Stochastic-Adversarial Channels: Online Adversaries With Feedback Snooping

Vinayak Suresh, Eric Ruzomberka and David J. Love
School of Electrical and Computer Engineering
Purdue University
Email: suresh20@purdue.edu, eruzombe@purdue.edu, djlove@purdue.edu

***Abstract*—The growing need for reliable communication over untrusted networks has caused a renewed interest in adversarial channel models, which often behave much differently than traditional stochastic channel models. Of particular practical use is the assumption of a *causal* or *online* adversary who is limited to causal knowledge of the transmitted codeword. In this work, we consider stochastic-adversarial mixed noise models. In the set-up considered, a transmit node (Alice) attempts to communicate with a receive node (Bob) over a binary erasure channel (BEC) or binary symmetric channel (BSC) in the presence of an online adversary (Calvin) who can erase or flip up to a certain number of bits at the input of the channel. Calvin knows the encoding scheme and has strict causal access to Bob's reception through *feedback snooping*. For erasures, we provide a complete capacity characterization with and without transmitter feedback. For bit-flips, we provide converse and achievability bounds.**

I. Introduction

A central endeavour in information theory is the study of capacity and strategies for reliable communication over different types of channels. Two different philosophies exist on how channels are modeled. Channels in the Shannon world are characterized by some stochastic process that injects errors independently of the communication scheme, while channels in the Hamming world are characterized by an adversary who injects worst-case errors. Historically, adversarial channels were studied under either full knowledge (*omniscient adversary*) or no knowledge (*oblivious adversary*) of the transmitted codeword. A number of recent works [2]–[9] instead consider coding against *online* or *causal* adversaries wherein at any point during the transmission, the adversary knows only part of the codeword transmitted thus far.

As noted in [5], the causal adversary model lies in between the stochastic and the omniscient adversary models. In this work, we further bridge together the Shannon and the Hamming worlds by studying a new model where both adversarial and random sources of error are present. Specifically, Alice attempts to send a message to Bob over a binary erasure channel BEC(q) or binary symmetric channel BSC(q) in the presence of a causal adversary Calvin who can erase or flip a certain number of bits at the input of the channel. This is depicted in Fig 1. Any transmission strategy must not only overcome the noise due to the random channel but also from the adversary. We also assume that Calvin has access to Bob's reception, which we refer to as *feedback snooping*. The ability to spy on both Alice and Bob aids Calvin in designing strong attacks. Our goal is to characterize the capacity of this channel.

When there is no random channel present, i.e., $q = 0$ in Fig. 1, the only source of noise is adversarial. A complete capacity characterization for this case is given in [5]–[7]. Our models differ from the ones considered previously in two ways:

- **Mixture of random and adversarial noise** - The noise in the received word is affected by the random channel BEC/BSC as well as the actions of Calvin who is erasing/flipping bits. For example in the erasure case, a bit not erased by Calvin can be erased by the BEC. Similarly, in the bit-flip case, a bit flipped by Calvin may be "unflipped" by the BSC. Conceptually, we think of the discrete memoryless channel (DMC) as the main channel through which Alice and Bob communicate, and Calvin as a malicious entity who attempts to disrupt the transmission.
- **Feedback to adversary** - In our setting, Calvin is allowed access to Bob's reception through *feedback snooping*. This becomes important due to the presence of the stochastic channel. The adversarial attacks described in [5], [6] if used directly do not provide the right distance bounds needed to establish our converse results. These are appropriately strengthened and crucially rely on Calvin's ability to snoop. Note that feedback snooping is unnecessary when $q = 0$.

Our contributions can be summarized as follows:

- We provide a complete characterization of capacity for the case of erasures. Our result implies that the presence of the random channel BEC(q) in addition to adversarial erasures simply scales the capacity of the $q = 0$ case by a multiplicative factor.
- For the case of erasures, we also characterize the capacity when Alice has causal access to Bob's reception and encoding is *closed-loop*. In this scenario, we show that Calvin gains no benefit from his ability to spy on Alice or Bob. In fact, he can do no better than making erasure decisions in an i.i.d. manner.
- Finally in the case of bit-flips, we prove non-trivial converse and achievability bounds.

There are other adversarial models intermediate between the

This work was funded in part by the National Science Foundation under grants CNS1642982, CCF1816013, and EEC1941529.

An extended version of this paper is available at [1].

2021 IEEE International Symposium on Information Theory (ISIT) | 978-1-5386-8209-8/21/$31.00 ©2021 IEEE | DOI: 10.1109/ISIT45174.2021.9517968

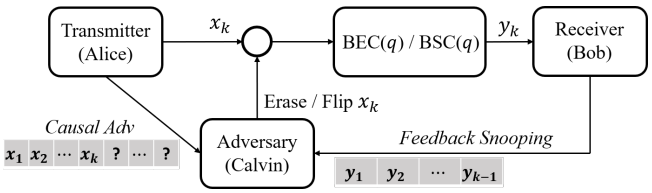

Fig. 1. Channel models considered in this work.

oblivious and omniscient models that have been considered in the literature which we do not pursue here (e.g., [10]–[15]). The problem of coding with feedback to the transmitter has been studied by several authors such as [16]–[18]. Finally, we note that our models can be cast under the more general framework of (non-state-deterministic) arbitrarily varying channels or AVCs [19], [20]. However, known results for AVCs do not directly imply the results of this paper.

II. Channel Models

Consider the channel depicted in Fig. 1. Alice (the transmitter) attempts to convey a message to Bob (the receiver) over a BEC(q), in the presence of a p-limited causal adversary (Calvin) where the terms will be clarified shortly. The input and output alphabets are $\mathcal{X} = \{0, 1\}$ and $\mathcal{Y} = \{0, 1, \Lambda\}$, respectively, where Λ denotes an erasure symbol. Encoding is done over n channel uses, and the size of the message set is 2^{nR}. We allow stochastic encoding and assume the presence of local randomness available only to Alice for this purpose. Denote $x_k \in \mathcal{X}$ to be the symbol selected by Alice at channel use k. At time k, Calvin makes a decision on whether to erase x_k based on his side-information to be specified later. If Calvin erases x_k, the received symbol at time k at the receiver is an erasure, i.e., $y_k = \Lambda$. If Calvin decides not to erase x_k, then $y_k = x_k$ with probability $1 - q$ and $y_k = \Lambda$ with probability q, i.e., x_k is erased with probability q.

We assume that Calvin knows the codebook used at the transmitter in the case of deterministic encoding or the distribution of codewords in the case of stochastic encoding. Calvin is assumed to be *causal*, i.e., at each channel use k, he knows only part of the codeword transmitted so far $(x_1, x_2, \cdots, x_k) \in \mathcal{X}^k$. Calvin is neither aware of the message nor future transmissions. However, he has access to Bob's reception $(y_1, y_2, \cdots, y_{k-1}) \in \mathcal{Y}^{k-1}$ through a delay-free and noise-free strictly causal feedback link as shown in Fig. 1.

A power constraint is further imposed by enforcing Calvin to be p-limited, meaning that he can erase at most a constant fraction p of the bits, i.e., if $\mathbf{a} \in \{0, \Lambda\}^n$ denotes the positions where Calvin decides to erase symbols from $(x_1, x_2, \cdots, x_n)$, we must have $weight(\mathbf{a}) \leq pn$. We refer to this model as *the BEC causal adversarial channel with feedback snooping* (or BEC(q)-ADV(p)-FS). Note that the BEC block in Fig. 1 is slightly different from the classical BEC. If Calvin erases x_k to an erasure symbol Λ, we have $y_k = \Lambda$, where Λ does not carry any information.

Our aim is to characterize the capacity of this channel, i.e., the largest value of R such that Alice can reliably convey one out of 2^{nR} possible messages to Bob. Precise definitions are given shortly. In Section IV, we also consider a related channel by replacing the BEC(q) with a BSC(q) and letting Calvin flip bits instead of erasing them, denoted henceforth as BSC(q)-ADV(p)-FS.

Notation and Definitions: In this work, we only consider fixed length encoding. The blocklength is denoted by n. The transmitted message is denoted by the random variable (r.v.) $\mathbf{U}$ chosen uniformly from the message set $\mathcal{U} = \{1, 2, 3, \cdots, 2^{nR}\}$. A deterministic code consists of a fixed encoder map $\Phi_d : \mathcal{U} \to \mathcal{X}^n$ and a decoder map $\Gamma_d : \mathcal{Y}^n \to \mathcal{U}$, where each message is associated to a unique codeword. In case of stochastic encoding, a codeword $\mathbf{x}$ is selected for a message u according to a chosen conditional distribution $\Phi(.|u)$ defined on $\mathcal{X}^n$. A stochastic code is fully specified by defining all conditional distributions $\{\Phi(.|u)\}_{u \in \mathcal{U}}$ and a decoder $\Gamma : \mathcal{Y}^n \to \mathcal{U}$. Without loss of generality, we assume in proving converse results that no two distinct messages map to the same codeword. The (maximum) probability of error is then

$$P_e = \max_{u \in \mathcal{U}} \max_{\text{ADV}(p)} \sum_{\mathbf{y}} \sum_{\mathbf{x}} P(\mathbf{y}|\mathbf{x}) \Phi(\mathbf{x}|u) \mathbb{1}(\Gamma(\mathbf{y}) \neq u) \quad (1)$$

where $\mathbb{1}(.)$ denotes the indicator function and ADV(p) denotes a feasible strategy chosen by Calvin. Note that $P(\mathbf{y}|\mathbf{x})$ in (1) is a function of both the stochastic channel and the chosen adversarial strategy. We say that $R > 0$ is achievable if for every $\delta > 0$ and every sufficiently large n, there is a code of rate R and blocklength n with $P_e < \delta$. The capacity is defined to be the supremum of all achievable rates. Let Ber(q) denote a Bernoulli r.v. with success probability q and $h_2(x) = -x \log_2(x) - (1-x) \log_2(1-x)$ be the binary entropy function. For $x, y \in [0, 1/2]$, let $x \star y = x(1 - y) + y(1 - x)$ and note that $x \star y = 1/2$ iff either $x = 1/2$ or $y = 1/2$ (or both). Denote by $d(\mathbf{x}, \mathbf{y})$ the Hamming distance between $\mathbf{x}$ and $\mathbf{y}$. For $\mathbf{s} = (s_1, s_2, \cdots, s_n)$, we let $\mathbf{s}_1 = (s_1, s_2, \cdots, s_\ell)$ and $\mathbf{s}_2 = (s_{\ell+1}, \cdots, s_n)$, where ℓ is specified when proving converse results.

III. Results for Erasures

A. No Transmitter Feedback

Denote by $C^E(p, q)$ the capacity of BEC(q)-ADV(p)-FS when Alice has no side-information, i.e., encoding is restricted to be *open-loop*. We prove the following result.

Theorem 1. *The capacity $C^E(p, q)$ of BEC(q)-ADV(p)-FS is given by*

$$C^E(p, q) = \begin{cases} (1 - 2p)(1 - q) & \text{for } 0 \leq p \leq \frac{1}{2},\ 0 \leq q \leq 1 \\ 0 & \text{otherwise} \end{cases}. \quad (2)$$

Remark. When there is no BEC, i.e., when $q = 0$, our model reduces to the one studied in [6], [7]. Our result implies that in the setting where both causal adversarial erasures and random erasures are present, the capacity simply scales by a factor of $1 - q$.

Proof (Sketch) of Converse : The proof is based on a *wait and snoop, then push* attack inspired by, but different from, an attack in [2], [6]. Let the transmitted codeword be denoted by $\mathbf{x}$. Fix $\epsilon > 0$. Let $R = (1-2p)(1-q) + \epsilon$.

- **Wait and Snoop**: Calvin waits and does not induce any erasures for the first $\ell = n\frac{R-\frac{\epsilon}{2}}{1-q}$ channel uses. Instead, Calvin simply snoops into Bob's reception to determine the erased/unerased bits and their positions. At the end of this phase, Bob receives $\mathbf{y}_1$ containing some erased and some unerased bits. Let $\{i_j\}_{j=1}^{m}$ be the indices of unerased symbols.
- **Push**: Calvin forms the set $\mathcal{B}_{\mathbf{y}_1}$ of codewords consistent with $\mathbf{y}_1$ as

$$\mathcal{B}_{\mathbf{y}_1} = \{\mathbf{v} \in \mathcal{X}^n : \exists \tilde{u} \in \mathcal{U} \text{ s.t. } \Phi(\mathbf{v}|\tilde{u}) > 0 \text{ and } v_{i_k} = x_{i_k}\ k = 1, 2, \cdots, m\}. \quad (3)$$

He then samples a codeword $\mathbf{x}'$ from $\mathcal{B}_{\mathbf{y}_1}$ according to the distribution $P_{\mathbf{X}|\mathbf{Y}_1=\mathbf{y}_1}(.|\mathbf{y}_1)$. In the push phase, Calvin simply erases bit x_i whenever $x_i \neq x'_i$. Recall that the total erasure budget is pn. Hence, if $\mathbf{x}$ and $\mathbf{x}'$ correspond to different messages u and u' and are sufficiently close such that $d(\mathbf{x}_2, \mathbf{x}'_2) < pn$, there is no way for Bob to distinguish between messages u and u' under Calvin's attack. The proof relies on showing that this indeed occurs with a positive probability *independent* of n.

Note that while the presence of the BEC(q) lowers the target rate, Calvin adds no erasures for approximately $n(1-2p)$ channel uses which from [6], [7] is optimal when there is no BEC(q). The main difference in attack when $q \neq 0$ is that even though Calvin knows the entire prefix of the transmitted codeword $\mathbf{x}_1 = (x_1, x_2, \cdots, x_\ell)$, he forms his set in (3) based only on the unerased bits. Thanks to feedback snooping, Calvin exploits the additional equivocation induced by the BEC(q) in the wait and snoop phase to pick a codeword that is sufficiently close to the transmitted codeword, and which corresponds to a message different from one that Alice chose. Note also that while we give Calvin full causal access to Bob's reception, an alternate model where Calvin is allowed *one-time block feedback* is sufficient - he would add no erasures for ℓ channel uses, retrieve through feedback the entire block $\mathbf{y}_1$ and then 'push'.

The steps in the proof closely follow [6] accounting for the addition of the BEC(q). In the push phase, let E_2 be the event $\{\mathbf{U} \neq \mathbf{U}'\}$ and E_3 be the event $\{d(\mathbf{X}_2, \mathbf{X}'_2) < pn\}$, and note that when both E_2 and E_3 occur simultaneously, Calvin's budget of pn erasures is enough to cause a decoding error with probability at least $1/2$. Thus, to finish the proof, we need only show a lower bound on $P(E_2, E_3)$. Using techniques from [6], it can be shown that

$$P(E_2, E_3) \geq \frac{\epsilon}{4}\frac{\epsilon}{8p}\left(\frac{\epsilon}{5}\right)^{\frac{9}{\epsilon}-1} = \epsilon^{O(1/\epsilon)}$$

which holds independent of n as required. Details are provided in [1].

Proof (Sketch) of Achievability: We resort to a random coding argument to claim existence of a stochastic code that achieves (2). Our construction is a modification of the encoder and decoder described in [7] which we first review. While reviewing, we provide key insights into how this decoder might fail once a BEC is added. Following the review, we use our insights to modify the decoder in order to account for the additional random noise when $q > 0$. Alice has a set of private secrets $\mathcal{S}$ she uses for (stochastic) encoding. Fix $\epsilon > 0$ and let $R = (1-2p-\epsilon)(1-q)$, $\theta = \frac{\epsilon}{4}$. The encoder and decoder of [7] is constructed as follows (here, $q = 0$):

- **Encoder:** A message u is mapped to several sub-codewords or chunks, each of size $n\theta$, which are concatenated together to form the transmitted codeword. Each chunk is obtained from a stochastic code where the secrets between chunks are chosen independently. Further technical details can be found in [7].
- **Decoder:** Decoding begins after Bob receives the entire n-symbol channel output $\mathbf{y}$. For some integer t^*, Bob partitions $\mathbf{y}$ into 2 strings: $\mathbf{y}_1 = (y_1, \ldots, y_{t^*})$ and $\mathbf{y}_2 = (y_{t^*+1}, \ldots, n)$. Decoding occurs in two sequential phases. In the first phase, Bob performs list decoding on $\mathbf{y}_1$ to create a list of messages $\mathcal{L}$. In the second phase, he refines the list by removing all messages in $\mathcal{L}$ that are not consistent with $\mathbf{y}_2$. Here, a message u' is said to be consistent with $\mathbf{y}_2$ iff some codeword corresponding to u' agrees with $\mathbf{y}_2$ on the unerased bits. If exactly one message, say $\hat{u}$, remains in $\mathcal{L}$ after refinement, the decoder outputs $\hat{u}$. If the refined list does not contain exactly one message, a decoding error is declared. Decoding is successful if $\hat{u} = u$.

Here, t^* is chosen as a function of the number of (purely adversarial since $q = 0$) erasures $\lambda^a_{t^*}$ observed in $\mathbf{y}$ up until time t^*. Specifically, Bob chooses t^* as the smallest integer

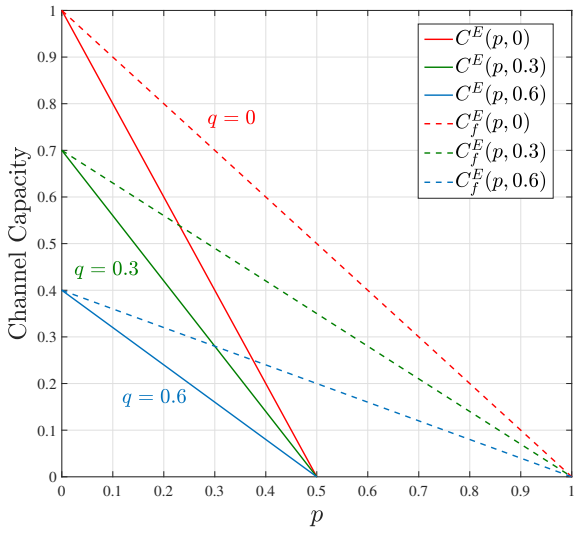

Fig. 2. Capacity of BEC(q)-ADV(p)-FS with ($C_f^E(p,q)$) and without transmitter feedback (resp. $C^E(p,q)$) as a function of p for $q = 0, 0.3, 0.6$. The cut-off value of p beyond which $C^E(p,q) = 0$ is $p = 1/2$ independent of q.

that satisfies the so-called *list-decoding condition*

$$\lambda_{t^*}^a \leq t^*(1-\theta) - ((1-2p)-\epsilon)n \tag{4}$$

and the *energy bounding condition*

$$np - \lambda_{t^*}^a \leq \frac{(n-t^*)(1-\theta)}{2}. \tag{5}$$

Condition (4) ensures the size of $\mathcal{L}$ is small (at most a constant) while condition (5) ensures the fraction of erasures that occur in $\mathbf{y}_2$ is small enough to perform list refinement.

Problems in this construction arise when $q > 0$. If the decoder assumes that all erasures that he sees are adversarial and performs decoding by selecting t^* according to (4) and (5), the maximum rate that can be achieved is $C^E(p+q-pq, 0) = C^E(p,q) - q$ which is strictly less than capacity. Therefore, simply counting erasures without knowing (or estimating) their source is no longer a viable strategy when $q > 0$. To circumvent this issue, we modify conditions (4) and (5) appropriately. Let λ_t denote the number of erasures observed by Bob up until time t, which includes contributions both from Calvin and the BEC(q). Then, Bob chooses t^* as the smallest integer that satisfies the modified list-decoding condition

$$\lambda_{t^*} - qt^* \leq t^*(1-q)(1-\theta) - Rn \tag{6}$$

and the modified list refinement condition

$$np(1-q) - (\lambda_{t^*} - qt^*) \leq \frac{(n-t^*)(1-q)(1-\theta)}{2}. \tag{7}$$

Note that if Calvin adds $\lambda_{t^*}^a$ erasures up until t^*, the total number of erasures λ_{t^*} that Bob observes is approximately $\lambda_{t^*} \approx \lambda_{t^*}^a + q(t^* - \lambda_{t^*}^a)$. On making this substitution we see that t^* satisfying (6) and (7) is nearly the same as that satisfying (4) and (5) i.e. it is sufficient to choose t^* only as a function of pure adversarial erasures. However, since Bob has no way of knowing this, he works with the quantity $\lambda_{t^*} - qt^*$ which is an estimate of the number of adversarial erasures that do not conincide with random erasures. Having selected t^*, Bob can then finish decoding using the two-phase decoding process of [7] to successfully recover the transmitted message. Further details of the proof are in [1].

B. With Transmitter Feedback

Suppose now that Alice in addition to Calvin has access to Bob's reception perfectly through a separate causal feedback link. This allows Alice to employ *closed-loop* encoding strategies where the input x_k at time k is possibly a function of both the message and Bob's reception thus far $(y_1, y_2, \cdots, y_{k-1})$, i.e.,

$$\mathbf{X}_k \sim f_k(\mathbf{U}, \mathbf{Y}_1, \mathbf{Y}_2, \cdots, \mathbf{Y}_{k-1}) \quad k = 1, 2, \cdots, n \tag{8}$$

where for each k, f_k is either deterministic or, more generally, a probabilistic map defining a conditional distribution $P_{\mathbf{X}|\mathbf{U},\mathbf{Y}_1,\mathbf{Y}_2,\cdots,\mathbf{Y}_{k-1}}$ over $\mathcal{X}$. Calvin is assumed to be causal. He does not know the message but knows the closed-loop encoding (possibly stochastic) maps $\{f_k\}_{k=1}^n$ used by Alice. Let the capacity in this case be denoted as $C_f^E(p,q)$. We have the following result.

Theorem 2. *The capacity $C_f^E(p,q)$ of BEC(q)-ADV(p)-FS with causal feedback to the transmitter is*

$$C_f^E(p,q) = (1-p)(1-q) \quad \forall\ 0 \leq p \leq 1, 0 \leq q \leq 1. \tag{9}$$

Remark. If Calvin were to simply erase each symbol with probability p, the rate is limited to[1] $(1-p)(1-q)$ which matches with the expression in (9). This implies that *the optimal attack for the adversary is to simply cause i.i.d. erasures. The knowledge of the (closed-loop) encoding scheme or the ability to snoop into Bob's reception does not buy Calvin any benefit.*

Proof (Sketch) of Converse : Fix $\epsilon > 0$. Calvin simply erases each symbol with probability $p - \frac{\epsilon}{1-q}$. By the Chernoff bound, the probability that Calvin will run out of his budget of pn erasures is at most $1 - 2^{-\Omega(\epsilon^2 n)}$. The combined effect of the adversary and the BEC(q) then is a BEC with erasure probability $s = p + q - pq - \epsilon$. Hence, $C_f^E(p,q) \leq 1 - s = (1-p)(1-q) + \epsilon$.

Proof (Sketch) of Achievability : Fix $\epsilon > 0$. The achievability scheme is essentially an ARQ scheme - transmit each of the k bits in the message repeatedly until it is successfully received. If e_Λ is the total number of erasures (a random quantity) that occur due to both the actions of Calvin and the BEC(q), Alice needs $n = k + e_\Lambda$ channel uses for this scheme to succeed. Note that at channel use t, since Calvin does not know whether the BEC(q) will introduce an erasure or not, we have that $P\left(e_\Lambda > ((p+q-pq)+\epsilon)n\right)$ is at most $1 - 2^{-\Omega(\epsilon^2 n)}$ and hence, $C_f^E(p,q) \geq (1-p)(1-q) - \epsilon$.

In Fig. 2, we plot $C^E(p,q)$ and $C_f^E(p,q)$ as a function of p for $q = 0, 0.3, 0.6$.

IV. Results for Bit-Flips

In this section, we assume that Calvin can attempt to flip up to pn bits and the random channel is a BSC(q) instead of a BEC(q). The input and output alphabets are $\mathcal{X} = \{0,1\}$ and $\mathcal{Y} = \{0,1\}$. At time k, Calvin produces $a_k \in \mathcal{A} = \{0,1\}$ based on his side information which is the same as before, i.e., he knows $(x_1, x_2, \cdots, x_k)$, the codebook or the codeword distribution, and $(y_1, y_2, \cdots, y_{k-1})$. The received symbol at time k at the receiver is $y_k = x_k \oplus a_k \oplus 1$ with probability q and $y_k = x_k \oplus a_k$ with probability $1-q$ where $\oplus$ denotes mod-2 addition and $q \in [0, 1/2]$. The constraint on the adversary can be expressed as $weight(a_1, a_2, \cdots, a_n) \leq pn$. In contrast to the erasure case, note that a flip-attempt of Calvin can now be undone by the BSC. No feedback to the transmitter is assumed. For this model denoted BSC(q)-ADV(p)-FS, we prove an upper bound and use the result of [7] to provide a simple achievable rate. The gap between the bounds gets larger when q gets larger. Eliminating this gap and proving a tight capacity characterization is left as future work.

[1]For a vanilla DMC such as the BEC, the capacity is the same under deterministic and stochastic encoding [20].

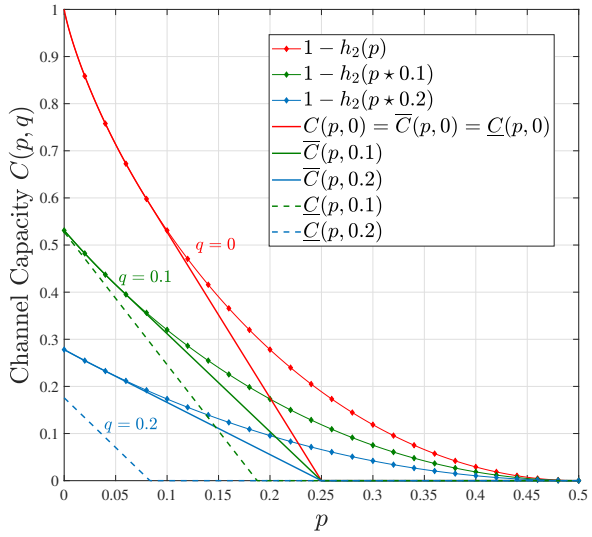

Fig. 3. Upper bound $\overline{C}(p,q)$ and lower bound $\underline{C}(p,q)$ on the capacity of BSC(q)-ADV(p)-FS as a function of p. The cut-off value of p beyond which $\overline{C}(p,q)=0$ is $p=1/4$ independent of q.

A. An Upper Bound $\overline{C}(p,q)$

Theorem 3. *The capacity $C(p,q)$ of BSC(q)-ADV(p)-FS is bounded as $C(p,q)\le\overline{C}(p,q)$ where*

$$\overline{C}(p,q)=\min_{\bar{p}:\bar{p}\in\mathcal{P}}\alpha(p,\bar{p},q)\left(1-h_2\left(\frac{\bar{p}}{\alpha(p,\bar{p},q)}\star q\right)\right), \quad (10)$$

$$\alpha(p,\bar{p},q)=1-4(p-\bar{p}) \quad , \quad \mathcal{P}=\{\bar{p}:0\le\bar{p}\le p\}$$

when $p<\frac{1}{4}$. When $p\ge\frac{1}{4}$, $C(p,q)=0$.

Remark. When $q=0$, i.e., there is no BSC, the channel model reduces to that considered in [5], [7], and the capacity expression (10) matches with the result proved in [5], [7].

Proof (Sketch): Fix any $\bar{p}\in[0,p]$ and $\epsilon>0$. Suppose that the transmitter attempts to communicate at a rate of $R=\alpha(p,\bar{p},q)\left(1-h_2\left(\frac{\bar{p}}{\alpha(p,\bar{p},q)}\star q\right)\right)+\epsilon$. We show a lower bound on the probability of error. Our proof is based on a *babble and snoop, then push* attack inspired in part from [5]. Let $\mathbf{x}$ and $\mathbf{y}$ denote the transmitted and received words.

- **Babble and Snoop**: For the first $\ell=(\alpha+\epsilon/2)n$ channel uses, Calvin injects random bit-flips and monitors Bob's reception - he flips each bit x_i independently with probability $\bar{p}n/\ell$. By the Chernoff bound, Calvin uses at most $\bar{p}n+\epsilon n/64$ flips with probability at least $1-e^{-\Omega(\epsilon^2 n)}$. Let this be event E_1. At the end of this phase, Calvin knows $\mathbf{x}_1$ and $\mathbf{y}_1$.
- **Push**: Calvin samples a codeword $\mathbf{x}'$ (corresponding to message u') according to the conditional distribution $P_{\mathbf{X}|\mathbf{Y}=\mathbf{y}_1}(.|\mathbf{y}_1)$. His goal is to confuse the receiver between $\mathbf{x}$ and $\mathbf{x}'$. At positions where $\mathbf{x}_2$ and $\mathbf{x}'_2$ agree, he does nothing. Positions j where $\mathbf{x}_2$ and $\mathbf{x}'_2$ disagree, he flips x_j with probability $1/2$. This way, Bob cannot distinguish between $\mathbf{x}$ and $\mathbf{x}'$ (even with the BSC(q)) due to the fact that $p(\mathbf{y}_2|\mathbf{x}_2)=p(\mathbf{y}_2|\mathbf{x}'_2)$. The proof relies on showing that with a small probability independent of n, u, u' are distinct and $\mathbf{x}_2$, $\mathbf{x}'_2$ are sufficiently close.

As was the case with erasures, the presence of the BSC(q) introduces additional equivocation at the receiver which Calvin is able to exploit thanks to his ability to snoop. Here also, *one-time block feedback* (of entire block $\mathbf{y}_1$) after the first ℓ channel uses is sufficient for the attack to succeed.

Conditioned on E_1, Calvin's remaining budget in the push phase is at least $(p-\bar{p})n-\epsilon n/64$. Letting $A_0=\left\{\mathbf{y}_1:H(\mathbf{U}\mid\mathbf{Y}_1=\mathbf{y}_1)>\frac{n\epsilon}{4}\right\}$, it can be shown that $P(\{\mathbf{Y}_1\in A_0\})\ge\epsilon/4$. Define the events $E_2=\{\mathbf{Y}_1\in A_0\}$, $E_3=\{\mathbf{U}\ne\mathbf{U}'\}$ and $E_4=\{d(\mathbf{X}_2,\mathbf{X}'_2)\le 2(p-\bar{p})n-\epsilon n/8\}$. Using techniques from Section A.2 of [6], we can show for $\mathbf{y}_1\in A_0$,

$$P(E_3,E_4\mid\{\mathbf{Y}_1=\mathbf{y}_1\})\ge\frac{\epsilon}{48}\left(\frac{\epsilon}{5}\right)^{\frac{12}{\epsilon}-1}=\epsilon^{\mathcal{O}(1/\epsilon)}. \quad (11)$$

The bound in (11) together with the bounds $P(E_1)\ge 1-e^{-\Omega(\epsilon^2 n)}$, $P(E_2)\ge\epsilon/4$ implies that the probability of error under the proposed attack strategy is at least $\epsilon^{O(1/\epsilon)}$ which is independent of n. Details are given in [1].

In Fig. 3, we plot $\overline{C}(p,q)$ as a function of p for $q=0,0.1,0.2$. For a fixed q, there is a $\tilde{p}_q$ such that for $p\le\tilde{p}_q$, $\overline{C}(p,q)$ is convex and equal to $(1-h_2(p\star q))$, which is the capacity when BSC(p) and BSC(q) are in cascade. Thus when $p\le\tilde{p}_q$, the babble, snoop, and push strategy outlined here provides no benefit over a simpler adversarial strategy of injecting i.i.d. Ber(p) bit-flips.

B. An Achievable Rate $\underline{C}(p,q)$

Theorem 4. *The capacity $C(p,q)$ of BSC(q)-ADV(p)-FS is at least $\underline{C}(p,q)=\overline{C}\left((p\star q),0\right)$.*

Proof (Sketch): From [5], [7], $\overline{C}(s,0)$ is a tight characterization of the capacity when there is no BSC present and Calvin has a total budget of sn bit-flips. Since at channel use k, Calvin does not know if the BSC will cause a bit-flip, it can be shown that the total number of bit-flips is at most $\left((p\star q)+\epsilon\right)n$ with probability at least $1-e^{-\Omega(n\epsilon^2)}$. If we now assume that all of the $\left((p\star q)+\epsilon\right)n$ flips are chosen in an adversarial manner by Calvin, a rate of $\overline{C}((p\star q)+\epsilon,0)$ is achievable.

In Fig. 3, we also plot achievable rates $\underline{C}(p,q)$ for $q=0,0.1,0.2$. As noted before, the gap between upper and lower bounds increases with q.

V. Conclusion

In this work, we considered communicating over a stochastic channel (BEC/BSC) in the presence of a powerful adversary who can spy on both communicating terminals and inject further erasures/bit-flips at the input of the channel. For erasures, we gave a complete capacity characterization and for bit-flips, we proved interesting converse and achievability bounds. Future work includes characterizing capacity tightly for bit-flips with and without transmitter feedback. Another interesting direction is to characterize capacity in the case where the adversary has no feedback snooping.

REFERENCES

[1] V. Suresh, E. Ruzomberka, and D. J. Love, "Stochastic-adversarial channels : Online adversaries with feedback snooping," 2021. [Online]. Available: https://arxiv.org/abs/2104.07194

[2] M. Langberg, S. Jaggi, and B. K. Dey, "Binary causal-adversary channels," in *2009 IEEE International Symposium on Information Theory*. IEEE, 2009, pp. 2723–2727.

[3] B. K. Dey, S. Jaggi, and M. Langberg, "Codes against online adversaries," in *2009 47th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2009, pp. 1169–1176.

[4] B. K. Dey, S. Jaggi, M. Langberg, and A. D. Sarwate, "Improved upper bounds on the capacity of binary channels with causal adversaries," in *2012 IEEE International Symposium on Information Theory Proceedings*, 2012, pp. 681–685.

[5] B. K. Dey, S. Jaggi, M. Langberg, and A. D. Sarwate, "Upper bounds on the capacity of binary channels with causal adversaries," *IEEE Transactions on Information Theory*, vol. 59, no. 6, pp. 3753–3763, 2013.

[6] R. Bassily and A. Smith, "Causal erasure channels," in *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 2014, pp. 1844–1857.

[7] Z. Chen, S. Jaggi, and M. Langberg, "A characterization of the capacity of online (causal) binary channels," in *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, 2015, pp. 287–296.

[8] Z. Chen, S. Jaggi, and M. Langberg, "The capacity of online (causal) q -ary error-erasure channels," *IEEE Transactions on Information Theory*, vol. 65, no. 6, pp. 3384–3411, 2019.

[9] B. K. Dey, S. Jaggi, and M. Langberg, "Codes against online adversaries: Large alphabets," *IEEE Transactions on Information Theory*, vol. 59, no. 6, pp. 3304–3316, 2013.

[10] B. K. Dey, S. Jaggi, and M. Langberg, "Sufficiently myopic adversaries are blind," *IEEE Transactions on Information Theory*, vol. 65, no. 9, pp. 5718–5736, 2019.

[11] B. K. Dey, S. Jaggi, M. Langberg, and A. D. Sarwate, "Coding against delayed adversaries," in *2010 IEEE International Symposium on Information Theory*. IEEE, 2010, pp. 285–289.

[12] ——, "A bit of delay is sufficient and stochastic encoding is necessary to overcome online adversarial erasures," in *2016 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2016, pp. 880–884.

[13] B. K. Dey, S. Jaggi, M. Langberg, A. D. Sarwate, and C. Wang, "The interplay of causality and myopia in adversarial channel models," in *2019 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2019, pp. 1002–1006.

[14] Q. E. Zhang, M. Bakshi, and S. Jaggi, "Covert communication over adversarially jammed channels," in *2018 IEEE Information Theory Workshop (ITW)*. IEEE, 2018, pp. 1–5.

[15] A. J. Budkuley and S. Jaggi, "Communication over an arbitrarily varying channel under a state-myopic encoder," in *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 616–620.

[16] E. R. Berlekamp, "Block coding with noiseless feedback," Ph.D. dissertation, Massachusetts Institute of Technology, 1964.

[17] V. S. Lebedev, "Coding with noiseless feedback," *Problems of Information Transmission*, vol. 52, no. 2, pp. 103–113, 2016.

[18] K. Zigangirov, "On the number of correctable errors for transmission over a binary symmetrical channel with feedback," *Problemy Peredachi Informatsii*, vol. 12, no. 2, pp. 3–19, 1976.

[19] I. Csiszar and J. Körner, *Information theory: coding theorems for discrete memoryless systems*. Cambridge University Press, 2011.

[20] A. Lapidoth and P. Narayan, "Reliable communication under channel uncertainty," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2148–2177, 1998.