Network Decoding and Packing Problems

Altan B. Kılıç

based on joint work with Allison Beemer and Alberto Ravagnani

Finite Geometries, 6th Irsee Conference

Irsee, 2022

EINDHOVEN UNIVERSITY OF TECHNOLOGY

Networks and Rules of the Game

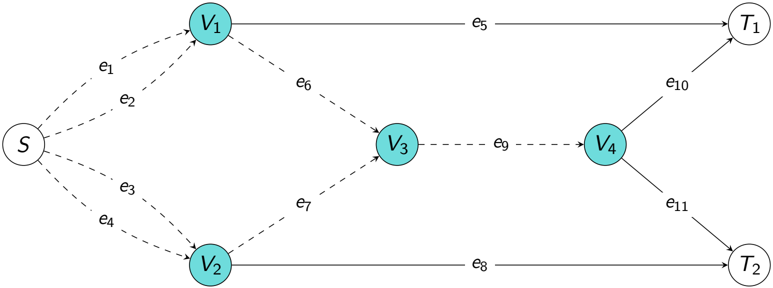

- Networks are finite directed acyclic multigraphs,
- The (single) source S wants to send symbols from a certain alphabet,
- All terminals want all the messages (multicast),
- Each edge can carry exactly one symbol from an alphabet $\mathcal{A}$,
- An adversary can corrupt up to t edges of the network (the dashed edges). t is called the **adversarial power**.

Some definitions

Let $\mathcal{A}$ be a finite alphabet (a finite set with $|\mathcal{A}| \geq 2$).

Definition

The vertices which are neither the source nor the terminals are called **intermediate nodes**.

Definition

A **network code (inner code)** $\mathcal{F}$ for a network $\mathcal{N}$ is a family of functions $\{\mathcal{F}_V \mid V \text{ is an intermediate node in } \mathcal{N}\}$, where

$$\mathcal{F}_V : \mathcal{A}^{|\text{in}(V)|} \to \mathcal{A}^{|\text{out}(V)|}.$$

Definition

An (**outer**) **code** $\mathcal{C}$ for a network $\mathcal{N}$ is a non-empty subset $\mathcal{C} \subseteq \mathcal{A}^{|\text{out}(S)|}$.

Some definitions

Let $\mathcal{A}$ be a finite alphabet (a finite set with $|\mathcal{A}| \geq 2$).

Definition

The vertices which are neither the source nor the terminals are called **intermediate nodes**.

Definition

A **network code (inner code)** $\mathcal{F}$ for a network $\mathcal{N}$ is a family of functions $\{\mathcal{F}_V \mid V \text{ is an intermediate node in } \mathcal{N}\}$, where

$$\mathcal{F}_V : \mathcal{A}^{|\text{in}(V)|} \to \mathcal{A}^{|\text{out}(V)|}.$$

Definition

An (**outer**) **code** $\mathcal{C}$ for a network $\mathcal{N}$ is a non-empty subset $\mathcal{C} \subseteq \mathcal{A}^{|\text{out}(S)|}$.

The adversary is omniscient!

Unambiguity

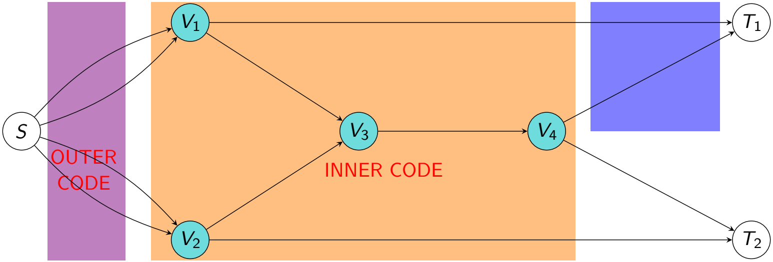

Unambiguity

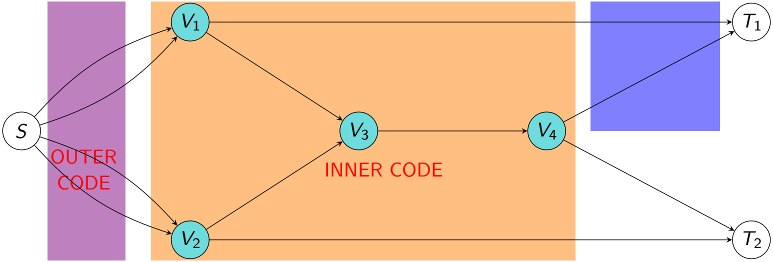

Intuitive definition

Given $x \in \mathcal{C}$, denote all possible outcomes that can appear in the blue region as $\Omega_1(x) \subseteq \mathcal{A}^{|\text{in}(T_1)|}$.

A pair $(\mathcal{C}, \mathcal{F})$ is called **unambiguous for** T_1 if $\Omega_1(x) \cap \Omega_1(x') = \emptyset$ for all $x, x' \in \mathcal{C}$ with $x \neq x'$.

It is called **unambiguous** if it is unambiguous for all terminals.

An Example

$\mathcal{A} = \mathbb{F}_2$, $t = 1$, $\mathcal{C} = \{000, 111\}$, $\mathcal{F}$ as in the picture.

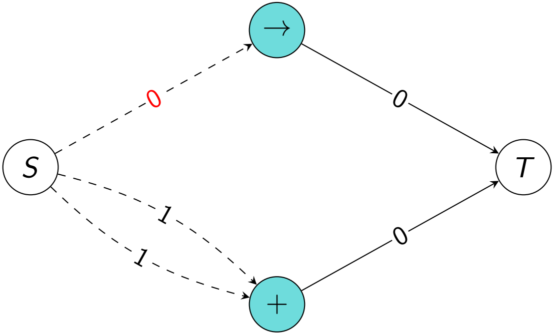

$\Omega(000) = \{00, 10, 01\}$ and $\Omega(111) = \{00, 11\}$. So,

$$\Omega(000) \cap \Omega(111) \neq \emptyset$$

and the pair $(\mathcal{C}, \mathcal{F})$ is not unambiguous.

1-shot capacity

Definition

The (**1-shot**) **capacity** of $\mathcal{N}$ is the largest real number α for which there exists an unambiguous pair $(\mathcal{C}, \mathcal{F})$ with $\alpha = \log_{|\mathcal{A}|} |\mathcal{C}|$.

1-shot capacity

Definition

The (**1-shot**) **capacity** of $\mathcal{N}$ is the largest real number α for which there exists an unambiguous pair $(\mathcal{C}, \mathcal{F})$ with $\alpha = \log_{|\mathcal{A}|} |\mathcal{C}|$.

Singleton Cut-Set Bound (Kschischang, Ravagnani '18)

Let $\mathcal{N}$ be a network, E be the set of edges in $\mathcal{N}$. If an adversary can act on $\mathcal{U} \subseteq E$ with adversarial power t, then

$$C_1(\mathcal{N}) \leq \min_{T_i} \min_{\mathcal{E}'} \left(|\mathcal{E}' \setminus \mathcal{U}| + \max\{0, |\mathcal{E}' \cap \mathcal{U}| - 2t\} \right),$$

where $\mathcal{E}' \subseteq E$ ranges over the edge-cuts between S and T_i.

1-shot capacity

Definition

The **(1-shot) capacity** of $\mathcal{N}$ is the largest real number α for which there exists an unambiguous pair $(\mathcal{C}, \mathcal{F})$ with $\alpha = \log_{|\mathcal{A}|} |\mathcal{C}|$.

Singleton Cut-Set Bound (Kschischang, Ravagnani '18)

Let $\mathcal{N}$ be a network, E be the set of edges in $\mathcal{N}$. If an adversary can act on $\mathcal{U} \subseteq E$ with adversarial power t, then

$$C_1(\mathcal{N}) \leq \min_{T_i} \min_{\mathcal{E}'} \left(|\mathcal{E}' \setminus \mathcal{U}| + \max\{0, |\mathcal{E}' \cap \mathcal{U}| - 2t\} \right),$$

where $\mathcal{E}' \subseteq E$ ranges over the edge-cuts between S and T_i.

Theorem (Silva, Kschischang, Kötter '08)

If $\mathcal{U} = E$ and [[assumptions on $\mathcal{A}$]], then the Singleton Cut-Set Bound is sharp.

1-shot capacity

Definition

The (**1-shot**) **capacity** of $\mathcal{N}$ is the largest real number α for which there exists an unambiguous pair $(\mathcal{C}, \mathcal{F})$ with $\alpha = \log_{|\mathcal{A}|} |\mathcal{C}|$.

Singleton Cut-Set Bound (Kschischang, Ravagnani '18)

Let $\mathcal{N}$ be a network, E be the set of edges in $\mathcal{N}$. If an adversary can act on $\mathcal{U} \subseteq E$ with adversarial power t, then

$$C_1(\mathcal{N}) \leq \min_{T_i} \min_{\mathcal{E}'} \left(|\mathcal{E}' \setminus \mathcal{U}| + \max\{0, |\mathcal{E}' \cap \mathcal{U}| - 2t\} \right),$$

where $\mathcal{E}' \subseteq E$ ranges over the edge-cuts between S and T_i.

Theorem (Silva, Kschischang, Kötter '08)

If $\mathcal{U} = E$ and [[assumptions on $\mathcal{A}$]], then the Singleton Cut-Set Bound is sharp. $\leftarrow$ *rank-metric codes*

Restricting the adversary ($\mathcal{U} \neq E$)

The Singleton Cut-Set Bound is not sharp in general, although often the best known bound.

Restricting the adversary ($\mathcal{U} \neq E$)

The Singleton Cut-Set Bound is not sharp in general, although often the best known bound.

The Diamond Network (with $t = 1$):

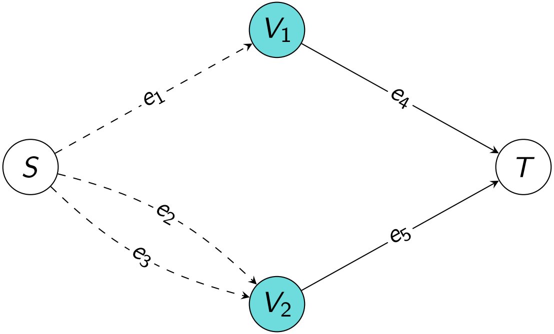

Restricting the adversary ($\mathcal{U} \neq E$)

The Singleton Cut-Set Bound is not sharp in general, although often the best known bound.
The Diamond Network (with $t = 1$):

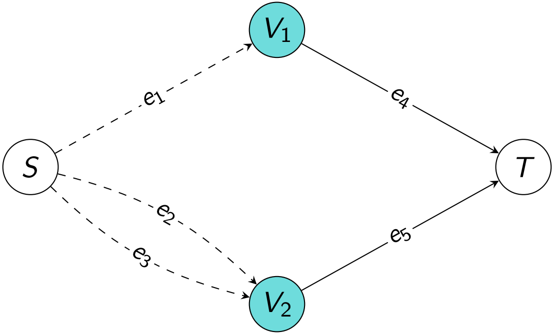

If $\mathcal{U} = \{e_1, e_2, e_3\}$, then the Singleton Cut-Set Bound is the best known upper bound and it gives $C_1(\mathcal{N}) \leq 1 = \log_{|\mathcal{A}|} |\mathcal{A}|$.

Restricting the adversary ($\mathcal{U} \neq E$)

The Singleton Cut-Set Bound is not sharp in general.

Main Interest

Find/Bound the value $\max\{|\mathcal{C}| : (\mathcal{C}, \mathcal{F})$ is unambiguous for $\mathcal{N}\}$.

Restricting the adversary ($\mathcal{U} \neq E$)

The Singleton Cut-Set Bound is not sharp in general.

Main Interest

Find/Bound the value $\max\{|\mathcal{C}| : (\mathcal{C}, \mathcal{F}) \text{ is unambiguous for } \mathcal{N}\}$.

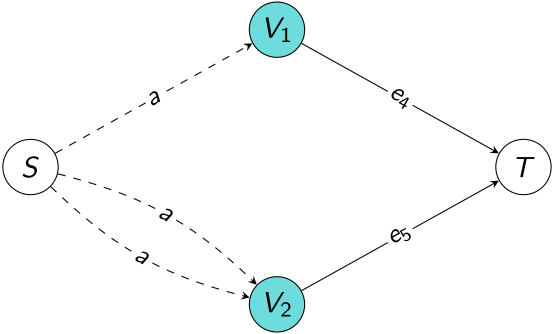

Natural candidate: $\mathcal{C} = \{(a, a, a) \mid a \in \mathcal{A}\}$.

Issue: One can globally <u>en</u>code, but not globally <u>de</u>code.

The Diamond Network

Theorem (Beemer, Ravagnani '21)

For the Diamond Network $\mathcal{N}$,

$$C_1(\mathcal{N}) = \log_{|\mathcal{A}|}(|\mathcal{A}| - 1).$$

The Diamond Network

Theorem (Beemer, Ravagnani '21)

For the Diamond Network $\mathcal{N}$,

$$C_1(\mathcal{N}) = \log_{|\mathcal{A}|}(|\mathcal{A}| - 1).$$

Capacity-achieving scheme:

The Diamond Network

Theorem (Beemer, Ravagnani '21)

For the Diamond Network $\mathcal{N}$,

$$C_1(\mathcal{N}) = \log_{|\mathcal{A}|}(|\mathcal{A}| - 1).$$

Capacity-achieving scheme:

- **sacrifice** an alphabet symbol $* \in \mathcal{A}$ for adversary detection
- $\mathcal{F}_{V_1}(a) = a$ for $a \in \mathcal{A}$
- $\mathcal{F}_{V_2}(a, b) = \begin{cases} a & \text{if} \;\; a = b \\ * & \text{if} \;\; a \neq b \end{cases}$
- T looks at the symbol on the outgoing edge of V_2. If it is not $*$, T decodes to that symbol. If it is $*$, then T decodes to the symbol on the outgoing edge of V_1.

Sacrifice an alphabet symbol $*$

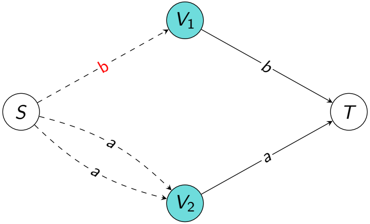

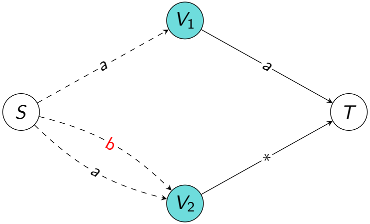

Simple 2-level networks

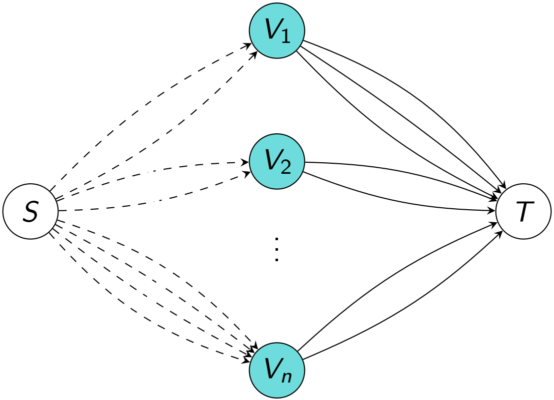

From examples to a more general theory

Theorem (Beemer, K., Ravagnani '22)

C_1(*any network*) $\leq C_1$(*simple 2-level induced from it*)

From examples to a more general theory

Theorem (Beemer, K., Ravagnani '22)

C_1(*any network*) $\leq$ C_1(*simple 2-level induced from it*)

Idea behind our approach

A possibly complicated network $\rightarrow$ 3-level network $\rightarrow$ 2-level network

From examples to a more general theory

Theorem (Beemer, K., Ravagnani '22)

$C_1(\textit{any network}) \leq C_1(\textit{simple 2-level induced from it})$

Idea behind our approach

A possibly complicated network $\rightarrow$ 3-level network $\rightarrow$ 2-level network

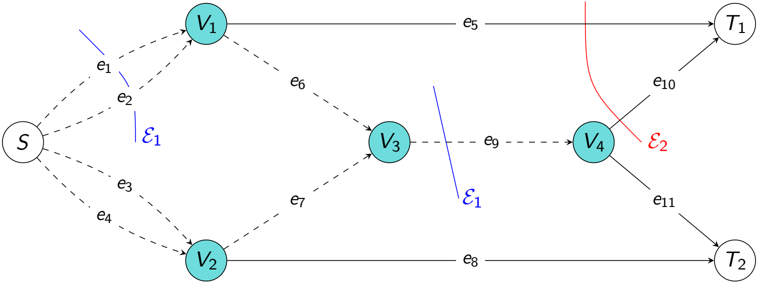

From examples to a more general theory

Theorem (Beemer, K., Ravagnani '22)

$C_1(\textit{any network}) \leq C_1(\textit{simple 2-level induced from it})$

Idea behind our approach

A possibly complicated network $\rightarrow$ 3-level network $\rightarrow$ 2-level network

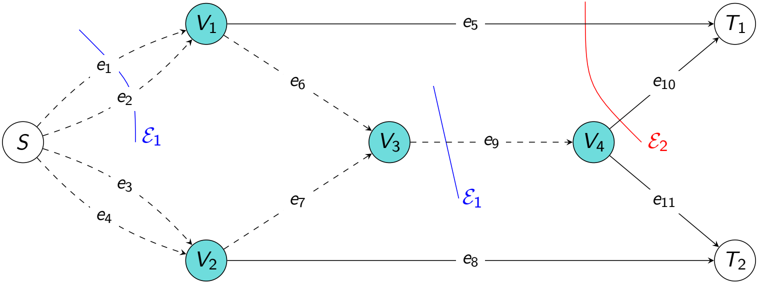

We prove a Double Cut-Set Bound (Beemer, K., Ravagnani '22)

Explaining via pictures

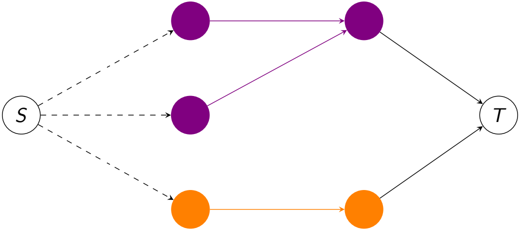

Explaining via pictures

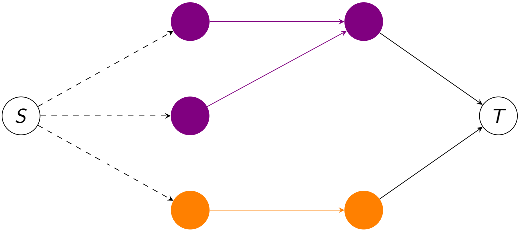

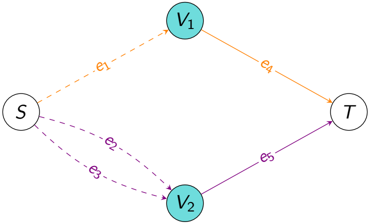

Explaining via pictures

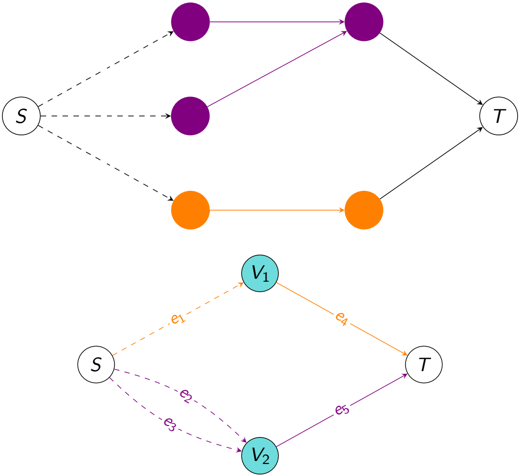

Observe that we end up in a simple 2-level network. We can now derive an upper bound for the capacity of the original more complicated network.

Simple 2-level networks

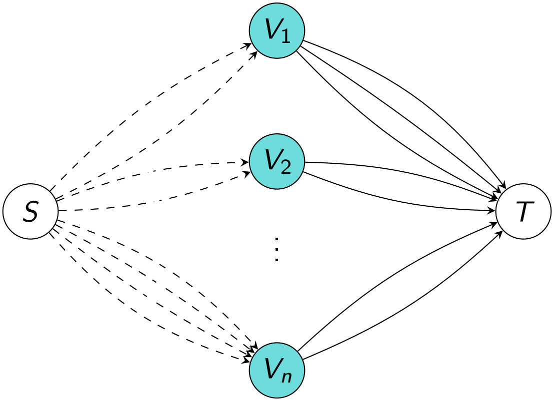

$a_i = |\text{in}(V_i)|$ and $b_i = |\text{out}(V_i)|$.

We also denote the network code by $\mathcal{F} = (\mathcal{F}_1, \ldots, \mathcal{F}_n)$.

First Packing Bound

Given an outer code $\mathcal{C} \subseteq \mathcal{A}^{a_1+a_2+\ldots+a_n}$, we let $\pi_i(\mathcal{C})$ be the projection of $\mathcal{C}$ onto the a_i coordinates corresponding to the edges to intermediate node V_i

Theorem (First Packing Bound) (Beemer, K., Ravagnani '22)

Consider a simple 2-level network with $a_i \leq b_i$ for all $1 \leq i \leq r$. Let $(\mathcal{C}, \mathcal{F})$ be unambiguous. Then

$$\sum_{\substack{t_1,\ldots,t_r \geq 0 \\ t_1+\ldots+t_r \leq t}} \prod_{i=1}^{r} \binom{a_i}{t_i} (|\mathcal{A}|-1)^{t_i} \sum_{x \in \mathcal{C}} \prod_{j=r+1}^{n} \left|\mathcal{F}_j\left(B_{t-(t_1+\ldots+t_r)}(\pi_j(x))\right)\right| \leq |\mathcal{A}|^{b_1+b_2+\ldots+b_n}$$

First Packing Bound

Given an outer code $\mathcal{C} \subseteq \mathcal{A}^{a_1+a_2+\ldots+a_n}$, we let $\pi_i(\mathcal{C})$ be the projection of $\mathcal{C}$ onto the a_i coordinates corresponding to the edges to intermediate node V_i

Theorem (First Packing Bound) (Beemer, K., Ravagnani '22)

Consider a simple 2-level network with $a_i \leq b_i$ for all $1 \leq i \leq r$. Let $(\mathcal{C}, \mathcal{F})$ be unambiguous. Then

$$\sum_{\substack{t_1,\ldots,t_r \geq 0 \\ t_1+\ldots+t_r \leq t}} \prod_{i=1}^{r} \binom{a_i}{t_i} (|\mathcal{A}|-1)^{t_i} \sum_{x \in \mathcal{C}} \prod_{j=r+1}^{n} \left| \mathcal{F}_j \left(B_{t-(t_1+\ldots+t_r)}(\pi_j(x)) \right) \right| \leq |\mathcal{A}|^{b_1+b_2+\ldots+b_n}$$

Proof Idea

- Whenever $a_i \leq b_i$, we can assume $a_i = b_i$ and take the corresponding function $\mathcal{F}_i$ to be identity (ignoring extraneous outgoing edges),
- $B_t(x) = \bigsqcup_{t_1+\ldots+t_n \leq t} [S_{t_1}(\pi_1(x)) \times \cdots \times S_{t_n}(\pi_n(x))]$,
- $\sum_{x \in \mathcal{C}} |\mathcal{F}(B_t(x))| \leq |\mathcal{A}|^{b_1+b_2+\ldots+b_n}$.

Simple 2-level Networks with $n = 2$

Corollary (Beemer, K., Ravagnani '22)

Consider a simple 2-level network with $n = 2$ and $a_1 \leq b_1$. Let $(\mathcal{C}, \mathcal{F})$ be unambiguous. Then,

$$\sum_{t_1=0}^{t} \binom{a_1}{t_1} (|\mathcal{A}| - 1)^{t_1} \sum_{x \in \mathcal{C}} |\mathcal{F}_2 \left(B_{t-t_1}(\pi_2(x)) \right)| \leq |\mathcal{A}|^{b_1 + b_2}.$$

Corollary of the above corollary

The Singleton Cut-Set Bound for the Diamond Network is not met.

Simple 2-level Networks with $n = 2$

Corollary (Beemer, K., Ravagnani '22)

Consider a simple 2-level network with $n = 2$ and $a_1 \leq b_1$. Let $(\mathcal{C}, \mathcal{F})$ be unambiguous. Then,

$$\sum_{t_1=0}^{t} \binom{a_1}{t_1} (|\mathcal{A}| - 1)^{t_1} \sum_{x \in \mathcal{C}} |\mathcal{F}_2 (B_{t-t_1}(\pi_2(x)))| \leq |\mathcal{A}|^{b_1+b_2}.$$

Corollary of the above corollary

The Singleton Cut-Set Bound for the Diamond Network is not met.

Using a similar idea, we can get a Hamming-type bound.

Lemma (Beemer, K., Ravagnani '22)

Let $(\mathcal{C}, \mathcal{F})$ be unambiguous for the simple 2-level network $\mathcal{N}$. Then, $\mathcal{F}^{-1}(\mathcal{F}(B_t(x))) \cap \mathcal{F}^{-1}(\mathcal{F}(B_t(x'))) = \emptyset$ for all distinct $x, x' \in \mathcal{C}$.

Second Packing Bound

Corollary (Beemer, K., Ravagnani '22)

Consider a simple 2-level network with $n = 2$ and $a_1 \leq b_1$. Let $(\mathcal{C}, \mathcal{F})$ be unambiguous. Then,

$$\sum_{t_1=0}^{t} \binom{a_1}{t_1} (|\mathcal{A}| - 1)^{t_1} \sum_{x \in \mathcal{C}} \left| \mathcal{F}_2^{-1}(\mathcal{F}_2 \left(B_{t-t_1}(\pi_2(x)) \right)) \right| \leq |\mathcal{A}|^{a_1+a_2}.$$

To be compared with:

Theorem (Hamming Bound)

Let $(\mathcal{C}, \mathcal{F})$ be unambiguous for a simple 2-level network with $n = 2$ and $a_1 \leq b_1$. Then, $|\mathcal{C}| \cdot \sum_{t_1=0}^{t} \binom{a_1+a_2}{t_1} (|\mathcal{A}| - 1)^{t_1} \leq |\mathcal{A}|^{a_1+a_2}$.

We expect the corollary to beat the Hamming bound for some classes of networks.

Future work

For example, compute the exact 1-shot capacity of all simple 2-level networks. Open:

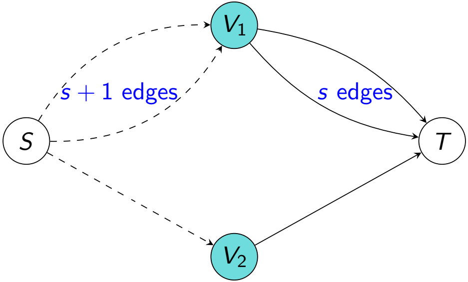

$t = 1 \implies s - 1 \leq C_1(\mathcal{N}) < s$, but what is the exact value?

$\rightarrow$ started collaboration at TU/e using combinatorial optimization.

Future work

For example, compute the exact 1-shot capacity of all simple 2-level networks. Open:

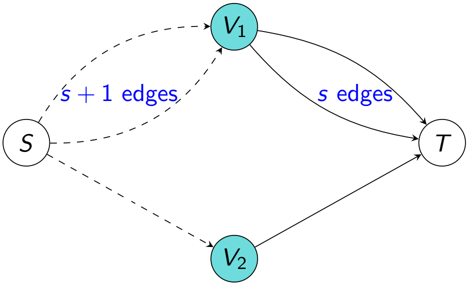

$t = 1 \implies s - 1 \leq C_1(\mathcal{N}) < s$, but what is the exact value?

$\rightarrow$ started collaboration at TU/e using combinatorial optimization.

Thank You!

https://arxiv.org/abs/2205.14655

Network Decoding

Allison Beemer[1], Altan B. Kılıç[*2], and Alberto Ravagnani[†3]

[1]Department of Mathematics, University of Wisconsin-Eau Claire, U.S.A.
[2,3]Department of Mathematics and Computer Science, Eindhoven University of Technology, the Netherlands

Abstract

We consider the problem of error control in a coded, multicast network, focusing on the scenario where the errors can occur only on a *proper subset* of the network edges. We model this problem via an adversarial noise, presenting a formal framework and a series of techniques to obtain upper and lower bounds on the network's (1-shot) capacity, improving on the best currently known results. In particular, we show that traditional cut-set bounds are not tight in general in the presence of a restricted adversary, and that the non-tightness of these is caused precisely by the restrictions imposed on the noise (and not, as one may expect, by the alphabet size). We also show that, in sharp contrast with the typical situation within network coding, capacity cannot be achieved in general by combining linear network coding with end-to-end channel coding, not even when the underlying network has a single source and a single terminal. We finally illustrate how network *decoding* techniques are necessary to achieve capacity in the scenarios we examine, exhibiting capacity-achieving schemes and lower bounds for various classes of networks.