The Curious Case of the Diamond Network

Allison Beemer and Alberto Ravagnani

University of Wisconsin-Eau Claire (beemera@uwec.edu)
Eindhoven University of Technology (a.ravagnani@tue.nl)

Abstract. We initiate the study of the one-shot capacity of communication networks with an adversary having access only to a proper subset of network edges. We introduce the Diamond Network as a minimal example to show that known cut-set bounds are not sharp in general, and that their non-sharpness comes precisely from restricting the action of the adversary to a region of the network. We also give a capacity-achieving scheme for the Diamond Network that implements an adversary detection strategy. Finally, we give a sufficient condition for tightness of the Singleton Cut-Set Bound in a family of two-level networks.

Keywords: Network coding, adversarial network, capacity, cut-set bound.

1 Introduction

As the prevalence of interconnected devices grows, vulnerable communication networks must be able to counter the actions of malicious actors; a unified understanding of the fundamental communication limits of these networks is therefore paramount. The correction of errors introduced by adversaries in networks has been studied in a number of previous works. Cai and Yeung give generalizations of some classical coding bounds to the network setting in [1, 2]. Other bounds and related code constructions for adversarial networks are presented in, e.g., [3–8]. The work most closely related to this paper is [8], where a unified combinatorial framework for adversarial networks and a method for porting point-to-point coding-theoretic results to the network setting are established. In contrast to works that address random errors in networks, or a combination of random and adversarial errors, [8] focuses purely on adversarial, or *worst-case*, errors. The results presented here assume the same model in a single-use regime.

Problem formulation. In contrast to most previous work, in this paper we concentrate on networks with an adversary who can possibly corrupt only a *proper subset* of the network edges. This paper is the first stepping stone of a long-term project aimed at understanding how the topology of the vulnerable region of a network determines (or at least affects) its capacity.

We focus on networks whose inputs are drawn from a finite alphabet and whose intermediate nodes may process information before forwarding. We assume that an omniscient adversary can corrupt up to some fixed number of alphabet symbols sent along a subset of network edges. The one-shot capacity of such an adversarial network measures the number of symbols that can be sent with zero error during a single transmission round. A universal approach to forming cut-set bounds, which are derived by reducing the capacity problem to a minimization across cut-sets of the underlying directed graph of the network, is presented in [8]. Any coding-theoretic bound may be ported to the networking setting, including the famous Singleton Bound.

Our contribution. In this paper, we exhibit a minimal example showing that known cut-set bounds for the one-shot capacity of a network subject to adversarial noise are not sharp in general. More precisely, we construct a network for which the Singleton Bound gives the best established upper bound on one-shot capacity, and show that it is not tight (regardless of the size of the network alphabet). The non-tightness of the bound comes precisely from limiting the adversary to operation on a certain region of the network. Our example, which we call the *Diamond Network*, requires that a single symbol be sacrificed to the task of locating the adversary within the network. Interestingly, this requirement results in a non-integer-valued one-shot capacity (which we are able to compute).

We note that the requirement that the receiver locate the adversary is related to the problem of authentication in networks (see, e.g. [9–11]). In our capacity-achieving scheme for the Diamond Network, one intermediate vertex must be able to either sound an alarm (if the adversary is detected), or decode correctly (when the adversary is absent). On the other hand, in our presented scheme for a modification of the Diamond Network, called the *Mirrored Diamond Network*, the way in which intermediate vertices sound the alarm must simultaneously serve as the way in which a particular alphabet symbol is transmitted. This interplay between authentication and correction is reminiscent of the work in [11], and the connection warrants further investigation.

Outline. This paper is organized as follows. In Section 2 we introduce necessary notation and background. Sections 3 and 4 together establish the exact one-shot capacity of the Diamond Network, proving that the Singleton Cut-Set Bound is not tight. In Section 5, we establish the (bound-achieving) one-shot capacity of the Mirrored Diamond Network. Section 6 expands our focus to the broader class of two-level networks, and gives a sufficient condition for a network in this class to meet the best cut-set bound. We conclude with future directions in Section 7.

2 Preliminaries

We introduce the terminology and notation for the remainder of the paper. We start by formally defining communication networks as in [8].

Definition 1. *A* ***(single-source communication) network*** *is a 4-tuple* $\mathscr{N} = (\mathscr{V}, \mathscr{E}, S, \mathbf{T})$*, where:*

(A) $(\mathscr{V}, \mathscr{E})$ *is a finite, directed and acyclic multigraph;*
(B) $S \in \mathscr{V}$ *is the* ***source****;*

(C) $\mathbf{T} \subseteq \mathscr{V}$ *is the set of* ***terminals****.*

We also assume the following:

(D) $|\mathbf{T}| \geqslant 1$ *and* $S \notin \mathbf{T}$*;*
(E) there exists a directed path from S *to any* $T \in \mathbf{T}$*;*
(F) for every $V \in \mathscr{V} \backslash (\{S\} \cup \mathbf{T})$ *there exists a directed path from* S *to* V *and from* V *to some terminal* $T \in \mathbf{T}$*.*

The elements of $\mathscr{V}$ *are called* ***vertices*** *or* ***nodes****, and those of* $\mathscr{E}$ *are called* ***edges****. The elements of* $\mathscr{V} \backslash (\{S\} \cup \mathbf{T})$ *are the* ***intermediate vertices/nodes****. The set of incoming and outgoing edges for a vertex* V *are denoted by* $\mathrm{in}(V)$ *and* $\mathrm{out}(V)$*, respectively. Their cardinalities are the* ***indegree*** *and* ***outdegree*** *of* V*, which are denoted by* $deg^{-}(V)$ *and* $deg^{+}(V)$*, respectively.*

Our communication model is as follows: all edges of a network $\mathscr{N}$ can carry precisely one element from a set $\mathscr{A}$ of cardinality at least 2, which we call the **alphabet**. The vertices of the network collect alphabet symbols over the incoming edges, process them according to functions, and send the outputs over the outgoing edges. Vertices are memoryless and transmissions are delay-free. We model errors as being introduced by an adversary $\mathbf{A}$, who can corrupt the value of up to t edges from a fixed set $\mathscr{U} \subseteq \mathscr{E}$. An alphabet symbol sent along one of the edges in $\mathscr{U}$ can be changed to any other alphabet symbol at the discretion of the adversary. In particular, the noise we consider is *not* probabilistic in nature, but rather worst-case: we focus on correcting *any* error pattern that can be introduced by the adversary. We call the pair $(\mathscr{N}, \mathbf{A})$ an **adversarial network**.

It is well-known that an acyclic directed graph $(\mathscr{V}, \mathscr{E})$ defines a partial order on the set of its edges, $\mathscr{E}$. More precisely, $e_1 \in \mathscr{E}$ **precedes** $e_2 \in \mathscr{E}$ (in symbols, $e_1 \preccurlyeq e_2$) if there exists a directed path in $(\mathscr{V}, \mathscr{E})$ whose first edge is e_1 and whose last edge is e_2. We may extend this partial order to a total order on $\mathscr{E}$, which we fix once and for all and denote by $\leqslant$. Important to note is that the results in this paper do not depend on the particular choice of $\leqslant$.

Definition 2. *Let* $\mathscr{N} = (\mathscr{V}, \mathscr{E}, S, \mathbf{T})$ *be a network. A* ***network code*** $\mathscr{F}$ *for* $\mathscr{N}$ *is a family of functions* $\{\mathscr{F}_V \mid V \in \mathscr{V} \backslash (\{S\} \cup \mathbf{T})\}$*, where* $\mathscr{F}_V : \mathscr{A}^{deg^{-}(V)} \to \mathscr{A}^{deg^{+}(V)}$ *for all* V*.*

A network code $\mathscr{F}$ describes how the vertices of a network $\mathscr{N}$ process the inputs received on the incoming edges. There is a unique interpretation for these operations thanks to the choice of the total order $\leqslant$.

Definition 3. *Let* $\mathscr{N} = (\mathscr{V}, \mathscr{E}, S, \mathbf{T})$ *be a network and let* $\mathscr{U}, \mathscr{U}' \subseteq \mathscr{E}$ *be non-empty subsets. We say that* $\mathscr{U}$ ***precedes*** $\mathscr{U}'$ *if every path from* S *to an edge of* $\mathscr{U}'$ *contains an edge from* $\mathscr{U}$*.*

Our next step is to define outer codes for a network and give necessary and sufficient conditions for decodability. We do this by introducing the notion of an adversarial channel as proposed in [8, Section IV.B].

Notation 1. *Let $(\mathscr{N},\mathbf{A})$ be an adversarial network with $\mathscr{N}=(\mathscr{V},\mathscr{E},S,\mathbf{T})$ and let $\mathscr{U},\mathscr{U}'\subseteq\mathscr{E}$ be non-empty such that $\mathscr{U}$ precedes $\mathscr{U}'$. Let $\mathscr{F}$ be a network code for $\mathscr{N}$. For $\mathbf{x}\in\mathscr{A}^{|\mathscr{U}|}$, we denote by*

$$\Omega[\mathscr{N},\mathbf{A},\mathscr{F},\mathscr{U}\to\mathscr{U}'](\mathbf{x})\subseteq\mathscr{A}^{|\mathscr{U}'|} \tag{1}$$

the set of vectors over the alphabet that can be exiting the edges of $\mathscr{U}'$ when:

- *the coordinates of $\mathbf{x}$ are the alphabet values entering the edges of $\mathscr{U}$,*
- *vertices process information according to $\mathscr{F}$,*
- *everything is interpreted according to the total order $\leqslant$.*

Note that (1) *is well-defined because $\mathscr{U}$ precedes $\mathscr{U}'$. Furthermore, $\mathscr{U}\cap\mathscr{U}'$ need not be empty. We refer to the discussion following [8, Definition 41]; see also [8, Example 42].*

Example 1. Let $(\mathscr{N},\mathbf{A})$ be the network in Figure 1, where the edges are ordered according to their indices. We consider an adversary capable of corrupting up to one of the dashed edges. At intermediate node V_1, let $\mathscr{F}_{V1}$ be the identity function; at V_2, let $\mathscr{F}_{V2}$ be the projection onto the second coordinate of the input pair. Then, for example, for $\mathbf{x}=(x_1,x_2,x_3)\in\mathscr{A}^3$ we have that

$$\Omega[\mathscr{N},\mathbf{A},\mathscr{F},\{e_1,e_2,e_3\}\to\{e_4,e_5\}](\mathbf{x})\subseteq\mathscr{A}^2$$

is the set of vectors $\mathbf{y}=(y_1,y_2)\in\mathscr{A}^2$ for which $d_{\mathrm{H}}((y_1,y_2),(x_1,x_3))\leqslant 1$, where d_{H} denotes the Hamming distance.

We now define error-correcting codes in the context of adversarial networks. Informally, codes are comprised of the alphabet vectors that may be emitted by the source.

Definition 4. *An **(outer) code** for a network $\mathscr{N}=(\mathscr{V},\mathscr{E},S,\mathbf{T})$ is a subset $C\subseteq\mathscr{A}^{\deg^+(S)}$ with $|C|\geqslant 1$. If $\mathscr{F}$ is a network code for $\mathscr{N}$ and $\mathbf{A}$ is an adversary, then we say that C is **unambiguous** (or **good**) for $(\mathscr{N},\mathbf{A},\mathscr{F})$ if for all $\mathbf{x},\mathbf{x}'\in C$ with $\mathbf{x}\neq\mathbf{x}'$ and for all $T\in\mathbf{T}$ we have*

$$\Omega[\mathscr{N},\mathbf{A},\mathscr{F},\mathrm{out}(S)\to\mathrm{in}(T)](\mathbf{x})\cap\Omega[\mathscr{N},\mathbf{A},\mathscr{F},\mathrm{out}(S)\to\mathrm{in}(T)](\mathbf{x}')=\varnothing.$$

The last condition in the above definition guarantees that every element of C can be uniquely recovered by every terminal, despite the action of the adversary. Finally, we define the one-shot capacity of an adversarial network.

Definition 5. *The **(one-shot) capacity** of an adversarial network $(\mathscr{N},\mathbf{A})$ is the maximum $\alpha\in\mathbb{R}$ for which there exists a network code $\mathscr{F}$ and an unambiguous code C for $(\mathscr{N},\mathbf{A},\mathscr{F})$ with $\alpha=\log_{|\mathscr{A}|}(|C|)$. We denote this maximum value by $\mathrm{C}_1(\mathscr{N},\mathbf{A})$.*

In [8], a general method was developed to "lift" bounds for Hamming-metric channels to the networking context. The method allows any classical coding bound to be lifted to the network setting. The next result states the lifted version of the well-known Singleton Bound. Recall that an edge-cut between source S and terminal T is a set of edges whose removal would separate S from T.

Theorem 2 (The Singleton Cut-Set Bound). *Let $\mathscr{N}$ be a network with edge set $\mathscr{E}$. Assume an adversary* $\mathbf{A}$ *can corrupt up to* $t \geqslant 0$ *edges from a subset* $\mathscr{U} \subseteq \mathscr{E}$. *Then*

$$\mathrm{C}_1(\mathscr{N}, \mathbf{A}) \leqslant \min_{T \in \mathbf{T}} \min_{\mathscr{E}'} \left(|\mathscr{E}' \backslash \mathscr{U}| + \max\{0, |\mathscr{E}' \cap \mathscr{U}| - 2t\} \right),$$

where $\mathscr{E}' \subseteq \mathscr{E}$ *ranges over all edge-cuts between* S *and* T.

3 The Diamond Network: Achievability

We present a minimal example of a network for which the *best* bound in [8], namely the Singleton Cut-Set Bound, is not sharp. The example will serve to illustrate the necessity of performing *partial* decoding at the intermediate nodes in order to achieve capacity.

Example 2 (The Diamond Network). The network $\mathscr{D}$ of Figure 1 has one source S, one terminal T, and two intermediate vertices V_1 and V_2. The vertices are connected as in the figure. We consider an adversary $\mathbf{A}_{\mathscr{D}}$ able to corrupt at most one of the dashed edges, and we call the pair $(\mathscr{D}, \mathbf{A}_{\mathscr{D}})$ the **Diamond Network**.

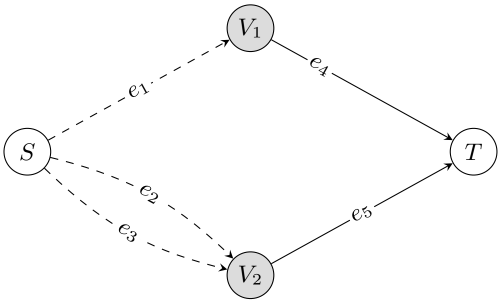

Fig. 1. The Diamond Network

For the Diamond Network, the best bound among those proved in [8] is the Singleton Cut-Set Bound, given below.

Corollary 1. *For the Diamond Network* $(\mathscr{D}, \mathbf{A}_{\mathscr{D}})$, $\mathrm{C}_1(\mathscr{D}, \mathbf{A}_{\mathscr{D}}) \leqslant 1$.

We will prove in this section and the next that the Diamond Network has capacity

$$\mathrm{C}_1(\mathscr{D}, \mathbf{A}_{\mathscr{D}}) = \log_{|\mathscr{A}|}(|\mathscr{A}| - 1). \tag{2}$$

In particular, this shows that the bounds of [8] are not sharp. The results presented in the remainder of the paper offer an intuitive explanation for why Corollary 1 is not sharp: in order to achieve the capacity of the Diamond Network, one alphabet symbol needs to be reserved to implement an adversary detection strategy. We will elaborate on this idea following the proof of achievability, given below.

Proposition 1. *For the Diamond Network* $(\mathscr{D}, \mathbf{A}_{\mathscr{D}})$ *we have*

$$\mathrm{C}_1(\mathscr{D}, \mathbf{A}_{\mathscr{D}}) \geqslant \log_{|\mathscr{A}|}(|\mathscr{A}| - 1).$$

Proof. We isolate a symbol $* \in \mathscr{A}$ and define $\mathscr{A}' = \mathscr{A} \backslash \{*\}$. Consider the scheme where the source S can send any symbol of $\mathscr{A}'$ via a three-times repetition code over its outgoing edges. Vertex V_1 simply forwards the received input, while vertex V_2 proceeds as follows: If the two received inputs coincide and are equal to $a \in \mathscr{A}'$, then it forwards a. Otherwise, it transmits $*$. It is not difficult to check that any symbol from $\mathscr{A}'$ can be uniquely decoded, showing that the proposed scheme is unambiguous. This concludes the proof. □

Notice that the communication strategy on which the previous proof is based reserves an alphabet symbol $* \in \mathscr{A}$ to pass information about the location of the adversary (more precisely, the symbol $*$ reveals whether or not the adversary is acting on the lower "stream" of the Diamond Network). The source is not allowed to emit the reserved symbol $*$, rendering $\log_{|\mathscr{A}|}(|\mathscr{A}| - 1)$ the maximum rate achievable by this scheme.

It is natural to then ask whether the reserved symbol $*$ can simultaneously be a part of the source's codebook, achieving a rate of $1 = \log_{|\mathscr{A}|}(|\mathscr{A}|)$ transmitted message per single channel use. In the next section, we will formally answer this question in the negative; see Proposition 2. In Section 5, we consider a modification of the Diamond Network and present a scheme where one symbol is reserved for adversary detection, but can nonetheless also be used as a message symbol.

4 The Diamond Network: The Converse

In this section, we establish an inequality for the cardinality of any unambiguous code C for the Diamond Network. The inequality is quadratic in the code's size and implies that $|C| \leqslant |\mathscr{A}| - 1$. Together with Proposition 1, this computes the exact capacity of $(\mathscr{D}, \mathbf{A}_{\mathscr{D}})$.

Proposition 2. *Let* $\mathscr{F}$ *be a network code for* $(\mathscr{D}, \mathbf{A}_{\mathscr{D}})$ *and let* $C \subseteq \mathscr{A}^3$ *be an outer code. If* C *is unambiguous for* $(\mathscr{D}, \mathbf{A}_{\mathscr{D}}, \mathscr{F})$*, then*

$$|C|^2 + |C| - 1 - |\mathscr{A}|^2 \leqslant 0.$$

In particular, we have $|C| \leqslant |\mathscr{A}| - 1$.

Proof. The argument is organized into various claims. We denote by $\pi : \mathscr{A}^3 \to \mathscr{A}$ the projection onto the first coordinate. The proofs for the following three claims are included in the full version of this work; all hinge on the assumption that C is unambiguous.

Claim A. *We have $|\pi(C)| = |C|$.*

Claim B. *The restriction of $\mathscr{F}_{V_1}$ to $\pi(C)$ is injective.*

To simplify notation, let $\Omega := \Omega[\mathscr{D}, \mathbf{A}_{\mathscr{D}}, \mathscr{F}, \{e_1, e_2, e_3\} \to \{e_5\}]$, which is well-defined because $\{e_1, e_2, e_3\}$ precedes e_5; see Definition 3.

Claim C. *There exists at most one codeword $\mathbf{x} \in C$ for which the cardinality of $\Omega(\mathbf{x})$ is 1.*

To simplify further, denote the transfer from S to T by

$$\Omega'' := \Omega[\mathscr{D}, \mathbf{A}_{\mathscr{D}}, \mathscr{F}, \{e_1, e_2, e_3\} \to \{e_4, e_5\}].$$

Since C is unambiguous, we have

$$\sum_{\mathbf{x} \in C} |\Omega''(\mathbf{x})| \leqslant |\mathscr{A}|^2. \tag{3}$$

For all $\mathbf{x} \in C$, write $\Omega''(\mathbf{x}) = \Omega''_1(\mathbf{x}) \cup \Omega''_2(\mathbf{x})$, where

$$\begin{aligned} \Omega''_1(\mathbf{x}) &= \{\mathbf{z} \in \Omega''(\mathbf{x}) \mid z_1 = \mathscr{F}_{V_1}(x_1)\}, \\ \Omega''_2(\mathbf{x}) &= \{\mathbf{z} \in \Omega''(\mathbf{x}) \mid z_2 = \mathscr{F}_{V_2}(x_2, x_3)\}. \end{aligned}$$

By definition, we have $|\Omega''(\mathbf{x})| = |\Omega''_1(\mathbf{x})| + |\Omega''_2(\mathbf{x})| - 1$. Summing over all $\mathbf{x} \in C$ and using Claims A, B and C we find

$$\begin{aligned} \sum_{\mathbf{x} \in C} |\Omega''(\mathbf{x})| &\geqslant 1 + 2(|C| - 1) + \sum_{\mathbf{x} \in C} |C| - |C| \\ &= 2|C| - 1 + |C|^2 - |C| \\ &= |C|^2 + |C| - 1. \end{aligned}$$

Combining this with (3), we find $|C|^2 + |C| - 1 \leqslant |\mathscr{A}|^2$, which is the desired inequality. □

We can now compute the capacity of the Diamond Network by combining Propositions 1 and 2.

Theorem 3. *For the Diamond Network $(\mathscr{D}, \mathbf{A}_{\mathscr{D}})$, $\mathrm{C}_1(\mathscr{D}, \mathbf{A}_{\mathscr{D}}) = \log_{|\mathscr{A}|}(|\mathscr{A}| - 1)$.*

The Diamond Network is admittedly a small example. However, we believe that it will provide valuable insight into the general behavior of the one-shot capacity of larger networks.

5 The Mirrored Diamond Network

It is interesting to observe that by adding a single vulnerable edge from S to V_1 in the Diamond Network (as in Figure 2), the capacity will be exactly the one predicted by the Singleton Cut-Set Bound of Theorem 2. We call this new network the **Mirrored Diamond Network**. The adversary can corrupt at most one edge from the four exiting S. The notation for the network-adversary pair is $(\mathscr{S}, \mathbf{A}_{\mathscr{S}})$.

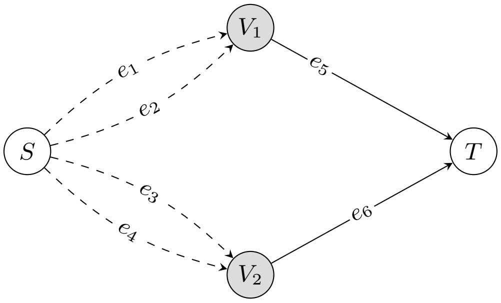

Fig. 2. The Mirrored Diamond Network.

Proposition 3. *We have* $\mathrm{C}_1(\mathscr{S}, \mathbf{A}_{\mathscr{S}}) = 1$.

Proof. By Theorem 2, $\mathrm{C}_1(\mathscr{S}, \mathbf{A}_{\mathscr{S}}) \leqslant 1$, so we need only prove achievability. Select $* \in \mathscr{A}$, and consider the scheme where the source S sends any symbol of $\mathscr{A}$ via a four-times repetition code. Vertices V_1 and V_2 both proceed as follows: If the two received inputs coincide and are equal to $a \in \mathscr{A}$, the vertex forwards a; otherwise it transmits $*$. At T, if the received symbols match and are equal to $a \in \mathscr{A}$, decode to a. Otherwise, decode to the symbol that is not equal to $*$. It is clear that any symbol from $\mathscr{A}$ can be uniquely decoded, including $*$, showing that the proposed scheme is unambiguous. This concludes the proof. □

Note that, as in the proof of Proposition 1, the above scheme uses an alphabet symbol to pass information about the location of the adversary. In strong contrast with the Diamond Network however, in the Mirrored Diamond Network this strategy comes at no cost, as the "reserved" alphabet symbol can be used by the source like any other symbol.

6 Two-Level Networks

We initiate a systematic study of communication with restricted adversaries. Since a global treatment is out of reach at the moment, we start by concentrating on a small but sufficiently interesting family of highly structured networks. These are defined as follows.

Definition 6. *A* ***two-level network*** *is a network* $\mathscr{N} = (\mathscr{V}, \mathscr{E}, S, \{T\})$ *with a single terminal* T *such that any path from* S *to* T *is of length* 2.

An example of a two-level network is given in Figure 3. By applying the Singleton Cut-Set Bound of Theorem 2 to two-level networks with vulnerable edges restricted to the first level, we establish the following bound.

Theorem 4. *Consider a two-level network* $\mathscr{N}$ *where the adversary* $\mathbf{A}$ *can act on up to* t *edges of the first level. Then,* $C_1(\mathscr{N}, \mathbf{A})$ *is upper bounded by the*

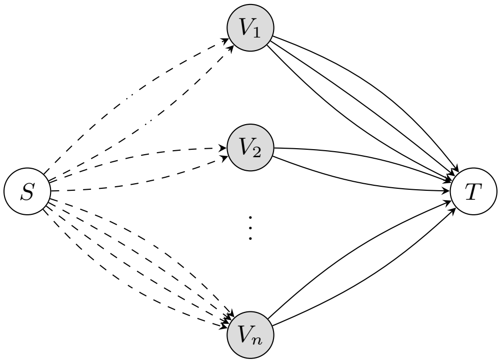

Fig. 3. An example of a two-level network where vulnerable edges are restricted to those in the first level. In general, there may be any number of edges between the source/sink and each intermediate node.

following value:

$$\min_{\mathscr{V}_1,\mathscr{V}_2}\left(\sum_{V_i\in\mathscr{V}_1} deg^+(V_i)+\max\left\{0,\sum_{V_i\in\mathscr{V}_2} deg^-(V_i)-2t\right\}\right),$$

where the minimum is taken over all 2-partitions $\mathscr{V}_1,\mathscr{V}_2$ of the set of intermediate vertices $\{V_1,\dots,V_n\}$.

To understand when the Singleton Cut-Set Bound is achievable in a two-level network, we introduce the following terminology.

Definition 7. *Consider a network where an adversary can act simultaneously on up to t edges. We call an intermediate vertex in the network **damming** if*

$$deg^+(V_i)+1\leqslant deg^-(V_i)\leqslant deg^+(V_i)+2t-1.$$

Notice that if the adversary can change at most one symbol, the above definition reduces to $\deg^-(V_i)=\deg^+(V_i)+1$; such a vertex is present in *both* the Diamond Network and the Mirrored Diamond Network.

We conclude the paper with a sufficient condition for the achievability of the Singleton Cut-Set Bound in a family of two-level networks. The proof will appear in the extended version of this work.

Theorem 5. *In a two-level network where an adversary can act on up to t edges of the first level, if no intermediate vertex is damming, then the Singleton Cut-Set Bound is achievable for sufficiently large alphabet size.*

The results of Section 5 demonstrate that the converse of Theorem 5 does not hold. Indeed, both intermediate vertices of the Mirrored Diamond Network are damming but its capacity is as predicted by the Singleton Cut-Set Bound; see Proposition 3.

7 Discussion and Future Work

We considered the problem of determining the one-shot capacity of communication networks with adversarial noise. In contrast with the typical scenario considered in the context of network coding, we allow the noise to affect only a subset of the network's edges. We defined the Diamond Network and computed its capacity, illustrating that previously known cut-set bounds are not sharp in general. We then studied the family of two-level networks, giving a sufficient condition under which the Singleton Cut-Set Bound is sharp over a sufficiently large alphabet.

Natural problems inspired by these results are the complete characterization of two-level networks for which cut-set bounds are sharp, and development of techniques to derive upper bounds for the capacity of more general adversarial networks. These will be the subject of future work.

References

1. R. W. Yeung and N. Cai, "Network error correction, I: Basic concepts and upper bounds," *Communications in Inf. & Systems*, vol. 6, no. 1, pp. 19–35, 2006.
2. N. Cai and R. W. Yeung, "Network error correction, II: Lower bounds," *Communications in Inf. & Systems*, vol. 6, no. 1, pp. 37–54, 2006.
3. S. Yang and R. W. Yeung, "Refined coding bounds for network error correction," in *IEEE Inf. Theory Workshop on Inf. Theory for Wireless Networks*, 2007, pp. 1–5.
4. S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, and M. Médard, "Resilient network coding in the presence of byzantine adversaries," in *26th IEEE Int'l Conference on Computer Communications.* IEEE, 2007, pp. 616–624.
5. R. Matsumoto, "Construction algorithm for network error-correcting codes attaining the singleton bound," *IEICE Trans. on Fundamentals of Electronics, Communications and Computer Sciences*, vol. 90, no. 9, pp. 1729–1735, 2007.
6. S. Yang, C. K. Ngai, and R. W. Yeung, "Construction of linear network codes that achieve a refined Singleton bound." in *IEEE Int'l Symp. on Inf. Theory*, 2007, pp. 1576–1580.
7. S. Yang, R. W. Yeung, and Z. Zhang, "Weight properties of network codes," *European Trans. on Telecommunications*, vol. 19, no. 4, pp. 371–383, 2008.
8. A. Ravagnani and F. R. Kschischang, "Adversarial network coding," *IEEE Trans. on Inf. Theory*, vol. 65, no. 1, pp. 198–219, 2018.
9. O. Kosut and J. Kliewer, "Network equivalence for a joint compound-arbitrarily-varying network model," in *IEEE Inf. Theory Workshop*, 2016, pp. 141–145.
10. N. Sangwan, M. Bakshi, B. K. Dey, and V. M. Prabhakaran, "Multiple access channels with adversarial users," in *IEEE Int'l Symp. on Inf. Theory.* IEEE, 2019, pp. 435–439.
11. A. Beemer, E. Graves, J. Kliewer, O. Kosut, and P. Yu, "Authentication and partial message correction over adversarial multiple-access channels," in *IEEE Conference on Communications and Network Security*, 2020, pp. 1–6.