

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: November 26, 2020

S. Yang
CUHK(SZ)
X. Huang
R.W. Yeung
CUHK
J.K. Zao
NCTU
May 25, 2020

BATS Coding Scheme for Multi-hop Data Transport
draft-yang-nwcr-g-bats-03

Abstract

This document describes a BATS coding scheme for communication through multi-hop networks. BATS code is a class of efficient linear network coding scheme with a matrix generalization of fountain codes as the outer code, and batch-based linear network coding as the inner code.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 26, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	3
2.	Procedures	3
2.1.	Introduction	3
2.2.	Data Delivery Procedures	4
2.2.1.	Source Node Data Partitioning and Padding	5
2.2.2.	Source Node Outer Code Encoding Procedure	5
2.2.3.	Recoding Procedures	7
2.2.4.	Destination Node Procedures	7
2.3.	Recommendation for the Parameters	8
2.4.	Example DDP Packet Format	8
2.4.1.	Packet Header	8
2.4.2.	Packet Payload	9
2.4.3.	Packet Footer	10
3.	BATS Code Specification	10
3.1.	Common Parts	10
3.2.	Outer Code Encoder	11
3.3.	Inner Code Encoder (Recoder)	12
3.4.	Belief Propagation Decoder	13
4.	IANA Considerations	13
5.	Security Considerations	13
5.1.	Provision of Confidentiality Protection	14
5.2.	Countermeasures against Pollution Attacks	14
6.	References	15
6.1.	Normative References	15
6.2.	Informative References	15
Appendix A.	Additional Stuff	16
	Authors' Addresses	16

[1. Introduction](#)

This document specifies a BATS code [[BATS](#)] scheme for data delivery in multi-hop networks. The BATS code described here includes an outer code and an inner code. The outer code is a matrix generalization of fountain codes (see also the RaptorQ code described in [RFC 6330](#) [[RFC6330](#)]), which inherits the advantages of reliability and efficiency and possesses the extra desirable property of being network coding compatible. The inner code, also called recoding, is formed by linear network coding for combating packet loss, improving the multicast efficiency, etc. A detailed design and analysis of BATS codes are provided in the BATS monograph [[BATSMonograph](#)].

The BATS coding scheme can be applied in multi-hop networks formed by wireless communication links, which are inherently unreliable due to interference. Existing network protocols like TCP/IP use end-to-end retransmission and store-and-forward at the relays, so that packet loss would accumulate along the way.

The BATS coding scheme can be used for various data delivery applications like file transmission, video streaming over wireless multi-hop networks, etc. Different from traditional forward error correcting (FEC) schemes that are applied either hop-by-hop or end-to-end, the BATS coding scheme combines the end-to-end coding (the outer code) with certain hop-by-hop coding (the inner code), and hence can potentially achieve better performance.

The coding scheme described here can be used in a network with multiple communication flows. For each flow, the source node encodes the data for transmission separately. Inside the network, however, it is possible to mix the packets from different flows for recoding. In this document, we describe a simple case where recoding is performed within each flow. Note that the same encoding/decoding scheme described here can be used with different recoding schemes as long as they follow the principle as we illustrate in this document.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

2. Procedures

2.1. Introduction

A BATS coding scheme includes an outer code encoder (also called encoder), an inner code encoder (also called recoder) and a decoder. The BATS coding scheme can be used for a single data flow that includes a single source and one or multiple destinations. Thus there exists only one encoder with multiple recoders and decoders. The BATS coding scheme described in this document can be used by a Data Delivery Protocol (DDP) with the following procedures.

Outer Code Encoding at a source node which has the data for transmission:

- * The DDP provides the data to be delivered and the related information to the BATS encoder.

- * The BATS encoder generates a sequence of batches, each consisting of a set of coded packets and the information pertaining to the batch.

The batches generated at the source node are further recoded before transmitting:

- * A BATS recoder generates recoded packets of a batch.
- * The DDP forms and transmits the DDP packets using the batches and the corresponding batch information.

Recoding at an intermediate node that does not need the data:

- * The DDP extracts the batches and the corresponding batch information from its received DDP packets.
- * A BATS recoder generates recoded packets of a batch.
- * The DDP forms and transmits DDP packets using the recoded packets and the corresponding batch information.

Decoding at a destination node that needs the data:

- * The DDP extracts the batches and the corresponding batch information from its received DDP packets.
- * A BATS decoder tries to recover the transmitted data using the received batches.
- * The DDP sends the decoded data to the application that needs the data.

2.2. Data Delivery Procedures

Suppose that the DDP has F octets of data for transmission. We describe the procedures of one BATS session for transmitting the F octets. There is a limit on F of a single BATS session. If the total data has more than the limit, the data needs to be transmitted using multiple BATS sessions. The limit on F of a single BATS session depends on the MTU (maximum transmission unit) of the network, which MUST be known by the DDP. We have F is no more than $(MTU-10)2^{16}-1$ octets.

2.2.1. Source Node Data Partitioning and Padding

The DDP first determines the following parameters:

- o Batch size (M): the number of coded packets in a batch.
- o Recoding field size (q): the number of elements in the finite field for recoding. q is 2 or 2^8
- o BATS payload size (T0): the number of payload octets in a BATS packet, including the coded data and the coefficient vector.

Based on the above parameters, the parameters T, 0 and K are calculated as follows:

- o 0: the number of octets of a coefficient vector, calculated as $0 = \text{ceil}(M \cdot \log_2(q)/8)$.
- o T: the number of data octets of a BATS packet, calculated as $T = T_0 - 0$.
- o K: number of source packets, calculated as $K = \text{floor}(F/T) + 1$.

The data MUST be padded to have $T \cdot K$ octets, which will be partitioned into K source packets $b[0], \dots, b[K-1]$, each of T octets. In our padding scheme, $b[0], \dots, b[K-2]$ are filled with data bits, and $b[K-1]$ is filled with the remaining data octets and padding octets. Let $P = K \cdot T - F$ denote the number of padding octets. We use $b[K-1, 0], \dots, b[K-1, T-P-1]$ to denote the T-P source octets and $b[K-1, T-P], \dots, b[K-1, T-1]$ to denote the P padding octets in $b[K-1]$, respectively. The padding process is shown in Figure 1.

```

Z = T - P
Let bl be the last source packet b[K-1]
for i = 1, 2, ... do
    if Z + i >= T - 1 do
        bl[Z...T-1] = i
        break
    bl[Z...Z+i-1] = i
    Z = Z + i

```

Figure 1: Data Padding Process

2.2.2. Source Node Outer Code Encoding Procedure

The DDP provides the BATS encoder with the following information:

- o Batch size (M): the number of coded packets in a batch.

- o Recoding field size (q): the number of elements in the finite field for recoding.
- o `MAX_DEG`: the size of DD.
- o The degree distribution (DD), which is an unsigned integer array of size `MAX_DEG+1`.
- o A sequence of batch IDs (j , $j = 0, 1, \dots$).
- o Number of source packets (K).
- o Packet size (T): the number of octets in a source packet.
- o The source packets ($b[i]$, $i = 0, 1, \dots, K-1$).

Using this information, the (outer code) encoder generates a batch for each batch ID. For the batch ID j , the encoder returns the DDP that contains

- o a sparse degree $d[j]$, and
- o M coded packets ($x[j,i]$, $i = 0, 1, \dots, M-1$), each containing T_0 octets.

The DDP will use the batches to form DDP packets to be transmitted to other network nodes towards the destination nodes. The DDP MUST deliver with each coded packet its

- o d : sparse degree
- o BID: batch ID

The DDP MUST deliver the following information to each recoder:

- o M : batch size M
- o q : recoding field size

The DDP MUST deliver the following information to each decoder:

- o M : batch size
- o q : recoding field size
- o K : the number of source packet
- o T : the number of octets in a source packet

The BATS payload size T_0 MUST be known by all the nodes.

2.2.3. Recoding Procedures

Both the source node and the intermediate nodes perform recoding on the batches before transmission. At the source node, the recoder receives the batches from the outer code encoding procedure. At an intermediate node, the DDP receives the DDP packets from the other network nodes, and should be able to extract coded packets and the corresponding batch information from these packets.

The DDP provides the recoder with the following information:

- o the batch size M ,
- o the recoding field size q ,
- o a number of received coded packets of the same batch, each containing T_0 octets, and
- o link statistics, e.g., packet loss rates.

For a received batch, the recoder determines a positive integer M_r , the number of recoded packets to be transmitted for the batch. The recoder uses the information provided by the DDP to generate M_r recoded packets, each containing T_0 octets. The DDP uses the M_r recoded packets to form the DDP packets for transmitting.

2.2.4. Destination Node Procedures

A destination node needs the data transmitted by the source node. At the destination node, the DDP receives DDP packets from the other network nodes, and should be able to extract coded packets and the corresponding batch information from these packets.

The DDP provides the decoder with the following information:

- o M : batch size,
- o q : recoding field size,
- o K : the number of source packets
- o T : the number of octets of a source packet
- o A sequence of batches, each of which is formed by a number of coded packets belonging to the same batch, with their corresponding batch IDs and degrees.

The decoder uses this information to decode the K source packets. If successful, the decoder returns the recovered K source packets to the DDP, which will use the K source packets to form the F octets data. The recommended padding process is shown as follows:

```
// this procedure returns the number P of padding octets
// at the end of b[K-1]
Let bl be the last decoded source packet b[K-1]
PL = bl[T-1]
if PL == 1 do
    return P = 1
WI = T - 1
while bl[WI] == PL do
    WI = WI - 1
return P = (1 + bl[WI]) * bl[WI] + T - WI - 1
```

Figure 2: Data Depadding Process

2.3. Recommendation for the Parameters

The recommendation for the parameters M and q is shown as follows:

- o When q=2, M=16,32,64
- o When q=256, M=8,16,32,64

It is RECOMMENDED that K is at least 128. However, the encoder/decoder SHALL support an arbitrary positive integer value less than 2^{16} .

2.4. Example DDP Packet Format

A DDP can form a DDP packet with a header (5 octets), a footer (3 octets) and a payload (T0 octets). A DDP packet has totally 8+T0 octets.

2.4.1. Packet Header

The BATS packet header has 40 bits (5 octets) and includes fields Packet_Count, Mq, Batch_ID, and Degree.

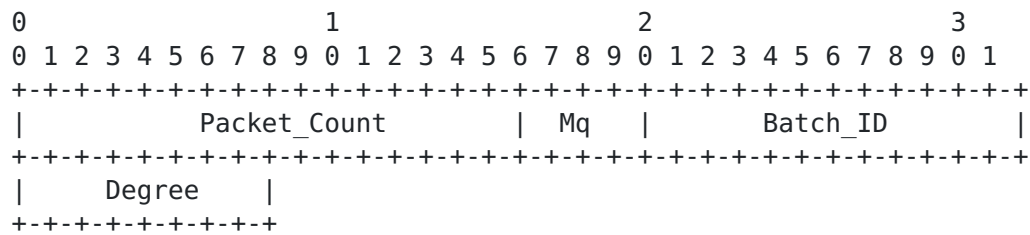


Figure 3: BATS packet header format.

- o Packet_Count: 16-bit unsigned integer, specifying the number K of packets of the BATS session.
- o Mq: 4-bit unsigned integer to specify the value of M and q as Table 1.
- o Batch_ID: 12-bit unsigned integer, specifying the batch ID BID of the batch the packet belonging to.
- o Degree: 8-bit unsigned integer, specifying the batch degree d of the batch the packet belonging to.

Mq	M	q	0
0010	16	2	2
0100	32	2	4
0110	64	2	8
0001	8	256	8
0011	16	256	16
0101	32	256	32
0111	64	256	64

Table 1: Values of Mq field

2.4.2. Packet Payload

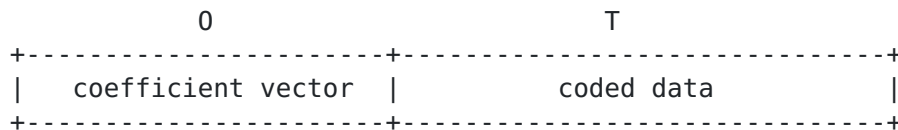


Figure 4: BATS packet payload format.

The payload has T0 octets, where the first 0 octets contain the coefficient vector and the remaining T octets contain the coded data.

Information in both fields MAY be encoded in JSON (ASCII) or protobuf (binary) formats.

- o coefficient vector: 0 octets. The range of the value of 0 is in Table 1.
- o coded data: T octets. T is at most MTU - 10, where 10 is the total of the header and footer length plus the minimum value of 0.

2.4.3. Packet Footer

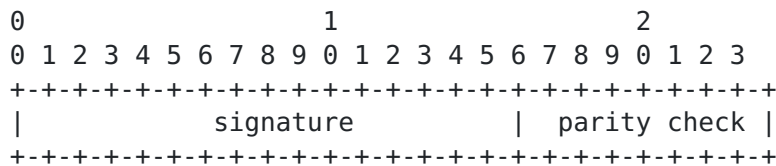


Figure 5: BATS packet footer format.

The footer has three octets.

- o signature: 2 octets. A signature of the individual packet to prevent pollution attack.
- o parity check: 1 octet. A parity check field used to verify the correctness of the packet.

3. BATS Code Specification

3.1. Common Parts

The T octets of a source packets are treated as a column vector of T elements in GF(256). Linear algebra and matrix operations over finite fields are assumed in this section.

Suppose that a pseudorandom number generator Rand() which generate an unsigned integer of 32 bits is shared by both encoding and decoding. The pseudorandom generator can be initialized by Rand_Init(S) with seed S. When S is not provided, the pseudorandom generator is initialized arbitrarily. One example of such a pseudorandom generator is defined in [RFC 8682](#) [RFC8682].

A function called BatchSampler is used in both encoding and decoding. The function takes two integers j and d as input, and generates an array idx of d integers and a d x M matrix G. The function first initializes the pseudorandom generator with j, sample d distinct integers from 0 to K-1 as idx, and sample d*M integers from 0 to 255 as G. See the pseudocode in Figure 6.


```

function BatchSampler(j,d)
    // initialize the pseudorandom generator by seed j.
    Rand_Init(j)
    // sample d distinct integers between 0 and K-1.
    for k = 0, ..., d-1 do
        r = Rand() % K
        while r already exists in idx do
            r = Rand() % K
        idx[k] = r

    // sample d x M matrix
    for r = 0, ..., d-1 do
        for c = 0,...,M-1 do
            G[r,c] = Rand() % 256

    return idx, G

```

Figure 6: Batch Sampler Function

3.2. Outer Code Encoder

Define a function called DegreeSampler that return an integer d using the degree distribution DD . We expect that the empirical distribution of the returning d converges to $DD(d)$ when $d < K$. One design of DegreeSampler is illustrated in Figure 7.

```

function DegreeSampler(j, DD)
    Let CDF be an array
    CDF[0] = 0
    for i = 1, ..., MAX_DEG do
        CDF[i] = CDF[i-1] + DD[i]
    Rand_Init()
    r = Rand() % CDF[MAX_DEG]
    for d = 1, ..., MAX_DEG do
        if r >= CDF[d] do
            return min(d,K)
    return min(MAX_DEG,K)

```

Figure 7: Degree Sampler Function

Let $b[0], b[1], \dots, b[K-1]$ be the K source packets. A batch with BID j is generated using the following steps.

- o Obtain a degree d by calling DegreeSampler with input j .
- o Obtain idx and $G[j]$ by calling BatchSampler with input j and d .

- o Let $B[j] = (b[idx[0]], b[idx[1]], \dots, b[idx[d-1]])$. Form the batch $X[j] = B[j]*G[j]$, whose dimension is $T \times M$.
- o Form the $T_0 \times M$ matrix $Xr[j]$, where the first 0 rows of $Xr[j]$ form the $M \times M$ identity matrix I with entries in $GF(q)$, and the last T rows of $Xr[j]$ is $X[j]$.

See the pseudocode of the batch generating process in Figure 8.

```
function GenBatch(j)
    d = DegreeSampler(j)
    (idx, G) = BatchSampler(j,d)
    B = (b[idx[0]], b[idx[i]], ..., b[idx[d-1]])
    X = B * G
    Xr = [I_M; X]
    return Xr
```

Figure 8: Batch Generation Function

3.3. Inner Code Encoder (Recoder)

The inner code comprises (random) linear network coding applied on the coded packets belonging to the same batch. At a particular network node, recoded packets are generated by (random) linear combinations of the received coded packets of a batch. The recoded packets have the same BID, sparse degree and coded packet length.

The number M_r of recoded packets for a batch is decided first by the recoder. M_r can be set as M . When the link statistics is known, the recoder can try to obtain the link packet loss rate e for the link to transmit the recoded batch, and set M_r to be $(1+e)M$.

Suppose that coded packets $xr[i]$, $i = 0, 1, \dots, r-1$, which have the same BID j , have been received at an intermediate node. Using the recommended packet format, it can be verified whether the corresponding packet headers of these coded packets are the same. Then a recoded packet can be generated by one of the following two approaches:

- o forwarding: when receiving $xr[i]$, directly use $xr[i]$ as a recoded packet.
- o linear combination recoding: (randomly) choose a sequence of coefficients $c[i]$, $i = 0, 1, \dots, r-1$ from $GF(q)$. Generate $c[0]xr[0]+c[1]xr[1]+\dots+c[r-1]xr[r-1]$ as a recoded packet.

A recoder can combine these two approaches to generate recoded packets. For example, the recoder will output $xr[i]$, $i = 0, 1, \dots$,

$r-1$ as r systematic recoded packets and generate $M-r$ recoded packets using linear combinations of randomly chosen coefficients.

3.4. Belief Propagation Decoder

The decoder receives a sequence of batches $Y_r[j]$, $j = 0, 1, \dots, n-1$, each of which is a $T \times d[j]$ matrix over $GF(256)$. The degree $d[j]$ of batch j is also known. Let $Y[j]$ be the submatrix of the last T rows of $Y_r[j]$. When $q = 256$, let $H[j]$ be the first M rows of $Y_r[j]$; when $q = 2$, let $H[j]$ be the matrix over $GF(256)$ formed by embedding each bit in the first $M/8$ rows of $Y_r[j]$ into $GF(256)$.

By calling BatchSampler with input j and $d[j]$, we obtain $idx[j]$ and $G[j]$. According to the encoding and recoding processes described in [Section 3.2](#) and [Section 3.3](#), we have the system of linear equations $Y[j] = B[j]G[j]H[j]$ for each received batch with ID j , where $B[j] = (b[idx[j],0], b[idx[j],1], \dots, b[idx[j],d-1])$ is unknown.

We describe a belief propagation (BP) decoder that can efficiently solve the source packets when a sufficient number of batches have been received. A batch j is said to be decodable if $\text{rank}(G[j]H[j]) = d[j]$ (i.e., the system of linear equations $Y[j] = B[j]G[j]H[j]$ with $B[j]$ as the variable matrix has a unique solution). The BP decoding algorithm has multiple iterations. Each iteration is formed by the following steps:

- o Decoding step: Find a batch j that is decodable. Solve the corresponding system of linear equations $Y[j] = B[j]G[j]H[j]$ and decode $B[j]$.
- o Substitution step: Substitute the decoded source packets into undecodable batches. Suppose that a decoded source packet $b[k]$ is used in generating an undecodable $Y[j]$. The substitution involves 1) removing the entry in $idx[j]$ corresponding to k , 2) removing the row in $G[j]$ corresponding to $b[k]$, and 3) reducing $d[j]$ by 1.

The BP decoder repeats the above steps until no batches are decodable during the decoding step.

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

Subsuming both Random Linear Network Codes (RLNC) and fountain codes, BATS codes naturally inherit both their desirable capability of

offering confidentiality protection as well as their vulnerability towards pollution attacks.

5.1. Provision of Confidentiality Protection

Since the transported messages are linearly combined with random coefficients at each recoding node, it is statistically impossible to recover the individual messages by capturing the coded messages at any one or small number of nodes. As long as the coding matrices of the transported messages cannot be fully recovered, any attempt of decoding is equivalent to randomly guessing the transported messages. Thus, with the use of BATS codes, information confidentiality throughout the data transport process is assured.

The only threat towards confidentiality exists in the form of eavesdropping on the initial encoding process, which takes place at the encoding nodes. In these nodes, the transported data are presented in plain text and can be read along their transfer paths. Hence, information isolation between the encoding process and all other user processes running on the node must be assured.

In addition, the authenticity and trustworthiness of the encoding, recoding and decoding program running on all the nodes must be attested by a trusted authority. Such a measure is also necessary in countering pollution attacks.

5.2. Countermeasures against Pollution Attacks

Like all network codes, BATS codes are vulnerable under pollution attacks. In these attacks, one or more compromised coding node(s) can pollute the coded messages or inject forged messages into the coding network. These attacks can prevent the receivers from recovering the transported data correctly. Although error detection mechanisms can be put in place to prevent the receivers from getting the wrong messages, detection and discard of the polluted messages still constitute a form of denial-of-service (DoS) attack.

The research community has long been investigating the use of various signature schemes (including homomorphic signatures) to identify the forged messages and stall the attacks (see Zhao07 [[Zhao07](#)], Yu08 [[Yu08](#)], Agrawal09 [[Agrawal09](#)]). Nevertheless, these countermeasures are regarded as being computationally too expensive to be employed in broadband communications. A practical approach to protect against pollution attacks consist of the following system-level countermeasures:

1. Attestation and Validation of all encoding, recoding and decoding nodes in the network. Remote attestation and repetitive

validation of a node based on valid public key certificates with proper authorization MUST be a pre-requisite for admitting that node to a network and permitting it to remain in that network.

2. Attestation of all encoding, recoding and decoding programs used in the coding nodes. All programs used to perform the BATS encoding, recoding and decoding processes MUST be remotely attested before they are permitted to run on any of the coding nodes. Reloading or alteration of programs MUST NOT be permitted during an encoding session. Programs MUST be attested or validated again when they are executed in new execution environments instantiated even in the same node.
3. Original Authentication of all coded messages using network or transport level secure protocols such as IP-sec or TLS/DTLS MUST be used to provide Peer or Message Origin Authentication to every coded message sent through the coding network.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8682] Saito, M., Matsumoto, M., Roca, V., Ed., and E. Baccelli, "TinyMT32 Pseudorandom Number Generator (PRNG)", [RFC 8682](#), DOI 10.17487/RFC8682, January 2020, <<https://www.rfc-editor.org/info/rfc8682>>.

6.2. Informative References

- [Agrawal09] Agrawal, S. and D. Boneh, "Homomorphic MACs: MAC-based integrity for network coding", International Conference on Applied Cryptography and Network Security , 2009.
- [BATS] Yang, S. and R. Yeung, "Batched Sparse Codes", IEEE Transactions on Information Theory 60(9), 5322-5346, 2014.
- [BATSMonograph] Yang, S. and R. Yeung, "BATS Codes: Theory and Practice", Morgan & Claypool Publishers , 2017.

- [RFC6330] Luby, M., Shokrollahi, A., Watson, M., Stockhammer, T., and L. Minder, "RaptorQ Forward Error Correction Scheme for Object Delivery", [RFC 6330](#), DOI 10.17487/RFC6330, August 2011, <<https://www.rfc-editor.org/info/rfc6330>>.
- [Yu08] Yu, Z., Wei, Y., Ramkumar, B., and Y. Guan, "An Efficient Signature-Based Scheme for Securing Network Coding Against Pollution Attacks", INFOCOM , 2008.
- [Zhao07] Zhao, F., Kalker, T., Medard, M., and K. Han, "Signatures for content distribution with network coding", ISIT , 2007.

[Appendix A](#). Additional Stuff

This becomes an Appendix.

Authors' Addresses

Shenghao Yang
CUHK(SZ)
Shenzhen, Guangdong
China

Phone: +86 755 8427 3827
Email: shyang@cuhk.edu.cn

Xuan Huang
CUHK
Hong Kong, Hong Kong SAR
China

Phone: +852 3943 8375
Email: 1155136647@link.cuhk.edu.hk

Raymond W. Yeung
CUHK
Hong Kong, Hong Kong SAR
China

Phone: +852 3943 8375
Email: whyeung@ie.cuhk.edu.hk

John K. Zao
NCTU
Hsinchu, Taiwan
China

Email: jkzao@ieee.org