

1 **LoRaWAN Fragmented Data Block Transport Specification v1.0.0**

2 Copyright © 2018 LoRa Alliance, Inc. All rights reserved.

3

4 **NOTICE OF USE AND DISCLOSURE**

5 Copyright © LoRa Alliance, Inc. (2018). All Rights Reserved.

6

7 The information within this document is the property of the LoRa Alliance (“The Alliance”) and its use and
8 disclosure are subject to LoRa Alliance Corporate Bylaws, Intellectual Property Rights (IPR) Policy and
9 Membership Agreements.

10

11 Elements of LoRa Alliance specifications may be subject to third party intellectual property rights, including
12 without limitation, patent, copyright or trademark rights (such a third party may or may not be a member of LoRa
13 Alliance). The Alliance is not responsible and shall not be held responsible in any manner for identifying or failing
14 to identify any or all such third party intellectual property rights.

15

16 This document and the information contained herein are provided on an “AS IS” basis and THE ALLIANCE
17 DISCLAIMS ALL WARRANTIES EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO (A) ANY
18 WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OF THIRD
19 PARTIES (INCLUDING WITHOUT LIMITATION ANY INTELLECTUAL PROPERTY RIGHTS INCLUDING
20 PATENT, COPYRIGHT OR TRADEMARK RIGHTS) OR (B) ANY IMPLIED WARRANTIES OF
21 MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE OR NON-INFRINGEMENT.

22

23 IN NO EVENT WILL THE ALLIANCE BE LIABLE FOR ANY LOSS OF PROFITS, LOSS OF BUSINESS, LOSS
24 OF USE OF DATA, INTERRUPTION OF BUSINESS, OR FOR ANY OTHER DIRECT, INDIRECT, SPECIAL OR
25 EXEMPLARY, INCIDENTAL, PUNITIVE OR CONSEQUENTIAL DAMAGES OF ANY KIND, IN CONTRACT OR
26 IN TORT, IN CONNECTION WITH THIS DOCUMENT OR THE INFORMATION CONTAINED HEREIN, EVEN IF
27 ADVISED OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

28

29

30 The above notice and this paragraph must be included on all copies of this document that are made.

31

32 LoRa Alliance™
33 5177 Brandin Court
34 Fremont, CA 94538
35 United States

36 *Note: All Company, brand and product names may be trademarks that are the sole property of their respective*
37 *owners.*



39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63

LoRaWAN Fragmented Data Block Transport Specification

Authored by the FUOTA Working Group of the LoRa Alliance Technical Committee

Technical Committee Chairs:

N.SORNIN (Semtech), A.YEGIN (Actility)

Working Group Chairs:

J.CATALANO (Kerlink), N.SORNIN (Semtech)

Editor:

J.CATALANO (Kerlink)

Contributors:

J.CATALANO (Kerlink), J-P.COUPIGNY (STMicroelectronics), N.SORNIN (Semtech),
J.STOKKING (The Things Network Foundation)

Version: v1.0.0

Date: September 10, 2018

Status: Final release

64 Contents

65	1	Conventions	5
66	2	Introduction	6
67	3	Downlink Fragmentation Transport Message Package	7
68	3.1	PackageVersionReq & Ans	8
69	3.2	FragSessionSetupReq & Ans	8
70	3.3	FragSessionDeleteReq & Ans	10
71	3.4	Downlink Data Fragment message	10
72	3.5	FragSessionStatusReq & Ans	11
73	4	File integrity check and authentication	13
74	5	Fragmentation algorithm	14
75		Appendix: Data block fragmentation forward error correction code proposal	15
76	6	Introduction	16
77	7	Fragment error coding	17
78	8	Fragment decoding and reassembling	20
79	9	Performance of the coding scheme	22
80	10	End-device memory requirement	24
81	11	Preliminary Matlab code	26
82	12	Glossary	28
83	13	Bibliography	29
84	13.1	References	29
85	14	NOTICE OF USE AND DISCLOSURE	30
86			

87 Tables

88	Table 1: Fragmentation Control messages summary	7
89	Table 2: PackageVersionAns	8
90	Table 3: FragSessionSetupReq	8
91	Table 4: FragSessionSetupReq FragSession field	8
92	Table 5: FragSessionSetupReq Control field	9
93	Table 6: FragSessionSetupAns	9
94	Table 7: FragSessionSetupAns StatusBitMask field	9
95	Table 8: FragSessionDeleteReq	10
96	Table 9: FragSessionDeleteReq Param bits	10
97	Table 10: FragSessionDeleteAns	10
98	Table 11: FragSessionDeleteAns Status bits	10
99	Table 12: Downlink Data Fragment payload	10
100	Table 13: Downlink Data Fragment Index&N field	10
101	Table 14: FragSessionStatusReq	11
102	Table 15: FragSessionStatusReq FragStatusReqParam field	11
103	Table 16: FragSessionStatusReq FragStatusReqParam field participant bit	11
104	Table 17: FragSessionStatusAns	12
105	Table 18: FragSessionStatusAns Received&index field	12
106	Table 19: FragSessionStatusAns Status field	12

107

108 Figures

109	Figure 1 : 26x52 parity check matrix.....	18
110	Figure 2 : 32x32 matrix A built during decoding process	21
111		

112 1 Conventions

113

114 The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD",
115 "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be
116 interpreted as described in RFC 2119.

117

118 The octet order over the air for all multi-octet fields is little endian (Least significant byte is
119 sent first).

120

121 **2 Introduction**

122

123 This document proposes an application layer messaging package running over LoRaWAN to
124 perform the following operations on a fleet of end-devices:

- 125 • Send a fragmented block of data to one or many end-devices

126

127 All messages described in this document are transported as application layer messages. As
128 such, all unicast messages (uplink or downlink) are encrypted by the LoRaWAN MAC layer
129 using the end-device's AppSKey. Downlink multicast messages are encrypted using a
130 multicast group McAppSKey common to all end-devices of the group. The setup of the group
131 is described in [RPD_Remote_Multicast_Setup].

132 The data block transported may be a firmware upgrade, but this document is not specific to
133 firmware upgrade. Any large (from 1kBytes to X Kbytes) data file may be sent to a (group of)
134 end-device using this protocol.

135

136 The “**fragmentation control**” package can be used to:

- 137 • Setup / report / delete fragmentation transport sessions
- 138 • Several fragmentation sessions MAY be supported simultaneously by an end-device
- 139 • Fragmentation can be used either over multicast or unicast
- 140 • Authenticate a data block once reconstructed (TBD)
- 141 • Report on the status of a fragmentation session

142

143 This package uses a dedicated port to separate its traffic from the rest of the applicative
144 traffic.

3 Downlink Fragmentation Transport Message Package

The identifier of the fragmentation transport package is 3. The version of this package is version 1.

This package supports all the commands necessary to transport reliably a large data block from a fragmentation server to an end-device (using unicast) or a group of end-device (if multicast is used over classB or classC). This package requires a dedicated port. The default port value is 201. Once declared, this port cannot be used for any other purposes.

All fragmentation related messages are exchanged on this port using application payload and encrypted using the end-device's AppSKey or the McAppSKey. All unicast or multicast control messages use the same format:

Command1	Command1 Payload	Command2	Command2 payload
----------	------------------	----------	------------------	------

A message MAY carry more than one command with the exception of the "Data fragment" command which MUST be the only command in a message's payload. The length of each command's payload is fixed and a function of the command. Commands are executed from first to last.

The following table summarizes the list of fragmentation control messages

CID	Command name	Transmitted by		Multicast (M) / Unicast (U)	Short Description
		End-device	server		
0x00	PackageVersionReq		x	U	Used by the AS to request the package version implemented by the end-device
0x00	PackageVersionAns	x		U	Conveys the answer to PackageVersionReq
0x01	FragStatusReq		x	U/M	Asks an end-device or a group of end-devices to send the status of a fragmentation session
0x01	FragStatusAns	x		U	Conveys answer to the FragSessionStatus request
0x02	FragSessionSetupReq		x	U	Defines a fragmentation session
0x02	FragSessionSetupAns	x		U	
0x03	FragSessionDeleteReq		x	U	Used to delete a fragmentation session
0x03	FragSessionDeleteAns	x		U	
0x08	DataFragment		x	U/M	Carries a fragment of a data block

Table 1: Fragmentation Control messages summary

The message marked "U/M" can be received using a unicast or multicast address. All other messages are exchanged only using the unicast end-device address.

171 3.1 PackageVersionReq & Ans

172

173 The *PackageVersionReq* command has no payload.

174 The end-device answers with a *PackageVersionAns* command with the following payload.

175

Field	PackageIdentifier	PackageVersion
Size (bytes)	1	1

176

Table 2: PackageVersionAns

177 *PackageIdentifier* uniquely identifies the package. For the “fragmentation transport package”
 178 this identifier is 3.

179 *PackageVersion* corresponds to the version of the package specification implemented by the
 180 end-device.

181 3.2 FragSessionSetupReq & Ans

182 This message is used to setup a DL fragmentation transport session.

183

Field	FragSession	NbFrag	FragSize	Control	Padding	Descriptor
Size (bytes)	1	2	1	1	1	4

184

Table 3: FragSessionSetupReq

185

186 *FragSession* identifies the fragmentation session and contains the following fields

187

FragSession Fields	RFU	FragIndex	McGroupBitMask
Size (bits)	2bits	2bits	4bits

188

Table 4: FragSessionSetupReq FragSession field

189 *FragIndex* [0 to 3] identifies one of the 4 simultaneously possible fragmentation sessions.

190

191 *McGroupBitMask* specifies which multicast group addresses are allowed as input to this
 192 defragmentation session. Bit number X indicates if multicast group with *McGroupID*=X is
 193 allowed to feed fragments to the defragmentation session. Unicast can always be used as a
 194 source for the defragmentation session and cannot be disabled. For example, 4'b0000
 195 means that only Unicast can be used with this fragmentation session. 4'b0001 means the
 196 defragmentation layer MAY receive packets from the multicast group with *McGroupID*=0 and
 197 the unicast address. 4'b1111 means that any of the 4 multicast groups or unicast may be
 198 used. If the end-device does not support multicast, this field SHALL be ignored.

199

200

201

202

203

204

205

206

207

208

Note: the *McGroupBitMask* is a mechanism allowing tying a defragmentation session to one or several specific multicast group addresses in a given end-device. For example, a street lighting controller end-device is part of 2 multicast groups, one used to control the lamps and one for firmware updates. Only data fragments coming from the second group shall be taken into account by the fragmented transport layer. The first group shall only transport ON/OFF lamp control packet and should not be allowed to transport firmware update data.

209 *NbFrag* (*Number of Fragments*) specifies the total number of fragments of the data block to
 210 be transported during the coming multicast fragmentation session. (example: 100 means
 211 that the data block that is going to be multicasted will be divided in 100 fragments)

212
 213 *FragSize* (fragment size) is the size in byte of each fragment. The data block size is
 214 therefore $NbFrag \times FragSize$

215
 216 *Control* consists of 2 fields
 217

Control Fields	RFU	FragAlgo	BlockAckDelay
Size (bits)	2bits	3bits	3bits

218 **Table 5: FragSessionSetupReq Control field**

219
 220 *FragAlgo* encodes the type of fragmentation algorithm used. This parameter is simply
 221 passed to the fragmentation algorithm. *FragAlgo* 0 corresponds to FEC fragmentation
 222 described in Appendix “data block fragmentation forward error correction code proposal”.

223
 224 *BlockAckDelay* encodes the amplitude of the random delay that end-devices have to wait
 225 between the reception of a downlink command sent using multicast and the transmission of
 226 their answer. This parameter is a function of the group size and the geographic spread and
 227 is used to avoid too many collisions on the uplink due to many end-devices simultaneously
 228 answering the same command. The actual delay SHALL be $rand().2^{BlockAckDelay+4}$ seconds
 229 where $rand()$ is a random number in the [0:1] interval.

230
 231 *Padding*: The binary data block size may not be a multiple of *FragSize*. Therefore, some
 232 padding bytes MUST be added to fill the last fragment. This field encodes the number of
 233 padding byte used. Once the data block has been reconstructed by the receiver, it SHALL
 234 remove the last “padding” bytes in order to get the original binary file.

235
 236 *Descriptor*: The descriptor field is a freely allocated 4 bytes field describing the file that is
 237 going to be transported through the fragmentation session. For example, this field MAY be
 238 used by the end-device to decide where to store the defragmented file, how to treat it once
 239 received, etc... If the file transported is a FUOTA binary image, this field might encode the
 240 version of the firmware transported to allow end-device side compatibility verifications. The
 241 encoding of this field is application specific.

242
 243 The end-device answers with a **FragSessionSetupAns** message with the following payload
 244

FragSessionSetupAns payload	StatusBitMask
Size (bytes)	1

245 **Table 6: FragSessionSetupAns**

246
 247

Bits	7:6	5:4	3	2	1	0
Status bits	FragIndex	RFU	Wrong Descriptor	FragSession index not supported	Not enough Memory	Encoding unsupported

248 **Table 7: FragSessionSetupAns StatusBitMask field**

249 If any of the bits [0:3] is set to 1 the **FragSessionSetup** command was not accepted.

250
 251 If a *FragSessionSetupReq* command with a *FragIndex* field corresponding to an already
 252 existing fragmentation session is received, the context of this session is cleared, and a new
 253 session is setup with the parameters of the new *FragSessionSetupReq* command.

254

255 3.3 FragSessionDeleteReq & Ans

256 This message is used to delete a fragmentation session. A fragmentation session MUST be
257 deleted before its index (FragIndex) can be reused for another one. The command payload
258 is:

259	Field	Param
260	Size (bytes)	1

261

262 **Table 8: FragSessionDeleteReq**

263 Where:

	Bits	7:2	1:0
Param bits	RFU	FragIndex	

264 **Table 9: FragSessionDeleteReq Param bits**

265 The end-device answers with FragSessionDeleteAns with payload:

266	Field	Status
267	Size (bytes)	1

268

269 **Table 10: FragSessionDeleteAns**

270

271 Where:

	Bits	7:3	2	1:0
Status bits	RFU	Session does not exist	FragIndex	

272 **Table 11: FragSessionDeleteAns Status bits**

273 If the bit 2 is set to 1 the **FragSessionDeleteReq** command was not accepted because the
274 fragmentation session corresponding to FragIndex did not exist in the end-device.

275 3.4 Downlink Data Fragment message

276 This message can be received by the end-device in a multicast or unicast downlink frame.
277 This message is used to carry a data block fragment.

278 The payload content is:

279	Size (bytes)	Index&N	0:MaxAppPI-3
	Payload	2	P_M^N

280 **Table 12: Downlink Data Fragment payload**

281	Index&N Fields	FragIndex	N
	Size (bits)	2bits	14bits

282 **Table 13: Downlink Data Fragment Index&N field**

283

284 If this message was received on a multicast address, the end-device MUST check that the
 285 multicast address used was enabled at the creation of the fragmentation session through the
 286 *McGroupBitMask* field of the **FragSessionSetup** command. If not, the frame SHALL be
 287 silently dropped.

288 Where P_M^N is the fragment N over M of the session.

289 More than M fragments MAY actually be transmitted to add redundancy and packet loss
 290 robustness. N is the index of the coded fragment transported.

291 M is equal to the *NbFrag* parameter.

292 Once the data block has been reconstructed the end-device SHALL drop any further
 293 message using that *fragIndex* until the fragmentation session is first deleted and a new
 294 fragmentation session is setup through the **FragSessionSetup** application command.

295 3.5 FragSessionStatusReq & Ans

296 This message can be transmitted by the server in a UNICAST or MULTICAST downlink
 297 frame.

298

299

Size (bytes)	1
FragParam Payload	FragStatusReqParam

300

Table 14: FragSessionStatusReq

301

302 Where:

303

bits	7:3	2:1	0
FragStatusReqParam field	RFU	FragIndex	Participants

304

Table 15: FragSessionStatusReq FragStatusReqParam field

305

306 Used by the fragmentation server to request receiver end-devices to report their current
 307 defragmentation status.

308

309 The receivers (in the case of multicast) SHOULD NOT answer this request all at the same
 310 time because this would potentially generate a lot of collisions. The receivers MUST
 311 therefore spread randomly their responses as specified by the *BlockAckDelay* field of the
 312 *FragSessionSetupReq* command.

313

314 The “participants” bit signals if all the fragmentation receivers should answer or only the
 315 ones still missing fragments.

316

Participant bit value	0	1
	Only the receivers still missing fragments MUST answer the request	All receivers MUST answer, even those who already successfully reconstructed the data block

317

Table 16: FragSessionStatusReq FragStatusReqParam field participant bit

318
319 The end-devices respond with a **FragSessionStatusAns** message. The message payload
320 is:
321

Size (bytes)	2	1	1
FragParam Payload	Received&index	MissingFrag	Status

322 **Table 17: FragSessionStatusAns**

323 Where:
324

bits	15:14	13:0
Received&index field	FragIndex	NbFragReceived

325 **Table 18: FragSessionStatusAns Received&index field**

bits	7:1	0
Status field	RFU	Not enough matrix memory. The defragmentation process was aborted because the number of missing fragments was greater than the available memory matrix storage capacity

326 **Table 19: FragSessionStatusAns Status field**

327
328 Used by the fragmentation receiver to report its defragmentation status for the fragmentation
329 session *FragIndex*.

330
331 *NbFragReceived* is the total number of fragments received for this fragmentation session
332 since the session was created.

333
334 *MissingFrag* is the number of independent coded fragments still required before being able
335 to reconstruct the data block. In the case where the block was already successfully
336 reassembled this field SHOULD be 0. If more than 255 fragments are missing, then
337 MissingFrag SHALL be set to 255.

338
339
340 As described in the "**FragSessionStatusReq**" command, the receivers MUST respond with
341 a pseudo-random delay as specified by the BlockAckDelay field of the
342 FragSessionSetupReq command.

343 **4 File integrity check and authentication**

344

345 The payloads transported by the fragmentation/defragmentation package are encrypted and
346 authenticated using the McAppSKey and McNwkSKey. However, those keys are identical in
347 all the end-devices of the multicast group. Because one of the group's end-devices might be
348 compromised (the end-device might have been physically destroyed and the keys
349 extracted), those keys cannot be considered safe except if a tamper-proof secure element is
350 used to store them in ALL the end-devices of the group.

351 If that is not the case (no secure element is used), then an additional file integrity and
352 authentication step SHOULD take place. The integrity/authentication check corresponds to
353 making sure that the block reconstructed is exactly what the fragmentation server wanted to
354 send to the end-device and that this block has not been modified in any way through the
355 transport process. This goal may be achieved by different means:

- 356 1. Public/private cryptography certificate
- 357 2. Unicast exchange protected by Symmetric key

358 Solution 1 does not require any additional exchange and MAY use a standard HASH +
359 certificate mechanism based on RSA or ECC cryptography. This solution is
360 RECOMMENDED when the fragmentation layer is used to transport a firmware upgrade file.
361 That solution increases the size of the file transported (cryptography/certificate overhead).
362 For ECC that overhead is typically around 100bytes.

363 Solution 2 relies on a unicast exchange between the end-device and the fragmentation
364 server. The messages exchanged contain a HASH of the reconstructed file and MUST be
365 protected by an end-device specific key. In that way even if the multicast keys are
366 considered unsafe, the final authentication is made on an end-device per end-device basis
367 and cannot be compromised. That solution has no file size overhead but requires an
368 additional unicast exchange between each end-device and the AS.

369 The choice between the two solutions is application specific and this is currently considered
370 out-of-scope of this specification. Following versions of this specification will provide a
371 recommended file integrity/authentication verification process.

372 5 Fragmentation algorithm

373 This section will contain a description of the fragmentation / coding / decoding /
374 defragmentation algorithm proposed.

375 The coding adds some redundancy to the fragment transmitted such that an end-device
376 missing some of the multicasted fragments can still reconstruct the complete data block.

377 The maximum ratio of lost fragment that can be tolerated is a parameter selected by the
378 “fragmentation server” preparing the multicast.

379

380

381 **APPENDIX: DATA BLOCK FRAGMENTATION**
382 **FORWARD ERROR CORRECTION CODE**
383 **PROPOSAL**

384
385

386 6 Introduction

387

388 This appendix proposes a simple Forward Error Correction (FEC) code to be used for
389 fragmented transport of large binary files over LoRaWAN. As all radio link, a LoRaWAN link
390 exhibit a certain ratio of lost frames. Adding FEC in the file fragmentation process allows an
391 end-device to autonomously recover the full file even in the presence of lost frames without
392 having to systematically request the missing fragments.

393 The transmitter of the fragmented binary file can select to add an arbitrary redundancy to the
394 transmission content through this FEC.

395 For example a 10% redundancy added by the fragmentation transmitter allows the receiver
396 performing the defragmentation to loose roughly 10% of the incoming frames and still be
397 able to reconstruct the binary file.

398

399

400 Fragmentation may be used for many different applications, for example:

401 • Broadcasting a firmware upgrade to a group of end-devices (Network -> end-devices
402 downlink multicast)

403 • Fragmenting a huge data block in several smaller messages before sending it up to
404 the network (end-device -> Network uplink) and make sure all the data has been
405 successfully received. Example: A sensor collects frequent data for a long time, and
406 then compresses it into one huge block that is sent using fragmentation, as soon as
407 the server is able to reconstruct the full block the end-device receives a notification
408 and stops transmitting.

409

410

411 The coding scheme proposed here is directly derived from the 1963 thesis of Robert
412 Gallager describing Parity-Check code: This thesis can be accessed at
413 <http://www.inference.phy.cam.ac.uk/mackay/gallager/papers/ldpc.pdf>
414

415 7 Fragment error coding

416
 417 The initial data block that needs to be transported must first be fragmented into M data
 418 fragments of arbitrary but equal length. The length of those fragments has to be chosen to
 419 be compatible with the maximum applicative payload size available.

420 The actual applicative payload length will be:

421
$$\text{fragLen} + 2 \text{ bytes}$$

422
 423 Where fragLen is the byte length of each fragment plus 2 bytes of fragmentation header
 424 (containing Index & N, see 3.4)

425
 426 Those original data fragments are named uncoded fragments and are noted B_n

427
 428 The full data block to be transported consists therefore of the concatenation of the B_n
 429 uncoded fragments [B₁ : B₂ : ... : B_m]

430
 431 The coded fragments are noted P_M^N and are derived from the uncoded fragments.

432 P_M^N is the Nth coded fragment from a fragmentation session containing M (B₁ to B_m)
 433 uncoded fragments. The coded fragments P_M^N all have exactly the same byte length than the
 434 uncoded (B_n) fragments.

435 To allow the original uncoded fragments reconstruction on the receiving end of the link even
 436 in presence of arbitrary packet loss, the transmitter (performing the fragmentation) adds
 437 redundancy. Therefore N might be greater than M, meaning that the sender may send more
 438 coded fragments than the total number of uncoded fragments to enable the reconstruction
 439 on the receiving end in presence of packet loss. The ratio between M and the number of
 440 actually sent coded fragments is called coding ratio (or redundancy factor) and noted CR

441
 442 The coded fragments P_M^N are constructed by performing a bit per bit Xor operation between
 443 different subset of the uncoded fragments. The Xor operator is noted +.

444
 445
 446 Each coded fragment is defined as :

447
$$P_M^N = C_M^N(1).B_1 + C_M^N(2).B_2 + .. + C_M^N(M).B_M$$

448
 449 Where C_M^N(i) is a function of M,N and i and whose value is either 0 or 1.

- 450 • 0. B₁ is a binary word of the same length than the fragment B₁ with all bits = 0.
- 451 • 1. B₁ is equal to B₁

452
 453 The binary vector C(N, M) = [C_M^N(1), C_M^N(2), ..., C_M^N(M)] of M bits is a function of M and N and
 454 is given by a function matrix_line(M,N) which will be described later.

455
 456 It is sufficient to know that this function generates either:

- 457 • If N<=M , a vector of length M with a single one at position N, all other bits = 0
- 458 • If N>M : a parity check vector containing statistically as many zeros as ones in a
 459 pseudo-random order.

460 The parity check matrix, as defined by Gallager in his 1963 thesis, is an MxN matrix
 461 containing the C_Nⁱ on column i and line N.

462
 463 The following picture illustrates an example of such a matrix generated by the proposed
 464 function for M=26 and with a coding ratio CR=1/2 , this matrix has 26/CR = 52 lines. Id, it
 465 allows the creation of 52 coded fragments from the 26 original uncoded fragments.

466




Figure 1 : 26x52 parity check matrix

467
468

469
470
471

We can see that the parity matrix consists first of an $M \times M$ identity matrix followed by a parity control $M \times M$ matrix.

The coded fragment P_M^N is therefore the bit-wise Xor of the uncoded fragment B_i such that $C_M^N(i)$ is non-zero. The coded fragments have exactly the same bit length than the uncoded fragments.

477
478

Step by Step encoding example:

The transmitter must send 2000 bytes allowing up to 50% packet error rate on the radio link. A coding ratio of $\frac{1}{2}$ is selected. For this purpose the 2000 bytes data block will be segmented in 100 fragments of 20 bytes each, each frame will transport 1 fragment. The transmitter will send $100/CR = 200$ coded fragments. We will see that the receiver will be able to decode as soon as it receives ~ 103 frames out of the 200 (depending on the exact combination of frames lost).

486

487 First split the 2000 bytes into 100 uncoded fragments of 20 bytes each, B1 to B100.
488
489 To generate the first coded fragment.
490 First generate the first line of the parity check matrix by calling $C = \text{matrix_line}(1, 100)$
491 Then perform a bitwise Xor operation between all the uncoded fragments corresponding to a
492 1 in the C parity check vector.
493
494 In this case $C = 1'b1$ followed by 99 zeros, the first coded fragment $P_{100}^1 = B1$
495
496 When the transmitter reaches the frame number 101 , we have for example:
497 $C = \text{matrix_line}(101, 100) = 100'b0110010\dots\dots$;
498 Therefore $P_{100}^{101} = B2 + B3 + B6 + \dots\dots$;
499 Where + is the bit-wise Xor operator
500

501 8 Fragment decoding and reassembling

502

503 The receiver of a fragmentation session must perform the following operations.

504

505 For each frame received extract the coded fragment and its index.

506

507 The receiver also needs to create a null binary $A = M \times M$ bit matrix structure in his memory.

508

509 Then process those fragments one by one.

510

511

512 1. For each new fragment P_M^N , first fetch the corresponding line of the parity check
 513 matrix : $C = \text{matrix_line}(N, M)$

514 2. Proceed from left to right along the C vector (*i varying from 1 to M*) : For each entry
 515 C_i equal to 1, check if the line i of the matrix A contains a 1 in row i . If yes, perform a
 516 XOR between line i of matrix A “A(i)” and the vector C and store the result in C. Also
 517 perform a xor between P_M^N and the coded fragment stored at position i in the
 518 fragment memory store S_i and update P_M^N with the result.

519 3. Once this process is finished there are two options:

520 a. Either C now contains only zeros, in that case just get rid of the coded
 521 fragment P_M^N ; it isn't bringing any new information

522 b. The vector C is non-null : write it in the matrix A at the line i corresponding to
 523 the first non-zero element of C. Also add the modified P_M^N fragment to the
 524 memory store at position i : S_i

525 4. Loop to 1 until all lines of the matrix A have been updated. The matrix A will have
 526 only 1's on its diagonal and will be a triangular matrix with only 0's on the lower left
 527 half. The fragment memory store will contain exactly M fragments.

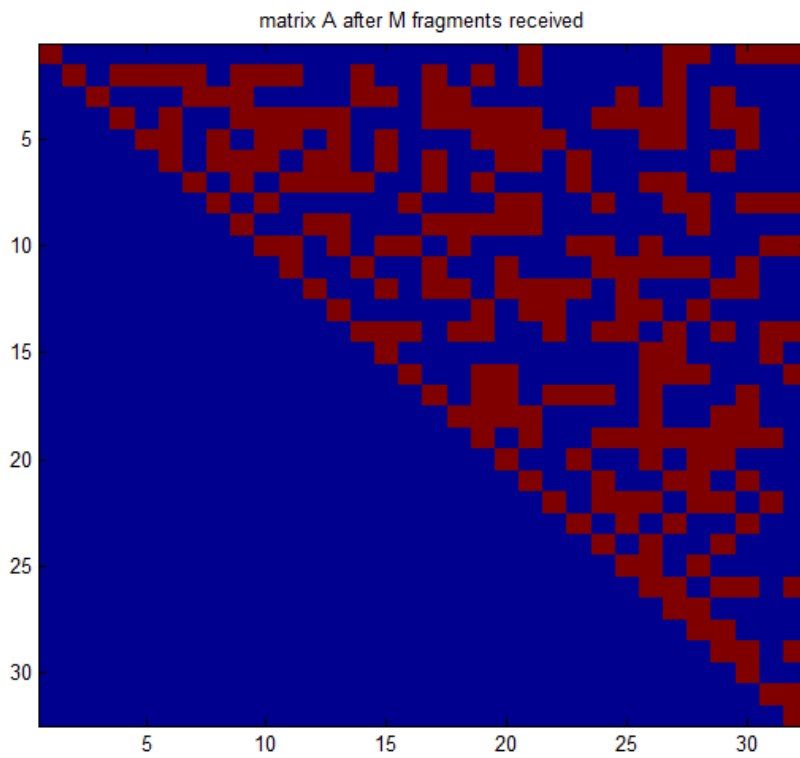
528 5. Starting from matrix line $i = M-1$ down to 1, fetch the i^{th} line of matrix A : A(i). The
 529 line A(i) has a 1 at position i and only zeros on the left. For any 1 at position $j > i$
 530 perform a xor between S_i and S_j and update S_i with the result.

531 6. The fragment memory store now contains the original uncoded fragments $S_i = B_i$

532 7. Reassemble the data block by concatenating all the uncoded fragments. If the
 533 fragment memory store is actually allocated as a continuous memory range, then this
 534 step is not even necessary, because the original data block consists of $S_1 : S_2 : \dots : S_M$
 535 where : represents the concatenation operator.

536

537



538
539

Figure 2 : 32x32 matrix A built during decoding process


540
541

542 **9 Performance of the coding scheme.**

543
544 As described in Gallager’s thesis Parity Check codes have a non-zero statistical overhead
545 independent of the coded word length. In our case the word length used is M . The actual
546 overhead depends on the way the parity check matrix is built. To be able to reconstruct the
547 uncoded fragments the receiver must receive at least M linearly independent coded
548 fragments. Said in another way, the parity check matrix reconstructed by the receiver based
549 on the fragments received must be of rank M .

550 This condition is fulfilled ideally as soon as M coded fragments have been received. But
551 sometimes, those M received first fragments are not all independent and the matrix resulting
552 rank is $<M$. in that case, more coded fragments need to be received until the rank of the
553 parity check matrix becomes M .

554
555 The following graph shows the probability of the matrix being of rank $<M$ with a number of
556 received fragment varying from M to $M+10$. The 5 curves corresponds to $M=32/40/48/56/64$
557 which are the number of uncoded fragments used in this proposal.
558
559



560 It can be seen that when the number of coded fragment received equals M , the matrix
561 cannot be inverted (is not of rank M) 70% of the cases. But this probability falls very rapidly
562 with a few additional received fragments. With $M+7$ fragments the matrix can be inverted in
563 99% of the occurrences. In takes in average $M+2$ coded fragments to recover the original
564 data block. The proposed fragmentation therefore works better (with a lower statistical
565 overhead) with a larger number of fragments. This coding scheme with a fixed overhead is
566 therefore close to ideal performance when the number of fragment is large (>100), because
567 adding 2 fragments on a 100 fragments session only represent 2% relative statistical
568 overhead. The overhead increases when the number of fragments is lower. This coding
569

570 scheme should not be used for less than 20 fragments (the average overhead is 10% in that
571 case).
572

573 **10 End-device memory requirement**

 574
 575 This coding scheme has been optimized to allow the transportation of large binary file with
 576 minimum memory overhead on the receiving end (performing the defragmentation).
 577 The following formula gives the decoding memory requirements for an optimized end-device
 578 implementation expressed in **Bytes** on **TOP** of the memory required to store the final
 579 reconstructed data block.

 580 **Example:** a 50kbytes data block must be received. The end-device
 581 requires 50kbytes of available memory + the value given by the
 582 following formula
 583

 584 For a data block fragmented using M fragments, let L be the number of coded fragments
 585 **lost** by the end-device amongst the first M coded fragments. (The fragments P_M^1 to P_M^M).
 586

 587 The required defragmentation memory overhead is a function of the maximum L value that
 588 the end-device is designed to tolerate. Interestingly, the memory overhead is **NOT** a function
 589 of the total number of fragments used to convey the data block. This is because the
 590 defragmentation implementation is optimized to use the fact that the first M fragments
 591 actually contain the original uncoded data block ($P_M^i = B_M^i$ for $i \leq M$).
 592

 593 Parity Matrix memory (bytes) = $l.(l+1)/2/8 + 2.l$, where l is the maximum value of L that the
 594 end-device tolerates (if $L > l$, the defragmentation fails and is aborted)

 595 The following table gives a few numerical examples.
 596

l	Parity matrix memory
32	130
40	183
48	243
56	312
64	388

 597
 598
 599 The total memory requirement is the sum of the full data block size and the parity matrix
 600 memory size.

 601 **Example:** a 50kbytes data block must be received, as 1000x50bytes
 602 fragments. The end-device is designed to be able to withstand the loss
 603 of 64 fragments out of the first 1000 transmitted ($l=64$). The end-device
 604 requires 50kbytes of available memory + 388bytes to store the parity
 605 matrix
 606
 607

 608 The data block can be reassembled in the memory directly at its final address (this means
 609 directly in the FLASH memory space for most end-devices), so there is no need for an
 610 additional “data block” swap memory space.

 611 The parity matrix memory space can be erased and reused once the defragmentation is
 612 finished.
 613

614 For the encoding process, there is no need to store the parity matrix so the required memory
615 simply corresponds to the data block to be segmented and transmitted plus a very little fixed
616 overhead for temporary calculations.
617
618

619 11 Preliminary Matlab code

620
621

622 This matlab code generates a list of coded fragments from an arbitrary binary file

623 Notations:

624 `w = 32;` %the number of fragments into which the binary file is split
 625 `fragment_size=10;` %the size of each fragment in bytes
 626 `DATA;` %a vector of bytes = the binary file to be sent. The length must be
 627 `w*fragment_size`

628

629

630

631 Encoding process in matlab

632

633

```
634 w = 32; %number of uncoded fragments
635 fragment_size=10; %nb of bytes per fragment
636 fprintf('\n number of uncoded fragments:%d, fragment size (bytes):%d bytes\n total
637 broadcast size %d bytes\n',w,fragment_size,w*fragment_size);
```

638

```
639 DATA = mod([0:w*fragment_size-1],256); %arbitrary binary file to be sent
```

640

```
641 % start of the fragmentation and encoding process
```

```
642 % UNCODED_F is an array of uncoded fragments
```

```
643 UNCODED_F = zeros(w,fragment_size);
```

```
644 for k=1:w
```

```
645     UNCODED_F(k,:) = DATA( (k-1)*fragment_size+1:k*fragment_size);
```

```
646 end
```

647

648

```
649 % now encode.
```

```
650 % we will create 2w CODED fragments, this number is arbitrary. Those can be
651 generated by the transmitter on the fly one by one.
```

```
652 CODED_F = []; % this will contain the array of 2w coded fragments
```

```
653 for y=1:w
```

```
654     CODED_F(y,:) = UNCODED_F(y,:); %the first w coded fragments are equal to the
655     uncoded fragments (binary data without coding)
```

```
656 end
```

```
657 %then we add w parity check fragments
```

```
658 for y=1:w
```

```
659     s=zeros(1,fragment_size);
```

```
660     A = matrix_line(y,w); %line y of w.w matrix
```

```
661     for x=1:w
```

```
662         if (A(x) == 1) %if bit x is set to 1 then xor the corresponding fragment
```

```
663             s = bitxor(s,UNCODED_F(x,:));
```

```
664         end
```

```
665     end
```

```
666     CODED_F = [CODED_F ; s]; % add the resulting coded fragment to the list
```

```
667 end
```

668

669

670 Matlab code of the matrix_line function generating a parity check vector:

```
671 %this funciton returns line N of the MxM parity matrix
```

```
672 function matrix_line = matrix_line(N,M)
```

```
673 matrix_line = zeros(1,M);
```

```
674 s=0;
```

675

```
676 % we must treat powers of 2 differently to make sure mtrix content is close
```

```
677 % to random . Powers of 2 tend to generate patterns
```

```
678 if (M == 2^floor(log2(M))) % if M is a power of 2
```

```
679     m=1;
680 else
681     m=0;
682 end
683
684
685 x= 1+1001*N; %initialize the seed differently for each line
686 nb_coeff=0;
687 while (nb_coeff<floor(M/2)) % will generate a line with M/2 bits set to 1 (50%)
688     r=2^16;
689     while (r>=M) %this can happen if m=1, in that case just try again with a
690 different random number
691         x=prbs23(x);
692         r=mod(x, M+m); %bit number r of the current line will be switched to 1
693     end
694
695     matrix_line(r+1) = 1; %set to 1 the column which was randomly selected
696     nb_coeff = nb_coeff + 1;
697 end
698
699
700
701 The prbs23() function implements a PRBS generator with 2^23 period.
702 %standard implementation of a 23bit prbs generator
703 function r=prbs23(start)
704 x= start;
705 b0 = bitand(x,1);
706 b1 = bitand(x,32)/32;
707 x = floor(x/2) + bitxor(b0,b1)*2^22;
708 r=x;
709
```

710 **12 Glossary**

711		
712	AES	Advanced Encryption Standard
713	AS	Application Server
714	FEC	Forward Error Correction
715		
716	TBD	To Be Done
717		

718 13 Bibliography**719 13.1 References**

- 720 [LoRaWAN 1.0.2]: LoRaWAN™ 1.0.2 Specification, LoRa Alliance, July 2016
- 721 [LoRaWAN 1.1]: LoRaWAN™ 1.1 Specification, LoRa Alliance, October 11, 2017
- 722 [LoRaWAN_Remote_Multicast_Setup]: LoRaWAN Remote Multicast Setup Specification
723 v1.0.0, LoRa Alliance, September 10, 2018
- 724

725 14 NOTICE OF USE AND DISCLOSURE

726 Copyright © LoRa Alliance, Inc. (2018). All Rights Reserved.

727 The information within this document is the property of the LoRa Alliance (“The Alliance”)
728 and its use and disclosure are subject to LoRa Alliance Corporate Bylaws, Intellectual
729 Property Rights (IPR) Policy and Membership Agreements.

730 Elements of LoRa Alliance specifications may be subject to third party intellectual property
731 rights, including without limitation, patent, copyright or trademark rights (such a third party
732 may or may not be a member of LoRa Alliance). The Alliance is not responsible and shall
733 not be held responsible in any manner for identifying or failing to identify any or all such third
734 party intellectual property rights.

735 This document and the information contained herein are provided on an “AS IS” basis and
736 THE ALLIANCE DISCLAIMS ALL WARRANTIES EXPRESS OR IMPLIED, INCLUDING
737 BUT NOT LIMITED TO (A) ANY WARRANTY THAT THE USE OF THE INFORMATION
738 HEREIN WILL NOT INFRINGE ANY RIGHTS OF THIRD PARTIES (INCLUDING WITHOUT
739 LIMITATION ANY INTELLECTUAL PROPERTY RIGHTS INCLUDING PATENT,
740 COPYRIGHT OR TRADEMARK RIGHTS) OR (B) ANY IMPLIED WARRANTIES OF
741 MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE OR
742 NONINFRINGEMENT.

743 IN NO EVENT WILL THE ALLIANCE BE LIABLE FOR ANY LOSS OF PROFITS, LOSS OF
744 BUSINESS, LOSS OF USE OF DATA, INTERRUPTION OF BUSINESS, OR FOR ANY
745 OTHER DIRECT, INDIRECT, SPECIAL OR EXEMPLARY, INCIDENTAL, PUNITIVE OR
746 CONSEQUENTIAL DAMAGES OF ANY KIND, IN CONTRACT OR IN TORT, IN
747 CONNECTION WITH THIS DOCUMENT OR THE INFORMATION CONTAINED HEREIN,
748 EVEN IF ADVISED OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

749 The above notice and this paragraph must be included on all copies of this document that
750 are made.

751 LoRa Alliance™
752 5177 Brandin Court
753 Fremont, CA 94538
754 United States

755 Note: All Company, brand and product names may be trademarks that are the sole property
756 of their respective owners.