BATS: Achieving the Capacity of Networks with Packet Loss

Raymond W. Yeung

Institute of Network Coding
The Chinese University of Hong Kong

WCI 2013

Joint work with Shenghao Yang (IIIS, Tsinghua U)

Outline

1. File Transmission in Packet Networks
2. Two Classes of Solutions
   - Fountain Codes in Networks
   - Chunked Codes
3. BATS Codes

Linear Network Coding

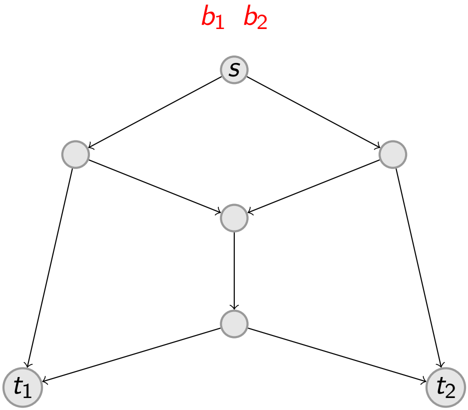

[ACLY00] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," IEEE Trans. Inform. Theory, 46(4):1204–1216, Jul 2000.

Linear Network Coding

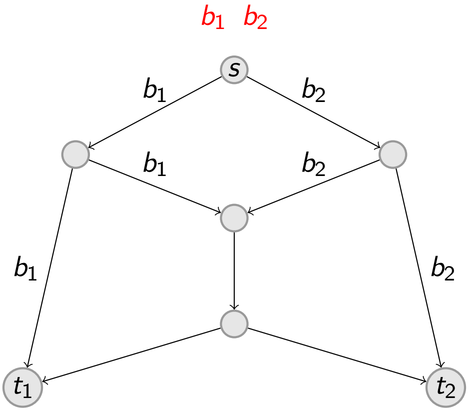

[ACLY00] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," IEEE Trans. Inform. Theory, 46(4):1204–1216, Jul 2000.

Linear Network Coding

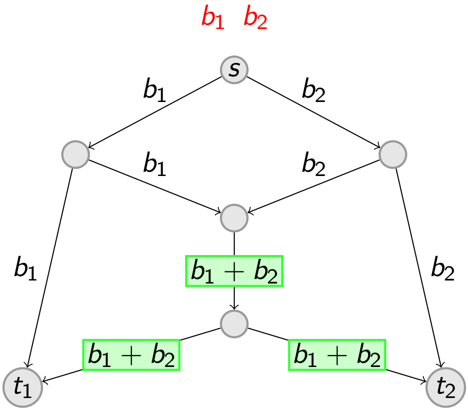

[ACLY00] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," IEEE Trans. Inform. Theory, 46(4):1204–1216, Jul 2000.

Transmission through Packet Networks (Erasure Networks)

One 20MB file ≈ 20,000 packets

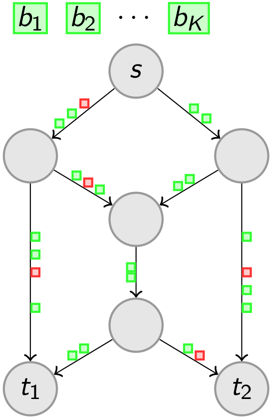

Transmission through Packet Networks (Erasure Networks)

One 20MB file ≈ 20,000 packets

A practical solution

- low computational and storage costs
- high transmission rate
- small protocol overhead

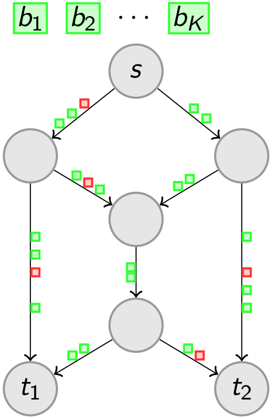

Routing Networks

Retransmission

- Example: TCP
- Not scalable for multicast
- Cost of feedback

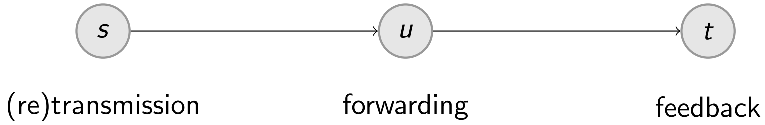

Routing Networks

Retransmission

- Example: TCP
- Not scalable for multicast
- Cost of feedback

Forward error correction

- Example: fountain codes
- Scalable for multicast
- Neglectable feedback cost

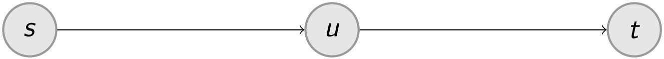

encoding forwarding decoding

Complexity of Fountain Codes with Routing

- K packets, T symbols in a packet.
- Encoding: $\mathcal{O}(T)$ per packet.
- Decoding: $\mathcal{O}(T)$ per packet.
- Routing: $\mathcal{O}(1)$ per packet and fixed buffer size.

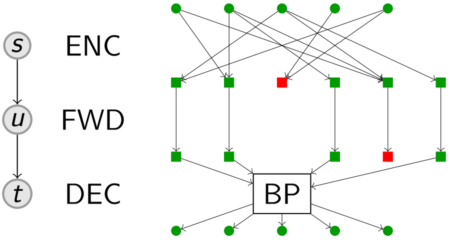

[Luby02] M. Luby, "LT codes," in Proc. 43rd Ann. IEEE Symp. on Foundations of Computer Science, Nov. 2002.

[Shokr06] A. Shokrollahi, "Raptor codes," IEEE Trans. Inform. Theory, vol. 52, no. 6, pp. 2551-2567, Jun 2006.

Achievable Rates

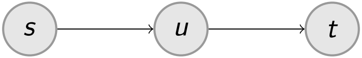

Both links have a packet loss rate 0.2.

The capacity of this network is 0.8.

Intermediate	End-to-End	Maximum Rate
forwarding	retransmission	0.64
forwarding	fountain codes	0.64

Achievable Rates

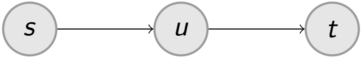

Both links have a packet loss rate 0.2.

The capacity of this network is 0.8.

Intermediate	End-to-End	Maximum Rate
forwarding	retransmission	0.64
forwarding	fountain codes	0.64
network coding	random linear codes	0.8

Achievable Rates: n hops

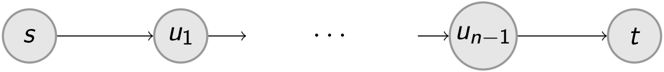

All links have a packet loss rate 0.2.

Intermediate Operation	Maximum Rate
forwarding	$0.8^n \to 0,\ n \to \infty$
network coding	0.8

An Explanation

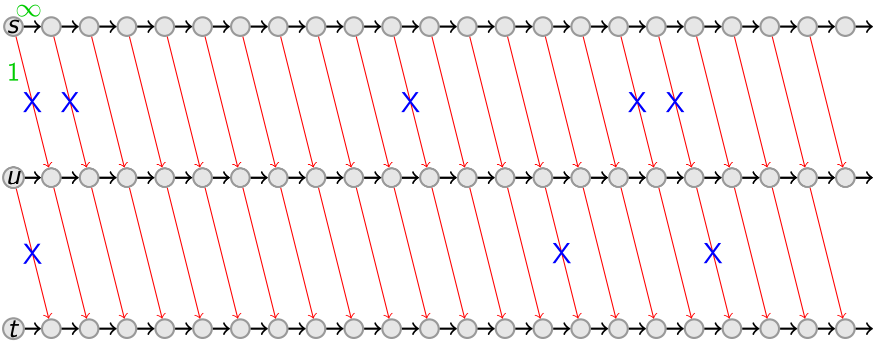

Multicast capacity of erasure networks

Theorem

Random linear network codes achieve the capacity of a large range of multicast erasure networks.

[Wu06] Y. Wu, "A trellis connectivity analysis of random linear network coding with buffering," in Proc. IEEE ISIT 06, Seattle, USA, Jul. 2006.

LMKE08] D. S. Lun, M. Médard, R. Koetter, and M. Effros, "On coding for reliable communication over packet networks," Physical Communication, vol. 1, no. 1, pp. 320, 2008.

Complexity of Linear Network Coding

- Encoding: $\mathcal{O}(TK)$ per packet.
- Decoding: $\mathcal{O}(K^2 + TK)$ per packet.
- Network coding: $\mathcal{O}(TK)$ per packet. Buffer K packets.

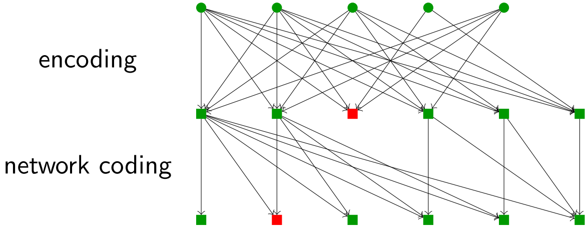

Quick Summary

Routing + fountain

- low complexity
- low rate

Network coding

- high complexity
- high rate

Outline

Fountain Codes with Coding in Intermediate Nodes

- Network coding changes the degree distribution of the received packets such that the low decoding complexity cannot be guaranteed.
- Works for special cases: P2P file sharing [CHKS09] and line networks [GS08].
  - Difficult to extend.
  - In the intermediate nodes, computational cost is $\mathcal{O}(TK)$ per packet and storage cost is K packets.

[CHKS09] M.-L. Champel, K. Huguenin, A.-M. Kermarrec, and N. L. Scouarnec. LT network codes. Research Report RR-7035, INRIA, 2009.

[GS08] R. Gummadi and R. Sreenivas. Relaying a fountain code across multiple nodes. In Proc. IEEE ITW 08, pages 149–153, May 2008.

Chunk Based Network Coding

- Using chunks to reduce complexity [CWJ03]
  - Encoding complexity: $\mathcal{O}(TKL)$
  - Decoding complexity: $\mathcal{O}(KL^2 + TKL)$
- Buffer requirement in the intermediate nodes?

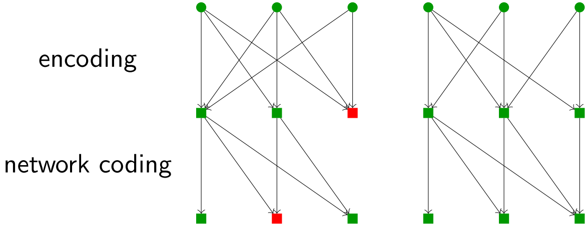

[CWJ03] P. A. Chou, Y. Wu, and K. Jain. Practical network coding. In Proc. Allerton Conf. Comm., Control, and Computing, Oct. 2003.

Scheduling of Chunks

- Sequential scheduling of chunks
  - Protocol overhead
  - Not scalable for multicast

Scheduling of Chunks

- Sequential scheduling of chunks
  - Protocol overhead
  - Not scalable for multicast
- Random scheduling of chunks [MHL06]
  - Intermediate network nodes cache K packets.
  - Less efficient when a major fraction of all the chunks have been decoded.

[MHL06] P. Maymounkov, N. J. A. Harvey, and D. S. Lun. Methods for efficient network coding. In Proc. Allerton Conf. Comm., Control, and Computing, Sept. 2006.

Scheduling of Chunks

- Sequential scheduling of chunks
  - Protocol overhead
  - Not scalable for multicast
- Random scheduling of chunks [MHL06]
  - Intermediate network nodes cache K packets.
  - Less efficient when a major fraction of all the chunks have been decoded.
  - Overlapped Chunks [SZK09] [HB10]
    - Improve the throughput of random scheduling
    - Cannot reduce the buffer size

[MHL06] P. Maymounkov, N. J. A. Harvey, and D. S. Lun. Methods for efficient network coding. In Proc. Allerton Conf. Comm., Control, and Computing, Sept. 2006.

[SZK09] D. Silva, W. Zeng, and F. R. Kschischang. Sparse network coding with overlapping classes. In Proc. NetCod 09, pages 74–79, 2009.

[HB10] A. Heidarzadeh and A. H. Banihashemi. Overlapped chunked network coding. In Proc. ITW 10, pages 1–5,2010.

Recent Progress in Chunked Codes

- Expander graph based approach [TYY12] and LDPC-like approach [MAB12] for chunked codes
  - For an erasure channel, the overhead is reduced to less than 7%.
  - But the performance for general networks is difficult to analyze.

[TYY12] B. Tang, S. Yang, Y. Yin, B. Ye, S. Lu, "Expander Graph Based Overlapped Chunked Codes", In Proc. IEEE Inte. Symp. on Information Theory ISIT '12, Cambridge, MA, USA, 2012.

[MAB12] K. Mahdaviani, M. Ardakani, H. Bagheri, and C. Tellambura, "Gamma codes: a low-overhead linear-complexity network coding solution," in Proceedings of 2012 International Symposium on Network Coding (NetCod), 2012, pp. 125130.

Outline

Batched Sparse (BATS) Codes

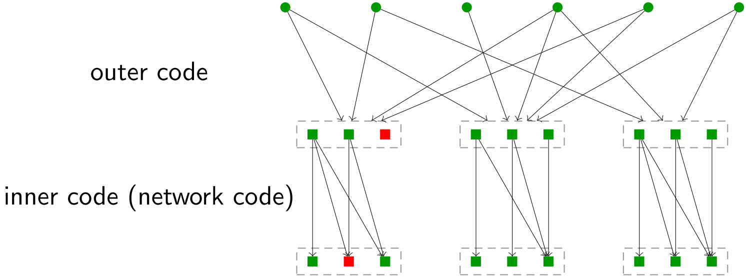

[YY11] S. Yang and R. W. Yeung. Coding for a network coded fountain. ISIT 2011, Saint Petersburg, Russia, 2011.

Encoding of BATS Code: Outer Code

- Apply a "matrix fountain code" at the source node:
  1. Obtain a degree d by sampling a degree distribution Ψ.
  2. Pick d distinct input packets randomly.
  3. Generate a batch of M coded packets using the d packets.
- Transmit the batches sequentially.

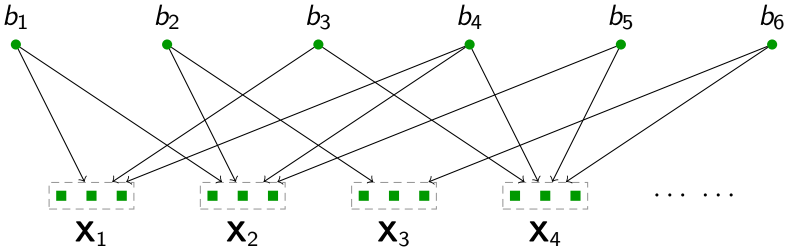

$$\mathbf{X}_i = \begin{bmatrix} b_{i1} & b_{i2} & \cdots & b_{id_i} \end{bmatrix} \mathbf{G}_i = \mathbf{B}_i \mathbf{G}_i.$$

Encoding of BATS Code: Inner Code

- The batches traverse the network.
- Encoding at the intermediate nodes forms the inner code.
- Linear network coding is applied in a causal manner within a batch.

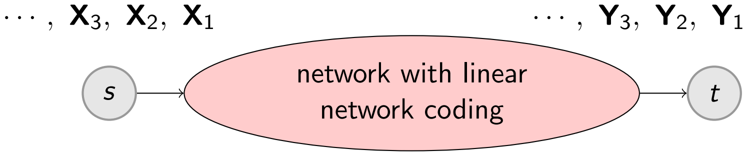

$$\mathbf{Y}_i = \mathbf{X}_i \mathbf{H}_i, \ i = 1, 2, \ldots.$$

Belief Propagation Decoding

1. Find a check node i with $\text{degree}_i = \text{rank}(\mathbf{G}_i\mathbf{H}_i)$.
2. Decode the ith batch.
3. Update the decoding graph. Repeat 1).

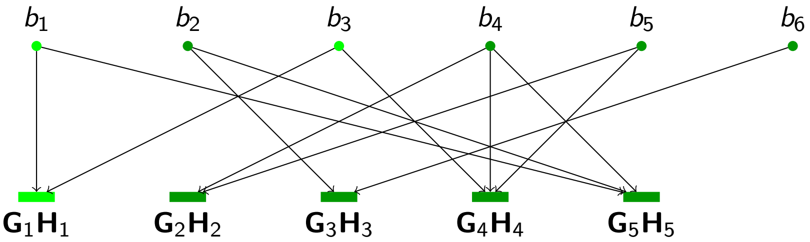

The linear equation associated with a check node: $\mathbf{Y}_i = \mathbf{B}_i\mathbf{G}_i\mathbf{H}_i$.

Precoding

- Precoding by a fixed-rate erasure correction code.
- The BATS code recovers $(1 - \eta)$ of its input packets.

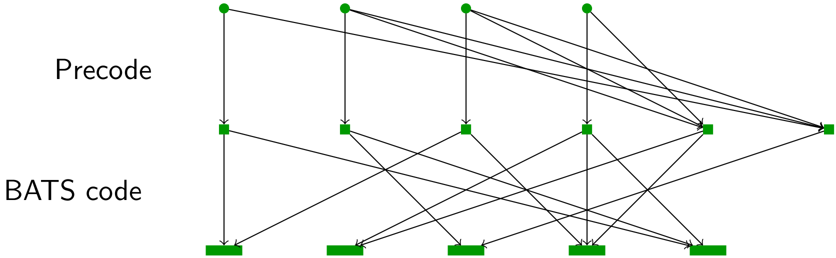

[Shokr06] A. Shokrollahi, Raptor codes, IEEE Trans. Inform. Theory, vol. 52, no. 6, pp. 25512567, Jun. 2006.

Degree Distribution

We need a degree distribution Ψ such that

1. The BP decoding succeeds with high probability.
2. The encoding/decoding complexity is low.
3. The coding rate is high.

A Sufficient Condition

Define

$$\Omega(x) = \sum_{r=1}^{M} h^*_{r,r} \sum_{d=r+1}^{D} d\Psi_d I_{d-r,r}(x) + \sum_{r=1}^{M} h_{r,r} r \Psi_r,$$

where $h^*_{r,r}$ is related to the rank distribution of H and $\mathrm{I}_{a,b}(x)$ is the *regularized incomplete beta function*.

Theorem

Consider a sequence of decoding graph BATS$(K, n, \{\Psi_{d,r}\})$ with constant $\theta = K/n$. The BP decoder is asymptotically error free if the degree distribution satisfies

$$\Omega(x) + \theta \ln(1 - x) > 0 \quad \text{for } x \in (0, 1 - \eta),$$

An Optimization Problem

$$
\begin{aligned}
\max \quad & \theta \\
\text{s.t.} \quad & \Omega(x) + \theta \ln(1-x) \geq 0, \quad 0 < x < 1-\eta \\
& \Psi_d \geq 0, \quad d = 1, \cdots, D \\
& \sum_d \Psi_d = 1.
\end{aligned}
$$

- $D = \lceil M/\eta \rceil$
- Solver: Linear programming by sampling some x.

Complexity of Sequential Scheduling

Source node encoding		$\mathcal{O}(TM)$ per packet
Destination node decoding		$\mathcal{O}(M^2 + TM)$ per packet
Intermediate Node	buffer	$\mathcal{O}(TM)$
	network coding	$\mathcal{O}(TM)$ per packet

- T: length of a packet
- K: number of packets
- M: batch size

Achievable Rates

Optimization

$$
\begin{aligned}
\max \quad & \theta \\
\text{s.t.} \quad & \Omega(x_k) + \theta \ln(1 - x_k) \geq 0, \quad x_k \in (0, 1 - \eta) \\
& \Psi_d \geq 0, \quad d = 1, \cdots, \lceil M/\eta \rceil \\
& \sum_d \Psi_d = 1.
\end{aligned}
$$

- The optimal values of θ is very close to $\mathrm{E}[\mathrm{rank}(H)]$.
- It can be proved when $\mathrm{E}[\mathrm{rank}(H)] = M \Pr\{\mathrm{rank}(H) = M\}$.

Simulation Result

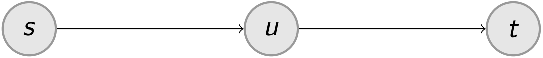

- Packet loss rate 0.2.
- Node s encodes K packets using a BATS code.
- Node u caches only one batch.
- Node t sends one feedback after decoding.

Coding rates obtained by simulation for $M = 32$

K	$q = 2$	$q = 4$	$q = 8$	$q = 16$
16000	0.5826	0.6145	0.6203	0.6248
32000	0.6087	**0.6441**	**0.6524**	**0.6574**
64000	0.6259	**0.6655**	**0.6762**	**0.6818**

M: batch size

K: number of packets

q: field size

Tradeoff

- $M = 1$: BATS codes degenerate to Raptor codes.
  - Low complexity
  - No benefit of network coding
- $M = K$ and $\mathrm{degree} \equiv K$: BATS codes becomes RLNC.
  - High complexity
  - Full benefit of network coding.
- Exist parameters with moderate values that give very good performance

Recent Developments

- Degree distribution optimization
  - Degree distribution depends on the rank distribution.
  - Robust degree distribution for different rank distributions.
  - Inactivation decoding alleviates the degree distribution optimization problem.
- Finite length analysis [3]
- Prototypes
  - 2-hop wireless transmission on 802.11
  - 802.11 based mesh network
  - Peer-to-peer network
- Other potential applications
  - Mobile ad hoc networks
  - Vehicular communication systems
  - Satellite delivery of digital films

Experiment setting

Experiment setting

Experiment setting

- Sender/receiver: a laptop with open source Atheros wireless drivers.
- Relay: one wireless router with Atheros chipset running OpenWrt (about 150HKD/20USD)
- WiFi 802.11 b/g/n at 2.4GHz
- Sender's rate is set to 1 Mb/s to reduce the effect of the router's low computation power.

Two modes of WiFi

- Normal 802.11
  - ACK and MAC layer retransmission (4 retries)
  - Benchmark
- Modified 802.11
  - without ACK and without MAC layer retransmission
  - Testing of BATS codes

Two modes of the router

- A normal wireless router
- Recoding: cache one batch and send random linearly coded packets

End-to-end packet loss

- Tested by iperf, without recoding
- Normal 802.11: five trials
  - end-to-end loss rate: 0%, 8.2%, 0%, 1%, 0%
- Modified 802.11: five trials
  - end-to-end loss rate: 30%, 26%, 55%, 24%, 28%
  - average capacity without recoding: 674 Kb/s
  - estimated link loss rate: 16%, 14%, 33%, 13%, 15%
  - estimated capacity with recoding: 818 Kb/s (assuming the loss rate on both links are the same)

BATS code parameters

- batch size $= 16$, finite field $= GF(2^8)$
- no. source packets $= 2{,}000$
- file size $= 2.835$ MByte
- Inactivation decoding with at most 256 inactivations
- coding overhead $= \frac{\text{no. received packets} - \text{no. source packets}}{\text{no. source packets}}$
- protocol overhead, including coding vectors (16 bytes per packet), batch IDs (2 bytes per packet): $< 2\%$

Results

- BATS codes with recoding in router (modified 802.11): three trials

	trial 1	trial 2	trial 3
rate (Kb/s)	588.45	600.62	589.495
overhead	5.3%	4.2%	6.6%

- BATS codes W/O recoding in router (modified 802.11): three trials

	trial 1	trial 2	trial 3
rate (Kb/s)	521.60	540.53	529.81
overhead	3.5%	7.9%	3.9%

Compare with estimated capacity

	Average rate	Est. capacity	Ratio
BATS w/ recoding	592.86 Kb/s	818 Kb/s	72.5%
BATS w/o recoding	530.65 Kb/s	674 Kb/s	78.7%

Compare with TCP

	Average rate (Kb/s)
BATS w/ recoding	592.86
BATS w/o recoding	530.65
TCP (normal 802.11)	420.33

Potential Improvements

- Better tuning of the degree distribution
- Use co-processors for encoding and decoding
- Larger file size would benefit more ($\sim$ 10MB)

Summary

- BATS codes provide a digital fountain solution with linear network coding:
  - Outer code at the source node is a matrix fountain code.
  - Linear network coding at the intermediate nodes forms the inner code.
  - Prevents BOTH packet loss and delay from accumulating along the way.
- The more hops between the source node and the sink node, the larger the benefit.
- Future work:
  - Proof of (nearly) capacity achieving
  - Design of intermediate operations to maximize the throughput and minimize the buffer size

References

S. Yang, R. W. Yeung,
"Batched Sparse Codes,"
submitted to *IEEE Trans. Inform. Theory*, 2012.

S. Yang, R. W. Yeung,
"Large File Transmission in Network-Coded Networks with Packet Loss – A Performance Perspective,"
ISABEL 2011, Barcelona, Spain, 2011.

T. C. Ng, S. Yang,
"Finite length analysis of BATS codes,"
NetCod 2013, Calgary, Canada, June, 2013.