Network Equivalence for a Joint Compound-Arbitrarily-Varying Network Model

Oliver Kosut and Jörg Kliewer

***Abstract*—We consider the problem of finding the capacity of noisy networks under the presence of Byzantine adversaries, modeled by a joint compound channel and arbitrarily varying channel (AVC) model. This extends our earlier work which considers these models only in isolation. The motivation for this setup is that typically the adversary first selects an arbitrary subset of edges from the network and then specifies adversarial transmissions to each of the selected edges. We show that in some cases equivalence between this network and another network holds in the sense that for a fixed selection of adversarial edges the noisy links can be replaced by noiseless bit-pipes with a capacity equal to the random coding capacity of the corresponding AVC. In particular, the capacity region for the noisy network can be outer bounded by the intersection of the individual capacity regions for the noiseless case, for each adversarial edge selection. Moreover, if the network is fully connected, we also show that this upper bound is equivalent to the capacity of the noisy network. We also provide necessary and sufficient condition for full connectivity, making use of a new condition for an AVC termed overwritability.**

I. Introduction

Recently, reduction results in network information theory have been increasingly utilized to establish connections between problems without explicitly knowing their individual solutions. This is in particular useful for the challenging problem of finding suitable bounds for the capacity of general noisy networks. For example, reductions between noisy and noiseless networks have been studied [1], [2] in the sense that, in order to estimate the capacity region of a given noisy network, upper and lower bounding noiseless network instances can be constructed which include the capacity region of the original noisy network.

As one fundamental problem in networked communication is to achieve robustness against Byzantine adversaries, one can ask the question whether such reductions can also be used to derive capacity results by relating an adversarial problem to an equivalent, often easier to analyze, non-adversarial setting. Although network error correction and capacity results for the adversarial noiseless case have been presented, e.g., in [3]–[5], our goal in this work is to address equivalence results for the *noisy* adversarial case. In this setting, the action of the adversary creates noisy channels with time-varying states for the adversarial set [6], i.e., a certain subset of links the adversary has access to.

Our previous work [7], [8] addressed network equivalence for the case where one specific channel in the network is either a compound channel (CC) [9], [10] or an arbitrarily varying channel (AVC) [11]–[13]. In the CC case, the selected state is fixed over the whole transmission of a codeword, whereas in the AVC case, the channel state may vary in an arbitrary manner. Specifically, for the AVC the deterministic capacity is either zero or equals the random coding capacity [11]. We have shown in [7], [8] that if there exists a parallel network path of fixed capacity $C > 0$ between two nodes connected by an AVC, the capacity region of a noisy network with AVC state is equivalent to the capacity region of a noiseless and stateless network, in which the AVC link is replaced by a noiseless bitpipe with a capacity equal to its random coding capacity. This is due to the fact that the parallel capacity link allows to share common randomness with negligible rate between the two nodes connected by an AVC and therefore, communication is possible with non-zero rate [11]–[13]. Thus *connectivity* becomes an important question; i.e. whether messages consisting of small number of bits may be transmitted throughout the network. This condition appears fundamental to all problems with adversarial state, as we find in this work as well.

In this work we extend our previous equivalence results obtained for an *isolated* AVC/CC channel in a network [7], [8] to a joint CC and AVC model, in which the adversary chooses a set $\mathcal{Z}$ of at most z links, fixed across the code block (as a CC state); it then specifies adversarial transmissions to each of the edges in $\mathcal{Z}$ (as an AVC state). Such a model is more general as it captures the selection of adversarial sets by the adversary in addition to the adversarial transmission or injection of symbols. Moreover, this model incorporates a separation of time scales, as the selection of the adversarial set happens on a much larger time scale than the adversarial transmission itself.

We show in the following that for an arbitrary selection of adversarial edges the capacity region for the noisy network can be outer bounded by the intersection of the individual capacity regions for the noiseless (and stateless) case, each generated for a specific edge selection. If the network is fully connected, we further show that this upper bound is equivalent to the capacity of the noisy network, thus providing a simple characterization of this quantity. We also provide necessary and sufficient conditions for network connectivity, making use of a new condition for an AVC termed *overwritability*.

O. Kosut is with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ 85287 (email: okosut@asu.edu

J. Kliewer is with the Department of Electrical and Computer Engineering, New Jersey Institute of Technology, Newark, NJ 07102 (email: jkliewer@njit.edu)

This material is based upon work supported by the National Science Foundation under grants CCF-1422358 and CNS-1526547.

II. Model

For the remainder of this paper we consider networks $\mathcal{N}$ consisting of

- nodes $\mathcal{V} = \{1, \dots, m\}$,
- for each pair of nodes $i, j \in \mathcal{V}$, a point-to-point channel with state from node i to node j given by
$$p(y^{(i\to j)}|x^{(i\to j)}, s^{(i\to j)}) \tag{1}$$
with input $x^{(i\to j)} \in \mathcal{X}^{(i\to j)}$, output $y^{(i\to j)} \in \mathcal{Y}^{(i\to j)}$, and state $s^{(i\to j)} \in \mathcal{S}^{(i\to j)}$,
- for each channel, a special state value $s_0^{(i\to j)} \in \mathcal{S}^{(i\to j)}$ that we refer to as the *null state*.

Note that we assume each pair of nodes has a channel connecting them (in both directions); this is without loss of generality, because the absence of a channel can be modeled by taking $\mathcal{X}^{(i\to j)} = \mathcal{Y}^{(i\to j)} = \mathcal{S}^{(i\to j)} = \emptyset$.

A rate vector $\mathcal{R}$ consists of multicast rates $R^{(v\to U)}$ from each source node v to each destination set $U \subseteq \mathcal{V}$. For each (v, U) pair, there is a message $W^{(v\to U)} \in \mathcal{W}^{(v\to U)} = [2^{nR^{(v\to U)}}]$. For a subset of nodes $V \in \mathcal{V}$, let $W^{(V\to *)}$ denote the vector of all messages originating at nodes $v \in V$, and let $\mathcal{W}^{(V\to *)}$ denote the corresponding message set. Also let W denote the vector of all messages.

A blocklength-n solution (i.e. code) $\mathsf{S}^{(n)}$ for network $\mathcal{N}$ consists of the following:

- for each pair $i, j \in \mathcal{V}$ and $t \in \{1, \dots, n\}$ a causal encoding function
$$X_t^{(i\to j)} : \mathcal{W}^{(\{i\}\to *)} \times \prod_{k\in\mathcal{V}\setminus\{i\}} (\mathcal{Y}^{(k\to i)})^{t-1} \to \mathcal{X}^{(i\to j)},$$
- for each pair (v, U) and each $i \in U$, a decoding function
$$\hat{W}^{(v\to U),i} : \mathcal{W}^{(\{i\}\to *)} \times \prod_{k\in\mathcal{V}\setminus\{i\}} (\mathcal{Y}^{(k\to i)})^{n} \to \mathcal{W}^{(v\to U)}.$$

We allow any of the above functions to be randomized, but use functional notation for convenience. The input symbol sent into channel $i \to j$ at time t is given by
$$X_t^{(i\to j)} = X_t^{(i\to j)}(W^{\{i\}\to *}, Y_{1:t-1}^{(k\to i)} : \forall k \in \mathcal{V} \setminus \{i\}).$$
The estimate of message $W^{(v\to U)}$ produced at node $i \in U$ is given by
$$\hat{W}^{(v\to U),i} = \hat{W}^{(v\to U),i}(W^{\{i\}\to *}, Y_{1:n}^{(k\to i)} : \forall k \in \mathcal{V} \setminus \{i\}).$$
Let $\hat{W}$ be the complete vector of message estimates, and $\{\hat{W} \neq W\}$ denotes the event that any estimate is incorrect.

The states are chosen by the adversary as follows. The adversary selects a subset of links $\mathcal{Z} \subset \mathcal{V} \times \mathcal{V}$ where $|\mathcal{Z}| \leq z$. This set is fixed across the coding block. For all links $(i, j) \notin \mathcal{Z}$, the state for the corresponding channel is set to the null state, i.e. $S_t^{(i\to j)} = s_0^{(i\to j)}$ for $t = 1, \dots, n$. For links $(i, j) \in \mathcal{Z}$, the states $S_t^{(i\to j)}$ may be selected arbitrarily from $\mathcal{S}^{(i\to j)}$ for all t. We assume that the adversary may know the code, but does not have any access to run-time information.

Given a blocklength-n solution $\mathsf{S}^{(n)}$, let $P_e(\mathsf{S}^{(n)})$ be the probability of error, i.e. $\mathbb{P}(\hat{W} \neq W)$, maximized over all possible adversary actions: that is, the choice of controlled links $\mathcal{Z}$, as well as the choices of state $S_t^{(i\to j)}$ for $(i, j) \in \mathcal{Z}$ and $t \in \{1, \dots, n\}$. We say a rate vector $\mathcal{R}$ is *achievable* if there exists a sequence of solutions $\mathsf{S}^{(n)}$ with rates $\mathcal{R}$ such that $P_e(\mathsf{S}^{(n)}) \to 0$ as $n \to \infty$. We also define the capacity region $\mathscr{R}(\mathcal{N})$ to be the closure of the set of all achievable rate vectors.

Definition 1: In a network $\mathcal{N}$, we say an ordered pair of nodes (u, v) are *connected* if there exists a sequence of solutions with vanishing probability of error and $|\mathcal{W}^{(i\to j)}| \geq 2^{n^\gamma}$ for some $\gamma \in (0, 1)$. We say a network $\mathcal{N}$ is *fully connected* if each pair of nodes is connected in both directions.

Note that connectivity, as defined above, is significantly weaker for than requiring even positive rate transmitted from u to v. However, transmitting this asymptotically-little data will be enough for our purposes.

III. Outer Bound

For each pair (i, j), define the ordinary channel capacity of link $i \to j$ with no adversary present (i.e. null state) as
$$C^{(i\to j)} = \max_{p(x^{(i\to j)})} I(X^{(i\to j)}; Y^{(i\to j)}|S^{(i\to j)} = s_0^{(i\to j)}).$$
Further define the random coding capacity for the channel viewed as an AVC, as
$$C_R^{(i\to j)} = \min_{p(s^{(i\to j)})} \max_{p(x^{(i\to j)})} I(X^{(i\to j)}; Y^{(i\to j)}) \tag{2}$$
where in the mutual information random variables $X^{(i\to j)}$ and $Y^{(i\to j)}$ are distributed according to
$$\sum_{s^{(i\to j)}} p(x^{(i\to j)})p(s^{(i\to j)})p(y^{(i\to j)}|x^{(i\to j)}, s^{(i\to j)}).$$
Given a subset of links $\mathcal{Z} \subset \mathcal{V} \times \mathcal{V}$, let $\tilde{\mathcal{N}}_\mathcal{Z}$ be a noiseless and stateless network in which each channel $i \to j$ is replaced by a noiseless link of capacity $C^{(i\to j)}$ if $(i, j) \notin \mathcal{Z}$ and it is replaced by a noiseless link of capacity $C_R^{(i\to j)}$ if $(i, j) \in \mathcal{Z}$. Note that determining $\mathscr{R}(\tilde{\mathcal{N}}_\mathcal{Z})$ is purely a network coding problem. The following is a straightforward outer bound.

Theorem 1: $\mathscr{R}(\mathcal{N}) \subseteq \bigcap_{\mathcal{Z}\subset\mathcal{V}\times\mathcal{V}:|\mathcal{Z}|\leq z} \mathscr{R}(\tilde{\mathcal{N}}_\mathcal{Z})$.

Proof: The capacity region is outer bounded by the resulting capacity region under any specific adversary action. In particular, suppose the adversary controls a set of links $\mathcal{Z}$ with $|\mathcal{Z}| \leq z$, and for each $(i, j) \in \mathcal{Z}$, the adversary generates a random i.i.d. state sequence a distribution $p(s^{(i\to j)})$ achieving the minimum in (2). This effectively turns the network into a state-less network of independent point-to-point links, where links in $\mathcal{Z}$ have capacity $C_R^{(i\to j)}$ and links not in $\mathcal{Z}$ have capacity $C^{(i\to j)}$. Thus, by standard network equivalence results [1], the capacity region of the network under this adversary action is precisely $\mathscr{R}(\tilde{\mathcal{N}}_\mathcal{Z})$. Hence $\mathscr{R}(\mathcal{N}) \subseteq \mathscr{R}(\tilde{\mathcal{N}}_\mathcal{Z})$. As this holds for all $\mathcal{Z}$, the theorem follows. ∎

IV. Achievability

Theorem 2: If network $\mathcal{N}$ is fully connected, then
$$\mathscr{R}(\mathcal{N}) = \bigcap_{\mathcal{Z}\subset\mathcal{V}\times\mathcal{V}:|\mathcal{Z}|\leq z} \mathscr{R}(\tilde{\mathcal{N}}_\mathcal{Z}). \tag{3}$$

Proof: We only need to prove achievability, as Theorem 1 gives the converse bound. The achievable scheme is given in

Algorithm 1, but we first build up some machinery to use in the algorithm.

The assumption that $\mathcal{N}$ is fully connected implies that for each pair of nodes (i, j) there is a solution that allows reliable transmission of low-rate messages from node i to node j. Thus, in constructing an achievable scheme to prove this theorem, any data consisting of a number of bits asymptotically negligible compared to the messages may be transmitted from node i to node j essentially for free. In our scheme, three kinds of asymptotically negligible data will be transmitted:

- hashes, which will be used to check whether previously transmitted data was received correctly,
- coordination data, which will be used to ensure that all nodes are using the same network-wide code,
- establishment of common randomness between terminals of a symmetrizable AVC, so as to transmit at its random code capacity.

The hash data requires the use of a hashing function. We assume the existence of a hashing function ψ, that takes as input a data sequence $x^n \in \{0,1\}^n$, and a key realization k consisting of q bits. This function $\psi(x^n, k)$, itself consisting of q bits, should have the property that, for any $x^n \neq x'^n$, and a key K chosen uniformly at random

$$\mathbb{P}(\psi(x^n, K) = \psi(x'^n, K)) \to 0 \text{ as } q \to \infty. \tag{4}$$

There are a number of different functions that could be used; e.g. functions based on finite field polynomials (see [3]).

Fix a rate vector $\mathcal{R}$ contained in the RHS of (3). For each $\mathcal{Z} \subset \mathcal{V} \times \mathcal{V}$ with $|\mathcal{Z}| \leq z$, we show that there exists a sequence of solutions on network $\mathcal{N}$ that achieves rate vector $\mathcal{R}$ assuming $\mathcal{Z}$ is the set of adversarial links. In particular, since $\mathcal{R} \in \mathscr{R}(\tilde{\mathcal{N}}_{\mathcal{Z}})$, there exists a sequence of rate-$\mathcal{R}$ solutions, which we denote $\tilde{\mathsf{S}}^{(n)}(\mathcal{Z}, \mathcal{R})$ for each length n, on network $\tilde{\mathcal{N}}_{\mathcal{Z}}$, with vanishing probability of error. We convert each of these solutions sequences to a solution sequence $\mathsf{S}^{(n)}(\mathcal{Z}, \mathcal{R})$ on network $\mathcal{N}$ as follows. All coding operations of $\tilde{\mathsf{S}}^{(n)}(\mathcal{Z}, \mathcal{R})$ are preserved in $\mathsf{S}^{(n)}(\mathcal{Z}, \mathcal{R})$, except that the signals sent across noiseless links of $\tilde{\mathcal{N}}_{\mathcal{Z}}$ are instead conveyed by channel codes over the noisy links of $\mathcal{N}$ as follows:

- For each $(i, j) \notin \mathcal{Z}$, an ordinary channel code of rate $C^{(i\to j)}$ is used on link (i, j) in network $\mathcal{N}$ to convey the signal sent across the noiseless link in $\tilde{\mathcal{N}}_{\mathcal{Z}}$.
- For each $(i, j) \in \mathcal{Z}$, an AVC code of rate $C_R^{(i\to j)}$ is used on link (i, j) in $\mathcal{N}$ to convey the signal sent across the noiseless link in $\tilde{\mathcal{N}}_{\mathcal{Z}}$. In the case that link $(i, j) \in \mathcal{Z}$ is symmetrizable, we need an AVC code with common randomness to achieve $C_R^{(i\to j)}$; the common randomness can be established via parallel reliable transmission from node i to node j. By Lemmas 12.8 and 12.10 of [13], the $\mathcal{O}(\log n)$ bits of common randomness are required, which may be transmitted using the assumed low-rate solution from i to j. (Recall the connectivity assumption allows transmitting n^γ bits.)

Note that, if the set $\mathcal{Z}$ is in fact the set of adversarial links, then $\mathsf{S}^{(n)}(\mathcal{Z}, \mathcal{R})$ achieves rate vector $\mathcal{R}$ with vanishing probability of error. However, there still remains the issue of finding this set $\mathcal{Z}$. The scheme to do so is provided in Algorithm 1, which also defines some of the quantities employed in the proof below. In Algorithm 1, a set $\widehat{\mathcal{Z}}$ is maintained throughout the network giving the current belief of the set of adversarial links. This set is updated via low-rate coordination signals sent through the network.

Algorithm 1 Network Scheme

For each (v, U) pair, break message $W^{(v,U)}$ into L blocks $W_k^{(v\to U)}$ for $k = 1, \ldots, L$
$k \leftarrow 1$, $\widehat{\mathcal{Z}} \leftarrow \emptyset$
while $k \leq L$ **do**
  Perform solution $\mathsf{S}^{(n)}(\widehat{\mathcal{Z}}, \mathcal{R})$ to transmit messages $W_k^{(v\to U)}$, yielding message estimates $\widehat{W}_k^{(v\to U),i}$ for all (v, U) and $i \in U$
  for all $(i, j) \in \mathcal{V} \times \mathcal{V}$ **do**
    Let $V_k^{(i\to j)}$ and $\hat{V}_k^{(i\to j)}$ be the encoded and decoded signals for the point-to-point code on channel $i \to j$ using in solution $\mathsf{S}^{(n)}(\widehat{\mathcal{Z}}, \mathcal{R})$.
    At node i, generate key $K_k^{(i\to j)}$ and hash $H_k^{(i\to j)} = \psi(V_k^{(i\to j)}, K_k^{(i\to j)})$
    Reliably transmit $(K_k^{(i\to j)}, H_k^{(i\to j)})$ from node i to node j
    Node j checks if $H_k^{(i\to j)} = \psi(\hat{V}_k^{(i\to j)}, K_k^{(i\to j)})$
    If hash does not match, node j floods network with a signal indicating a hash violation on channel $i \to j$ All nodes update $\widehat{\mathcal{Z}} \leftarrow \widehat{\mathcal{Z}} \cup \{(i, j)\}$.
  if no hash violations occurred **then**
    For all (v, U) and $i \in U$, node i declares $\widehat{W}_k^{(v\to U),i}$ to be its estimate of message block $W_k^{(v,U)}$
    $k \leftarrow k + 1$

Proof of correctness: To prove that the solution described in Algorithm 1 achieves rate vector $\mathcal{R}$, we must show that the probability of error is arbitrarily small, and that the rate vector is arbitrarily close to $\mathcal{R}$. We do this via a series of claims as follows. In the following, we write "w.h.p." (with high probability) meaning with probability approaching 1 as the blocklength goes to infinity.

Claim 1: If $(i, j) \notin \mathcal{Z}$, then w.h.p. $\hat{V}_k^{(i\to j)} = V_k^{(i\to j)}$. This follows by the reliability of the ordinary channel code used on link $i \to j$.

Claim 2: If $(i, j) \in \mathcal{Z}$ and $(i, j) \in \widehat{\mathcal{Z}}$, then w.h.p. $\hat{V}_k^{(i\to j)} = V_k^{(i\to j)}$. This follows by the reliability of the AVC code used for link $i \to j$, in conjunction with common randomness sent on a parallel low-rate link.

Claim 3: If $(i, j) \in \mathcal{Z}$ and $\hat{V}_k^{(i\to j)} \neq V_k^{(i\to j)}$, then w.h.p. $H_k^{(i\to j)} \neq \psi(\hat{V}_k^{(i\to j)}, K_k^{(i\to j)})$. We show this by contradiction. If the adversary controls link $i \to j$, then it chooses state sequence $S_{1:n}^{(i\to j)}$ arbitrarily. However, this choice is made without knowledge of runtime events, and in particular is independent of the choice of key $K_k^{(i\to j)}$. Dropping superscripts $(i \to j)$ and subscripts k for convenience, we have

$$\mathbb{P}(H = \psi(\hat{V}, K) | \hat{V} \neq V) \tag{5}$$

$$\leq \max_{s_{1:n}} \mathbb{P}\left(\psi(V, K) = \psi(\hat{V}, K) \middle| S_{1:n} = s_{1:n}, \hat{V} \neq V\right) \tag{6}$$

$$= \max_{s_{1:n}} \sum_{v,\hat{v}:\hat{v}\neq v} \mathbb{P}(V = v, \hat{V} = \hat{v} | S_{1:n} = s_{1:n}, \hat{V} \neq V)$$

$$\cdot \mathbb{P}(\psi(v,K)=\psi(\hat{v},K)) \tag{7}$$

$$\leq \max_{v,\hat{v}:v\neq\hat{v}} \mathbb{P}(\psi(v,K)=\psi(\hat{v},K)) \tag{8}$$

where in (6) we have used the fact that $H_k^{(i\to j)} = \psi(V_k^{(i\to j)}, K_k^{(i\to j)})$. By the assumption on the hash function (4) and again the fact that the key $K_k^{(i\to j)}$ is independent of all other variables, (8) can be made arbitrarily small.

Claim 4: If $\hat{V}_k^{(i\to j)} \neq V_k^{(i\to j)}$ for any $(i,j) \in \mathcal{Z}$, then w.h.p. (i,j) will be added to $\widehat{\mathcal{Z}}$ in Step 12 of Algorithm 1. From Claim 3, if $\hat{V}_k^{(i\to j)} \neq V_k^{(i\to j)}$ then w.h.p. the corresponding hash will fail to match, meaning node j will flood the network with a hash violation signal. By the assumption of the existence of reliable low-rate transmission between any two nodes, these hash violation signals are received w.h.p. and all nodes agree that (i,j) is added to $\widehat{\mathcal{Z}}$.

Claim 5: If no hash violations occur, w.h.p. the messages are decoded correctly. Consider round k. If no hash violations occur, by Claim 4, for all $(i,j) \in \mathcal{Z}$, w.h.p. $\hat{V}_k^{(i\to j)} = V_k^{(i\to j)}$. Moreover, by Claim 1, for any $(i,j) \notin \mathcal{Z}$, w.h.p. $\hat{V}_k^{(i\to j)} = V_k^{(i\to j)}$. Thus, the network successfully simulates the noiseless network $\tilde{\mathcal{N}}_{\widehat{\mathcal{Z}}}$. By the reliability of solution $\tilde{\mathsf{S}}^{(n)}(\widehat{\mathcal{Z}},\mathcal{R})$, the round k messages are decoded correctly.

Claim 6: If a hash violation occurs on link (i,j), then w.h.p. $(i,j) \in \mathcal{Z}\setminus\widehat{\mathcal{Z}}$. By Claim 1, w.h.p. hash violations do not occur on links $(i,j) \notin \mathcal{Z}$. By Claim 2, w.h.p. hash violations do not occur on links $(i,j) \in \mathcal{Z}\cap\widehat{\mathcal{Z}}$. This leaves only links in $\mathcal{Z}\setminus\widehat{\mathcal{Z}}$.

Claim 7: For sufficiently large L and n, the achieved rate is arbitrarily close to $\mathcal{R}$. By Claim 6, hash violations occur in at most z rounds. By Claim 5, this means that w.h.p. all messages are eventually decoded correctly. Moreover, if L/z is sufficiently large, the $\leq z$ rounds with hash violations constitute an asymptotically negligible fraction of rate loss. Similarly, all hashes, coordination data, and common randomness signals take asymptotically negligible time for sufficiently large n. ■

V. Connectivity Conditions

Recall the definition from [12] that an AVC $p(y|x,s)$ is *symmetrizable* if there exists a distribution $p(s|x)$ such that

$$\sum_s p(y|x,s)p(s|x') = \sum_s p(y|x',s)p(s|x) \text{ for all } x,x',y.$$

It is shown in [12] that, in a symmetrizable AVC, the adversary can choose a state sequence that confuses the receiver between two messages, and that non-symmetrizable AVCs have deterministic capacity equal to their random code capacity.

Even though symmetrizable channels have 0 deterministic capacity, for most such channels the decoder can still detect the adversary. We define the notion of *overwritability* as the condition under which it is impossible to even detect the adversary. We say an AVC $p(y|x,s)$ with null state s_0 is *overwritable* if there exists a distribution $p(s|\tilde{x})$ such that

$$\sum_s p(s|\tilde{x})p(y|x,s) = p(y|\tilde{x},s_0) \text{ for all } x,\tilde{x},y. \tag{9}$$

That is, in an overwritable channel, the adversary can make it appear that it is not present, and that the input is whatever it chooses. Moreover, the following lemma asserts that if a channel is non-overwritable, then adversarial corruption can always be detected, if not corrected. This lemma will be used in the network setting to detect adversaries on non-overwritable channels.

Lemma 3: Consider a non-overwritable AVC. There exists a length-n code with $\binom{\sqrt{n}}{\sqrt{n}/2}$ messages such that, w.h.p., (i) if no adversary is present, the message is decoded correctly, and (ii) if an adversary is present, then either the message is decoded correctly or the decoder declares an error.

The proof of this lemma is given in Appendix A. Note that for any $\gamma < 1/2$, $\binom{\sqrt{n}}{\sqrt{n}/2} \geq 2^{n^\gamma}$ for sufficiently large n, so the code satisfies the connectivity condition.

The following proposition is not hard to prove, so the proof is omitted for space.

Proposition 4: The following properties are ordered from strongest to weakest (i.e. each implies the next):

1) An AVC has zero no-adversary capacity.
2) An AVC is overwritable.
3) An AVC's random code capacity is zero.
4) An AVC is symmetrizable.

We now develop necessary and sufficient conditions for connectivity. Let $\mathcal{E} = \mathcal{V}\times\mathcal{V}$ be the complete set of point-to-point links. We divide this set into four categories:

1) Links $i\to j$ that are non-symmetrizable.
2) Links $i \to j$ that are symmetrizable but non-overwritable.
3) Links $i\to j$ that are overwritable but have positive no-adversary capacity $C^{(i\to j)}$.
4) Links $i\to j$ with $C^{(i\to j)} = 0$.

Let $\mathcal{E}_1,\mathcal{E}_2,\mathcal{E}_3,\mathcal{E}_4$ be the set of pairs (i,j) satisfying each of the above conditions respectively. Let

$$w_1 = \infty, \quad w_2 = 1, \quad w_3 = \tfrac{1}{2}, \quad w_4 = 0$$

and define a weight function on $\mathcal{E}$ as $w(i,j) = w_\ell$ if $(i,j) \in \mathcal{E}_\ell$. The motivation for this weight function is that, for channels in $\mathcal{E}_1$, the adversary cannot corrupt no matter what it does; in $\mathcal{E}_2$, adversarial corruption can be detected, so it effectively causes erasures; in $\mathcal{E}_3$, the adversary cannot be detected, so it can cause arbitrary errors; in $\mathcal{E}_4$, the capacity is zero even without the adversary. Thus, while an adversary must control a non-overwriteable channel in order to avoid a message getting through, it only needs to control half of the overwritable channels to do the same, because it can overwrite the legitimate messages with counterfeits, thereby confusing the receiver about which message is the true one. This is formalized in the following theorem.

Theorem 5: The pair (u,v) is connected if and only if $\text{min-cut}(u,v) > z$, where $\text{min-cut}(u,v)$ is the min-cut from u to v on the weighted directed graph $(\mathcal{V},\mathcal{E},w)$.

Proof: **Converse:** Suppose $\text{min-cut}(u,v) \leq z$. We prove that no solution can have vanishing probability of error. Let $\mathcal{Y}$ be a set of links crossing a min-cut. Note $\mathcal{Y}\cap\mathcal{E}_1 = \emptyset$, or else the min-cut would be infinite. Let $J = |\mathcal{Y}\cap\mathcal{E}_3|$. Thus $|\mathcal{Y}\cap\mathcal{E}_2| = \text{min-cut}(u,v) - \frac{J}{2} \leq z - \frac{J}{2}$. Thus there exists a set $\mathcal{Z}$ of z links containing $\mathcal{Y}\cap\mathcal{E}_2$ and $J/2$ links in $\mathcal{Y}\cap\mathcal{E}_3$. The adversary may control these links and do the following.

Since it knows the code, it generates an independent copy of all signals transmitted in the network. On links in $\mathcal{E}_2$, it uses symmetrizability to confuse the receiver about which scenario is the real one. On links in $\mathcal{E}_3$ in uses overwritability to make it appear that the false scenario is the real one. To all nodes on the down-stream side of the cut, it is impossible to tell which scenario is real and which is false, resulting in an error probability of at least $1/2$.

Achievability: Suppose $\text{min-cut}(u,v) > z$. Note that $\text{min-cut}(u,v)$ must be a multiple of $1/2$. Thus there exists a set of $2\,\text{min-cut}(u,v)$ paths from u to v, where the number of paths passing through (i,j) is at most $2w(i,j)$. Also, no paths contain links in $\mathcal{E}_4$, as they have zero weight. Let J be the number of paths containing a link in $\mathcal{E}_3$. The remaining $2\,\text{min-cut}(u,v) - J$ paths, containing only links in $\mathcal{E}_1 \cup \mathcal{E}_2$, may be assumed to be doubled-up. That is, we may assume there are $\text{min-cut}(u,v) - J/2$ paths, each of which appears twice, in addition to the J paths containing edges in $\mathcal{E}_3$. We claim that, for any set of z links controlled by the adversary, either (i) one of the $\text{min-cut}(u,v) - J/2$ paths through $\mathcal{E}_1 \cup \mathcal{E}_2$ are not controlled, or (ii) more than half of the J paths containing links in $\mathcal{E}_3$ are uncontrolled. This is because, to control all $\text{min-cut}(u,v) - J/2$ paths through $\mathcal{E}_1 \cup \mathcal{E}_2$ leaves only $z - (\text{min-cut}(u,v) - J/2) < J/2$ links.

We use the following strategy: Transmit a message simultaneously on each of these paths from u to v, using the codes from Lemma 3 for links in $\mathcal{E}_1 \cup \mathcal{E}_2$, and standard channel codes assuming no adversary for links in $\mathcal{E}_3$. For any path in $\mathcal{E}_1 \cup \mathcal{E}_2$, adversarial corruption can be detected by Lemma 3, because all these channels are non-overwritable. If the message on any path gets through without a detection flag, node v can decode. Otherwise, node v uses majority rule on the remaining paths with links in $\mathcal{E}_3$. As the adversary can influence strictly less than half of these links, majority rule is guaranteed to yield the correct message. ∎

References

[1] R. Koetter, M. Effros, and M. Medard, "A theory of network equivalence—Part I: Point-to-point channels," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 972–995, 2011.

[2] ——, "A theory of network equivalence—Part II: Multiterminal channels," *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 3709–3732, July 2014.

[3] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, M. Medard, and M. Effros, "Resilient network coding in the presence of Byzantine adversaries," *IEEE Trans. Inf. Theory*, vol. 54, no. 6, Jun. 2008.

[4] S. Kim, T. Ho, M. Effros, and S. Avestimehr, "Network error correction with unequal link capacities," in *Proc. 47th Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2009, pp. 1387–1394.

[5] O. Kosut, L. Tong, and D. N. C. Tse, "Polytope codes against adversaries in networks," *IEEE Transactions on Information Theory*, vol. 60, no. 6, pp. 3308–3344, June 2014.

[6] A. Lapidoth and P. Narayan, "Reliable communication under channel uncertainty," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2148–2177, Oct. 1998.

[7] O. Kosut and J. Kliewer, "Equivalence for networks with adversarial state," in *Proc. IEEE Int. Sympos. on Inform. Theory*, Honolulu, HI, Jun. 2014, pp. 2401–2405.

[8] ——, "Equivalence for networks with adversarial state," [Online] arXiv.org, arXiv:1404.6701, Jan. 2015.

[9] D. Blackwell, L. Breiman, and A. J. Thomasian, "The capacity of a class of channels," *The Annals of Mathematical Statistics*, vol. 30, no. 4, pp. 1229–1241, Dec. 1959.

[10] J. Wolfowitz, *Coding Theorems of Information Theory*. Springer Verlag, 1978.

[11] R. Ahlswede, "Elimination of correlation in random codes for arbitrarily varing channels," *Probability Theory and Related Fields*, vol. 44, no. 2, pp. 159–175, 1978.

[12] I. Csiszár and P. Narayan, "The capacity of the arbitrarily varying channel revisited: positivity, constraints," *IEEE Trans. Inf. Theory*, vol. 34, no. 2, pp. 181–193, 1988.

[13] I. Csiszár and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, 2011.

Appendix A
Proof of Lemma 3

By the assumption of non-overwritability, there exist two letters $x_1, x_2 \in \mathcal{X}$ such that for all $p(s)$,

$$\sum_s p(s)p(y|x_1,s) \neq p(y|x_2,s_0) \text{ for some } y. \qquad (10)$$

Note that $x_1 \neq x_2$, because if they were equal we could simply take $p(s) = \mathbf{1}(s = s_0)$ and (10) would not hold.

Codebook: The codebook is given by the set of all $\sqrt{n}$-length sequences of letters in $\{x_1, x_2\}$, where x_1 and x_2 occur equally often. (Assume $\sqrt{n}$ is an even integer.) Denote the mth codeword $x^{\sqrt{n}}(m)$ for $m \in \{1, \ldots, \binom{\sqrt{n}}{\sqrt{n}/2}\}$.

Encoding: Given message m, transmit $x_1(m)$ $\sqrt{n}$ times, followed by $x_2(m)$ $\sqrt{n}$ times, and so on.

Decoding: Divide output sequence Y^n into $\sqrt{n}$ subsequences of length $\sqrt{n}$. For each $i \in \{1, \ldots, \sqrt{n}\}$, let $p_i(y)$ be the type of the ith length-$\sqrt{n}$ subsequence. Let $\hat{m}$ be the smallest message such that

$$\|p_i(y) - p(y|x_i(\hat{m}), s_0)\| \leq \epsilon \text{ for } i = 1, \ldots, \sqrt{n}$$

where $\|\cdot\|$ denotes infinity-norm on the joint distribution. If there is no such message, declare an error.

Probability of error analysis: Suppose no adversary is present. By large deviation results, there exists a positive constant D such that

$$\mathbb{P}(\|p_i(y) - p(y|x_i(m), s_0)\| > \epsilon) \leq 2^{-D\sqrt{n}}.$$

Thus, the probability of making an error is at most $\sqrt{n}2^{-D\sqrt{n}}$ which can be made arbitrarily small for sufficiently large n.

Now suppose an adversary is present and chooses state sequence $S^n = s^n$. We need to show that, if the decoder does not declare an error, w.h.p. the message is decoded correctly. Let $p_i(s)$ be the type of the ith block of s^n of length $\sqrt{n}$. Thus

$$\mathbb{P}\Big(\Big\|p_i(y) - \sum_s p_i(s)p(y|x_i(m),s)\Big\| > \epsilon\Big) \leq 2^{-D\sqrt{n}}.$$

Thus, if no error is declared, there must be some $\tilde{x}_i \in \{x_1, x_2\}$ such that

$$\Big\| \sum_s p_i(s)p(y|x_i(m),s) - p(y|\tilde{x}_i, s_0)\Big\| \leq 2\epsilon.$$

However, by (10), for sufficiently small ϵ, this cannot hold if $x_i(m) = x_1$ and $\tilde{x}_i = x_2$. Thus, all letters where $x_i(m) = x_1$ are decoded correctly w.h.p. Moreover, if even one letter where $x_i(m) = x_2$ is decoded incorrectly, then by the assumption that all codewords have the same type, the estimated sequence will not correspond to any codeword, so an error is declared.