Coding for Noisy Networks

Abbas El Gamal

Stanford University

ISIT Plenary, June 2010

Introduction

- Over past 40+ years, there have been many efforts to extend Shannon's information theory to noisy networks
- Although we may be far from a complete network information theory, several coding schemes that are optimal or close to optimal for some important classes of networks have been developed
- My talk is about these schemes
- Focus is on two recently developed coding schemes:
  - Noisy network coding
  - Compute–forward
- And how they relate and compare to better known schemes:
  - Decode–forward
  - Compress–forward
  - Amplify–forward
  - Network coding and its extensions

Noisy Network Model

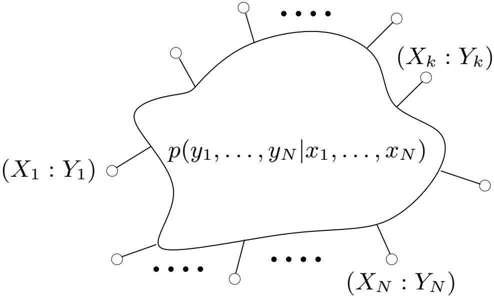

- Consider N-node discrete memoryless network (DMN)
- Allows for noise, interference, multi-access, broadcast, relaying, multi-way communication, ...
- Includes noiseless, erasure, and deterministic networks
- Can be modified to include Gaussian networks and networks with state

Noisy Network Model

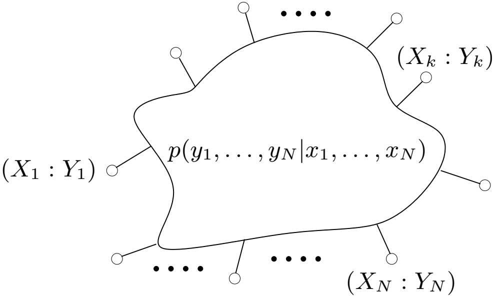

- Each source node has independent message and wishes to send it to a set of destination nodes
- The problem is to find the capacity region of the network

  The coding scheme that achieves it?

Noisy Multicast Network

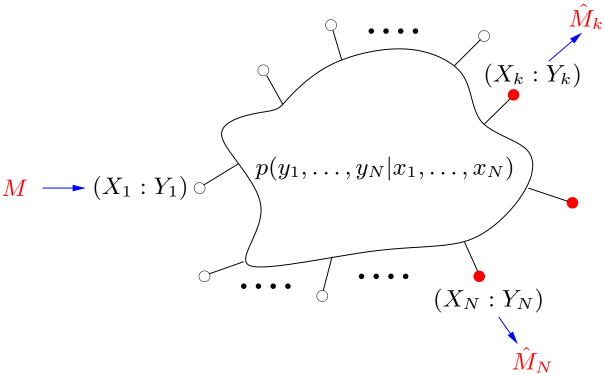

- Source node 1 wishes to send message M to destination nodes $\mathcal{D}$

Noisy Multicast Network

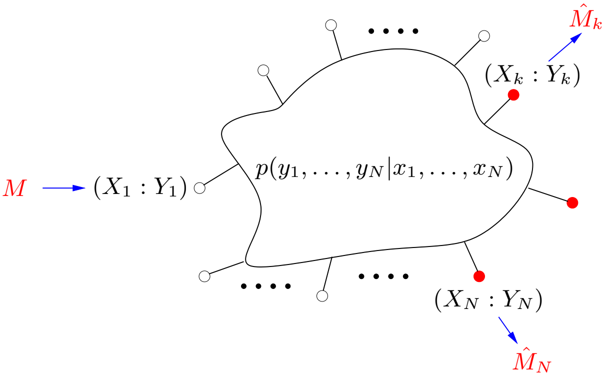

- A $(2^{nR}, n)$ code for the DM-MN:
  - Encoder: $x_{1i}(m, y_1^{i-1})$ for every $m \in [1 : 2^{nR}]$ and y_1^{i-1}, $i \in [1 : 2^{nR}]$
  - Relay encoders: $x_{ji}(y_j^{i-1})$ for every y_j^{i-1}, $i \in [1 : n]$, $j \in [2 : N]$
  - Decoders: $\hat{m}_k(y_k^n)$ for every y_k^n, $k \in \mathcal{D}$
- Average probability of error $P_e^{(n)} = \mathrm{P}\{\hat{M}_k \neq M \text{ for some } k \in \mathcal{D}\}$

Noisy Multicast Network

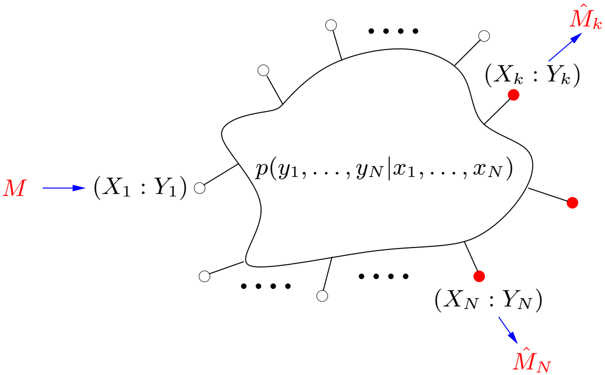

- Rate R achievable if there exists a sequence of codes with $P_e^{(n)} \to 0$
- The capacity C of the DM-MN is supremum of achievable rates
- Capacity is not known in general
- There are upper and lower bounds that coincide in some special cases

Cutset Upper Bound

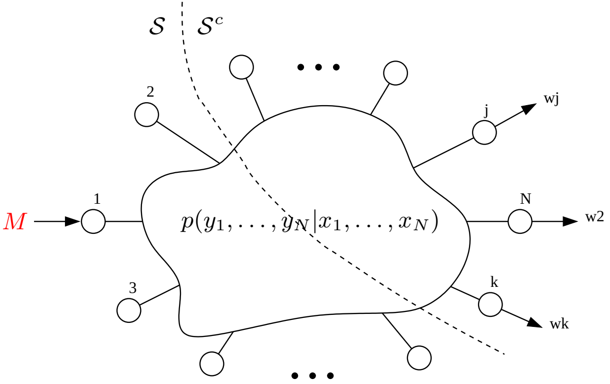

- $X(\mathcal{S})$ inputs in $\mathcal{S}$; $X(\mathcal{S}^c)$, $Y(\mathcal{S}^c)$ inputs/outputs in $\mathcal{S}^c$

Cutset upper bound (EG 1981)

$$C \le \max_{p(x^N)} \min_{k \in \mathcal{D}} \min_{\mathcal{S}: 1 \in \mathcal{S},\, k \in \mathcal{S}^c} I(X(\mathcal{S}); Y(\mathcal{S}^c) | X(\mathcal{S}^c))$$

Cutset Bound Is Sometimes Tight

- Point-to-point channel (Shannon 1948)
- Noiseless unicast network (Ford, Fulkerson 1956)
- Relay channel:
  - Degraded, reversely degraded (Cover, EG 1979)
  - Semi-deterministic (EG, Aref 1982), (Cover, Kim 2007)
  - Orthogonal sender components (EG, Zahedi 2005)
- Noiseless multicast networks (Ahlswede, Cai, Li, Yeung 2000)
- Erasure multicast networks (Dana et al. 2006)
- Deterministic multicast networks:
  - No interference (Aref 1980; Ratnakar, Kramer 2006)
  - Finite-field (Avestimehr, Diggavi, Tse 2007)
- Tight within constant gap for some Gaussian networks (Etkin, Tse, Wang 2006; Avestimehr, et al. 2007)
- Cutset bound is not tight in general (Zhang 1988; Aleksic et al. 2007)
- Bound can be tightened for multiple sources (Kramer, Savari 2006)

Coding Schemes: Outline

- Decode–forward: relay channel (Cover, EG 1979)
  Extension to DMNs (Aref 1980; Kramer, Gastpar, Gupta 2005)
- Compress–forward: relay channel (Cover, EG 1979)
  Extension to DMNs (Kramer, Gastpar, Gupta 2005)
- Amplify–forward: Gaussian networks (Schein, Gallager 2000)
- Network coding: noiseless networks (Ahlswede, Cai, Li, Yeung 2000)
  Extension to deterministic networks (Avestimehr, Diggavi, Tse 2007)
  And erasure networks (Dana, Gowaikar, Palanki, Hassibi, Effros 2006)
- Noisy network coding: DMNs (EG, Kim 2009)
  Extensions to multi-source networks (Lim, Kim, Chung, EG 2010)
- Compute–forward: Gaussian networks (Nazer, Gastpar 2007)

Relay Channel

- The relay channel (van der Meulen 1971) is a 3-node DMN

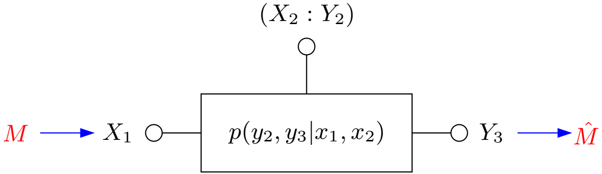

- Node 1 wishes to send M to node 3 with help of node 2 (relay)
- Capacity is not known in general
- Cutset upper bound simplifies to (Cover, EG 1979)

$$C \leq \max_{p(x_1,x_2)} \min \{ I(X_1, X_2; Y_3), I(X_1; Y_2, Y_3 | X_2) \}$$

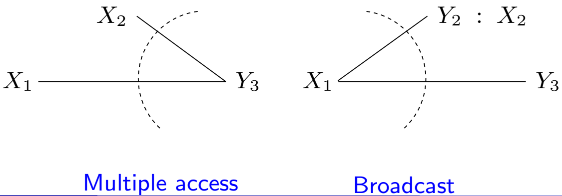

Decode–Forward Scheme

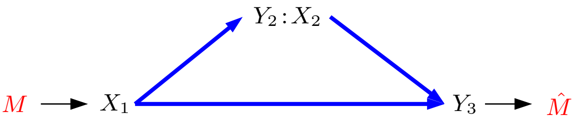

- Decode–forward can be viewed as a "digital-to-digital" interface

  Relay decodes message and coherently cooperates with sender to transmit it to receiver
- Use a block Markov scheme to send $b-1$ messages over b blocks

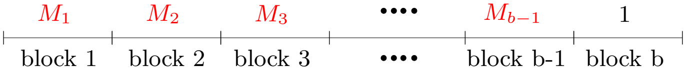

- At end of block $j \in [1:b-1]$, the relay decodes M_j
- Receiver decodes messages backwards after all b blocks are received (Willems, van der Meulen 1985)

Decode–Forward Lower Bound

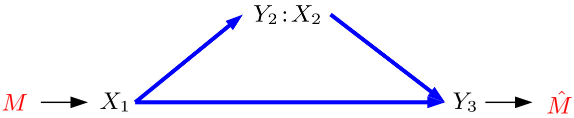

Decode–forward lower bound (Cover, EG 1979)

$$C \geq \max_{p(x_1,x_2)} \min\{I(X_1, X_2; Y_3), I(X_1; Y_2|X_2)\}$$

- Cutset upper bound

$$C \leq \max_{p(x_1,x_2)} \min\{I(X_1, X_2; Y_3), I(X_1; Y_2, Y_3|X_2)\}$$

- Bounds coincide when relay channel is physically degraded

Compress–Forward Scheme

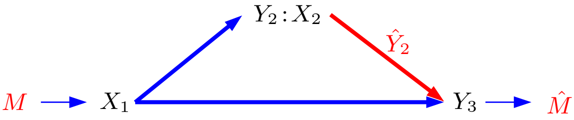

- Compress–forward can be viewed as an "analog-to-digital" interface

  Relay compresses its received signal and forwards it to receiver
- Node 1 transmits $b-1$ messages M_j, $j \in [1:2^{nR}]$, over b blocks
- At the end of block j:
  - Relay chooses reproduction sequence $\hat{y}_2^n(j)$ of $y_2^n(j)$
  - It uses Wyner–Ziv binning to reduce rate necessary to send $\hat{y}_2^n(j)$
  - It sends Compression bin index to receiver in block $j+1$ via $x_2^n(j+1)$
- At end of block $j+1$: Receiver decodes M_j sequentially
  1. It decodes compression bin index from which it finds $\hat{y}_2^n(j)$
  2. It then decodes m_j from $\hat{y}_2^n(j), y_3^n(j)$

Compress–Forward Lower Bound

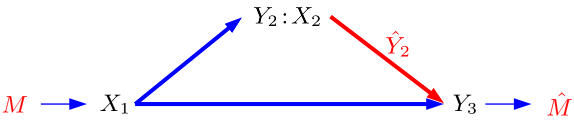

Compress–forward lower bound (Cover, EG 1979)

$$C \geq \max_{p(x_1)p(x_2)p(\hat{y}_2|y_2,x_2)} I(X_1; \hat{Y}_2, Y_3 | X_2)$$

subject to $I(X_2; Y_1) \geq I(Y_2; \hat{Y}_2 | X_2, Y_3)$

- Cutset bound:

$$C \leq \max_{p(x_1,x_2)} \min \{ I(X_1, X_2; Y_3), I(X_1; Y_2, Y_3 | X_2) \}$$

Amplify–Forward Scheme (Schein, Gallager 2000)

- Scheme for AWGN relay channel:

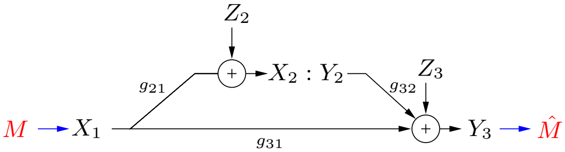

  g_{21}, g_{31}, g_{32} are channel gains

  Z_2, Z_2 are $\mathrm{N}(0,1)$

  Assume power constraint P on each sender

- Amplify–forward is an analog-to-analog interface

  The relay sends scaled version of its previously received symbol:
  $X_{2i} = aY_{2,i-1}$ for $i \in [1:n]$
  Amplification factor a picked to satisfy relay sender power constraint

- Obtain an ISI channel with known capacity

Comparison Between Schemes

- Consider AWGN relay channel

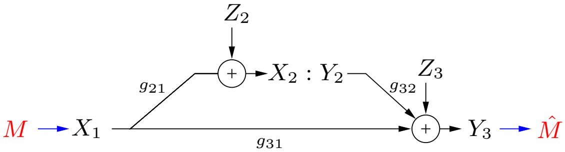

- Decode–forward: within 1/2 bit of cutset bound
- Compress–forward: within 1/2 bit of cutset (Chang et al. 2008)
- Amplify–forward: within 1 bit of cutset (Chang et al. 2008)
- Compress–forward always outperforms amplify–forward
- Compress–forward outperforms decode–forward if:
  $g_{21}^2 < g_{31}^2$ or $g_{21}^2 << g_{32}^2$
  Decode–forward is better, otherwise

Extensions to Multicast Networks

- Network decode–forward (Aref 1980; Xie, Kumar (2005); Kramer, Gastpar, Gupta 2005)
  - Decode–forward along a path
  - Bound tight for physically degraded network (Aref, EG 1981)
- Network compress–forward (Kramer, Gastpar, Gupta 2005)
  - Again use Wyner-Ziv binning and sequential decoding
  - Scheme extended by decode–forward of compression bin indices
- Amplify–forward can also be extended to Gaussian networks
  - Each relay sends a scaled version of its previously received symbol

Noiseless Multicast Network

- Consider noiseless network modeled by weighted graph

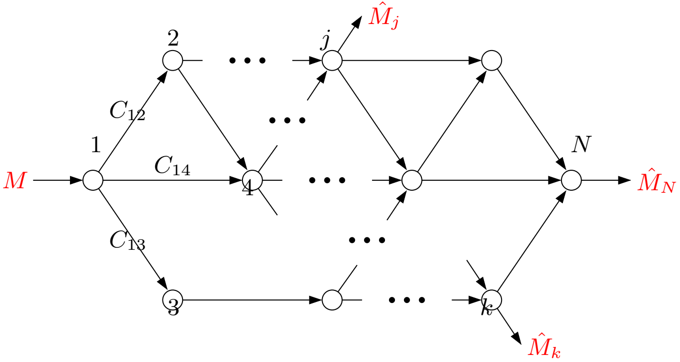

- Node 1 wishes to send message M to set of destination nodes $\mathcal{D}$
- Capacity coincides with cutset bound

Network Coding Theorem (Ahlswede, Cai, Li, Yeung 2000)

$$C = \min_{k\in\mathcal{D}} \min_{\mathcal{S}:\, 1\in\mathcal{S},\, k\in\mathcal{S}^c} C(\mathcal{S})$$

Outline of Proof: Acyclic Network

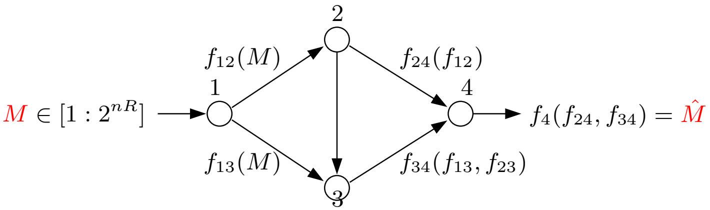

- Wolog assume zero node delay
- Use block coding (assume C_{jk} are integer valued)
- Random codebook generation:

  $f_{jk} \in [1:2^{nC_{jk}}]$, $(j,k) \in \mathcal{E}$, and f_4 are randomly and independently generated, each according to uniform pmf
- Key step: If $R < \min_{\mathcal{S}} C(\mathcal{S})$, $f_4(m)$ is one-to-one with high prob.
- Cutset bound can be achieved with zero error using linear network coding (Li, Yeung, Cai 2003; Koetter, Medard 2003)

Outline of Proof: Cyclic Network

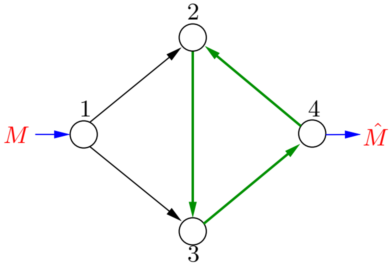

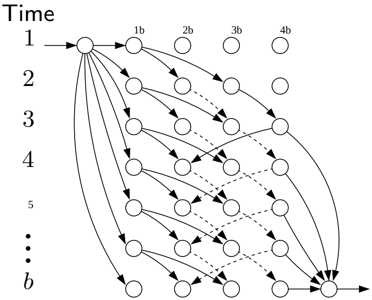

- Cannot assume zero delay nodes. Assume unit delay at each node
- Unfold to time extended (acyclic) network with b blocks
- Key step: Min-cut capacity of the new network is $\approx bC$ for b large
- By result for acyclic case, cutset for new network is achievable
- Key insight: Send same message b times using independent mappings

Deterministic Multicast Network

- Generalizes noiseless multicast network with broadcast, interference

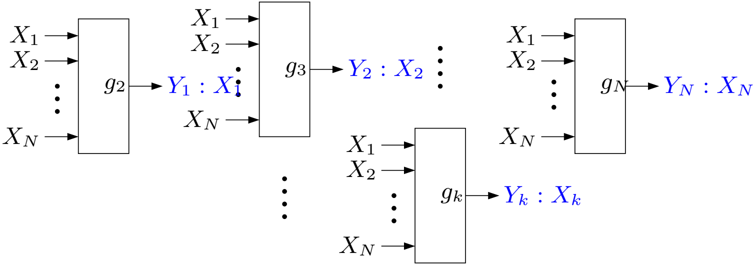

- Node 1 wishes to send message to subset of nodes $\mathcal{D}$
- Capacity is not known in general
- Cutset upper bound reduces to

$$C \leq \max_{p(x^N)} \min_{k \in \mathcal{D}} \min_{\mathcal{S}: 1 \in \mathcal{S},\, k \in \mathcal{S}^c} H(Y(\mathcal{S}^c) | X(\mathcal{S}^c))$$

Deterministic Multicast Network

Lower bound on capacity (Avestimehr, Diggavi, Tse 2007)

$$C \geq \max_{\prod_{j=1}^{N} p(x_j)} \min_{k \in \mathcal{D}} \min_{\mathcal{S}: 1 \in \mathcal{S},\, k \in \mathcal{S}^c} H(Y(\mathcal{S}^c) | X(\mathcal{S}^c))$$

- Cutset bound:

$$C \leq \max_{p(x^N)} \min_{k \in \mathcal{D}} \min_{\mathcal{S}: 1 \in \mathcal{S},\, k \in \mathcal{S}^c} H(Y(\mathcal{S}^c) | X(\mathcal{S}^c))$$

- Bounds coincide for:

  No interference (Ratnakar, Kramer 2006):

$$Y_k = (y_{k1}(X_1), \ldots, y_{kN}(X_N)),\ k \in [2:N]$$

  Finite-field network (Avestimehr, Diggavi, Tse 2007):

$$Y_k = \sum_{j=1}^{N} g_{jk} X_j \text{ for } g_{jk}, X_j \in \mathbb{F}_q,\ j \in [1:N],\ k \in [2:N]$$

  Used to approximate capacity of Gaussian networks in high SNR

Outline of Proof

- Layered networks:

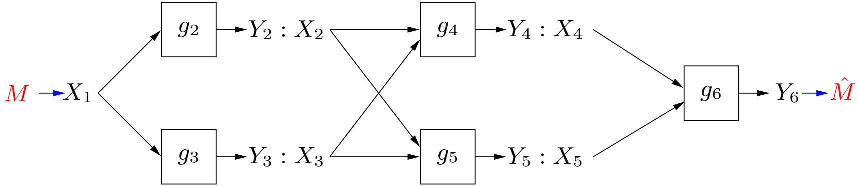

  - Random codebook generation:
    Randomly and independently generate $x_j^n(y_j^n)$ for each sequence y_j^n
  - Key step: If R satisfies lower bound, end-to-end mapping is one-to-one with high probability

- Non-layered network:
  - Construct time extended (layered) network with b blocks
  - Key step: If R satisfies lower bound, end-to-end mapping is one-to-one with high probability
  - Again send the same message b times using independent mappings

Noisy Network Coding Scheme

- Alternative characterization of compress–forward lower bound:

Compress–forward lower bound (EG, Mohseni, Zahedi 2006)

$$C \geq \max_{p(x_1)p(x_2)p(\hat{y}_2|y_2,x_2)} \min\{I(X_1, X_2; Y_3) - I(Y_2; \hat{Y}_2 | X_1, X_2, Y_3), \\ I(X_1; \hat{Y}_2, Y_3 | X_2)\}$$

- Original compress–forward lower bound (Cover, EG 1979):

$$C \geq \max_{p(x_1)p(x_2)p(\hat{y}_2|y_2,x_2)} I(X_1; \hat{Y}_2, Y_3 | X_2)$$

subject to $I(X_2; Y_3) \geq I(Y_2; \hat{Y}_2 | X_2, Y_3)$

- Cutset bound:

$$C \leq \max_{p(x_1,x_2)} \min\left\{I(X_1, X_2; Y_3), I(X_1; Y_2, Y_3 | X_1)\right\}$$

Noisy Network Coding Scheme

- The alternative characterization of compress–forward lower bound for relay channel generalizes naturally to noisy multicast networks

Theorem (EG, Kim Lecture on NIT 2009)

$$C \geq \max \min_{k \in \mathcal{D}} \min_{\substack{\mathcal{S} \subseteq [1:N] \\ 1 \in \mathcal{S},\, k \in \mathcal{S}^c}} \big(I(X(\mathcal{S}); \hat{Y}(\mathcal{S}^c), Y_k | X(\mathcal{S}^c)) - I(Y(\mathcal{S}); \hat{Y}(\mathcal{S}) | X^N, \hat{Y}(\mathcal{S}^c), Y_k) \big),$$

where the maximum is over $\prod_{k=1}^{N} p(x_k) p(\hat{y}_k | y_k, x_k)$

- Includes as special cases:
  - Capacity of noiseless multicast networks
  - Lower bound on deterministic multicast networks
  - Capacity of wireless erasure muticast networks (Dana, Gowaikar, Palanki, Hassibi, Effros 2006)
- Simpler and more general proof (deals directly with cyclic networks)

Outline of Proof

- Source node sends same message b times; relays use compress–forward; decoders use simultaneous decoding
- No Wyner–Ziv binning

  Do not require decoding compression indices correctly!
- For simplicity, consider proof for relay channel

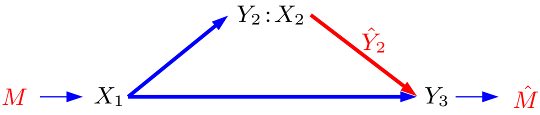

- The relay uses independently generated compression codebooks:

$$\mathcal{B}_j = \{\hat{y}_2^n(l_j|l_{j-1}) : l_j, l_{j-1} \in [1:2^{nR_2}]\},\ j \in [1:b]$$

  l_j is compression index of $\hat{Y}_2^n(j)$ sent by relay in block $j+1$
- The senders use independently generated transmission codebooks:

$$\mathcal{C}_j = \{(x_1^n(j,m), x_2^n(l_{j-1})) : m \in [1:2^{nbR}],\ l_{j-1} \in [1:2^{nR_2}]\}$$

Outline of Proof

Block	1	2	3	…	$b-1$	b
X_1	$x_1^n(1,m)$	$x_1^n(2,m)$	$x_1^n(3,m)$	…	$x_1^n(b-1,m)$	$x_1^n(b,m)$
Y_2	$\hat{y}_2^n(l_1\|1)$	$\hat{y}_2^n(l_2\|l_1)$	$\hat{y}_2^n(l_3\|l_2)$	…	$\hat{y}_2^n(l_{b-1}\|l_{b-2})$	$\hat{y}_2^n(l_b\|l_{b-1})$
X_2	$x_2^n(1)$	$x_2^n(l_1)$	$x_2^n(l_2)$	…	$x_2^n(l_{b-2})$	$x_2^n(l_{b-1})$
Y_3	$\emptyset$	$\emptyset$	$\emptyset$	…	$\emptyset$	$\hat{m}$

- Decoding: After receiving all blocks $y_3^n(j)$, $j \in [1:2^{nR}]$, the receiver finds unique $\hat{m}$ such that:

$$(x_1^n(j,\hat{m}), \hat{y}_2^n(l_j|l_{j-1}), x_2^n(l_{j-1}), y_3^n(j)) \in \mathcal{T}_\epsilon^{(n)}$$

for all $j \in [1:b]$ and for some $l_1, l_2, \ldots, l_b$

Extension: Noisy Multi-source Multicast Network

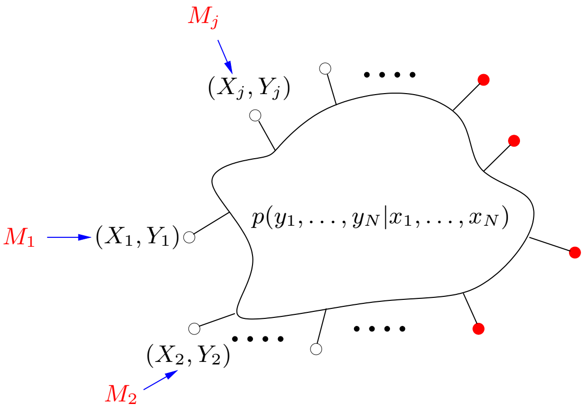

- Noisy network coding generalizes to this case (Lim et al. 2010)
- Includes results on erasure, deterministic networks (Dana, Gowaikar, Palanki, Hassibi, Effros 2006; Perron 2009) as special cases

Extension: Multi-source Multicast Gaussian Networks

- Channel model: $Y^N = GX^N + Z^N$
  - G is network gain matrix
  - Z^N is i.i.d. $\mathrm{N}(0, 1)$
  - Power constraint P on every sender X_k, $k \in [1:N]$
- Noisy network coding can be extended to this case:
  - Extend scheme to DMN with input cost
  - Apply discretization procedure in (EG, Kim LN-NIT 2009)
- Optimal distribution on inputs X_ks and $\hat{Y}$s is not known
- Assume $X_j \sim \mathrm{N}(0, P)$, and $\hat{Y}_j = Y_j + \hat{Z}_j$, $\hat{Z}_j \sim \mathrm{N}(0, 1)$

Extension: Multi-source Multicast Gaussian Networks

- Noisy network coding bound with these choices yields

$$\sum_{j \in \mathcal{S}} R_j < \frac{1}{2} \log \left| I + \frac{P}{2} G(\mathcal{S}) G(\mathcal{S})^T \right| - \frac{|\mathcal{S}|}{2}$$

- Cutset bound is upper bounded as

$$\sum_{j \in \mathcal{S}} R_j \leq \frac{1}{2} \log \left| I + G(\mathcal{S}) K(\mathcal{S}) G(\mathcal{S})^T \right|,$$

  $K(\mathcal{S})$ is covariance matrix of $X(\mathcal{S})$
- By loosening cutset bound further, can show that noisy network coding is within $(N/2)\log 6 \approx 1.3N$ bits/trans of cutset bound
- This improves previous results (Avestimehr, Diggavi, Tse 2007; Perron 2009)

Example: AWGN Two-Way Relay

- AWGN two-way relay channel is a 3-node DMN

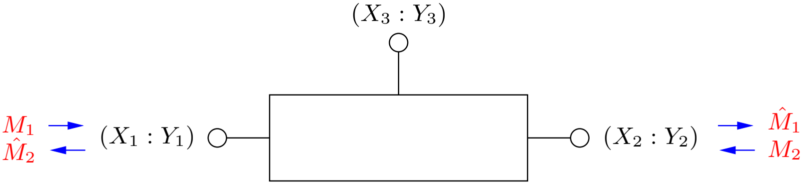

- $Y_k = \sum_{j \neq k} g_{kj} X_j + Z_k$ for $k = 1, 2, 3$,
  $Z_k \sim \mathrm{N}(0, 1)$
  Assume power constraint P on every sender
- Node 1 wishes to send message M_1 to node 2
  Node 2 wishes to send message M_2 to node 1

Example: AWGN Two-Way Relay Channel

- Extensions of decode–forward, compress–forward, and amplify–forward compared (e.g., Rankov, Wittneben 2006; Katti, Maric, Goldsmith, Katabi, Medard 2007)

Node 1 to 2 distance: 1; node 1 to 3 distance: $d \in [0, 1]$; $g_{13} = g_{31} = d^{-3/2}$, $g_{23} = g_{32} = (1 - d)^{-3/2}$

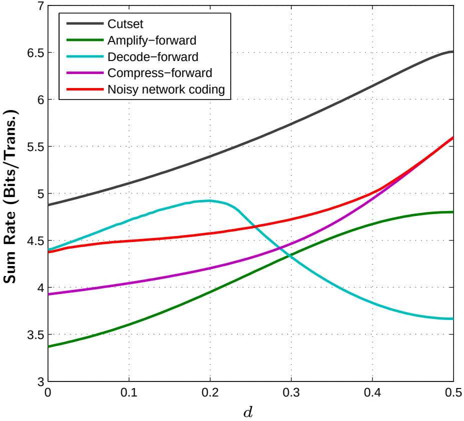

Example: AWGN Two-Way Relay Channel

- NNC sum rate within 1.5 bit of cutset bound
- Gap unbounded (as $P \to \infty$) for all other schemes

$g_{12} = g_{21} = 0.1,\ g_{13} = g_{32} = 0.5,\ g_{23} = g_{31} = 2$

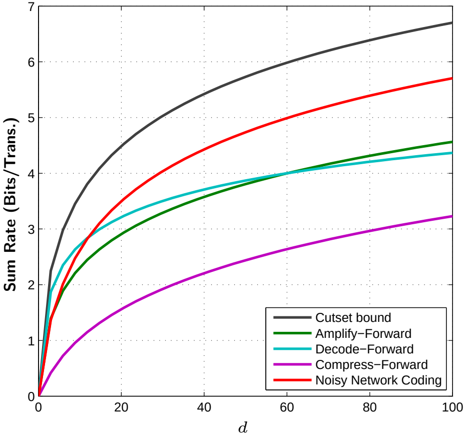

Extension: Multi-unicast Networks

- For example, consider an N-node DMN:
  node 1 wishes to send message M_1 to node 3
  node 2 wishes to send message M_2 to node 4
- Using noisy network coding, can view network as interference channel with senders X_1 and X_2 and respective receivers $(Y_3, \hat{Y}_5, \hat{Y}_6, \ldots, \hat{Y}_N)$ and $(Y_4, \hat{Y}_5, \hat{Y}_6, \ldots, \hat{Y}_N)$
- Can use coding strategies for interference channel:
  - Each receiver decodes only its message (treats interference as noise)
  - Each receiver decodes both messages
  - One receiver uses the former strategy, the other uses the later
- Presentation by Lim on Friday morning in Gaussian Networks session

Example: AWGN Interference Relay Channel

- AWGN interference channel with a relay is a 5-node network:

$$Y_k = g_{k1}X_1 + g_{k2}X_2 + Z_k \text{ for } k = 3, 4, 5$$

$Z_k \sim \mathrm{N}(0, 1)$ and power constraint P on each sender

Noiseless broadcast link with capacity C_0 from node 5 to nodes 3,4

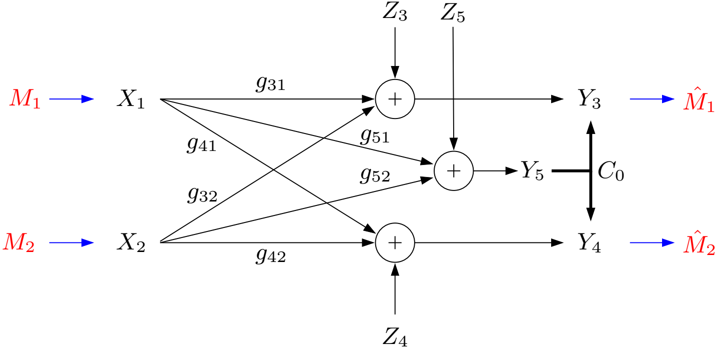

Example: AWGN Interference Relay Channel

- Extensions of compress–forward and hash–forward (Cover, Kim 2007) compared (Razaghi, Yu 2010)

$$g_{31} = g_{42} = 1,\ g_{41} = g_{32} = g_{51} = 0.5,\ g_{52} = 0.1,\ C_0 = 1$$

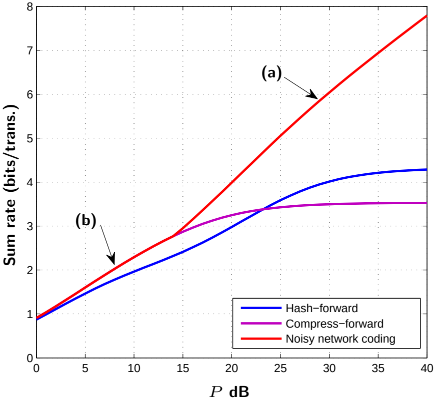

(a) treating interference as noise (b) decoding both messages

Noisy Network Coding: Summary

- Noisy network coding generalizes network coding and its extensions and compress–forward to noisy networks
- Performance comparison:
  - Noisy network coding strictly outperforms network compress–forward
  - Can outperform network decode–forward and amplify–forward
  - Achieves tightest known gap to cutset bound for Gaussian multi-source multicast networks

    Other schemes have arbitrarily large gap (2-way relay)
- Take aways:
  - Proving achievability for DMN can be easier and cleaner than for special cases (noiseless, deterministic, Gaussian)
  - Simultaneous decoding is more powerful than sequential decoding
- To learn more: Paper at: http://arxiv.org/abs/1002.3188, Friday talk

Compute–forward Scheme

- Idea spurred by distributed computing of mod-2 sum of DSBS (Körner, Marton 1979):

  $X_1, X_2 \sim \mathrm{Bern}(1/2)$; $X_1 \oplus X_2 = Z \sim \mathrm{Bern}(p)$ separately encoded; decoder wishes to find Z losslessly

- Slepian–Wolf achieves sum rate $H(X_1, X_2) = 1 + H(p)$

- Körner, Marton showed that minimum sum rate is $2H(p)$

  Achieved using same random linear code

- Scheme builds on previous work on structured codes

- Zamir's plenary talk on Friday

Compute–forward Scheme

- Consider AWGN 2-way relay with no direct links

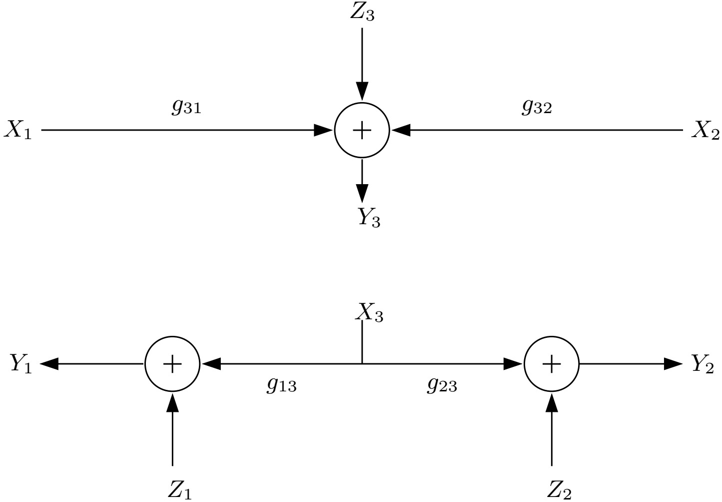

- In decode–forward, the relay decodes both messages
- Compute–forward: Decode–forward $g_{31}X_1^n + g_{32}X_2^n$ instead

Random Code Is Not Good

- Using Gaussian random codes, compute–forward ≡ decode–forward
  Sum of two codewords uniquely determines individual codewords

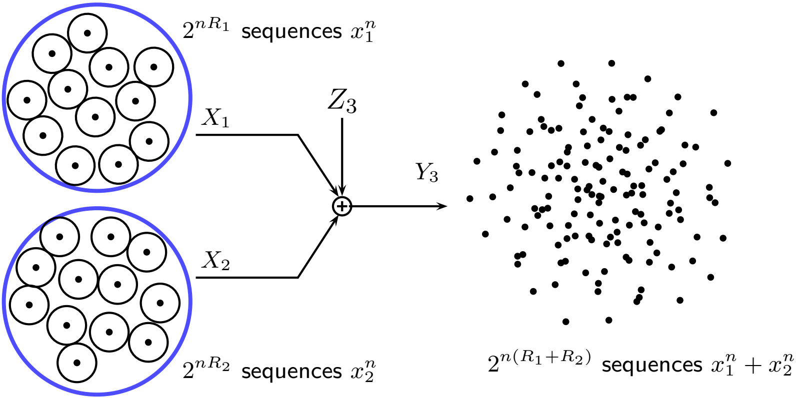

Compute–forward Scheme

- Senders use same lattice code (Nazer, Gastpar 2007; Narayanan, Wilson, Sprintson 2007)

  Sum of two codewords is a codeword

  Allows us to increase transmission rates

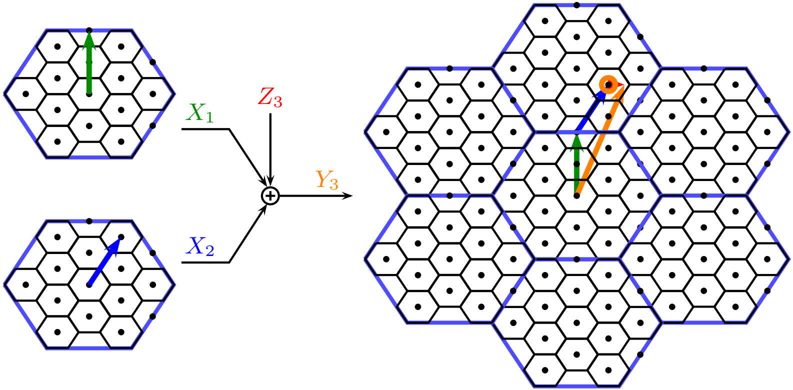

Compute–forward Can Outperform Other Schemes

- Sum rate within 0.58 bit of cutset bound
  Noisy network coding within 1 bit

$g_{31} = g_{32} = 1$ and $g_{13} = g_{23} = 2$

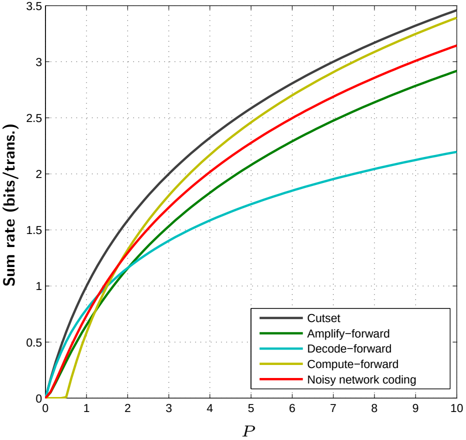

But Not Always

$$g_{31} = g_{32} = 0.5 \text{ and } g_{13} = g_{23} = 1$$

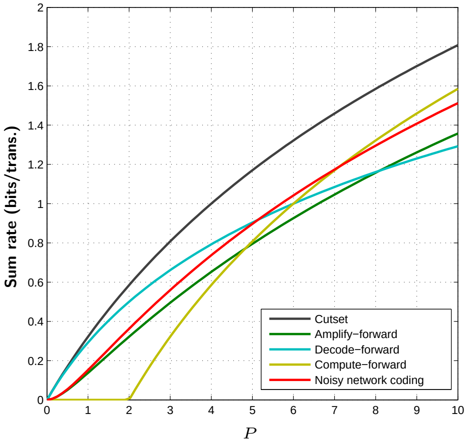

Conclusion

- Presented two recently developed coding schemes:
  - Noisy network coding: General scheme for noisy networks
  - Compute–forward: Specialized scheme for Gaussian networks
- Compared them to other schemes
- Many open questions:
  - How to apply noisy network coding to wireless networks
  - How to generalize compute–forward to arbitrary Gaussian networks
  - Compute–forward may be viewed as "structured" decode–forward
    How about "structured" noisy network coding
  - How to combine decode–forward and noisy network coding
  - How about combining compute–forward and noisy network coding
  - . . .
- We may be long way from finding optimal coding scheme for noisy networks

  But we may be closer than we think . . .

Acknowledgments

- Young-Han Kim
- Sung Hoon Lim
- Sae-Young Chung
- Michael Gastpar
- Bernd Bandemer
- Partial support from DARPA ITMANET