Network Decoding Against Restricted Adversaries

Allison Beemer * **Altan B. Kılıç** ** **Alberto Ravagnani** ***

* *University of Wisconsin-Eau Claire, WI 54701, USA (e-mail: beemera@uwec.edu)*
** *Eindhoven University of Technology, 5612 AZ Eindhoven, The Netherlands (e-mail: a.b.kilic@tue.nl)*
*** *Eindhoven University of Technology, 5612 AZ Eindhoven, The Netherlands (e-mail: a.ravagnani@tue.nl)*

Abstract: We initiate the study of the one-shot capacity of communication (coded) networks with an adversary having access only to a proper subset of the 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 give a capacity-achieving scheme for the Diamond Network that implements an adversary detection strategy. We also show that linear network coding does not suffice in general to achieve capacity, proving a strong separation result between the one-shot capacity and its linear version. We then give a sufficient condition for tightness of the Singleton Cut-Set Bound in a family of two-level networks. Finally, we discuss how the presence of nodes that do not allow local encoding and decoding does or does not affect the one-shot capacity.

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 Yeung and Cai (2006); Cai and Yeung (2006). Other bounds and related code constructions for adversarial networks are presented in, e.g., Yang and Yeung (2007); Jaggi et al. (2007); Matsumoto (2007); Yang et al. (2007, 2008); Ravagnani and Kschischang (2018). The work most closely related to this paper is Ravagnani and Kschischang (2018), 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, Ravagnani and Kschischang (2018) 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 Ravagnani and Kschischang (2018). 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

* A. B. K. is supported by the Dutch Research Council through grant VI.Vidi.203.045. A. R. is supported by the Dutch Research Council through grants VI.Vidi.203.045, OCENW.KLEIN.539, and by the Royal Academy of Arts and Sciences of the Netherlands.

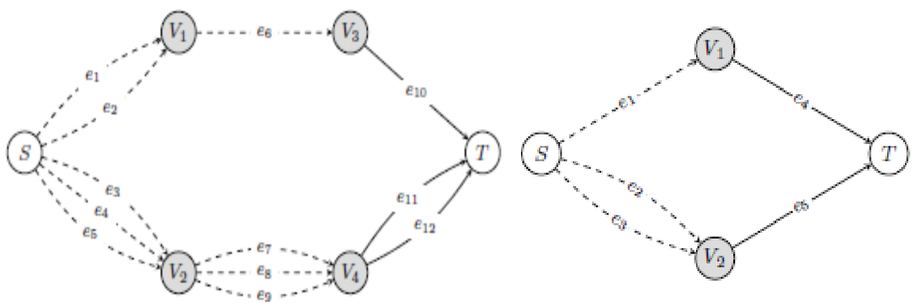

Allison Beemer

Assistant Professor
Mathematics

Co-Authors:
Altan B. Kilic &
Alberto Ravagnani

Proceedings of the 25th International Symposium of Mathematical Theory of Networks and Systems, Sept. 2022

Network Decoding Against Restricted Adversaries

Adversarial networks model communication scenarios where data is transmitted through a web of links that may be corrupted by a malicious actor. The network capacity measures the largest amount of information that can be reliably communicated through a given network. In this work, we initiate the study of network capacity when an adversary has access to some, but not all, network links. The problem differs in surprising (and interesting) ways from previously-solved cases; we illustrate these differences using a minimal example called the Diamond Network (see lower figure). A full version, Network Decoding, is under review and posted on arXiv.