On Networks with Side Information

A. Cohen, S. Avestimehr and M. Effros

California Institute of Technology

Motivation

Side information is data available within the network, which is correlated with the main source we wish to transmit. Side information can significantly reduce the rates required throughout the network.

The model is interesting in cases where:

- Not enough capacity from the source.
- Side information is "close" or "less expensive".
- Collaborative networks (sensor, organizational, army).

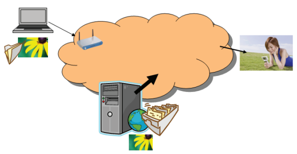

Problem Statement

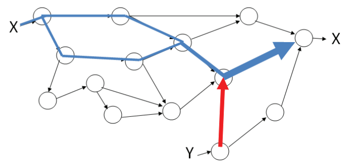

- Network is modeled as a directed acyclic graph.
- One source node, one side information node, possibly several terminals requiring the source (with arbitrarily small error).
- Data processing is allowed at all nodes.

How much data do we have to send?

Goal: compute the amount of data required - the *Rate Region*

For a network with nodes $\mathcal{V}$, edges $\mathcal{E}$, source at $s \in \mathcal{V}$, side information at $z \in \mathcal{V}$ and terminals $T \subset \mathcal{V}$, characterize $\mathcal{R}(\mathcal{V}, \mathcal{E}, s, z, T) \subseteq \mathbb{R}^{|\mathcal{E}|}$: the set of achievable rates.

Converse:

- What are the minimal link capacities required?

Achievability:

- Which X descriptions to create?
- Which Y descriptions to create?
- How to decode X?

Complexity:

- What are the required operations at the nodes?

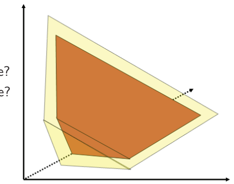

Related Work

- Ahlswede et al. 00: Network information flow. Coding: Li et al. Ho et al. 03 (linear, random); Koetter & Medard 03 (algebraic); Jaggi et al. 05 (polynomial time).

- Ho et al. 05: correlated sources

- Bakshi & Effros 08: multicast with side information at the Terminals

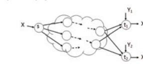

- Ahlswede & Korner 75: Coded side information

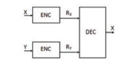

Cut Sets

The majority of network coding related work is on cases where the cut-set bounds are tight!

Generalizing Point-to-Point Networks and Beyond

Shannon 48:

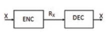

Ford-Fulkerson 56:

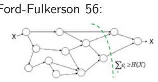

So we can easily also do:

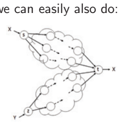

Yet, this generalization of point-to-point is sub-optimal in general.

Example 1: one terminal

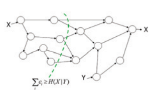

Example 2: more than one terminal

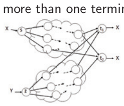

But we know how to solve more than point-to-point networks!

Characterizing the Smallest Required Rates (Outer Bound)

$\mathcal{E}_{XY} \subseteq \mathcal{E}$ denotes the set of edges for which there is a directed path from s to $o(e)$. $\mathcal{E}_Y$ denotes the set $\mathcal{E} \setminus \mathcal{E}_{XY}$. Given any non-intersecting sets $A, B \subset \mathcal{V}$, $\mathcal{V}_{A;B}$ denotes a cut with $A \subseteq \mathcal{V}_{A;B}$ and $B \cap \mathcal{V}_{A;B} = \emptyset$. Let $\mathcal{C}(\mathcal{V}_{A;B})$ be the set of edges $e \in \mathcal{E}$ for which $o(e) \in \mathcal{V}_{A;B}$ and $d(e) \notin \mathcal{V}_{A;B}$.

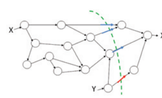

Theorem

Given a side information network $(\mathcal{V}, \mathcal{E}, s, z, T)$, if $(c(e) : e \in \mathcal{E}) \in \mathcal{R}(\mathcal{V}, \mathcal{E}, s, z, T)$, then for each $t \in T$ and each cut $\mathcal{V}_{s,z;t}$ there exists a random variable $U \in \mathcal{U}$ such that $U \leftrightarrow Y \leftrightarrow X$, $|\mathcal{U}| \leq |\mathcal{Y}|$, and

$$\sum_{e \in \mathcal{E}_{XY} \cap \mathcal{C}(\mathcal{V}_{s,z;t})} c(e) \geq H(X|U)$$

$$\sum_{e \in \mathcal{E}_{Y} \cap \mathcal{C}(\mathcal{V}_{s,z;t})} c(e) \geq I(Y;U).$$

Comparison to the Cut-Set Bound

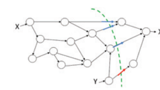

The cut-set bound is:

$$\sum_{e \in \mathcal{C}(\mathcal{V}_{s,z;t})} c(e) \geq H(X),$$

while the new bound gives

$$\sum_{e \in \mathcal{C}(\mathcal{V}_{s,z;t})} c(e) \geq H(X|U) + I(Y;U).$$

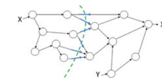

The cut-set bound is:

$$\sum_{e \in \mathcal{C}(\mathcal{V}_{s;z,t})} c(e) \geq H(X|Y),$$

while the new bound gives

$$\sum_{e \in \mathcal{C}(\mathcal{V}_{s;z,t})} c(e) \geq \sum_{e \in \mathcal{E}_{XY} \cap \mathcal{C}(\tilde{\mathcal{V}}_{s,z;t})} c(e) \geq H(X|U).$$

Tightness

Since $U \leftrightarrow Y \leftrightarrow X$, the new bound is at least as tight as the cut-set bound.

Two Sinks - Arbitrary Network from Y

Theorem

Assume both X and Y are binary symmetric. The outer bound is tight for the "Y" network, and $(c(e) : e \in \mathcal{E}) \in \mathcal{R}(\mathcal{V}, \mathcal{E}, s, z, \{t_1, t_2\})$ iff $\exists\, U_1 \in \mathcal{U}_1$ and $U_2 \in \mathcal{U}_2$ such that $U_1 \leftrightarrow Y \leftrightarrow X$, $U_2 \leftrightarrow Y \leftrightarrow X$, $|\mathcal{U}_1| \leq |\mathcal{Y}|$, $|\mathcal{U}_2| \leq |\mathcal{Y}|$, and

$$\sum_{e \in \mathcal{C}(\mathcal{V}^*_{s;t_1})} c(e) \geq H(X|U_1)$$

$$\sum_{e \in \mathcal{C}(\mathcal{V}^*_{z;t_1})} c(e) \geq I(Y;U_1)$$

$$\sum_{e \in \mathcal{C}(\mathcal{V}^*_{s;t_2})} c(e) \geq H(X|U_2)$$

$$\sum_{e \in \mathcal{C}(\mathcal{V}^*_{z;t_2})} c(e) \geq I(Y;U_2)$$

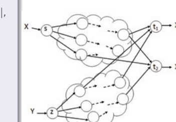

Cut-sets

Cut-set bounds are loose, but the new generalization is tight!

Multi-Resolution codes are Necessary

- A source may have a high rate to one terminal and a low rate to the other.
- As a result, two descriptions may be required.

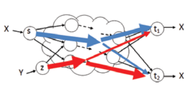

Multi-resolution codes:

- Each point represents a sequence in $\{0,1\}^n$.
- $R = \log(\#\text{ of regions})$.
- Ideally, the high rate description *refines* the coarse one.

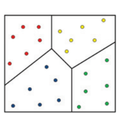

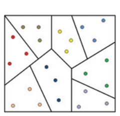

K Sinks with Direct Links from Y - Result

Theorem

$(c(e) : e \in \mathcal{E}) \in \mathcal{R}(\mathcal{V}, \mathcal{E}, s, z, \{t_i\}_{i=1}^K)$ iff for any $1 \leq i \leq K$ there exist $U_i \in \mathcal{U}_i$ such that $U_i \leftrightarrow Y \leftrightarrow X$, $|\mathcal{U}_i| \leq |\mathcal{Y}|$ and

$$\sum_{e \in \mathcal{C}(\mathcal{V}^*_{s;t_i})} c(e) \geq H(X|U_i)$$

$$c((z, t_i)) \geq I(Y;U_i)$$

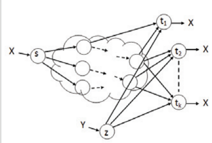

The main idea: randomly bin the source sequences

Each node sends as much independent equations on the bin index as it can.

Routing is sub-optimal

To send the two descriptions (at the lowest possible rates) we must use network coding.

Proof Sketch - random binning and independent equations

The figure below represents binning to 8 bins and its binary representation:

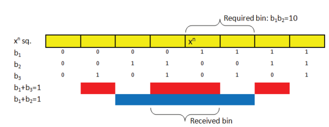

The bin we received is not the one we "intended" to send, but:

- It has the correct size.
- It contains the true x^n.
- It is just as random.

The power of random binning

We can refine the random binning without having a true incremental multicast network code.

Conclusion

- We derived inner and outer bounds on the rate region, and identified scenarios in which they are tight.
- The outer bound is at least as tight as the cut-set, and tighter in most non-trivial cases.
- While the cut-set bounds are loose, we extended the range of scenarios for which cut analysis describes the rate region.
- Our methods inherit the desirable properties of the building blocks.
- We identified an interesting connection between coding for networks with side information and successive refinement.
- In fact, this method might yield interesting results for other source coding problems as well. In a sense, this is a way to rigorously formulate the intuition from canonical problems to larger networks.