



Original software publication

Fonsure 1.0: An erasure code library with efficient repair and update features

Şuayb Ş. Arslan

MEF University, Maslak, Istanbul, Turkey



ARTICLE INFO

Article history:

Received 8 September 2020

Received in revised form 11 January 2021

Accepted 13 January 2021

Keywords:

Distributed storage

Erasure coding

Fountain coding

Single instruction multiple data (SIMD)

OpenMP

Reliability

ABSTRACT

Fonsure is an open-source software library that implements a multi-dimensional graph-based erasure coding entirely based on fast exclusive OR (XOR) logic. Its implementation utilizes compiler optimizations and multi-threading to generate the right assembly code for the given multi-core CPU architecture with vector processing capabilities. Fonsure possesses important features that shall find various applications in modern data storage, communication, and networked computer systems, in which the data needs protection against device, hardware, and node failures. As data size reached unprecedented levels, these systems have become hungry for network bandwidth, computational resources, and average consumed power. To address that, the proposed library provides a three-dimensional design space that trades off the computational complexity, coding overhead, and data/node repair bandwidth to meet different requirements of modern distributed data storage and processing systems. Fonsure library enables efficient encoding, decoding, repairs/rebuilds, and updates while all the required data storage and computations are distributed across the network nodes.

© 2021 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Code metadata

Current code version	1.0
Permanent link to code/repository used for this code version	https://github.com/ElsevierSoftwareX/SOFTX-D-20-00043
Legal Code License	LGPLv3
Code versioning system used	git
Software code languages, tools, and services used	C, Python, OpenMP, etc.
Compilation requirements, operating environments & dependencies	Linux Runtime Environment and a C compiler ($\geq 4.8.4$). Python 2.3 or above.
If available Link to developer documentation/manual	https://github.com/suaybarslan/fonsure/blob/master/tests/Fonsure_1_0_User_Manual.pdf
Support email for questions	arslans@mef.edu.tr

1. Motivation and significance

Erasure coding is a fault tolerance mechanism that provides data protection and high availability in distributed data storage and processing systems [1]. Reed–Solomon (RS) codes are the conventional option for constructing erasure codes based on overhead-optimal design. In other words, they use the storage space as efficiently as information-theoretically possible [2]. As the modern data storage systems evolved to possess different requirements, the set of constraints on the design of erasure codes has dramatically changed. For instance, previous research work on erasure coding such as RS focused on optimizing the

coding overhead i.e., minimization of storage space for a given target data reliability [2,3]. Moreover, some of the most popular designs considered pure eXclusive Or (XOR) operations to provide durability, and efficient computation [4]. More recently, locally repairable codes have attracted attention due to their efficient utilization of network resources and eventually achieve better overall reliability [5] at the expense of suboptimal coding overhead. Besides the proprietary implementations of advanced erasure coding algorithms, many open-source implementations with different mathematical constructions become available online [6, 7]. Previous works provided overhead optimal and fast/efficient erasure coding library functions. However, they did not take into

E-mail address: arslans@mef.edu.tr.

<https://doi.org/10.1016/j.softx.2021.100662>

2352-7110/© 2021 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

account the peculiarities of the harnessed network, scarce computational resources, and data/node regeneration in a distributed setting.

In this study, we propose Founsure library which utilizes binary operations on a three-dimensional bipartite graph to construct a multi-functional erasure code. The design space includes computational complexity, coding overhead, and repair bandwidth as different dimensions of optimization. If one prefers to have an overhead optimal design with a reasonable complexity performance, then it might be advisable to use the well known Jerasure 2.0 [7] erasure coding library, which is now fully supported by the *RedHat Ceph* community [8]. Unfortunately, this library and the likes (e.g. *zfec* [9]) do not provide a sufficient code structure to address modern problems of distributed storage systems such as degraded reads, data repair/regeneration bandwidth, data security, etc. On the other hand, the main objective and purpose of developing the Founsure library have been to provide different operating points in the three-dimensional design space based on the requirements of the storage applications through a set of parameter configurations. Therefore, Founsure can be shown to demonstrate the huge potential for distinct storage applications using smart/guided configuration steps. For instance, Founsure is shown to be configured to a baseline deduplication engine within an archival scenario [10].

The encoding process of Founsure begins by defining a two-dimensional conventional bipartite graph which leads to a non-systematic low-density generator matrix (LDGM) code. A version of this class of codes, with appropriate input distributions, are generally known as fountain codes [11,12]. The degree distribution of Founsure's LDGM code is specially selected to meet a good trade-off operating point between computational complexity, coding overhead, and repair bandwidth. The current version (1.0) supports Robust Soliton Distribution (RSD) [12] (if one prefers good coding overhead design), as well as all possible finite max-degree distributions (if one prefers different operating points) including the one in [13] by default. However, we note that the referred distributions are optimized for the minimum coding overhead criterion only.

One of the building pillars of Founsure is genuine symbol check relationships. The terminology of check symbols is quite common in Low Density Parity Check (LDPC) coding community. What check symbols do is that they provide a mathematical relationship between a subset of data symbols to check a certain condition, usually a simple binary sum or equivalent Galois field operation. The check idea is also quite beneficial for local data repairs/rebuilds [14]. So on top of the two-dimensional bipartite graph of Founsure, the encoding engine generates check nodes for "data-only" (referred hereafter as check #1), "data & coding" (referred as check #2) and "coding-only" (referred as check #3) chunks/symbols. As can be imagined, these check nodes (mathematical relationships) can be added to the two-dimensional bipartite graph to give it a three-dimensional look. This new data representation shall be used to provide advanced decoding, repair, and update features of the library. Throughout the document, nodes typically contain multiple chunks and chunks typically contain multiple symbols.

Founsure uses Belief Propagation (BP) [15] (a.k.a. message passing) algorithm to resolve or decode the user data, to repair the encoded data, or update the encoded data. Sticking to BP as a design criterion is to ensure a low-complexity decoding process and allow fast/efficient operation. Library also supports register-level parallelism through compiler optimizations as well as multi-threading using the open standard *openMP* primitives with its encode, decode, repair and update functions. The multi-threading feature, once properly configured and used, allows parallel processing and ensures acceleration for shared-memory architectures. By reducing the processing time, Founsure

secures quick responses to the common read/write requests of any generic distributed storage system.

With the current release, the original user data does not appear in the output files. Instead, all output files are a mathematical function of the user data due to the so-called non-systematic encoding. In other words, one cannot read off data from the encoder output without any further decoding. Therefore, using Founsure, one can think of the user data encrypted automatically after encoding operation. Note that we use pseudo-random number generators (based on linear congruential generator) and seed (integers of long type) to build edges of the underlying bipartite graph. So without the seed number (we can treat them as keys in an encryption context), there is no way to recover original user data simply because the underlying graph generation is contingent upon the seed availability. Therefore, the Founsure software package also provides a user-configurable lightweight built-in encryption feature. Unlike systematic codes in which the data is explicit at the encoder output, Founsure comes with the non-systematic format as mentioned before. This, however, provides data security in return in addition to data protection. As long as decoding is fast using various novel techniques we discuss in this study, non-systematic codes should not be an overall performance bottleneck for the rest of the system.

This paper shall briefly describe the details of the software architecture, a set of functionalities provided with the library as well as some of the associated advanced features. The source code, a comprehensive user guide, few test results, and all related documentation are available also from [github](https://github.com) and the web link <http://www.suaybarslan.com/founsure.html>.

2. Software description and architecture

2.1. Software functionalities

Founsure has the following three executable main components that implement four important functionalities.

- `founsureEnc`: Encoder engine that generates s number of data chunks (to be stored in s different failure domains) under a local *Coding* directory and a metadata file that includes information about the file, coding parameters, and the seed information.
- `founsureDec`: Decoder engine that requires a local *Coding* directory with enough number of files, a valid file name, and an associated metadata file to run multiple Belief Propagation (BP) passes in order to decode the user data.
- `founsureRep`: Repair engine that also requires a *Coding* directory with sufficient number of files and
 - fixes/repairs one or more data chunks should they have been erased, corrupted, or flagged as unavailable.
 - generates extra coding chunks should a code update has been requested. The system update is triggered if data reliability is decreased/degraded overtime or increased due to equipment replacements.

These functions are used to execute encoding, decoding, repair, and update operations. There are also utility functions of Founsure used to help system admins to make correct design choices on degree distributions, required reliability, desired complexity, and storage space efficiency. We also use utility functions to trigger update functionality as will be demonstrated later. One of the distinctive features of utility functions is that they do not directly process user data, instead, they help us configure right parameters for the main functions to modify and process the user data properly. The current version supports two utility functions as listed below.

- **simDisk**: This function can be used to exhaust all possible combinations of disk failures for a given set of coding parameters. In other words, this function checks whether the provided coding parameters are sufficient to achieve a user-defined reliability goal. Therefore, running this function can help us design target-policy erasure codes by configuring degree distributions for achieving various system-level goals besides reliability.
- **genChecks**: This utility function is crucial for two different important functionalities: (1) fast/efficient repair/rebuild of data and (2) seamless on-the-fly update. For the repair process, it generates two types of checks: check #2 and check #3 and registers them into a `<testfile>_check.data` file using a format described within this document. In case of an update, it modifies the metadata and `<testfile>_check.data` files so that the coding chunks can be updated by running `founsureRep` function.

Next, we provide the details of Founsure encoding, decoding, repair and update operations, particularly the implementation details of `founsureEnc`, `founsureDec` and `founsureRep` functions. For more details on underlying theory and computational complexity, we refer the reader to the appropriate documents such as [16].

2.2. Implementation details of encoding/decoding operations

In graph theory terminology, nodes (sometimes referred to as equations) are represented by graph vertices and node relationships by edges of the graph. There are three types of nodes in a 3-D bipartite graph; *data* nodes, *coding* nodes and *check* nodes. The coding nodes represent a set of linear combinations of data nodes generated through a predetermined mathematical function such as XOR logic operation. Check nodes represent all the local sets of data and coding nodes for which a certain mathematical relationship is satisfied such as even or odd *parity*. A simple mathematical function used by Founsure is the region XOR operation that operates over multiple data blocks of the same size and generates a single block of information. We use f flag to indicate the file name, k to indicate the total number of data nodes/symbols where b of these are the original user data nodes/symbols, n to indicate the total number of coding nodes/symbols, and t to indicate the number of bytes to store per node/symbol.

In `founsureEnc` function, data file with $filesize$ bytes is partitioned into multiple $b \times t$ bytes and each partition is encoded independently as shown in Fig. 1. With the current version, *partition coupling* is not supported between distinct partitions i.e., partitions are processed independently of each other. This technique is currently under investigation and might have interesting performance improvements to our design/implementation in analogy to spatially-coupled LDPC codes [17]. However, the coupling may have different effects for partial disk failures and may eventually lead to non-uniform decoding performances across partitions.

If $filesize$ is not a multiple of $b \times t$ bytes, then we use zero padding to make $filesize$ a multiple. On the other hand, `founsureEnc` also checks whether s divides n ($s|n$). If not, the least largest n is selected automatically to satisfy $s|n$. Such a requirement enables us to store exact same amount of information bytes across different failure domains. This is particularly relevant to a balanced system design where the underlying storage devices are of the same type, quality and equally resilient against various types of failures.

Encoding proceeds as follows. First, a memory space (a buffer) worth $(k + n)t$ bytes is allocated, and the buffer content is initialized to zero. Next, check # 1 equations are generated by an

efficient array LDPC encoding [18,19]. The choice of array LDPC as the precode is to enable efficient encoding operation and fast processing. As a result of this operation, an extra $k - b$ chunks are created to make up a total of k chunks of data. The precoding process is shown as ① in Fig. 1. Later, a total of n coding chunks are generated from the whole set of k data chunks based on an LDGM base code with a configured “FiniteDist” degree and pseudo-random selection distributions. This process is shown as ② in Fig. 1. Finally, n coding chunks are distributed (striped) equally across distinct output files for allocation on s number of drives. We repeat this process for each data partition in a loop and append coding chunks at the end of the corresponding output files. For a given `<filename>.ext` file, we use `<filename>_disk0.Oi.ext` to refer to the i th output file. The number of zeros that appear in the name of output files is set by the “parameter.h” variable `DISK_INDX_STRNG_LEN`.

In Founsure implementation, we have distinct object definitions for encoding, decoding, and repair operations. These objects have the trailer “*Obj” in common and include the same set of parameters in their object fields. For instance, both encoding and/or decoding functions accept `EncoderObj` and/or `DecoderObj` constructs as inputs. Similarly, b and k variables can be accessed using the standard way `EncoderObj.sizesb` and `EncoderObj.sizek`.

Each encoding/decoding object is associated with a seed value (`EncoderObj.seed`¹) from which other seed values and the local sets of data chunks are pseudo-randomly generated. Each coding chunk within `EncoderObj` and `DecoderObj` has their own unique ID. These IDs are used to identify the erased coding chunks. The seed value is used by the pseudorandom generator to create a sequence of integers. These integers form the basis of coding chunk degree number assignment and the selected data chunks for coding chunk computations. These numbers are stored as part of the object and can be regenerated using the same initial seed number followed by the regular recurrence relationship. Let us assume we have s number of output files (failure domains), then we use the default value `EncoderObj.seed + i` as the seed of the i th output file with $0 \leq i < s$.

Each check node c is associated with a degree number c_d (chosen according to an appropriate degree distribution $\Omega(x) = \sum_i \Omega_i$ where Ω_i is the probability of choosing degree i) and c_d data node neighbors are selected to be involved in final symbol computation. The degree distribution $\Omega(x)$ is typically selected to minimize the coding overhead. For instance, the following degree distribution is proposed for Raptor codes [13]

$$\begin{aligned} \Omega(x) = & 0.007969x + 0.49357x^2 + 0.16622x^3 \\ & + 0.072646x^4 + 0.082558x^5 \\ & + 0.056058x^8 + 0.037229x^9 + 0.05559x^{19} \\ & + 0.025023x^{64} + 0.003135x^{65} \end{aligned}$$

where $\Omega_1 = .007969$, $\Omega_2 = .49357$, $\Omega_3 = 0.16622$, $\Omega_4 = .072646$, $\Omega_5 = .082558$, $\Omega_8 = .056058$, $\Omega_9 = .037229$, $\Omega_{19} = .05559$, $\Omega_{64} = .025023$, $\Omega_{65} = .003135$.

However, Founsure does not necessarily minimize overhead. It may optimize overhead, repair bandwidth, and complexity at the same time. We recommend choosing degree distributions that will give us a good trade-off point between these three objectives. A systematic optimization procedure to achieve the desired operating point is the subject of further investigation. Although there is no optimal point for all applications, Founsure is designed to be highly configurable to fit in different requirements and sensitivities of modern storage ecosystems.

¹ The default value selected for the seed is 1389488782 which is experimentally observed to give good recovery performance.

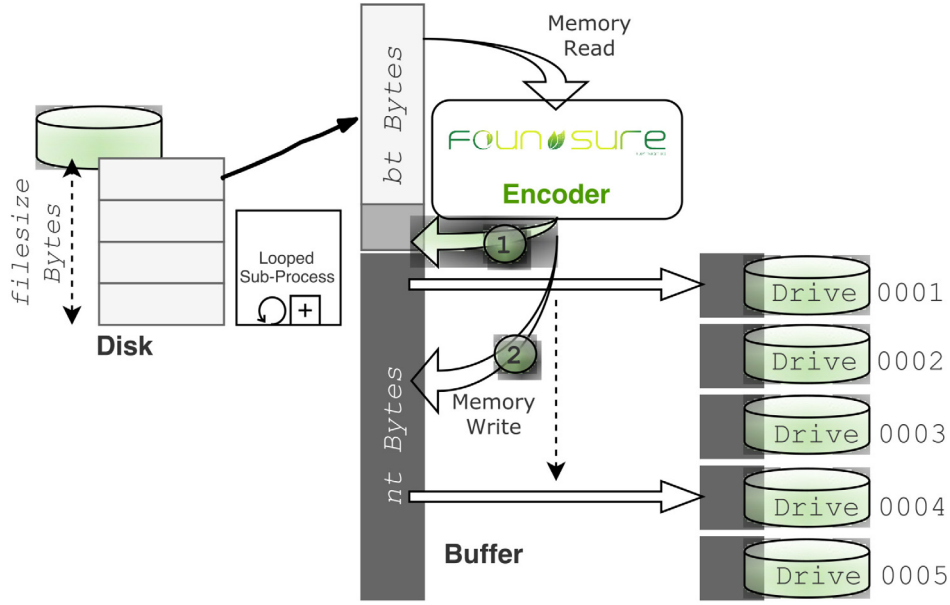


Fig. 1. The function `founsoreEnc` allocates $(n + k)t$ bytes of buffer, generates coding chunks, and write them to a shaded area on memory before they are written to distinct drives. The file on disk is striped and processed in a looped subprocess as shown. The number of stripes is stored in the `readin` variable and written to the metadata file.

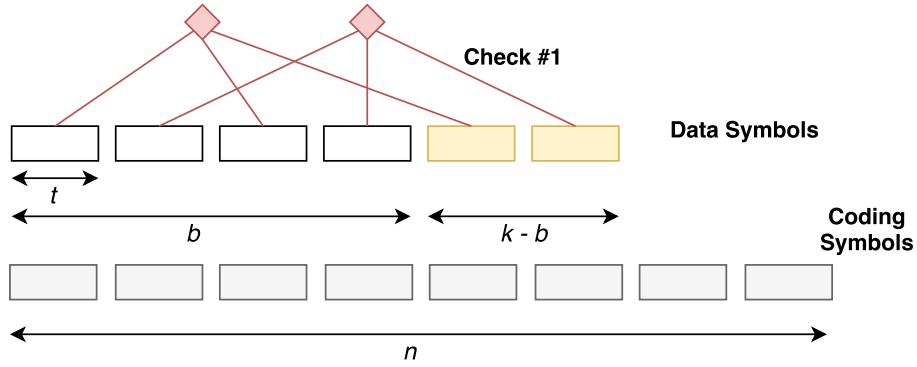


Fig. 2. Founsure symbols are t bytes each and a precode is applied to find check #1 equations. We add $k - b$ extra symbols to satisfy check equations as shown. Founsure encoding engine generates n coding chunks from k chunks as shown.

Algorithm 1 Belief Propagation (BP) Algorithm (*Skeleton*)

<pre> 1: procedure BP($\mathbf{B}, \mathcal{E}, maxit$) 2: $\mathcal{N} = \{1, \dots, n\}$ 3: $\mathcal{K} = \{1, \dots, k\}$ 4: $\mathbf{C} \leftarrow \mathbf{b}_{:, \mathcal{N} \setminus \mathcal{E}}$ 5: $\mathcal{F} \leftarrow \mathcal{K}$ 6: while $\mathcal{F} \neq \emptyset$ and $i < maxit$ do 7: for $g = 1, \dots, \mathcal{N} \setminus \mathcal{E}$ do 8: $\mathbf{c}_{\mathcal{K} \setminus \mathcal{F}, g} \leftarrow \mathbf{0}$ 9: if $weight(\mathbf{c}_{:, g}) = 1$ then 10: $\mathcal{F} \leftarrow \mathcal{F} - \{f : \mathbf{c}_{f, g} = 1 \text{ for } f \in \mathcal{F}\}$ 11: return \mathcal{F} </pre>	<p>▷ Inputs: $\mathbf{B} \in \mathbb{F}_2^{k \times n}$, $\mathcal{E} \subset \mathcal{N}$: The set of erasures.</p> <p>▷ Initialize indexes.</p> <p>▷ Initialize indexes.</p> <p>▷ Find survival matrix.</p> <p>▷ Initialization: \mathcal{F} holds the unrecovered indexes.</p> <p>▷ Zero-out rows for decoded symbols.</p> <p>▷ Find degree-1 coding symbols.</p> <p>▷ Return unrecovered indexes.</p>
--	--

We run `founsoreDec` when we want to collect a subset of output data files and recover the input data file. Decoder is based on belief propagation algorithm a summary of which is provided in Algorithm 1. BP function admits `DecoderObj`, indexes of erasures $\mathcal{E} \subset \mathcal{N} = \{1, 2, \dots, n\}$, the generator matrix of the base code $\mathbf{B} \in \mathbb{F}_2^{k \times n}$ and a maximum number of iterations `maxit`. In Algorithm 1, we use $\mathbf{b}_{i,:}$ to refer to the i th row of \mathbf{B} and $\mathbf{b}_{:,i}$ to refer to the i th column of \mathbf{B} . Additionally, $\mathbf{b}_{\mathcal{A}, \mathcal{B}}$ refers to a matrix

whose rows and columns are given by the rows and columns of \mathbf{B} indexed by the sets \mathcal{A} and \mathcal{B} . The decoder utilizes the information contained in metadata file to generate (prepare) the contents of `DecoderObj`, particularly the underlying coding graph. It works in a similar fashion to `founsoreEnc` i.e., it reads the striped coding chunks, loads the buffer, and runs BP algorithm at most twice (once for the outer graph code and if need be, additional one for the inner Array LDPC precoder) and recovers the bt bytes at each

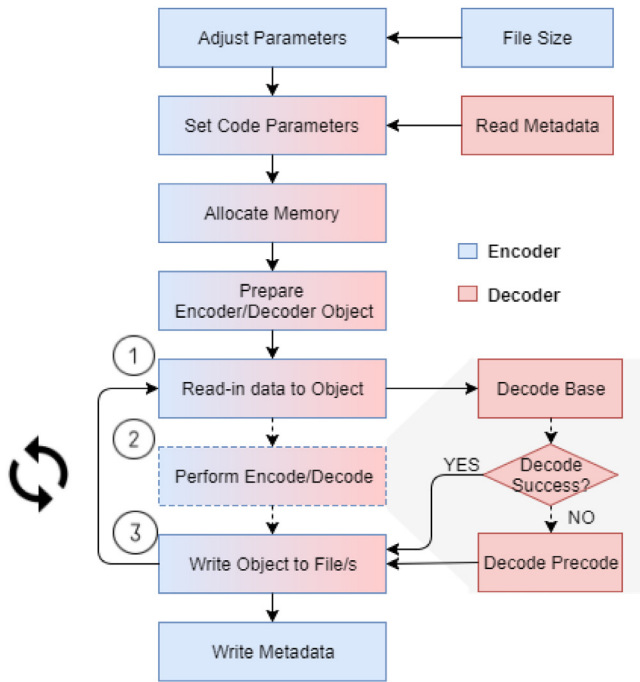


Fig. 3. Summary of software blocks of Founsure Encoder and Decoder.

turn. Finally, these bytes are written to decoded/recovered data file by calling standard kernel I/O commands.

To summarize the software architecture of Founsure’s encoding and decoding schemes, we provide Fig. 3 to illustrate the order of software building blocks that take place to encode/decode the user data. The repair operation’s architecture resembles a lot to this figure and hence omitted to save space. As can be seen, we used different colors for blocks to differentiate the encoding and decoding processes. Some of these blocks are used only by one of the functions. On the other hand, some blocks are used by both the encoder and decoder processes.

Having all computation based on pseudorandomly selected chunks and carrying out these computations solely in terms of simple XOR logic has the cost of making the code non-optimal in terms of overhead (though it might be near-optimal through choosing appropriate degree distributions). If n coding symbols are distributed over s drives and when one of the drives fail, a subset of coding symbols are lost. To find what fraction of f -failure combinations can be tolerated for a given degree distribution, we provide a utility function `simDisk` that exhausts all possible combinations of failures to report which failures cases are tolerable by the code and which are not. Such a utility function is extremely useful for determining the reliability of the data protected by the Founsure library.

2.3. Check equations and the data repair process

Given the three-dimensional Founsure graph representation as shown in Fig. 4, we have three types of check nodes as mentioned before. We provide the details of such a checking process in this subsection.

Checks #1: These check equations are defined by the precode of the Founsure (for version 1.0, we selected an Array LDPC code family [18] for efficient processing as given in Algorithm 2). Based on the selection of good precodes, the mathematical and coding parameter selections etc., the graph connections are automatically determined. Founsure includes a precode support based on a binary array LDPC code. Future releases of the library




Fig. 4. Founsure graph code is three dimensional bipartite graph with three sets of check equations.

shall include external precode support which can be provided by the user using a preformatted input file. Please see the precoding subsection to find more information about the construction of these check equations.

Checks #2: These check equations are generated as given in Algorithm 4. One of the special features of these checks is that only one neighbor is selected from the data nodes and the rest of the neighbors of the check node are from the coding nodes. This special feature can be used to partially decode the input data without running the complete decoder and reconstruct the unnecessary parts of the input data. An application of this could be securely stored multimedia source in which the Region of Interest (ROI) can be directly reconstructed using this type of check equations.

Checks #3: These check equations are generated as given in Algorithm 4. These checks form the local groups based on the coding nodes. These checks are primarily used to repair the permanently erased, long-time unavailable, or unresponsive coding nodes in case of hardware, software, and network failures.

2.4. Precoding process - generation of check #1 equations

As shown in Fig. 2, a (b, k, n) Founsure code takes b data symbols (a total of bt bytes) and initially generates $k - b$ check #1 parity symbols based on the binary array LDPC encoding [18]. This special choice of array LDPC codes enables efficient encoding operation (linear with blocklength) and improves the complexity performance of the overall Founsure library.

The procedure outlined in Algorithm 2 uses a generic function `largest_prime_factor(.)` which chooses the largest prime factor of the argument. The rate of the array LDPC is defined as $r_{LDPC} = b/k$. The user can choose any k, n and r_{LDPC} and hence we can calculate the appropriate $b = \lfloor kr_{LDPC} \rfloor$. Let $p = \text{largest_prime_factor}(k)$, we may not be able to get the quantity $\lfloor kr_{LDPC} \rfloor / p$ equal to an integer. We can use the floor function to get an estimate of j' . However, the array LDPC code performance is heavily dependent on k' and j' values and there is no array LDPC code for all (k, r_{LDPC}) pairs. If (j', k') pair are small, the code performance is observed to be pretty bad. For this reason, we provide an algorithm that reasonably chooses a good performing array LDPC code and satisfies (within some error margin) the user-provided parameters k, n and r_{LDPC} at the same time. One can see the chosen parameters by adding “-v” flag at the end of Founsure main functions.

Let us define the following system parameters. After that, we shall formally provide the algorithm that determines the closest good-performing array LDPC code for the user-provided parameters k, n and r_{LDPC} . These system parameters with their default values are defined in “parameter.h” file and can easily be modified.

- DIFF_TH: Allowed error threshold between the estimated and user provided b values.

Algorithm 2 Array LDPC Checks (check # 1)

```

1: procedure CHECKONE( $b, k$ )
2:    $\mathbf{L}^{(1)} \leftarrow \mathbf{0}_{k \times k - b}$ 
3:    $p \leftarrow \text{largest\_prime\_factor}(k)$ 
4:    $k' \leftarrow k/p$ 
5:    $j' \leftarrow k' - b/p$ 
6:   for  $j = 0; j < j'; j++$  do
7:     for  $i = 0; i < p; i++$  do
8:       for  $m = 1; m < k' - j' + 1; m++$  do
9:          $\mathbf{L}_{i+jp, m-1}^{(1)} \leftarrow k'p - (j' - j + m - 1)p - (m(j' - j - 1) - i - 1 + p) \pmod{p} - 1$ 
10:  return  $\mathbf{L}^{(1)}$ 

```

▷ Inputs: b, k
▷ Initial zero matrix
▷ $\mathbf{L}^{(1)}$: check #1 sets.

- RRATE_TH: Allowed error threshold between the estimated and user provided precode rate.
- RED_BYTE_TH: Allowed redundant zero bytes to be appended at the end of the file for parameter consistency.
- RAND_WIN_MAX: Random number search window maximum value.
- RAND_WIN_MIN: Random number search window minimum value.
- ARRAY_MIN_JJ: Minimum Array LDPC “j” parameter value.
- ARRAY_MIN_KK: Minimum Array LDPC “k” parameter value.
- TRIES_TH: Threshold on the number of tries before incrementing DIFF_TH and RRATE_TH.
- DELTA_DIFF_TH: Step size increment for DIFF_TH
- DELTA_RRATE_TH: Step size increment for RRATE_TH.

Next, we provide the algorithm that returns the estimated values of b, k , and the redundant number of zeros (*redundantzeros*) that need to be appended to the input user data. The algorithm admits four inputs, namely r_{LDPC} , *filesize*, b and t . Initial values of system parameters shall be set by “parameter.h” file and are changed locally within the function implementing Algorithm 3.

2.5. Generating information for efficient repair

To efficiently repair the lost data, we need to extract repair information from the underlying graphical content of the generated Founsure code. We observe that check #3 nodes are the most suitable node type for the repair process since it establishes a direct relationship between the coded symbols. It is not hard to see having more of these node types (created independently or dependently) gives alternative ways of repairing a given node in case of different combinations of node failures happening in the communication network. In other words, the more of these check types we find, the more potential we have for the regeneration of the lost coded chunks. With regard to this observation, we can use two techniques to increase the number of check #3 type information based on check #1 and check #2 equations:

- The first method is as follows. We identify the coded nodes with degree one (say we have M of those nodes). Identify their data node neighbors. There are M check #2 equations that connect these data nodes with the coding nodes. Since the corresponding M coding nodes carry the same information, we can use these check #2 type check equations as additional check #3 check equations. Note that since check #2 and check #3 equations are derived from the same base graph, this technique is likely to generate already existent local recovery groups or local groups that can be derived from existent local groups for the coding nodes.
- The second method is as follows. Note that check #1 is user-defined although subject to a predefined structure. This defines local recovery groups over the data nodes. Since

Algorithm 3 Adjust Parameters with Precode (APP)

```

1: procedure APP( $r_{LDPC}, filesize, b, t$ )
2:    $\hat{b} \leftarrow b, b \leftarrow 0, k \leftarrow 1, k' \leftarrow 1e9, j' \leftarrow 1e9, p \leftarrow 2, blocks \leftarrow 1, tries \leftarrow 0, iter \leftarrow 0$ 
3:   while  $(|b - \hat{b}| > \text{DIFF\_TH} \parallel |b/k - r_{LDPC}| > \text{RRATE\_TH} \parallel k' > p \parallel j' > p$ 
4:      $\parallel k' < j' \parallel |blocks \times t \times b - filesize| > \text{RED\_BYTE\_TH}$  do
5:      $k \leftarrow \lfloor \hat{b}/r_{LDPC} \rfloor + \text{rand}() \pmod{(\text{RAND\_WIN\_MAX} + \text{RAND\_WIN\_MIN}) - \text{RAND\_WIN\_MIN}}$ 
6:      $p \leftarrow \text{largest\_prime\_factor}(k)$ 
7:      $k' \leftarrow k/p, j' \leftarrow k' - \hat{b}/p$ 
8:     if  $j' < \text{ARRAY\_MIN\_JJ}$  then
9:        $j' \leftarrow 2$ 
10:     $b \leftarrow (k' - j')p$ 
11:    if  $b > 0$  then
12:       $blocks \leftarrow filesize / tb + 1$ 
13:       $tries++$ 
14:      if  $tries > \text{TRIES\_TH}$  then
15:         $\text{DIFF\_TH} \leftarrow \text{DIFF\_TH} + \text{DELTA\_DIFF\_TH}$ 
16:         $\text{RRATE\_TH} \leftarrow \text{RRATE\_TH} + \text{DELTA\_RRATE\_TH}$ 
17:         $tries \leftarrow 0$ 
18:       $iter++$ 
19:      if  $iter > 1e7$  then
20:        print error and exit.
21:     $blocks \leftarrow filesize / (tb) + 1$ 
22:     $redundantzeros \leftarrow bt \times blocks - filesize$ 
23:  return  $b, k, redundantzeros$ 

```

each data node is linked to local recovery groups of coded nodes through check #2, we can use this relationship to derive check #3 type check equations. For example, suppose we have the following check #1 local recovery group defined for data nodes D_0, D_1 and D_2 : (D_0, D_1, D_2) . Also, suppose that we have the following check #2 equations:

$$\Rightarrow D_0 = C_0 \oplus C_1 \oplus C_2 \quad (1)$$

$$\Rightarrow D_1 = C_1 \oplus C_{12} \oplus C_{17} \oplus C_{20} \oplus C_{99} \quad (2)$$

$$\Rightarrow D_2 = C_0 \oplus C_{21} \oplus C_{99} \quad (3)$$

Thus, we can find a check #3 type equation given by $(C_2, C_{12}, C_{17}, C_{20}, C_{21})$ by observing the following equivalence,

$$C_2 \oplus C_{12} \oplus C_{17} \oplus C_{20} \oplus C_{21} = D_0 \oplus D_1 \oplus D_2 \quad (4)$$

Note that since this technique uses check #1 equations, it is likely to generate distinct check #3 local recovery groups and help improve repair performance dramatically. These additional check #3 local recovery groups (for instance the operation of Eq. (4)) are efficiently computed by the set union function *setXOR* given in “encoder.c” file.

2.6. Algorithm for jointly generating check #2 and check #3 equations

We propose a heuristic algorithm to generate check #2 and check #3 equations at the same time for efficiency. This algorithm uses XOR operation (\oplus) to sparsify the generator matrix \mathbf{B} . If \mathbf{B} is full rank, i.e., $\text{rank}(\mathbf{B}) = k$ then the algorithm is guaranteed to converge successfully. This is due to elementary matrix row operations shall generate $n - k$ zero columns for a full rank \mathbf{B} and hence the algorithm will leave the main while loop and generating a modified \mathbf{B} with column weights k ones and $n - k$ zeros. In mathematical terms at the end, we should be able to find a permutation matrix \mathbf{P} such that $\mathbf{BP} = [\mathbf{I}_{k \times k} \mid \mathbf{0}_{k \times n-k}]$. Also, $\tilde{\mathbf{L}}^{(2,3)} = \mathbf{L}^{(2,3)}\mathbf{P}$ shall hold all the local recovery sets i.e., check #2 equations in the first k columns and check #3 equations in the last $n - k$ columns. We can express different types of checks as the union of all check #2 and check #3 equations as given by

$$\bigcup_{i < k} (D_i, \{C_s : \tilde{\mathbf{L}}_{s,i}^{(2,3)} \neq 0\}) \cup \bigcup_{i \geq k} (\{C_s : \tilde{\mathbf{L}}_{s,i}^{(2,3)} \neq 0\}) \quad (5)$$

where $\tilde{\mathbf{L}}_{s,i}^{(2,3)}$ denote sth row and ith column entry of $\tilde{\mathbf{L}}^{(2,3)}$. In Algorithm 4 we provide the details of the algorithm using pseudocode. We use a simple function `zero_columns(.)` that finds the number of nonzero columns of the matrix in the argument. Since \mathbf{B} is typically sparse, we use a sparse representation of matrices in the library implementation for efficient memory utilization. We also order the local recovery groups based on their cardinality i.e., the set with the smallest cardinality comes first. Such an arrangement helps us reduce the repair/update complexity since the repair function processes local groups in sequential order.

2.7. Management of check #2 and check #3 equations and the generation of `<filename>_check.data` file for efficient repair/update process

Check #1 equations are determined through a binary array LDPC code as explained before. The user-defined number of equations are selected from a set of precode rates based on the reliability imposed by the application. The graph connections are deterministic and given by the constraints of the array code.

Unlike check #1, check #2 and check #3 are determined pseudo-randomly by the Founsure base code. Based on the generator matrix of the code \mathbf{B} , Algorithm 4 is run to determine n equations. If the algorithm converges, then we should have k equations for check #2 type and $n - k$ equations for check #3 type. The algorithm produces a correct set of local equations (sets) but does not guarantee those equations to be independent. In generating those equations, we do not employ any matrix inversions (which is quite costly for large size matrices) to find check equations and hence we trade off the efficiency by performance. The function that generates check #2 and check #3 local recovery equations is the utility function `genChecks`.

The function `genChecks` assumes that a metadata file is already generated by a previous run of the encoder `founsureEnc`. Hence `genChecks` generates check groups and modifies the `meta_data` file (appends the size of check data in terms of `sizeof(int)` bytes at the end of the metadata file if “-m” flag parameter is True). The check information is stored in another binary file called `<filename>_check.data`. This file stores an integer array with a specific format. The reason for introducing a format is to use bulk read/write capabilities of `fread` and `fwrite` C library functions which will make kernel’s I/O performance acceptable.

The proposed format in this study is pretty straightforward and can be improved. We use flag bits to differentiate between the two distinct check equations. Thus, the integer value of the first `sizeof(int)` bytes in `<filename>_check.data` is either 0 or 1.

- If it is 1 (Check #2), then the next integer value (next `sizeof(int)` bytes) gives the data symbol index which is involved within a local recovery group whose degree is given by the following integer (next `sizeof(int)` bytes). This degree also indicates the next “degree” number, i.e., the number of integers to be read as part of one local recovery group for the coded symbols.
- If it is 0 (Check #3), then the next integer value (next `sizeof(int)` bytes) gives the degree number i.e., the number of integers to be read as part of one local recovery group for coding symbols.

The nice thing about Algorithm 4 is that if it converges, then all of the data symbols are covered exactly by one particular Check #2 local recovery equation. Let us provide an example to illustrate the working principle and suppose that we have the following integer array stored in `<filename>_check.data`:

```
0 4 13 56 17 66 1 19 2 11 13 0 2 39 88 ...
```

If we decode this integer array, we will be able to say that the first local set is of type check #3 and this set has four elements. In other words, 13th, 56th, 17th, and 66th coding symbols form a local recovery group i.e., their binary sum should produce all-zero content. The next local set belongs to check #2 and the associated data symbol index is 19. This data symbol along with 11th and 13th coded symbols (two coding symbols) forms a local recovery group. This way we can decode the whole integer array stored in `<filename>_check.data`. If the algorithm converges, there should be k leading 1’s and $n - k$ leading 0’s in the integer array not necessarily written in sequential order. Note that the total number of integers contained in the array is given by

$$N = \sum_{c=0}^{n-1} L_c + k + 2n \quad (6)$$

where L_c is the total number of elements in check #2 (excluding the data symbols) and check #3 indexed by c . Note that even if the algorithm does not converge, the maximum memory occupancy possible is $N \times \text{sizeof(int)}$ bytes. So it is sufficient to allocate the size of memory given by Eq. (6) for the file without encountering a segmentation fault.

Fig. 5 summarizes how different functions of Founsure interact with each other, what other metadata is generated/used, and what type of read/write permissions are granted to each of these functions for the proper operation of the Founsure library.

2.8. Reading/formatting the contents of `<filename>_check.data` file

When the repair process is initiated, memory allocation, and repair object (`RepairObj`) preparation, begins. The main repair engine shall look for `<filename>_check.data` under `/Coding` directory. If it finds one and if the metadata is appropriately formatted (after a successful format check), it will read-in the metadata and format the check # 2 and check # 3 equations for the preparation of `RepairObj`. A bulk read kernel call is performed and all the content is transferred to memory (inside the buffer pointed by `content2read`). Since Founsure’s decoding, repair, and update operations are solely based on the BP algorithm, it sequentially searches only one unknown over the available local sets in a loop. To reduce the computation and bandwidth, the repair/decode process must use small size check # 3 equations first so that we do not have to run through the end of the loop to complete the overall repair process. Founsure implementation extracts check # 3 equations from the buffer (`content2read`) using the standard `qsort(.)` function and then fills in the appropriate fields of `RepairObj`. The ordering can be enabled or disabled for check # 2 and # 3 equations using parameters `ORDER_CHECK_2` and `ORDER_CHECK_3` in “parameter.h” file.

Algorithm 4 Check Generation Algorithm (CGA)

```

1: procedure CGA(B)
2:    $\mathbf{L}^{(2,3)} \leftarrow \mathbf{I}_{n \times n}$ 
3:   while zero_columns(B) <  $n - k$  do
4:     temp  $\leftarrow 0$ 
5:     for j = 0; j < n; j++ do
6:       for i = 0; i < n; i++ do
7:         if  $j \neq i$  and  $2w(\mathbf{b}_{:,j}) > w(\mathbf{b}_{:,i})$  then
8:           if  $w(\mathbf{b}_{:,j} \oplus \mathbf{b}_{:,i}) < w(\mathbf{b}_{:,j})$  then
9:              $\mathbf{b}_{:,j} \leftarrow \mathbf{b}_{:,j} \oplus \mathbf{b}_{:,i}$ 
10:             $\mathbf{l}_{:,j}^{(2,3)} \leftarrow \mathbf{l}_{:,j}^{(2,3)} \oplus \mathbf{l}_{:,i}^{(2,3)}$ 
11:            temp  $\leftarrow 1$ 
12:     if temp = 0 then
13:       break;
14:   return B,  $\mathbf{L}^{(2,3)}$ 

```

▷ Input matrix $\mathbf{B} \in \mathbb{F}_2^{k \times n}$
 ▷ Identity matrix
 ▷ Number of zero-columns is checked.
 ▷ **B**:Check type, $\mathbf{L}^{(2,3)}$:Local sets.

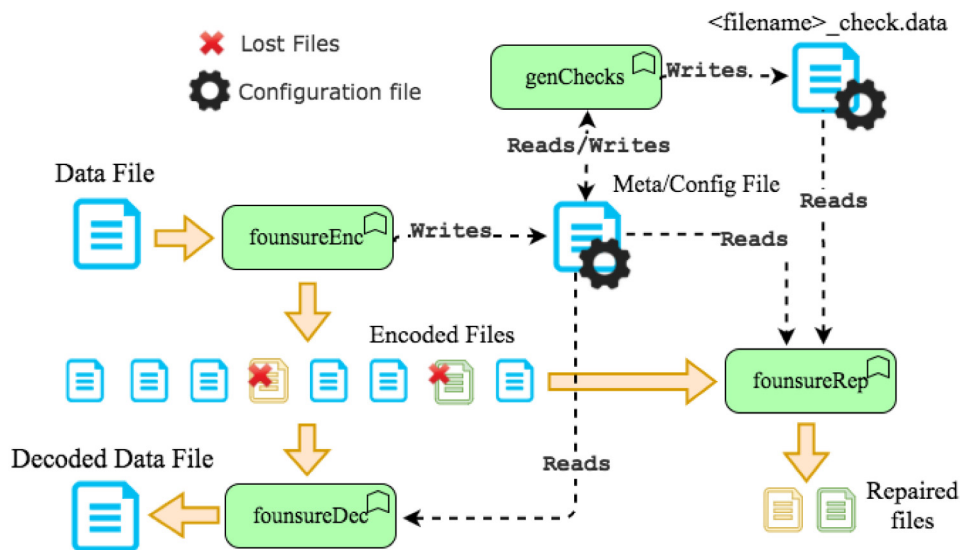


Fig. 5. Summary of interactions between different functions of Fountainsure.

2.9. Update process

An update process is about making the existing Fountainsure code stronger or weaker by either generating more redundancy (in case of increased failures or wear-out) or taking away unwanted redundancy (in case of using more reliable devices for storing information). If we would like to make the existing code weaker, it would not be hard. We just need to modify the metadata accordingly and erase the redundancy manually. Fountainsure does not automatically erase files and leave them to upper layer software management. So for the rest of this section, updating the code structure would mean making the code stronger.

The desirable features of a generic update process can be listed as follows.

- An update process should minimize the modification of the data generated by the encoding operation.
- An update process should generate extra redundancy consistent with the encoded data with minimum processing effort.
- An update process should have a minimum limit on the extent of extra redundancy that can be generated.
- An update process should not violate the rules set by the encoding and decoding processes.

Fountainsure's update mechanism poses no modification changes to the already encoded data. In that sense, its update process

functions as ideal as possible. Fountainsure update process is tightly related to the repair process. This is mainly because updating a code is about repairing the missing blocks of information to help increase the reliability of data. We call genChecks to update the current code using the flag '-e'. There must be valid metadata associated with the code at the time genChecks is called. The code update process will rewrite n , the number of bytes used for the integer array due to check #2 and check #3 equations and update <filename>_check.data. Hence the repair process uses the metadata (the rule set) generated by the previous runs of the encoder/decoder pair. This series of modifications do not make any changes to the existent data/coding chunks. To trigger/sync changes with data, we finally need to call founsureRep function with the appropriate file name. Since the existent data is only read and we use minimum cardinality, local recovery groups, while updating, the processing effort is minimized. We finally note that since Fountainsure is based on fountain-like codes, there is no practical limit to the number of coding symbols that can be generated. As can be seen, the update functionality of Fountainsure is designed and implemented in the observation of desirable features listed above.

3. Advanced features

Although many advanced features of Founsure are described in previous sections, we have two more important implementation-specific advanced features that make Founsure's performance stand out.

3.1. Shared memory parallelism

In shared-memory multiprocessor architectures, threads can be used to implement parallelism. The shared-memory standard openMP is a high-level and portable interface that makes it easier to use multi-threading capability and obtain satisfactory performance improvements. Many erasure coding libraries such as Jersure 2.0 [7] has encoding/decoding engines that comprise independent "for" loop iterations and hence possess huge potential for multi-threaded processing. Multi-threaded implementations of Jersure 2.0 are studied in [20] and [21].

As can be seen in Fig. 3 for founsureEnc, the software consists of three stages executed in a loop. However, two of these stages, namely reading the data into the *EncoderObj* and *DecoderObj* that stay on memory and writing the object contents to the persistent storage devices require kernel I/O calls. As a result of these calls, the performance will be inhibited by the throughput performance of the underlying storage devices I/O bandwidth. Thus, our focus is essentially the second stage, namely the pure encoding and decoding in which the data traverses only between CPU caches and the main memory.

Founsure utilizes the shared memory approach standard OpenMP library directives to help use multiple threads to handle the workload of encoding, decoding, repair, and update operations in parallel. However, to use OpenMP directives effectively, we needed to implement encoding/decoding operations differently. We use '-m' flag to set the number of threads in the main functions of Founsure. This parameter can independently be assigned but we provide recommendations for picking out the right number of threads for each main function because selecting the wrong number could result in degraded performance. For instance, we recommend it to be equal to the number of failure domains (disks for instance) for founsureEnc so that each data block is generated by a different thread. Considering each output file gets written to a different disk or storage node, we can maximize the overall throughput of the system if '-m' and '-s' flags are set to the same number provided that the underlying CPU architecture supports that many concurrent threads.

In founsureDec, remember that we use Algorithm 1 to resolve the user data in an iterative manner. Suppose that for a maximum of convergent t_m iterations, we decode a set of data symbols (also known as the ripple size) \mathcal{G}_i at the i th iteration for $i \in 0, 1, \dots, t_m$. However, we note that a decoded data symbol $s \in \mathcal{G}_i$ might be using another symbol $h \in \mathcal{G}_i$ while decoding, which results in intra-iteration decoding dependency. If we let j th iteration to use only decoded symbols in $\cup_{i=0}^{j-1} \mathcal{G}_i$, this would lead to another decoded set sequence $\bar{\mathcal{G}}_0, \bar{\mathcal{G}}_1, \dots, \bar{\mathcal{G}}_{\bar{t}_m}$ where $\bar{\mathcal{G}}_0 = \bar{\mathcal{G}}_0$ and $\bar{t}_m > t_m$. Note that upon convergence, we should have $\cup_{i=0}^{t_m} |\mathcal{G}_i| = \cup_{i=0}^{\bar{t}_m} |\bar{\mathcal{G}}_i|$. Although this new delayed BP will converge late compared to original version with single thread, this observation is not necessarily true with multi-threads as $\bar{\mathcal{G}}_i$ s can be computed by multiple threads because data symbols in $\bar{\mathcal{G}}_i$ are decoded completely independent of each other and decoding process for each only share data for read operations eliminating potential race conditions. We note that for a given k , if we increase the block length n we would need less number of iterations i.e., smaller t_m and \bar{t}_m with larger ripple sizes in each iteration. Finally, we note that we have many calls of Algorithm 1 for decoding independent partitions of the user data. Using a

shared memory approach, we use multi-threading to compute \mathcal{G}_i in parallel in each iteration. Thus, the larger is ripple size, the better would become the performance of our implementation. It is recommended to use more threads as the number of coding blocks n increases. We also recommend testing the best number of threads for a given n to find the optimal value because this number is heavily dependent on the degree distribution. Finally, the multi-threaded implementation of repair/update operations is similar to decoding since in both cases, we resolve the repaired/updated data using the BP algorithm. For founsureRep function, the recommended number of threads equal to either the number of repaired coding blocks or the number of extra coding blocks generated out of an update operation.

We have two functions that take advantage of the multi-core multi-threaded systems and carry out the main operations of Founsure in parallel. These functions are *EncodeComputeFast_mt* for performing the data encoding and generating output files simultaneously and *runBP_mt* which runs the BP algorithm as parallel as possible. In our revised BP implementation, we remove the intra-iteration decoding dependency (see also Algorithm 5) by

- allowing BP to proceed using only the decoded symbols in $\cup_{i=0}^{j-1} \mathcal{G}_i$ in j th iteration or decoding step,
- allowing only one particular (lowest-degree) coding symbol to decode each source symbol in \mathcal{G}_j ,

where the latter eliminates the possibility of race conditions (double writes by different threads) and optimizes the complexity performance by reducing the number of XOR operations. Note that as the number of failures increases, the number of coded symbols decreases, and hence finding the lowest degree coding symbol will not usually end up with much performance improvement. We finally note that, since different threads deal with different levels of workloads, we use dynamic scheduling of threads in openMP.

3.2. Optimal decoding path generation

In this section, we assume that node degree and selection distributions of Founsure is determined by different requirements of the system. Also, \mathcal{DP} represents the set of source and coding symbol pairs in which the coding symbol is used to decode the paired up source symbol. Thus, in the case of convergence of BP, we should expect $|\mathcal{DP}| = 2k$. The elements of \mathcal{DP} is found according to Algorithm 5. A careful look at the algorithm reveals that the following line does the local optimization of finding the lowest-degree coding symbol that decodes a specific source symbol.

$$\begin{aligned} \mathcal{DP} \leftarrow \mathcal{DP} \cup \{(f, g) : \text{For each } f \in \mathcal{F}, g \in \bar{\mathcal{G}}_i \\ \text{s.t. } \mathbf{c}_{f,g} = 1 \text{ and } g_d \text{ is minimum.}\} \end{aligned} \quad (7)$$

One another note about Algorithm 5 is that by keeping unrecovered symbols in \mathcal{F} , we do not allow the same source symbol to be decoded more than once. This leads to suboptimality but helps us with multi-threaded implementation since it will save us from dealing with race conditions that would otherwise be handled with time-consuming locks. Also comparing it with Algorithm 1, we can observe that symbol decodings are done one iteration at a time and hence symbols that are decoded at a given iteration do not help with other symbols that could have potentially be decoded within the same iteration. This approach is adapted to help with the shared memory implementation of the previous subsection.

Algorithm 5 Decoding Path Generation (DPG) Algorithm.

```

1: procedure DPG(B,  $\mathcal{E}$ , maxit)
2:    $\mathcal{N} = \{1, \dots, n\}$ 
3:    $\mathcal{K} = \{1, \dots, k\}$ 
4:    $i \leftarrow 0$ ,  $\mathcal{DP} \leftarrow \emptyset$ 
5:    $\mathbf{C} \leftarrow \mathbf{b}_{\mathcal{N} \setminus \mathcal{E}}$ 
6:    $\mathcal{F} \leftarrow \mathcal{K}$ 
7:   while  $\mathcal{F} \neq \emptyset$  and  $i < \text{maxit}$  do
8:      $\mathbf{c}_{\mathcal{K} \setminus \mathcal{F},:} \leftarrow \mathbf{0}$ 
9:      $\bar{\mathcal{G}}_i \leftarrow \{g : \text{weight}(\mathbf{c}_{:,g}) = 1 \text{ for } g = 1, \dots, |\mathcal{N} \setminus \mathcal{E}|\}$ 
10:     $\mathcal{DP} \leftarrow \mathcal{DP} \cup \{(f, g) : \text{For each } f \in \mathcal{F}, g \in \bar{\mathcal{G}}_i \text{ s.t. } \mathbf{c}_{f,g} = 1 \text{ and } g_d \text{ is minimum.}\}$ 
11:     $\mathcal{F} \leftarrow \{f : \text{weight}(\mathbf{c}_{f,g_i}) = 0 \text{ for } f \in \mathcal{F}\}$ 
12:     $i \leftarrow i + 1$ 
13:  return  $\mathcal{F}$ 

```

▷ Inputs: $\mathbf{B} \in \mathbb{F}_2^{k \times n}$, $\mathcal{E} \subset \mathcal{N}$: The set of erasures.
 ▷ Initialize indexes.
 ▷ Initialize indexes.
 ▷ Find survival matrix.
 ▷ Initialization: \mathcal{F} holds the unrecovered indexes.
 ▷ Zero-out rows for decoded symbols.
 ▷ Find degree-1 coding symbols.
 ▷ Update \mathcal{F}
 ▷ Return unrecovered indexes.

4. Illustrative examples

In this section, we provide a set of commands to use encoding, decoding, repair, and update features of Founsure. Note that Founsure comes with man pages or you can always use “-h” flag command for immediate help when you call Founsure functions. Moreover, these examples are also included in the GitHub page along with several performances and unit tests.

The following command will encode a test file *testfile.txt* with $k = 500$ data chunks with each chunk occupying $t = 512$ bytes. The encoder generates $n = 1000$ coding chunks using $d = \text{‘FiniteDist’}$ degree distribution and $p = \text{‘ArrayLDPC’}$ precoding. Finally, generated chunks are striped/written to $s = 10$ distinct files for default disk/drive allocation under */Coding* directory.² The flag “-v” is used to output parameter information used during the encoding operation. Founsure encoder also generates a metadata file with critical coding parameters which will later be useful for decoding, repair, and update operations. Without appropriate metadata, Founsure cannot operate on files.

```
founsureEnc -f testfile.txt -k 500 -n 1000 -t 512
-d 'FiniteDist' -p 'ArrayLDPC' -s 10 -v
```

Now, let us erase one of the coding chunks and run the Founsure decoder. The decoder shall generate a decoded file *test_file_decoded.txt* under */Coding* directory. You can use “diff” command to compare this file with the original.

```
rm -rf Coding/testfile_disk0007.txt
founsureDec -f testfile.txt -v
```

One of the things we notice about founsureDec function is that it does not recover the lost drive data *Coding/testfile_disk0007.txt*, because this function is responsible only for the original data recovery process. In storage systems, however, we need to recover lost data to maintain acceptable data reliability. In Founsure, it is extremely easy to initiate the repair (current version only supports exact repair at the moment) process by running the following command.

```
founsureRep -f testfile.txt -v
```

This would trigger the conventional repair operation and first shall decode the entire data and then re-run partial encoding to generate the lost chunks. In addition, founsureRep outputs the pure computation speed as well as the bandwidth consumed due to repair. We observe that conventional repair is a heavily time and bandwidth-consuming operation. In fact, due to non-optimal overhead, the number of bytes that need to be transferred for the

conventional repair is a little larger than the size of the original user file. Founsure supports fast and efficient repair as well. In order to use this feature, one needs to modify the metadata file and create an extra helping data/file called *testfile_check.data* which shall contain information for fast repair. Details can be found later in the document. To make these changes, we primarily run genChecks function. Finally, we can re-run the repair function as before and you will realize from the comments pointed out that the function will be able to recognize that there is available information for fast/efficient repair and will run that process instead of switching to conventional repair. You should be able to observe the reduced bandwidth consumed by the repair operation.

```
genChecks -f testfile.txt -m 1 -v
founsureRep -f testfile.txt -v
```

We can also use genChecks to trigger ‘update’ functionality. For example, let us assume that the system reliability is degraded due to drive wear and we want to generate an extra two drive-worth information, in addition to already generated 10 drive-worth information. We use ‘-e’ flag to modify metadata file as well as *testfile_check.data* for the update operation. This shall change the code and all its related parameters. However, in order to apply it to encoded data, we shall use founsureRep to generate new coding chunks and output files. Alternatively, you can erase drive info as well by supplying negative values for ‘-e’ flag. In this case, you do not need to call founsureRep because there is nothing to generate. You can simply erase corresponding drive chunks after you scale the system down.

```
genChecks -f testfile -m 1 -v -e 2
founsureRep -f testfile.txt -v
```

5. Numerical results and impact

To check the set-up, accuracy, functionality of the library, several tests are included with the software package. Besides, to measure the encoding/decoding speed and bandwidth consumed in case of a data repair, we included several performance tests as well. Our performance tests are run on a server system the details of which are given in Table 1. To be able to draw a summary of the library performance, we provide Table 2 for quantification of Encoding/Decoding speed. In our test, we used a 64MiB file, and encoded data are spread across 10 disks equally. We have used multi-threading support and set -m parameter set to 12. While decoding, we have removed 2, 3, and 4 disk worth of information before running the decoder. The -t parameter is judiciously chosen in powers of two to enable hardware-friendly operation. No exhaustive optimization is carried out. As can be

² This directory naming is conventional and maintained for the legacy of Jerasure erasure code library.

Table 1
Server System CPU features.

Property	Name: Intel Xeon CPU E5-2620	
	Value	Explanation
Socket	2	
Core	12	6 in each socket
Clock speed	2.4	GHz (Max.)
Threads	24	2 in each core
Arc.	IA64	X86_64
L1 cache	32KiB	
L2 cache	256KiB	
L3 cache	15360KiB	
Main memory	≈ 8GiB	

Table 2
Encoder/Decoder Performance in MB/sec.

k	n	Rate	t	Encoder performance (MB/sec)	Failure #	Decoder performance (MB/sec)
1036	2180	0.4752	1024	2722	2	3291
					3	2852
					4	2340
10246	21800	0.47	512	1691	2	2668
					3	2496
					4	2356

seen, with almost a half code rate, we could achieve super-fast encoding and decoding speeds with the current implementation.

While the execution performance of the library is quite attractive, we can also show that it is also bandwidth friendly when the data is repaired. We considered the case $k = 10246$ and $n = 21800$ with a $100\text{MiB} = 104,857,600$ bytes file while all the rest of the parameters are the same as before. If we consider double disk failures, the conventional repair method requires us to transfer $108,267,520$ bytes of data for the repair to be successful. This is a little over 100MiB as expected due to overhead sub-optimality of the code used in Founsure. On the other hand, if we use the improved repair scheme suggested in this study, we can achieve a maximum of $65,952,320$ bytes of transfer for successful recovery, which is almost $2\times$ more efficient use of bandwidth over that of the conventional method. Note again that no optimization is performed in terms of degree distribution $\Omega(x)$ and advanced graph partitioning to minimize repair bandwidth.

Finally in Table 3, we have run a simple test to compare the performance of Founsure against two of the most efficient and heavily used erasure coding libraries, namely Jerasure 2.0³ and ISA-L of Intel,⁴ based on RS codes (using Cauchy and Vandermonde Matrix constructions). This time, we used a $1\text{GiB} = 1,073,741,824$ bytes. We realize that the parameters k and n cannot be selected too large for these libraries, due to their algebraic constructions which makes it extremely complex to deal with such low-level sub-packetizations. We set $k = 10$ and $n = 20$ to simulate half code-rate RS codes which are closest to Founsure's 0.47 code rate previously selected. Note that with this selection, the number of disks used to store the generated content can be 20 at most (Parameter n also characterizes the number of disks). Having more disks to distribute data would make the presented encoding/decoding performances worse. On the other hand, the number of disks does not change the performance of Founsure. Parameters of Jerasure and ISA-L libraries are selected to give the best performance on the same system defined in Table 1. The decoding of these libraries is set to decode 8 blocks (4 disks

Table 3
Performance comparisons between different erasure code libraries.

Library name	Worst/Best Time complexity	Encoder performance (MB/sec)	Decoder performance (MB/sec)	Repair BW (MB)
ISA-L	$O(n^2)/O(n \log^2(n))$ [22]	1420	1390	1073,7
Jerasure 2.0	$O(n^2)/O(n \log^2(n))$ [22]	780	870	1073,7
Founsure 1.0	$O(n \log(n))/O(n)$ [16]	1620	2245	675,35

worth information). We also know that there is no simple repair mechanism defined for these RS-based constructions in publicly available libraries and hence adopted the conventional decoding-based repair while we compare the bandwidth consumed for repairing a single failed disk content in Table 3. Finally, we also provided approximate best and worst-case time complexity of encoding/decoding for RS and fountain codes as available from other research works in literature [22].

To the best of my knowledge, Founsure is the most flexible erasure coding library that is open source and can be configured based on the requirements of the application. With the current software architecture, many more functionalities can be integrated such as partial user data construction and advanced error detection for failure localization. Also, as numerical results suggest even if many more optimizations are possible to make the performance better, the current version's performance in terms of execution speed and repair bandwidth still stand out. Founsure has highly parallel architecture and lends itself to parallel programming. Unlike Founsure, existing research is mostly focused on overhead optimal designs using the inherent nature of the parallel hardware. Originally, Founsure is developed for data storage systems, it can simply be adapted to packet-switched networks in which the underlying channel is erasure channel or sporadic erasure channels [23]. With advanced features such as error detection, erased content reconstruction, multi-threaded support, advanced decoding, Founsure can further be used for error correction which could open up more areas of applications such as image reconstruction and data protection over noisy communication channels. Consequently, we believe that Founsure could be a strong candidate to be used for any system that secures data protection and recovery in one way or another.

6. Conclusions

In this work, we have developed and presented an erasure coding library that can be used to operate on various points of the trade-off between computational complexity, coding overhead, and repair bandwidth. For example, through the right selection of coding parameters, the Founsure library can be used to save storage space and minimize the data storage overhead. On the other hand, by allowing some overhead again through tweaking parameters, Founsure can reduce the data repair bandwidth. Unlike previous software packages, such freedom of parameter selections makes Founsure library more application-centric and configurable for future generation reliable system design.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

³ <https://github.com/ceph/jerasure>.

⁴ <https://github.com/intel/isa-l>.

Acknowledgments

This research is funded by The Scientific and Technological Research Council of Turkey (TUBITAK) under grant numbers 115C111, 119E235 and supported by Quantum Corporation in reaching out the necessary hardware platforms on time for extensive testing.

References

- [1] Nisbet B. FAS storage systems: Laying the foundation for application availability, Network Appliance white paper. Available online: <http://www.netapp.com/us/library/analyst-reports/ar1056.html>, February. 2008.
- [2] Reed IS, Solomon G. Polynomial codes over certain finite fields. *J Soc Ind Appl Math* 1960;8:300–4.
- [3] Blomer J, Kalfane M, Karpinski M, Karp R, Luby M, Zuckerman D. An XOR-Based Erasure-Resilient Coding Scheme. Technical Report TR-95-048, International Computer Science Institute; 1995, August.
- [4] Blaum M, Brady J, Bruck J, Menon J. EVENODD: An efficient scheme for tolerating double disk failures in RAID architectures. *IEEE Trans Comput* 1995;44(2):192–202, February.
- [5] Asteris M, Papailiopoulos D, Dimakis AG, Vadali R, Chen S, Borthakur D. XORing elephants: Novel erasure codes for big data. *Proc VLDB Endow (PVLDB)* 2013;6(5).
- [6] Partow A. Schifra reed-solomon ECC library. 2000-2007, Open source code distribution: <http://www.schifra.com/downloads.html>.
- [7] Plank JS, Simmerman S, Schuman CD. Jerasure: A Library in C/C++ Facilitating Erasure Coding for Storage Applications - Version 1.2. Tech. Rep. CS-08-627, University of Tennessee; 2008, August.
- [8] CEPH software-defined storage. Available online: <http://www.ceph.com>.
- [9] Wilcox-O'Hearn Z. Zfec 1.4.0. 2008, Open source code distribution: Available online: <http://pypi.python.org/pypi/zfec>.
- [10] Arslan SS, Goker T, Wideman R. A joint dedupe-fountain coded archival storage. In: *IEEE ICC'17*. Paris, France. 2017. p. 1–7.
- [11] MacKay DJC. Fountain codes. *IEE Proc Commun* 2005;152:1062–8, Dec.
- [12] Luby M. LT-codes. In: *Proc. 43rd annu. IEEE symp. foundations of computer science*. Vancouver, BC, Canada. Nov. 2002. p. 271–80.
- [13] Shokrollahi A. Raptor codes. *IEEE Trans Inform Theory* 2006;52(6):2551–67.
- [14] Lee D, Park H, Moon J. Reducing repair-bandwidth using codes based on factor graphs. In: *IEEE international conference on communications*. Kuala Lumpur, Malaysia. 2016. p. 1–6.
- [15] Pearl J. Reverend Bayes on inference engines: A distributed hierarchical approach. In: *Second national conference on artificial intelligence*. Menlo Park, California: AAAI Press; 1982, p. 133–6.
- [16] Arslan SS. Incremental redundancy, fountain codes and advanced topics. 2014, Available online: [arXiv:1402.6016](https://arxiv.org/abs/1402.6016).
- [17] Felstrom AJ, Zigangirov KS. Time-varying periodic convolutional codes with low-density parity-check matrix. *IEEE Trans Inform Theory* 1999;45(6):2181–90, Sep.
- [18] Eleftheriou E, Olcer S. Low-density parity-check codes for digital subscriber lines. In: *Proc. IEEE int. conf. communications*, vol. 3. 2002. p. 1752–57.
- [19] Fan JL. Array codes as low-density parity-check codes. In: *Proc. 2nd international symposium on turbo codes and related topics*. Brest, France. Sept. 2000. p. 543–46.
- [20] Arslan SS. Implementation of multi-threaded erasure coding under multi-processing environments. In: *24th Signal processing and communication application conference*. Zonguldak. 2016. p. 1773–76.
- [21] Arslan SS, Le H, Landman J, Goker T. OpenMP and POSIX thread implementation of Jerasure 2.0. In: *IEEE international black sea conference on communications and networking*. Istanbul, Turkey. 2017. p. 1–5.
- [22] Chen N, Yan Z. Complexity analysis of Reed–Solomon decoding over $G F(2^m)$ without using syndromes. *EURASIP J Wireless Commun Networking* 2008;2008(16):1–11.
- [23] Liva G, Paolini E, Matuz B, Chiani M. A decoding algorithm for LDPC codes over erasure channels with sporadic errors. In: *48th Annual Allerton conference on communication, control, and computing*. Allerton, IL. 2010. p. 458–65.