexity, $T(M_k) \leq T(F'_k)$. Combining the above inequalities, $T(M_k) \leq T(F_{\tilde{k}})$.

Case ii-c) : $x > y$. Construct $\lambda(F'_{\tilde{k}}) = (\lambda(F_{\tilde{k}-1}), \lambda_{\tilde{k}} - x + y)$ and $\lambda(F'_k) = (\lambda(F'_{\tilde{k}}), 0, \ldots, 0)$. That is, add one non-zero component so that the two lists have the same sum, and append any required zero-value components so that the two lists have the same length. Although the lists have the same sum load, i.e., $\Lambda(M_k) = \Lambda(F'_{\tilde{k}}) = \Lambda(F'_k)$, because $\lambda_{\tilde{k}} > \lambda_{\tilde{k}} - x + y$ and $\Lambda(F_{\tilde{k}-1}) < 1$, it follows that $T(F'_{\tilde{k}}) \leq T(F_{\tilde{k}})$ (Cor. 2). Moreover, $\lambda(M_k) \prec \lambda(F'_k)$ and, by Schur convexity, $T(M_k) \leq T(F'_k) = T(F'_{\tilde{k}})$. Combining the above inequalities, $T(M_k) \leq T(F_{\tilde{k}})$.

Thus, the forward-packed critically overloaded set has a higher throughput than that of any modal set, and hence maximizes throughput over all subsets of $[n]$. ∎

Theorem 5: For offered loads λ and $k \in [n]$, the throughputs of all k-sets i.e., $\{T(A) : A \subseteq [n], |A| = k\}$ are bounded as:

1) $T(R_k) \leq T(A_k) \leq T(F_k), \ k \leq \tilde{k}$
2) $T(F_k) \leq T(A_k) \leq T(R_k), \ k \geq \hat{k}$

Proof: For $k \leq \tilde{k}$, all three of R_k, A_k, F_k are underloaded, i.e., $R_k, A_k, F_k \in \mathcal{A}_u$. Hence by Fact 1 and Lem. 4, $T(R_k) \leq T(A_k) \leq T(F_k)$. For $k \geq \hat{k}$, all three of R_k, A_k, F_k are pure overloaded, i.e., $R_k, A_k, F_k \in \mathcal{A}_{\text{op}}$, and hence by Fact 1 and Cor. 3, $T(F_k) \leq T(A_k) \leq T(R_k)$. ∎

Recall that Thm. 5 gives lower and upper bounds on the throughput, in terms of the offered loads λ, as a function of the cardinality k, in the regimes $k \leq \tilde{k}(\lambda)$ and $k \geq \hat{k}(\lambda)$ (Def. 9). By contrast, Cor. 1 gives bounds for a particular throughput sequence $T(\lambda(\sigma))$ (31), corresponding to permutation σ of

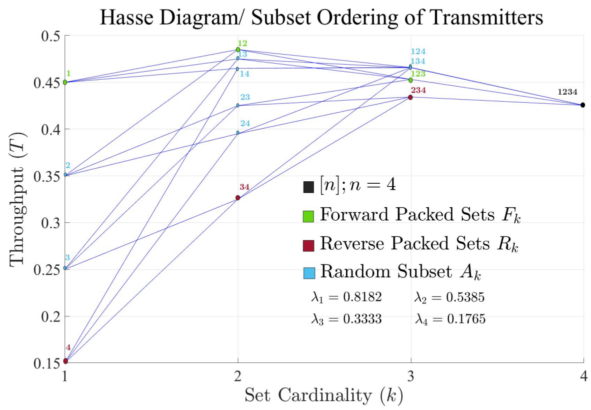

Fig. 2. Sum throughput $T(A)$ (y-axis) vs. subset cardinality k (x-axis) for various subsets $A \subseteq [n]$, for $n = 4$ and $\lambda \approx (0.818, 0.539, 0.333, 0.177)$. The forward-packed set F_k and reverse-packed set R_k for each k are shown in green and red, respectively.

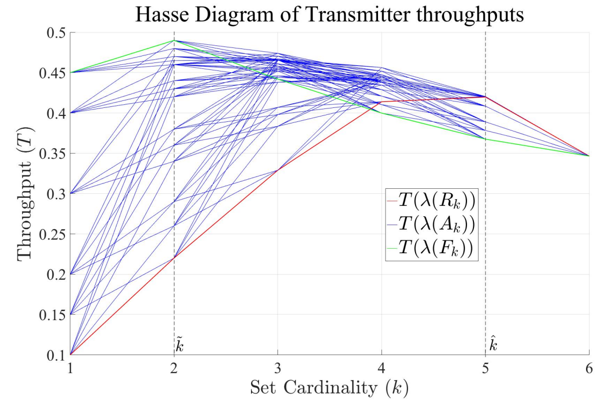

Fig. 3. Sum throughput $T(A)$ (y-axis) vs. subset cardinality k (x-axis) for various subsets $A \subseteq [n]$, for $n = 6$. The critical cardinalities $\tilde{k}(\lambda)$ and $\hat{k}(\lambda)$ (Def. 9) and the throughput bounds (Thm. 5) provided by the forward- and reverse-packed sets, i.e., $T(F_k)$ and $T(R_k)$, for $k \leq \tilde{k}$ and $k \geq \hat{k}$, are evident.

$[n]$, for the intermediate regime $k \in [\tilde{k} : \hat{k}]$, using the index $k'(\lambda, \sigma)$ (Def. 9). The proof follows mutatis mutandis from that of Thm. 5.

Corollary 1: For a given set of radios with offered loads λ, and permutation σ of $[n]$, the throughput sequence $T(\lambda(\sigma))$ (31) is lower bounded as follows:

1) $T(R_k) \leq T(A_k),\ k \in [\tilde{k}(\lambda) : k'(\lambda, \sigma)]$
2) $T(F_k) \leq T(A_k) \leq T(R_k),\ k \in [k'(\lambda, \sigma) : \hat{k}(\lambda)]$

VII. Numerical results

Numerical results are presented for $n \in \{4, 6, 50, 100\}$ in Figures 2, 3, 4, and 5, respectively. The smaller values of n, i.e., $n \in \{4, 6\}$, are such that the throughput of each possible subset of $[n]$ may be seen. For larger n, i.e., $n \in \{50, 100\}$, the trends in both the bounds and the "typical" throughput sequence are evident. Several points merit mention.

First, the critical indices $\tilde{k}(\lambda)$ and $\hat{k}(\lambda)$ (Def. 9) are shown on several of the plots. Recall from Thm. 5 that the forward- and reverse-packed sets provide throughput bounds, $T(F_k), T(R_k)$ vs. k, for both $k \leq \tilde{k}(\lambda)$ and $k \geq \hat{k}(\lambda)$.

Second, the numerical results demonstrate that the forward-packed subset $F_{\tilde{k}}$ from Thm. 4 is in fact the throughput-optimal subset. Yet, also evident from several of the plots is the fact that near-optimal throughput can be obtained by using significantly larger values of k than $\tilde{k}$. From a throughput–fairness tradeoff perspective, sacrificing a small throughput in order to enable participation by a much larger cardinality subset of radios may be worthwhile.

Finally, the 100 randomly selected throughput sequences (§VI-C) in Fig. 5 demonstrate that, although the forward-packed and reverse-packed subset throughput bounds are tight from the perspective of all $n!$ possible throughput trajectories, they are quite loose from a "typical" trajectory perspective. This suggests that it would be useful to characterize the expectation and variance of the random throughput of a randomly selected k-set, as a function of k, and use these to derive *statistical* throughput bounds.

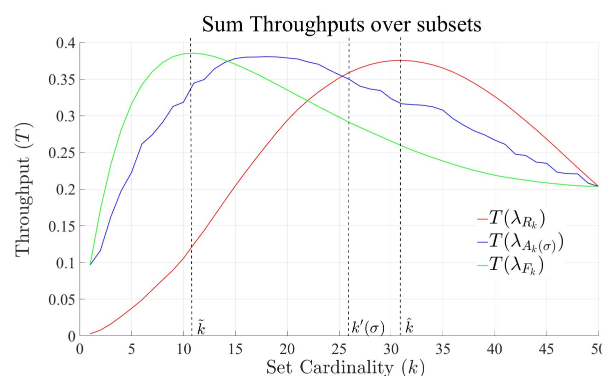

Fig. 4. Sum throughput $T(A)$ (y-axis) vs. subset cardinality k (x-axis) for various subsets of a single random permutation $\sigma \in \Sigma$ and $A \equiv \sigma([n])$, for $n = 50$. The critical cardinalities $\tilde{k}(\lambda)$ and $\hat{k}(\lambda)$ and k' (Def. 9) and the throughput bounds (Thm. 5) provided by the forward- and reverse-packed sets, i.e., $T(F_k)$ and $T(R_k)$, for $k \leq \tilde{k}$ and $k \geq k'$, are evident.

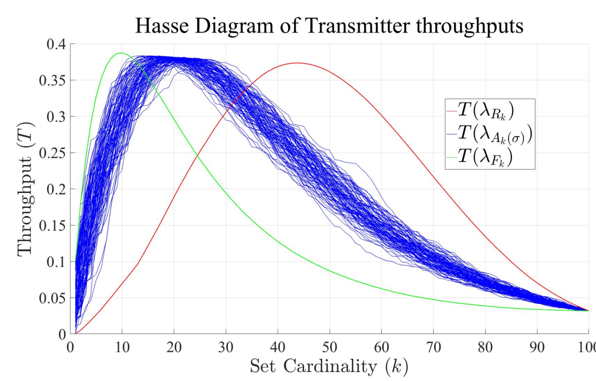

Fig. 5. Sum throughput $T(A)$ (y-axis) vs. subset cardinality k (x-axis) for various subsets $A \subseteq [n]$, for $n = 100$. Shown are the throughput sequences (§VI-C) for 100 randomly selected permutations σ, i.e., $T(\lambda(\sigma))$.

References

[1] V. Gkatzelis and S. Weber, "Participation incentives on a wireless random access erasure collision channel," in *NetEcon Workshop*, Cambridge, MA, June 2017.

[2] B. Yang, G. Feng, and X. Guan, "Random access in wireless ad hoc networks for throughput maximization," in *IEEE ICARCV*, Singapore, December 2006.

[3] R. Ma, V. Misra, and D. Rubenstein, "An analysis of generalized slotted-Aloha protocols," *IEEE/ACM Transactions on Networking*, vol. 17, no. 3, pp. 936–949, March 2009.

[4] G. Miao, Y. Li, and A. Swami, "Channel-aware distributed medium access control," *IEEE/ACM Transactions on Networking*, vol. 20, no. 4, pp. 1290–1303, April 2012.

[5] H. Dhillon, H. Huang, H. Viswanathan, and R. Valenzuela, "Fundamentals of throughput maximization with random arrivals for M2M communications," *IEEE Transactions on Communications*, vol. 62, no. 11, pp. 4094–4109, November 2014.

[6] Y. Li and L. Dai, "Maximum sum rate of slotted Aloha with capture," *IEEE Transactions on Communications*, vol. 64, no. 2, pp. 690–705, February 2016.

[7] L. Barletta, F. Borgonovo, and I. Filippini, "The S-Aloha capacity: Beyond the e^{-1} myth," in *IEEE International Conference on Computer Communications (INFOCOM)*, San Francisco, CA, April 2016.

[8] D. Malak, H. Huang, and J. Andrews, "Throughput maximization for delay-sensitive random access communication," *IEEE Transactions on Wireless Communications*, vol. 18, no. 1, pp. 709–723, January 2019.

[9] N. Xie and S. Weber, "On the Aloha throughput-fairness tradeoff," *IEEE Transactions on Information Theory*, vol. 63, no. 9, pp. 6085–6112, September 2017.

[10] A. Marshall, I. Olkin, and B. Arnold, *Inequalities: Theory of Majorization and Its Applications, 2nd Edition*. Springer, 2011.

APPENDIX

Let A, B, C be subsets of $[n]$, with $A \cap B = A \cap C = \emptyset$, and define $\Delta_T(A,B,C) \equiv T(A \cup B) - T(A \cup C)$; we interpret $\Delta_T(A,B,C)$ as the change in throughput when modifying the set of participating radios from $A \cup C$ to $A \cup B$, where $\Delta_T(A,B,C) > 0$ indicates this change is advantageous. Observe *any* pair of sets D, E may be decomposed into A, B, C as above with $A = D \cap E$, $B = D \setminus A$, and $C = E \setminus A$.

Lemma 2: $\Delta_T(A,B,C)$, defined above, equals

$$\frac{1}{\Pi(A)}\left[\Lambda(A)\left(\frac{1}{\Pi(B)} - \frac{1}{\Pi(C)}\right) + T(B) - T(C)\right]. \quad (41)$$

Proof: $\Delta_T(A,B,C)$

$$\begin{aligned} &= \frac{\Lambda(A \cup B)}{\Pi(A \cup B)} - \frac{\Lambda(A \cup C)}{\Pi(A \cup C)} \\ &= \frac{\Lambda(A) + \Lambda(B)}{\Pi(A)\Pi(B)} - \frac{\Lambda(A) + \Lambda(C)}{\Pi(A)\Pi(C)} \end{aligned} \quad (42)$$

■

The following results are immediate corollaries of Lem. 2.

Corollary 2: The condition for $\Delta_T(\emptyset, B, C) \geq 0$ in the cases below are:

#	A	B	C	$A \cup B$	$A \cup C$	$\Delta_T(\emptyset, B, C) \geq 0$
1)	A	i	$\emptyset$	$A \cup i$	A	$\Lambda(A) \leq 1$
2)	$A \setminus i$	$\emptyset$	i	$A \setminus i$	A	$\Lambda(A \setminus i) \geq 1$

(43)

The interpretation of these results is immediate:

1) Adding i to A improves throughput if A is underloaded; adding i to A degrades throughput if A is overloaded;
2) Removing i from A improves throughput if $A \setminus i$ is overloaded; removing i from A degrades throughput if $A \setminus i$ is underloaded.
3) The magnitude of the change in throughput is proportional to the value of the load added or removed.

Proof: Specializing (41) in Lem. 2 to the cases above:

$$\begin{aligned} \Delta_T(A, i, \emptyset) &= \frac{1}{\Pi(A)}\left[\frac{(1-\Lambda(A))\lambda_i}{(1+\lambda_i)}\right] \\ \Delta_T(A \setminus i, \emptyset, i) &= \frac{1}{\Pi(A \setminus i)}\left[\frac{\lambda_i(\Lambda(A \setminus i) - 1)}{1+\lambda_i}\right] \end{aligned} \quad (44)$$

We can see that Δ_T is proportional to $\frac{\lambda_i}{1+\lambda_i}$ which is an increasing function of λ_i. ■

Fix λ and $A \subseteq [n]$. The throughput gradient is denoted $\nabla T(\lambda(A)) \equiv \left(\frac{\partial}{\partial \lambda_i} T(\lambda(A)), i \in A\right)$.

Lemma 3: The throughput gradient components are:

$$\frac{\partial}{\partial \lambda_k} T(\lambda(A)) = \frac{1 - \Lambda(\lambda(A \setminus k))}{(1+\lambda_k)\Pi(\lambda(A))}, \ k \in A. \quad (45)$$

Proof:

$$\begin{aligned} \frac{\partial}{\partial \lambda_k} T(\lambda(A)) &= \frac{\partial}{\partial \lambda_k} \frac{\Lambda(\lambda(A))}{\Pi(\lambda(A))} \\ &= \frac{\Pi(\lambda(A)) \cdot 1 - \Lambda(\lambda(A))\Pi(\lambda(A \setminus k))}{\Pi(\lambda(A))^2} \\ &= \frac{1 + \lambda_k - \Lambda(\lambda(A))}{(1+\lambda_k)\Pi(\lambda(A))} \end{aligned} \quad (46)$$

■

Let A, B be subsets of $[n]$, with common cardinality denoted by $k = |A| = |B|$, and recall $\lambda(A), \lambda(B)$ are subvectors of λ, indexed by A, B, respectively. Define the *direction* vector $\delta^{(\lambda)}(A,B) \equiv (\delta_i^{(\lambda)}(A,B), i \in [k])$, with $\delta_i^{(\lambda)}(A,B) \equiv \lambda(B)_i - \lambda(A)_i$, for $i \in [k]$. The throughput *directional derivative* in direction $\delta^{(\lambda)}(A,B)$ is defined as $\nabla_{\delta^{(\lambda)}(A,B)} T(\lambda(A)) \equiv \nabla T(\lambda(A)) \cdot \delta^{(\lambda)}(A,B)$.

Lemma 4: If A, B are subsets of equal cardinality ($|A| = |B|$) obeying product order $A \leq_{\text{po}} B$, and A is underloaded or critically overloaded ($A \in \mathcal{A}_{\text{u}} \cup \mathcal{A}_{\text{oc}}$), then $T(A) \geq T(B)$.

Proof: Write $\delta = \delta^{(\lambda)}(A,B)$, set $k = |A| = |B|$, let $A = (a_1, \ldots, a_k)$, and $B = (b_1, \ldots, b_k)$. Thus $\lambda(B)_i = \lambda_{b_i}$ and $\lambda(A)_i = \lambda_{a_i}$, and so, using the results of Lem. 3,

$$\begin{aligned} \nabla_\delta T(\lambda(A)) &\equiv \nabla T(\lambda(A)) \cdot \delta^{(\lambda)}(A,B) \\ &= \sum_{i \in [k]} \frac{\partial}{\partial \lambda_{a_i}} T(\lambda(A))(\lambda_{b_i} - \lambda_{a_i}) \\ &= \frac{1}{\Pi(\lambda(A))} \sum_{i \in [k]} \frac{1 - \Lambda(\lambda(A \setminus a_i))}{1 + \lambda_{a_i}} (\lambda_{b_i} - \lambda_{a_i}) \end{aligned} \quad (47)$$

By assumption, $A \in \mathcal{A}_{\text{u}} \cup \mathcal{A}_{\text{oc}}$, and as such $\Lambda(\lambda(A \setminus a_i)) < 1$ for each $i \in [k]$. Moreover, also by assumption, $A \leq_{\text{po}} B$, and as such $\lambda_{b_i} \leq \lambda_{a_i}$ for each $i \in [k]$. It follows that $\nabla_\delta T(\lambda(A)) < 0$, and thus $T(A) \geq T(B)$. ■

The following result is an immediate corollary of Lem. 4.

Corollary 3: If A, B are subsets of equal cardinality ($|A| = |B|$) obeying product order $A \leq_{\text{po}} B$, and A is pure overloaded ($A \in \mathcal{A}_{\text{op}}$), then $T(A) \leq T(B)$.

Proof: Consider (47) with δ the direction from $\lambda(A)$ to $\lambda(B)$. As $A \in \mathcal{A}_{\text{op}}$, thus $\Lambda(\lambda(A \setminus a_i)) > 1$ for each $i \in [k]$. It follows that $\nabla_\delta T(\lambda(A)) > 0$, and thus $T(A) \leq T(B)$. ■