Channel Capacity for Adversaries with Computationally Bounded Observations

Eric Ruzomberka, Chih-Chun Wang and David J. Love
School of Electrical and Computer Engineering, Purdue University, West Lafayette, USA
email: {eruzombe,chihw,djlove}@purdue.edu

***Abstract*—We study reliable communication over point-to-point adversarial channels in which the adversary can observe the transmitted codeword via some function that takes the n-bit codeword as input and computes an rn-bit output for some given $r \in [0,1]$. We consider the scenario where the rn-bit observation is *computationally bounded* – the adversary is free to choose an arbitrary observation function as long as the function can be computed using a polynomial amount of computational resources. This observation-based restriction differs from conventional channel-based computational limitations, where in the later case, the resource limitation applies to the computation of the (adversarial) channel error. For all $r \in [0, 1 - H(p)]$ where $H(\cdot)$ is the binary entropy function and p is the adversary's error budget, we characterize the capacity of the above channel. For this range of r, we find that the capacity is identical to the completely obvious setting ($r = 0$). This result can be viewed as a generalization of known results on myopic adversaries and channels with active eavesdroppers for which the observation process depends on a fixed distribution and fixed-linear structure, respectively, that cannot be chosen arbitrarily by the adversary.**

I. Introduction

Beginning with Shannon's seminal paper [1], early channel coding research observed that fundamental coding limits are highly sensitive to channel modeling assumptions. This sensitivity is demonstrated by a gap in capacity between the two classical models: the *Shannon model* in which channel errors follow a known random distribution and the *Hamming model* in which error patterns are worst-case for some fixed number of bit errors. In the design of robust codes, the more conservative Hamming model is particularly attractive as it makes no assumptions about the channel distribution and thus any resulting conclusion is *robust* against a wide variety of channel imperfections. The downside of the Hamming model, however, is that it admits a smaller capacity than the Shannon model. In many cases, the gap in capacity is large [2].

A. Closing the gap

Recent research efforts have made progress in closing this gap by considering settings in between the two classical models. Ideally, the following two properties hold for a good channel model:

> **Property 1:** The channel is *mild* in the sense that its capacity coincides with the Shannon model capacity.
>
> **Property 2:** The channel inherits conservative aspects of the Hamming model. In particular, the channel may vary in an arbitrary manner unknown to the communicating parties.

In the following Section I-B, we focus on two different approaches which have had success towards producing good channel models. These approaches have been to 1) bound the channel's computing power (i.e., computational complexity) [3], [4] and 2) bound the information known to the channel about the communication scheme [5]–[13].

B. Complexity Bounded Channels and Oblivious Channels

Consider a transmitter Alice who wishes to communicate a message m from a set of M possible messages over a noisy channel to a receiver Bob. To protect the message from noise corruption, Alice encodes m into an n bit codeword $\boldsymbol{x}$ of rate $R = (1/n)\log M$ and transmits $\boldsymbol{x}$ over the channel. The channel adds an n-bit error vector $\boldsymbol{e}$ to $\boldsymbol{x}$, and Bob receives the binary channel output $\boldsymbol{y} = \boldsymbol{x} \oplus \boldsymbol{e}$. The channel is controlled by an *adversary* who chooses $\boldsymbol{e}$ to prevent reliable (unique) decoding by Bob. For an error budget $p \in (0, 1/2)$, the adversary can only induce pn bit flips, i.e., the Hamming weight of $\boldsymbol{e}$ must be bounded above by pn.

In the *computationally bounded model* (first proposed by Lipton [3]), the adversary computes $\boldsymbol{e}$ using limited computational resources, e.g., via an algorithm that takes a finite number of computational steps. This model has the appeal of sufficiently describing practical channels, including channels with memory and channels governed by natural, but unknown processes. However, the computationally bounded model can be *severe* – an impossibility result of Guruswami and Smith [4] is that the model's capacity can be less than the Shannon capacity, and can even be 0 when the latter is positive. Thus, Property 1 does not hold for the computationally bounded model.

Another existing approach is the *partially oblivious model*, where the adversary chooses $\boldsymbol{e}$ using incomplete side-information about the transmitted codeword $\boldsymbol{x}$. This model includes myopic channels, e.g., [5]–[7], causal channels, e.g., [8]–[10], channels with active eavesdroppers, e.g., [14], and some arbitrarily varying channels (AVCs), e.g., [11], [12]. Although the model can vary between works, the setting usually has the following general structure: for $r \in [0,1]$ and some (deterministic) observation function $f_n : \{0,1\}^n \to \{0,1\}^{rn}$, the adversary makes an rn-bit observation $f_n(\boldsymbol{x})$ of codeword

This work was supported in part by the Office of Naval Research under ONR Grant N00014-21-1-2472, by NSF Grants CCF-1618475, CCF-1816013, CCF-2008527, CNS-2107363, and also by National Spectrum Consortium (NSC) under grant W15QKN-15-9-1004.

$\boldsymbol{x}$ prior to choosing $\boldsymbol{e}$.[1] The special cases $r = 0$ and $r = 1$ correspond to no information (i.e., completely oblivious) and perfect information (i.e., omniscient), respectively.

Property 1 can hold for the oblivious model when r is sufficiently small.[2] However, Property 2 does not hold for many oblivious channels in the literature. For example, in the myopic channel model, the adversary randomly draws f_n from a known distribution.[3] *For Property 2 to hold, however, we must allow f_n to be arbitrarily chosen and require Alice and Bob to devise their communication scheme without knowledge of f_n.* This is equivalent to the adversary choosing a *worst-case* f_n for a fixed r – a model studied by Langberg [13] under the name of the $(1-r)$-*oblivious channel*. The capacity of the $(1-r)$-oblivious channel remains an open problem, where the best known lower bound is given by [13] and will be summarized in Section I-D.

C. This Work

In this paper, we consider a channel model that has qualities of both the computationally bounded model and the partially oblivious model. We do so by requiring the adversary to observe $\boldsymbol{x}$ via an rn-bit observation function f_n that is computationally bounded.

Specifically, for fixed positive integers c and s, the adversary chooses a sequence of observation functions $f_n(\cdot)$, $\forall n \geq 1$ that belongs to $\mathsf{CPX}(r, cn^s)$ – the set of observation functions with n input bits and rn output bits that can be computed by a Boolean circuit with at most cn^s gates. We allow the choice of f_n to be unknown to Alice or Bob. On the other hand, the f_n chosen by the adversary can depend on the codebook of Alice but cannot depend on the actual message being sent.[4] Using the observation function f_n of its choice, the adversary observes $f_n(\boldsymbol{x})$ and chooses $\boldsymbol{e}$ with no computational bound. We refer to the above adversary as a $\mathsf{CPX}(r, cn^s)$-*oblivious adversary*. By construction, Property 2 holds for a channel controlled by a $\mathsf{CPX}(r, cn^s)$-oblivious adversary due to f_n being unknown to Alice or Bob..

Our imposed computational restriction is practical and sufficiently models realistic adversarial channels. A channel controlled by a $\mathsf{CPX}(r, cn^s)$-oblivious adversary closely approximates a $(1-r)$-*oblivious channel* (i.e., a channel controlled by a $\mathsf{CPX}(r, \infty)$-oblivious adversary) without weakening the power of the adversary too much. Indeed, the adversary is quite strong. To illustrate its strength, if for a sequence of functions $\{f_n\}_{n=1}^{\infty}$ and for $c, s \geq 1$ there exists a finite n_0 such that for all $n \geq n_0$ $f_n \notin \mathsf{CPX}(r, cn^s)$, then the sequence is widely regarded to be an infeasible computation [15]. The technical value of the computational constraint is to bound the number of observation functions that the adversary can choose from.

[1]The error vector depends *non-causally* on the entire observation $f_n(\boldsymbol{x})$ such that the adversary begins choosing $\boldsymbol{e}$ after observing all rn bits.

[2]This fact is an analog to a channel being *sufficiently myopic* (see [6]).

[3]In the myopic channels studied in [5]–[7], the adversary observes $\boldsymbol{x}$ through a discrete memoryless channel (DMC), not through a function f_n. Regardless, the important thing to note is that the adversary's observation does in fact depend on a distribution known to Alice and Bob.

[4]If the adversary's choice of f_n can depend on the actual message, it is as if it knows the entire message, which defeats the purpose of limiting the observation to rn bits.

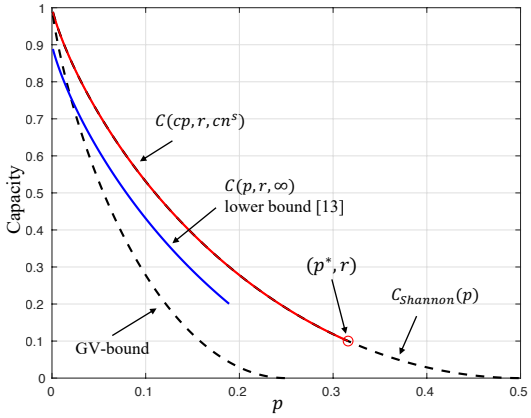

Fig. 1. Capacity when $r = 0.1$. Herein the value p^* satisfies $C_{Shannon}(p^*) = r = 0.1$.

D. Results

We assume that Alice uses deterministic encoding and we consider capacity under the *diminishing average error probability* criterion in which the probability of decoding error is averaged over the message set. Under the above model, the Shannon capacity is $C_{Shannon}(p) = 1 - H(p)$ where $H(p) = -p \log p - (1-p)\log(1-p)$ is the binary entropy function [12], [13]. We remark that $C_{Shannon}(p)$ is achievable in our model if the $\mathsf{CPX}(r, cn^s)$-oblivious adversary ignores its observation $f_n(\boldsymbol{x})$ and naively chooses $\boldsymbol{e}$ randomly from the set all possible error vectors with Hamming weight pn.

For $p \in (0, 1/2)$, $r \in [0, 1]$ and positive integers c, s, let $C(p, r, cn^s)$ denote the capacity of a channel controlled by a $\mathsf{CPX}(r, cn^s)$-oblivious adversary with error budget p. Similarly, let $C(p, r, \infty)$ denote the capacity of $(1-r)$-oblivious channel. The following result shows that Property 1 holds for our model for a wide range of r.

Theorem 1. *For $p \in (0, 1/2)$, $r \in [0, C_{Shannon}(p))$, and positive integers c and s, $C(p, r, cn^s) = C(p, 0, cn^s) = C(p, 0, \infty) = C_{Shannon}(p)$.*

We share a few remarks on the above theorem. When $r < C_{Shannon}(p) = 1 - H(p)$, Theorem 1 implies that the adversary can do no better than to *ignore* its side-information $f_n(\boldsymbol{x})$ and choose $\boldsymbol{e}$ randomly from the set of all n-bit vectors with Hamming weight pn. Additionally, we note that the largest known lower bound on $C(p, r, \infty)$ is $1 - r - H(p)$ for $r \in [0, \frac{1-H(p)}{3})$ [13]. Since $C(p, r, \infty)$ is a lower bound to $C(p, r, cn^s)$, Theorem 1 significantly sharpens the best known lower bound of $C(p, r, cn^s)$ to an exactly tight characterization. For $r > C_{Shannon}(p)$, an immediate lower bound of $C(p, r, cn^s)$ is given by the Gilbert-Varshamov (GV) bound (i.e. $C(p, r, \infty) \geq 1 - H(2p)$) [16], [17]. All results discussed thus far are summarized in Fig. 1.

Theorem 1 generalizes a few known results on myopic channels and channels with active eavesdroppers. For $r < C_{Shannon}$, $C_{Shannon}(p)$ is known to be the capacity of a

myopic channel where the adversary a) non-causally views $\boldsymbol{x}$ through a binary erasure channel with erasure probability $1-r$ (BEC$(1-r)$) then b) injects pn bit errors [6, Theorem III.12]. It is clear that this result is generalized by Theorem 1 after observing that a $\mathsf{CPX}(r, cn^s)$-oblivious adversary can choose f_n *randomly* in a way that simulates a BEC$(1-r)$. Although the results of [6] generalize to a number of channel models, it remains an open problem whether the proof techniques of [6] can be used to develop statements similar to Theorem 1 of this work. Similarly, for $r < C_{Shannon}(p)$, $C_{Shannon}(p)$ is known to be the capacity of a *wiretap channel with an active eavesdropper* where the adversary a) chooses rn indices in $\{1, \ldots, n\}$ and observes rn-bits of $\boldsymbol{x}$ at the chosen indices then b) injects pn bit errors [14, Theorem 4.2]. It is clear that this result is generalized by Theorem 1 after observing that a $\mathsf{CPX}(r, cn^s)$-oblivious adversary can choose $f_n(\boldsymbol{x})$ to output rn-bits of $\boldsymbol{x}$. A detailed proof of Theorem 1 can be found in the extended version of this paper [18].

II. Channel Model

A. Notation

All vectors are in bold notation. Let $d(\boldsymbol{z}, \boldsymbol{z}')$ denote the Hamming distance between two binary vectors $\boldsymbol{z}$ and $\boldsymbol{z}'$. For $t > 0$ and $\boldsymbol{z} \in \{0,1\}^n$ we define $\mathcal{B}_t(\boldsymbol{z}) = \{\boldsymbol{z}' \in \{0,1\}^n : d(\boldsymbol{z}, \boldsymbol{z}') \leq t\}$ to be the Hamming ball of radius t centered around $\boldsymbol{z}$. The functions $\log(\cdot)$ and $\ln(\cdot)$ denote the base 2 and base e logarithm, respectively. For a number $K \geq 1$, let $[K]$ denote the set $\{1, \ldots, \lfloor K \rfloor\}$. For a blocklength n and rate $R \in (0,1]$, a $[n, M]$ codebook $\mathcal{C}_n$ is a function $\mathcal{C}_n : [M] \to \{0,1\}^n$. When useful, we will think of $\mathcal{C}_n = \{\mathcal{C}_n(1), \ldots, \mathcal{C}_n(M)\}$ as a subset of $\{0,1\}^n$ and the ith codeword $\mathcal{C}_n(i)$ is a vector in $\{0,1\}^n$.

B. Channel Model

Alice's Encoding: A transmitter Alice communicates over a noisy channel with a receiver Bob in the following manner. For a rate $R \in (0,1]$ and blocklength n, Alice randomly draws a message m_0 uniformly from a message set $[M] = [2^{Rn}]$. For a $[n, M]$ codebook $\mathcal{C}_n$, Alice encodes m_0 into a codeword $\boldsymbol{x} \in \{0,1\}$ by computing $\boldsymbol{x} = \mathcal{C}_n(m_0)$. Since $\boldsymbol{x} = \mathcal{C}_n(m_0)$ is a deterministic function of m_0, we say that Alice is using a *deterministic encoding*. Following encoding, Alice transmits $\boldsymbol{x}$ into the channel.

Bob's Decoding: At the channel output, Bob receives $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{e}$ where $\boldsymbol{e} \in \{0,1\}^n$ is an error vector added by the channel. For $p \in (0, 1/2)$, we restrict $\boldsymbol{e}$ to have a Hamming weight bounded above by pn, i.e., we restrict $\boldsymbol{e} \in \mathcal{B}_{pn}(\boldsymbol{0})$. Bob performs list decoding by creating a list $\mathcal{L} \subseteq [M]$ of all messages whose corresponding codewords are contained in the ball $\mathcal{B}_{pn}(\boldsymbol{y})$. If $\mathcal{L}$ contains exactly one message, then Bob sets $\hat{m}$ equal to that message. Otherwise, Bob sets $\hat{m}$ equal to an *error symbol* (i.e., some symbol outside the set $[M]$). We say that a decoding error occurs if $\hat{m} \neq m_0$. Note that Bob is using the min-distance decoder.

Adversary: The channel is controlled by an adversary who has side-information about Alice's and Bob's communication scheme but not exact knowledge of the actual message m_0. In particular, the adversary knows Alice's codebook $\mathcal{C}_n$ and is *partially oblivious* to the transmitted codeword $\boldsymbol{x}$. By partially oblivious, we mean that for $r \in [0,1]$ and some function $f_n : \{0,1\}^n \to \{0,1\}^{rn}$, the adversary observes a realization $\boldsymbol{\psi}$ of the random variable $\boldsymbol{\Psi} = \boldsymbol{\Psi}(m_0) = f_n(\mathcal{C}_n(m_0)) = f_n(\boldsymbol{x})$.[5] Due to the adversary's computational bound, for positive integers c, s, the adversary chooses f_n from the set $\mathsf{CPX}(r, cn^s)$ (we provide a rigorous definition of $\mathsf{CPX}(r, cn^s)$ is Section II-C) using its knowledge of $\mathcal{C}_n$ but not the realization of m_0.[6] The chosen function f_n is not revealed to Alice or Bob. Finally, the adversary chooses $\boldsymbol{e} \in \mathcal{B}_{pn}(\boldsymbol{0})$ based on the knowledge of the codebook $\mathcal{C}_n$ and the observation $\boldsymbol{\Psi}(m_0)$. We refer to the above adversary as the $\mathsf{CPX}(r, cn^s)$-oblivious adversary with error budget p.

C. Adversary's Complexity Constraint

For $r \in [0,1]$ and positive integers c, s, we precisely define the set $\mathsf{CPX}(r, cn^s)$. Let $\mathcal{F}_{n,r}$ denote the set of all Boolean functions of the form $f_n : \{0,1\}^n \to \{0,1\}^{rn}$. To define $\mathsf{CPX}(r, cn^s)$, we first define the circuit complexity of a function $f_n \in \mathcal{F}_{n,r}$.

A Boolean circuit B_n is an acyclic directed graph where each node is either an input node (with in-degree 0) or a logic gate (with in-degree 2). All nodes in B_n have out-degree 1 with unbounded fan-out and each logic gate computes an arbitrary Boolean function from $\{0,1\}^2$ to $\{0,1\}$. The *size* of B_n is the total number of nodes in B_n (input nodes and logic gates). Note that an observation function $f_n \in \mathcal{F}_{n,r}$ can be computed by some Boolean circuit that takes n bits as input and produces rn bits as output. The *circuit (size) complexity* of an observation function $f_n \in \mathcal{F}_{n,r}$ is the size of the smallest size Boolean circuit B_n that can compute f_n. We define $\mathsf{CPX}(r, cn^s)$ to be the set of all functions $f_n \in \mathcal{F}_{n,r}$ with a circuit complexity of at most cn^s. In modern complexity theory, the study of circuit complexity is a common approach to proving lower bounds on the complexity of certain problems [15].

D. Capacity

For a fixed $[n, M]$ codebook $\mathcal{C}_n$, the (average) probability of decoding error $\bar{P}_e(\mathcal{C}_n)$ is defined as the maximum over all $f_n \in \mathsf{CPX}(r, cn^s)$ of the quantity

$$\mathbb{E}_{\boldsymbol{\Psi}}\left[\max_{\boldsymbol{e} \in \mathcal{B}_{pn}(\boldsymbol{0})} \mathbb{P}_{m_0}(\hat{m}(\boldsymbol{e}, m_0) \neq m_0 | \boldsymbol{\Psi}(m_0) = \boldsymbol{\psi})\right]$$

where the probability measure $\mathbb{P}_{m_0}(\cdot)$ is w.r.t. the distribution of m_0, and the expectation $\mathbb{E}_{\boldsymbol{\Psi}}[\cdot] = \sum_{\boldsymbol{\psi} \in \{0,1\}^{rn}} (\cdot) \mathbb{P}(\boldsymbol{\Psi}(m_0) = \boldsymbol{\psi})$. Given the above channel model, we can define achievable rate in the usual way.

Definition 1 (Achievable Rate)**.** *For $p \in (0, 1/2)$, $r \in [0,1]$, and positive integers c, s, a rate $R \in (0,1]$ is said to be (c,s)-achievable if for any $\epsilon_e > 0$, there exists an n_0 such that*

[5]The fact that $\boldsymbol{\Psi}$ is a random variable follows from its dependency on the random variable m_0.

[6]However, the adversary knows that m_0 is drawn uniformly from $[M]$.

for all $n \geq n_0$, there exists an $[n, M]$ codebook $\mathcal{C}_n$ such that $\bar{P}_e(\mathcal{C}_n) \leq \epsilon_e$.

For $p \in (0, 1/2)$, $r \in [0, 1]$, and positive integers c, s, we define the capacity $C(p, r, cn^s)$ as the supremum of (c, s)-achievable rates.

III. Proof Outline, Overview of Proof Technique

In this section, we outline the proof of Theorem 1 and discuss an overview of our proof technique. A detailed proof of Theorem 1 can be found in the extended version of this paper [18].

A. Achievability Scheme

For our proof of Theorem 1, we construct a specific $\mathcal{C}_n$.

Encoder Construction: Alice's $[n, M]$ codebook $\mathcal{C}_n$ is constructed as follows. Let $\rho \in (R, C_{Shannon}(p))$. Codebook $\mathcal{C}_n$ is a concatenation of two codebooks: an *outer* $[\rho n, M]$ codebook $\mathcal{C}_{out} : [M] \to \{0,1\}^{\rho n}$ and for $N = 2^{\rho n}$, an *inner* $[n, N]$ codebook $\mathcal{C}_{in} : \{0,1\}^{\rho n} \to \{0,1\}^n$. Encoding proceeds as follows. First, Alice encodes m_0 with $\mathcal{C}_{out}$ where we denote the resulting codeword as $\mathcal{C}_{out}(m_0)$. Subsequently, Alice encodes $\mathcal{C}_{out}(m_0)$ with $\mathcal{C}_{in}$ where we denote the resulting codeword as $\mathcal{C}_n(m_0) = \mathcal{C}_{in}(\mathcal{C}_{out}(m_0))$. After encoding, Alice transmits the codeword $\boldsymbol{x} = \mathcal{C}_n(m_0)$ over the channel.

Decoder Construction: Bob's list decoder is constructed as follows. Given the channel output $\boldsymbol{y}$, Bob first performs list decoding by forming a list $\mathcal{L}_{in}(m_0, \boldsymbol{e}, \mathcal{C}_{in})$ of all words $\boldsymbol{w} \in \{0,1\}^{\rho n}$ such that $\mathcal{C}_{in}(\boldsymbol{w})$ is contained in the ball $\mathcal{B}_{pn}(\boldsymbol{y})$. After list decoding, Bob refines the list (i.e., Bob performs disambiguation) by removing all words $\boldsymbol{w} \in \mathcal{L}_{in}$ that are *inconsistent* with $\mathcal{C}_{out}$: we say that a word $\boldsymbol{w}$ is inconsistent with $\mathcal{C}_{out}$ if $\boldsymbol{w} \neq \mathcal{C}_{out}(m)$ for any $m \in [M]$.

Denote the refined list as $\mathcal{L}_{out}$. After $\mathcal{L}_{in}$ is refined to $\mathcal{L}_{out}$, a decoding decision is made according to the following rules. If $|\mathcal{L}_{out}| = 1$, then we have exactly one $m \in [M]$ s.t. $C_{out}(m) \in \mathcal{L}_{out}$, and the decoder outputs $\hat{m} = m$. If $\mathcal{L}_{out}$ is empty or $|\mathcal{L}_{out}| > 1$, then the decoder declares an error by setting $\hat{m}$ to an error symbol. We say that a decoding error occurs if $\hat{m} \neq m_0$. However, by the list decoding logic, $\mathcal{C}_{out}(m_0)$ is guaranteed to be in $\mathcal{L}_{out}$, and so the only non-trivial decoding error event occurs when $|\mathcal{L}_{out}| > 1$.

Probability of Error: Given the above construction, the probability of decoding error $\bar{P}_e(\mathcal{C}_{out}, \mathcal{C}_{in})$ can be written as

$$\max_{f_n \in \mathsf{CPX}(r, cn^s)} \mathbb{E}_{\boldsymbol{\Psi}} \left[\max_{\boldsymbol{e} \in \mathcal{B}_{pn}(\boldsymbol{0})} \mathbb{P}_{m_0}(|\mathcal{L}_{out}| > 1 | \boldsymbol{\Psi}(m_0) = \boldsymbol{\psi}) \right]$$

which in turn is equal to the maximum over all $f_n \in \mathsf{CPX}(r, cn^s)$ of the quantity

$$\mathbb{E}_{\boldsymbol{\Psi}} \left[\max_{\boldsymbol{e} \in \mathcal{B}_{pn}(\boldsymbol{0})} \mathbb{P}_{m_0}\Big(\bigcup_{i=1}^{|\mathcal{L}_{in}|} \{\boldsymbol{w}_i \in \mathcal{I}_{m_0}\} | \boldsymbol{\Psi}(m_0) = \boldsymbol{\psi} \Big) \right] \quad (1)$$

where for $i = 1, \ldots, N$ we define[7]

$$\boldsymbol{w}_i(m_0, \boldsymbol{e}, \mathcal{C}_{in}) = \arg \min_{\boldsymbol{w} \in \mathcal{L}_{in} \setminus \{\boldsymbol{w}_1, \ldots, \boldsymbol{w}_{i-1}\}} d(\boldsymbol{y}, \mathcal{C}_{in}(\boldsymbol{w}))$$

to be the word corresponding to the ith closest codeword in $\mathcal{C}_{in}$ to $\boldsymbol{y}$, and we define

$$\mathcal{I}_{m_0} = \{\boldsymbol{z} \in \{0,1\}^{\rho n} : \exists m' \in [M], m' \neq m_0, \boldsymbol{z} = \mathcal{C}_{out}(m')\} \quad (2)$$

to be the set of words that are not inconsistent with $\mathcal{C}_{out}$ and do not correspond to the true message m_0.

B. Outline of the proof of Theorem 1

Random Coding Argument: In the sequel, for $p \in (0, 1/2)$, $r \in [0, 1 - H(p))$, positive integers c, s and for any $\epsilon_\rho, \epsilon_R > 0$, we set the outer-code rate $\rho = C_{Shannon}(p) - \epsilon_\rho$ and the inner-outer combined code rate $R = \rho - \epsilon_R$. Thus, rate R can be arbitrarily close to $C_{Shannon}$. To prove Theorem 1, we show that the rate R is (c, s)-achievable. We show this by using a *random-coding argument* in conjunction with the code construction presented in Section III-A. The argument states that for any fixed $[\rho n, M]$ outer code $\mathcal{C}_{out}$ that is a 1:1 function and for any $\epsilon_e > 0$, if there exists some n_0 such that for all $n \geq n_0$ there exists some non-empty set $\mathcal{G}$ of $[n, N]$ codebooks such that for any $\mathcal{C}_{in} \in \mathcal{G}$ we have $\bar{P}_e(\mathcal{C}_{in}, \mathcal{C}_{out}) \leq \epsilon_e$, then Theorem 1 holds.

Setup: In the sequel, we drop the dependency on $\mathcal{C}_{out}$ from all notation due to the outer codebook being fixed. To apply the random-coding argument, we first apply a simple union bound to $\bar{P}_e(\mathcal{C}_{in})$ and bound it above by $\bar{P}_e^{ub}(\mathcal{C}_{in})$ defined as

$$\max_{f_n \in \mathsf{CPX}(r, cn^s)} \sum_{i=1}^{N} \mathbb{E}_{\boldsymbol{\Psi}} \left[\max_{\boldsymbol{e} \in \mathcal{B}_{pn}(\boldsymbol{0})} q_i(f_n, \boldsymbol{\psi}, \boldsymbol{e}, \mathcal{C}_{in}) \right] \quad (3)$$

where for $f_n \in \mathsf{CPX}(r, cn^s)$, $\boldsymbol{\psi} \in \{0,1\}^{rn}$, $\boldsymbol{e} \in \mathcal{B}_{pn}(\boldsymbol{0})$ and $i \in [N]$ we define $q_i(f_n, \boldsymbol{\psi}, \boldsymbol{e}, \mathcal{C}_{in}) = \mathbb{P}(\boldsymbol{w}_i \in \mathcal{L}_{in}, \boldsymbol{w}_i \in \mathcal{I}_{m_0} | \boldsymbol{\Psi}(m_0) = \boldsymbol{\psi})$ to be the probability that $\boldsymbol{w}_i$ results in a decoding error.

To bound $\bar{P}_e^{ub}$, we leverage the list decodable properties of the inner codebook.

Definition 2. *For $\ell > 0$, an $[n, N]$ codebook $\mathcal{C}_{in}$ is said to be $[\ell, p]$ list decodable if $|\mathcal{C}_{in} \cap \mathcal{B}_{pn}(\boldsymbol{y})| \leq \ell$ for every $\boldsymbol{y} \in \{0,1\}^n$.*

For each n large enough, to show that there exists an $[n, N]$ codebook $\mathcal{C}_{in}$ such that $\bar{P}_e^{ub}(\mathcal{C}_{in}) \leq \epsilon_e$, it is sufficient to show that for some number $L = O(1/\epsilon_\rho)$ (independent of n), the probability (over the choice of inner codebook) that $\mathcal{C}_{in}$ is *not* $[L, p]$ list decodable or $q_i(f_n, \boldsymbol{\psi}, \boldsymbol{e}, \mathcal{C}_{in}) > \epsilon_e / L$ for *some* $f_n \in \mathsf{CPX}(r, cn^s)$, $\boldsymbol{\psi} \in \{0,1\}^{rn}$, $\boldsymbol{e} \in \mathcal{B}_{pn}(\boldsymbol{0})$ and $i \in [L]$ (excluding, possibly, a few $\boldsymbol{\psi} \in \{0,1\}^{rn}$ that have a vanishing probability over m_0 of being observed by the adversary) is strictly less than 1.[8] We note the probability that $\mathcal{C}_{in}$ is not $[L, p]$ list decodable is vanishing as $n \to \infty$ following known results on the list decodability of random codes (e.g., [8, Claim A.15]), . Hence, following a union bound, we only need to show that with probability bounded away from 1, we have $q_i(f_n, \boldsymbol{\psi}, \boldsymbol{e}, \mathcal{C}_{in}) > \epsilon_e / L$ for some parameters as described above. When confusion can be avoided, we drop the notated

[7] We address the scenario where $\boldsymbol{w}_i$ is not unique in the extended version [18].

[8] Note that a $[L, p]$ list decodable $\mathcal{C}_{in}$ implies that $q_i(\cdot, \cdot, \cdot, \mathcal{C}_{in}) = 0$ for $L \leq i \leq N$.

dependencies and subscripts of $q_i(f_n, \boldsymbol{\psi}, \boldsymbol{e}, \mathcal{C}_{in})$ and simply write $q(\mathcal{C}_{in})$ to emphasize the dependency on $\mathcal{C}_{in}$.

Analysis of $q(\mathcal{C}_{in})$: We fix $i \in [L]$, $f_n \in \mathsf{CPX}(r, cn^s)$, $\boldsymbol{\psi} \in \{0,1\}^{rn}$ and $\boldsymbol{e} \in \mathcal{B}_{pn}(\mathbf{0})$ and for $n = 1, 2, \ldots$ we study the *concentration of measure* of $q(\mathcal{C}_{in})$ around its expectation $\mathbb{E}_{\mathcal{C}_{in}}[q]$ (here, the expectation is w.r.t. the distribution of $\mathcal{C}_{in}$). We do so by deriving concentration inequalities via the logarithmic Sobolev inequalities, e.g., [19]. This derivation is also known as the entropy method.

An example of a common inequality derived via the entropy method is as follows. Define the *variation* of $q(\mathcal{C}_{in})$ as $V(\mathcal{C}_{in}) = \sum_{j=1}^{N} \mathbb{E}_{\boldsymbol{z}} |q(\mathcal{C}_{in}) - q(\mathcal{C}_{in}(j, \boldsymbol{z}))|^2$ where codebook $\mathcal{C}_{in}(j, \boldsymbol{z})$ is equal to $\mathcal{C}_{in}$ with the jth codeword replaced with the codeword $\boldsymbol{z}$ uniformly distributed in $\{0,1\}^n$. The quantity $V(\mathcal{C}_{in})$ captures how smoothly $q(\mathcal{C}_{in})$ varies for incremental changes to $\mathcal{C}_{in}$. For $a > 0$, we say that q is *a-smooth* if for all $[n, N]$ codebooks $\mathcal{C}_{in}$ we have $V(\mathcal{C}_{in}) \leq a$.

Proposition 1 ([19, Corollary 3])**.** *Suppose there exists an $a = a(n) > 0$ such that q is a-smooth. For $\lambda > 0$,*

$$\mathbb{P}_{\mathcal{C}_{in}}(q - \mathbb{E}_{\mathcal{C}_{in}}[q] > \lambda) \leq \exp\left\{\frac{\lambda^2}{4a}\right\} \tag{4}$$

where the probability $\mathbb{P}_{\mathcal{C}_{in}}$ is w.r.t. the distribution of $\mathcal{C}_{in}$.[9]

To apply Proposition 1, one can first set $\lambda = \epsilon_e/L - \mathbb{E}_{\mathcal{C}_{in}}[q]$, and then find a value of a small enough such that q is a-smooth and the R.H.S. of the inequality (4) approaches 0 in the limit $n \to \infty$. In practice, this value of a is difficult to find. The difficulty arises from the fact that for a small subset of $[n, N]$ codebooks $\mathcal{C}_{in}$, the variation $V(\mathcal{C}_{in})$ is large. Thus, Proposition 1 cannot be used directly.

Handling large variation: To address this issue, we take the following bootstrapping approach. We first approximate $q(\mathcal{C}_{in})$ with a function $q'(\mathcal{C}_{in})$ that is equal to $q(\mathcal{C}_{in})$ with high probability over the choice of $\mathcal{C}_{in}$ and has a small variation $V'(\mathcal{C}_{in})$ for all $[n, N]$ codebooks $\mathcal{C}_{in}$. We then imply the concentration of q by showing that the approximation q' is concentrated via a entropy-method-type concentration inequality that resembles Proposition 1.

To approximate $q(\mathcal{C}_{in})$, we first define a set $\mathcal{T}$ of *typical* $[n, N]$ *codebooks* such that $\mathbb{P}_{\mathcal{C}_{in}}(\mathcal{C}_{in} \notin \mathcal{T})$ is small. If $\mathcal{C}_{in} \in \mathcal{T}$, we define $q'(\mathcal{C}_{in}) = q(\mathcal{C}_{in})$. If otherwise $\mathcal{C}_{in} \notin \mathcal{T}$, we define $q'(\mathcal{C}_{in})$ such that for sufficiently small $a > 0$, q' is a-smooth.

Lastly, we briefly discuss the reasoning behind our choice of code construction. We remark that our construction and the resulting definition of q help us to show that $V'(\mathcal{C}_{in})$ is sufficiently small when $\mathcal{C}_{in} \notin \mathcal{T}$. Concatenated coding and list decoding/refinement allow us to isolate for $i = 1, \ldots, N$ the effect of the ith codeword of $\mathcal{C}_{in}$ on $\bar{P}_e(\mathcal{C}_{in})$ and show $V'(\mathcal{C}_{in}) = O(i)$ $\forall \mathcal{C}_{in} \notin \mathcal{T}$. Without the construction, for non-typical codebooks, the function under analysis may have a variation equal to $O(2^{Rn})$ which is too large to apply our concentration inequalities.

[9]Proposition 1 requires that each codeword of $\mathcal{C}_{in}$ be independently and uniformly drawn from $\{0,1\}^n$.

Prior work: Our above approach is inspired by Langberg's framework [13] to study concentration of measure when the function under analysis is non-smooth. The main technical contribution of [13] is to carefully define the typical set $\mathcal{T}$ based on the codebooks' list decodable properties in way where one can then apply Vu's martingale-type concentration inequalities for non-smooth functions [20]. We follow Langberg's framework by also defining typicality in terms of list decodability. However, we use entropy-method-type concentration inequalities.

The major technical difference between our work and [13] lies at the definition of smoothness. For $a > 0$, reference [13] defines smoothness in terms of *a-Lipschitz*: for the quantity $W(\mathcal{C}_{in}) = N \max_{j \in [N], \boldsymbol{z} \in \{0,1\}^n} |q(\mathcal{C}_{in}) - q(\mathcal{C}_{in}(j, \boldsymbol{z}))|^2$, q is a-Lipschitz if for all $[n, N]$ codebooks $\mathcal{C}_{in}$ we have $W(\mathcal{C}_{in}) \leq a$.[10] We remark that for $a > 0$, q is a-smooth if q is a-Lipschitz, and thus a-Lipschitz is a stronger notion of smoothness than the notion used in our work.

The advantage to characterizing smoothness using a-smooth as opposed to a-Lipschitz is apparent in the following observation. Suppose that for $a_L > 0$, q is a_L-Lipschitz. Then one can usually find a smaller value $a_S \in (0, a_L)$ such that q is a_S-smooth, and in turn, leverage the a_S-smooth criterion to apply tighter concentration inequalities than those reliant on the a_L-Lipschitz criterion. Indeed, we take this approach to show that the probability that q is not concentrated is at most $2^{-2^{\Omega(n)}}$, whereas bounds on the order of $2^{-poly(n)}$ are obtainable using a_L-Lipschitz together with the framework of [13]. The cost of this approach is that finding a smaller value a_S can require significant effort compared to finding a_L. Indeed, our proof of Theorem 1 invests significant effort into finding a_S.

Computational Bound: Lastly, we briefly discuss the role that the adversary's computational bound plays in our random-coding argument. Recall that to show a rate R is (c, s)-achievable, it is sufficient to show that there exists one $[n, N]$ codebook $\mathcal{C}_{in}$ such that quantity (1) is small for *all* $f_n \in \mathsf{CPX}(r, cn^s)$. Showing the existence of this codebook becomes difficult if the set $\mathsf{CPX}(r, cn^s)$ contains many functions. We simplify our search for this codebook by bounding the number of functions in the set $\mathsf{CPX}(r, cn^s)$. We remark that the set $\mathsf{CPX}(r, cn^s)$ can be shown to have $2^{poly(n)}$ functions, and therefore, $\mathsf{CPX}(r, cn^s)$ is much smaller than the set of functions $\mathsf{CPX}(r, \infty) = \mathcal{F}_{n,r}$ with unbounded circuit complexity which has 2^{rn2^n} functions.

IV. Conclusion

In this work, we study the capacity of adversarial channels in which the adversary can observe the transmitted codeword via some computationally bounded process. We characterize the capacity for certain parameters under deterministic encoding and average probability of error criterion.

[10]In the definition of $W(\mathcal{C}_{in})$, we include constant N and square the absolute difference term to normalize $W(\mathcal{C}_{in})$ with respect to $V(\mathcal{C}_{in})$.

REFERENCES

[1] C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, vol. 27, no. 3, 1948.

[2] R. J. McEliece, E. R. Rodemich, H. Rumsey, and L. R. Welch, "New Upper Bounds on the Rate of a Code via the Delsarte—MacWilliams Inequalities," *IEEE Transactions on Information Theory*, 1977.

[3] R. J. Lipton, "A new approach to information theory," in *Annual Symposium on Theoretical Aspects of Computer Science*, 1994.

[4] V. Guruswami and A. Smith, "Optimal rate code constructions for computationally simple channels," *Journal of the ACM*, vol. 63, no. 4, 2016.

[5] A. D. Sarwate, "Coding against myopic adversaries," in *2010 IEEE Information Theory Workshop, ITW 2010 - Proceedings*, 2010.

[6] B. K. Dey, S. Jaggi, and M. Langberg, "Sufficiently Myopic Adversaries Are Blind," *IEEE Transactions on Information Theory*, 2019.

[7] A. J. Budkuley, B. K. Dey, S. Jaggi, M. Langberg, A. D. Sarwate, and C. Wang, "Symmetrizability for Myopic AVCs," in *IEEE International Symposium on Information Theory - Proceedings*, vol. 2020-June, 2020.

[8] Z. Chen, S. Jaggi, and M. Langberg, "A characterization of the capacity of online (causal) binary channels," in *ACM symposium on Theory of Computing*, 2015, pp. 287–296.

[9] B. K. Dey, S. Jaggi, M. Langberg, and A. D. Sarwate, "A bit of delay is sufficient and stochastic encoding is necessary to overcome online adversarial erasures," in *IEEE International Symposium on Information Theory - Proceedings*, vol. 2016-August, 2016.

[10] V. Suresh, E. Ruzomberka, and D. J. Love, "Stochastic-Adversarial Channels: Online Adversaries with Feedback Snooping," in *IEEE International Symposium on Information Theory - Proceedings*, vol. 2021-July, 2021.

[11] I. Csiszár and P. Narayan, "Capacity and Decoding Rules for Classes of Arbitrarily Varying Channels," *IEEE Transactions on Information Theory*, vol. 35, no. 4, 1989.

[12] I. Csiszar and P. Narayan, "The Capacity of the Arbitrarily Varying Channel Revisited: Positivity, Constraints," *IEEE Transactions on Information Theory*, vol. 34, no. 2, 1988.

[13] M. Langberg, "Oblivious communication channels and their capacity," *IEEE Transactions on Information Theory*, 2008.

[14] C. Wang, "On the capacity of the binary adversarial wiretap channel," in *54th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2016*, 2017.

[15] M. Sipser, *Introduction to the theory of computation*, 2nd ed. Boston: Thompson Course Technology, 2006.

[16] E. N. Gilbert, "A comparison of signaling alphabets," *Bell System Technical Journal*, vol. 31, no. 3, p. 504–522, 1952.

[17] R. R. Varshamov, "Estimate of the number of signals in error correcting codes," *Dokl. Acad. Nauk*, vol. 117, no. 739–741, 1957.

[18] E. Ruzomberka, C.-C. Wang, and D. J. Love, "Channel capacity for adversaries with computationally bounded observations," *ArXiv preprint*.

[19] S. Boucheron, G. Lugosi, and P. Massart, "Concentration inequalities using the entropy method," *Annals of Probability*, vol. 31, no. 3, 2003.

[20] V. H. Vu, "Concentration of Non-Lipschitz Functions and Applications," in *Random Structures and Algorithms*, 2002.