Aalborg Universitet

AALBORG UNIVERSITY

DENMARK

On one-round reliable message transmission

Christensen, René Bødker

Link to publication from Aalborg University

Accepted Manuscript

On one-round reliable message transmission

René Bødker Christensen

Highlights

- Attaining constant transmission rate for constant-size messages is impossible.
- Application of the Guruswami-Sudan algorithm allows increased message sizes.
- An RMT-protocol for smaller field sizes is proposed.

On one-round reliable message transmission

René Bødker Christensen

Department of Mathematical Sciences, Aalborg University, Skjernvej 4A, 9220 Aalborg Øst, Denmark

Abstract

In this paper, we consider one-round protocols for reliable message transmission (RMT) when t out of $n = 2t + 1$ available channels are controlled by an adversary. We show impossibility of constructing such a protocol that achieves a transmission rate of less than $\Theta(n)$ for constant-size messages and arbitrary reliability parameter. In addition, we show how to improve two existing protocols for RMT to allow for either larger messages or reduced field sizes.

Keywords: reliable message transmission, cryptography

1. Introduction

The concept of secure message transmission was first introduced in [1], and the term comprises a model where a sender and a receiver are connected via n channels. Up to t of these channels are controlled by a computationally unbounded active adversary who can read and alter the symbols sent across these t channels. More specifically, we consider the setting where $n = 2t + 1$. In keeping with cryptographic tradition, we will call the sender 'Alice', the receiver 'Bob', and the adversary 'Eve'. The challenge is to devise a strategy that allows Alice and Bob to communicate securely and reliably in a limited number of transmission rounds. We focus on one-round protocols.

In the original setting of [1], the protocols are required to be perfectly secure, meaning that no matter what Eve might attempt, she will gain no information about the message. They are also required to be perfectly reliable such that Bob will always recover the correct message. Later, [2] relaxed these conditions to allow some small failure probabilities for both security and reliability. Taking this idea even further, [3] considers protocols where the security of the message delivery is *not* required, but only reliable transmission

---

is of interest. They call this *unconditionally reliable message transmission*, but we will omit 'unconditionally' and write RMT instead.

To asses the efficiency of a message transmission-protocol, it is common to use the *transmission rate* defined as the total number of transmitted bits divided by the bit-length of the message. Hence, a low transmission rate is preferable. As shown in [3, Theorem 3], we cannot do better than $\Omega(1)$ for RMT, and this bound is tight. In Section 3, however, we show that this transmission rate is not achievable for messages of a constant size.

1.1. Related work

RMT has also been studied in [3, 4]. The protocol in [4] is based on list-decoding of folded Reed-Solomon codes, but although it attains the optimal transmission rate, the computational cost for the receiver to recover the message is exponential in the number of channels. The work [3] contains bounds and constructions for both the secure and the reliable-only settings. In addition, they achieve this while tolerating a mixed adversary, giving more fine-grained control of the adversarial assumptions.

Although this paper is only concerned with RMT, we also direct the reader to related works on secure message transmission; that is, protocols that also offer privacy. This additional guarantee comes at a cost. As shown by [1], perfect security for $n = 2t+1$ requires at least two rounds, and a single-round protocol can only offer security in the case $n \geq 3t + 1$. In the former setting, Agarwal et al. [5] gave a perfectly secure two-round protocol that achieves optimal performance asymptotically, albeit at a high computational cost. A computationally efficient protocol was subsequently achieved by Kurosawa and Suzuki [6] using the concept of pseudobases. This idea was also taken up by [7], who obtained further improvements, reducing the minimally required message size from $\mathcal{O}(n^2 \log n)$ to $\mathcal{O}(n \log n)$.

The setting where privacy is perfect, but reliability is not, was initially handled by [2] under the assumption that channels support multicast. The proposed solution, however, was inefficient for certain values of t and n. This was rectified in [8], where an efficient protocol for these values was given.

2. Preliminaries

2.1. Model assumptions

We assume that Alice and Bob are connected via $n = 2t + 1$ *simple* channels, meaning that the channels allow both Alice and Bob to transmit data, but no additional functionality is assumed. Before the protocol begins, Eve chooses t of these to be under her control. In other words, the adversary in our model is *static* and *active*.

For simple channels, [2] showed that $2t \geq n$ leads to a probability of failure of at least $1/4$. Hence, the setting where $n = 2t + 1$ has the maximal number of corruptions that we can hope to overcome. Since a majority of the channels are honest – i.e. not controlled by the adversary – a naive solution to the RMT-problem is to broadcast the message across all n channels. This leads to a transmission rate of n, but gives perfect reliability. Thus, this is the benchmark performance.

2.2. Universal hash families

The methods we present rely on the concept of ε-almost universal hash families as introduced by [9].

Definition 2.1. Let $\mathcal{H}$ be a family of hash functions from $\mathcal{M}$ to A, and let $\varepsilon \in \mathbb{R}_+$. Then $\mathcal{H}$ is called ε-almost universal if for any $m \neq m' \in \mathcal{M}$,

$$\Pr_{h \leftarrow \mathcal{H}}[h(m) = h(m')] \leq \varepsilon.$$

In particular, we use a hash family based on polynomial evaluation similar to the one used in [10], but generalized to evaluate in several points.

Definition 2.2. Let $\mathbb{F}$ be a finite field, and $\mathcal{K} \subseteq \mathbb{F}$. For every pair of positive integers $\eta \leq a$, define the map $\mathsf{PEval}^\eta \colon \mathbb{F}^a \times \mathcal{K}^\eta \to \mathbb{F}^\eta$ by

$$\mathsf{PEval}^\eta(\mathbf{m}, \mathbf{k}) = (f_{\mathbf{m}}(k_1), f_{\mathbf{m}}(k_2), \ldots, f_{\mathbf{m}}(k_\eta)),$$

where $f_{\mathbf{m}}(x) = \sum_{i=1}^a m_i x^i$. We use the notation $\mathsf{PEval}_{\mathbf{k}}^\eta(\mathbf{m}) = \mathsf{PEval}^\eta(\mathbf{m}, \mathbf{k})$.

It may be shown that the family $\mathcal{H}_{\mathsf{PEval}}^\eta = \{\mathsf{PEval}_{\mathbf{k}}^\eta \colon \mathbb{F}^a \to \mathbb{F}^\eta\}_{\mathbf{k} \in \mathcal{K}^\eta}$ of hashes is $(a/|\mathcal{K}|)^\eta$-almost universal.

3. Constant-size messages

One could hope that the overall optimal transmission rate $\Theta(1)$ is achievable for constant-size messages. As we show in Proposition 3.2, however, this is not possible for arbitrary reliability parameters. The proof of the proposition relies on the following result from [2, Theorem 5.1].

Theorem 3.1. *Assume that $n \leq 2t$, and denote by $\mathcal{M}$ the message space. Then any reliable message transmission protocol fails with probability at least $\frac{1}{2}(1 - 1/|\mathcal{M}|)$.*

Proposition 3.2. *Let $n = 2t + 1$, and consider the RMT-problem for a message of size $\Theta(1)$ bits. Then it is impossible to construct a protocol attaining a transmission rate lower than $\Theta(n)$ for arbitrary reliability parameters.*

3

PROOF. Assume for contradiction that $\mathcal{P}$ is such a protocol. We show the existence of an adversarial strategy such that $\mathcal{P}$ will fail with a probability greater than a constant.

Note that if all n available channels are used, at least n bits will be transmitted during the protocol. Hence, $\mathcal{P}$ can use at most $n-1$ channels. Let $X \in \{1, 2, \ldots, n\}$ be a random variable describing the unused channel. No assumptions are made about the probability distribution of X; it simply depends on $\mathcal{P}$. Consider an adversarial strategy where the corrupt channels are chosen uniformly at random. Equivalently, we can assume that the honest channels are given by the set $\{I_1, I_2, \ldots, I_{t+1}\}$, where each I_j is chosen uniformly at random in $\{1, 2, \ldots, n\}$ under the condition that $I_j \neq I_{j'}$ for $j \neq j'$. It may be shown that in fact $\Pr[I_j = a] = 1/(2t+1)$ for any $j \in \{1, 2, \ldots, t+1\}$ and $a \in \{1, 2, \ldots, n\}$.

Denote by E the event that Alice leaves out one of the honest channels when following $\mathcal{P}$; that is, $X = I_j$ for some $j \in \{1, 2, \ldots, t+1\}$. Since Alice does not know the outcomes of $I_1, I_2, \ldots, I_{t+1}$, it follows that X is independent from these variables. Using this fact and the fact that the events $X = I_1, X = I_2, \ldots, X = I_{t+1}$ are disjoint, we obtain that

$$\Pr[E] = \Pr[X = I_1 \vee \cdots \vee X = I_{t+1}] = \sum_{j=1}^{t+1} \Pr[X = I_j]$$

$$= \sum_{j=1}^{t+1} \sum_{k=1}^{n} \Pr[X = k] \Pr[I_j = k] = \sum_{j=1}^{t+1} \frac{1}{2t+1} \sum_{k=1}^{n} \Pr[X = k] = \frac{t+1}{2t+1}.$$

If E occurs, it follows from Theorem 3.1 that the probability of protocol failure is at least $\frac{1}{2}(1 - 1/|\mathcal{M}|)$, where $\mathcal{M}$ is the message space. Otherwise, the protocol $\mathcal{P}$ gives a contradiction to Theorem 3.1 since for $n = 2t$, we could introduce a 'dummy channel', discard it, and then mimic protocol $\mathcal{P}$ to obtain a lower probability of failure.

By applying the law of total probability, we obtain

$$\Pr[\mathcal{P} \text{ fails}] = \Pr[\mathcal{P} \text{ fails} \mid E] \Pr[E] + \Pr[\mathcal{P} \text{ fails} \mid \bar{E}] \Pr[\bar{E}]$$
$$\geq \Pr[\mathcal{P} \text{ fails} \mid E] \Pr[E]$$
$$\geq \frac{1}{2} \left(1 - \frac{1}{|\mathcal{M}|}\right) \frac{t+1}{2t+1} > \frac{1}{4} \left(1 - \frac{1}{|\mathcal{M}|}\right).$$

In conclusion, it is not possible to obtain arbitrarily levels of reliability with a transmission rate of less than $\Theta(n)$ for constant size messages. □

It is worth pointing out that this result is true for any RMT-protocol; not only one-round ones.

4

4. A method based on list-decoding

As part of a protocol for robust secret sharing, [10] introduced the notion of a 'robust distributed storage'. Their method for achieving this can easily be converted to a one-round protocol for RMT. In brief, the idea is to encode the message using a list-decodable code – e.g. a Reed-Solomon code – and transmit each position of the resulting codeword across the corresponding channel. In addition, each channel will deliver a key/tag-pair from an ε-almost universal hash family. The receiver can then use these tags to recover the intended message from the list of potential messages returned by the list-decoding algorithm.

However, since the original authors only need the asymptotical performance, they base their method on the list-decoding algorithm of Sudan [11], and use messages of size at most $\lfloor n/8 \rfloor + 1$. This may be increased to $\lfloor n/5 \rfloor + 1$ with no penalty in reliability by applying the Guruswami-Sudan algorithm [12] instead. This protocol has optimal transmission rate when the message has size $\Theta(n)$.

5. A method based on erasure decoding

In the following, we will describe the one-round RMT-protocol given in [3] in the language of Reed-Solomon codes and hash families. In this representation, the original authors are essentially relying only on the erasure correcting capabilities of the codes. We show that a careful choice of parameters allows correction of errors as well, causing the required field size to be quadratic rather than cubic in n.

The message we consider is an $a \times b$-matrix M over a finite field $\mathbb{F}$. Each row of this message is encoded by means of an $[n, b]$ Reed-Solomon code, yielding an $a \times n$-matrix S where each row is a codeword. Across the i'th channel, Alice sends the i'th column s_i of S. Since Bob needs to determine if Eve modified some of these columns during transmission, Alice also computes n verification tags $\{v_{i1}, v_{i2}, \ldots, v_{in}\}$ for each s_i by applying uniformly sampled hash functions from some family $\mathcal{H}$. Denote the keys of these functions by $\{k_{i1}, k_{i2}, \ldots, k_{in}\}$. Across the i'th channel, Alice then sends $\{s_i\} \cup \{k_{ji}, v_{ji}\}_{j=1}^n$. That is, each channel will transmit the codeword entries s_i, and a key/tag-pair (k_{ji}, v_{ji}) for every channel j.

When Bob receives the possibly modified values $\{s_i'\} \cup \{k_{ji}', v_{ji}'\}_{j=1}^n$, he will check the integrity of s_i' by computing the hash value $h_{k_{ij}'}(s_i')$ and comparing the result with the received tag v_{ij}'. He will do so for each received key/tag-pair, and if more than t tags disagree with the computed values, Bob will mark s_i' as modified and treat it as an erasure when recovering the message.

5

1. The message is represented as a matrix $M \in \mathbb{F}^{a \times b}$ and each row is encoded using an $[n, b]$ Reed-Solomon code over $\mathbb{F}$.

2. For each column s_i of the resulting codewords, Alice samples uniformly and independently n keys $\{k_{i1}, k_{i2}, \ldots, k_{in}\}$ and computes $v_{ij} = h_{k_{ij}}(s_i)$ for each $j \in \{1, 2, \ldots, n\}$.

3. Across the i'th channel, Alice transmits $\{s_i\} \cup \{k_{ji}, v_{ji}\}_{j=1,2,\ldots,n}$.

4. Bob receives the possibly modified values $\{s_i'\} \cup \{k_{ji}', v_{ji}'\}_{j=1,2,\ldots,n}$ for $i = 1, 2, \ldots, n$. For each i, he compares the tag v_{ij}' received from the j'th channel to the hash value $h_{k_{ij}'}(s_i')$. If these disagree for more than t channels, he will mark s_i as modified.

5. For each row in S', Bob computes the syndrome to check if it contains errors. Depending on the result, he proceeds with one of the three following steps.

    (a) **The syndrome is zero:** S' contains no errors, meaning that Bob can simply use polynomial interpolation to recover the message.

    (b) **The syndrome is nonzero, and S' contains at least $t - e$ erased columns:** Bob uses a decoding algorithm for Reed-Solomon codes to correct the erasures and errors, hereby recovering the message.

    (c) **The syndrome is nonzero, and S' contains less than $t - e$ erased columns:** Too many modified channels have passed the integrity checks. The protocol has failed.

Figure 1: This protocol allows Alice to reliably send ab symbols of a finite field $\mathbb{F}$ to Bob in one round by using $n = 2t + 1$ channels, t of which may be controlled by an adversary. Beforehand, Alice and Bob have agreed upon a parameter $e \in \mathbb{N}$, which satisfies $e \leq t + 1 - b$. Additionally, they agree on an ε-almost universal hash family $\mathcal{H} = \{h_k \colon \mathbb{F}^a \to \mathbb{F}^\eta \mid \mathbb{F}^\eta\}$.

With large probability, these checks performed by Bob reveal a considerable part of the corrupt channels delivering erroneous information. This causes a number of columns in S' to be marked as erasures. However, some small number e of corrupted channels may have passed the checks, meaning that the remaining entries in S' may still contain errors. In fact, each row of S' may contain up to $t - e$ erasures and e errors. If the parameter b agreed upon by Alice and Bob is sufficiently small, Bob may nevertheless correct these erasures and errors in S'. Since the rows of S' are codewords of an $[n, b]$ Reed-Solomon code which has minimal distance $n - b + 1$, Bob can recover the correct message if

$$2e + t - e < n - b + 1 \quad \Longrightarrow \quad b \leq n - (t + e) = t + 1 - e.$$

Thus, after verifying the received values, Bob can determine if the message can be recovered by simply counting the number of non-erased columns and

computing syndromes. The complete description of our protocol is given in Figure 1. The correctness of the protocol follows from essentially the same arguments as used by [3], albeit with the following modification.

Lemma 5.1. *If at least $t - e$ columns of S' are marked as erasures in step 4 of the protocol, Bob will recover the correct message.*

PROOF. Let $u \geq t - e$ be the number of erased columns, meaning that each row of S' contains at most $t - u$ errors. The minimal distance of the code is $d = n - b + 1$, which means that u erasures and $t - u$ errors can be corrected if $2(t - u) + u < d$. This is true because

$$2(t - u) + u = 2t - u \leq t + e \leq n - b,$$

where the last inequality follows from the requirement $e \leq t + 1 - b$ given in the protocol specification. □

5.1. Protocol reliability

Under the assumption that the hash family $\mathcal{H}$ applied in the protocol is ε-almost universal, we can bound the probability that Bob cannot recover the correct message.

Proposition 5.2. *If $\mathcal{H}$ is an ε-almost universal family of hash functions, then*

$$\Pr[\text{The protocol fails}] \leq \frac{t(t + 1)\varepsilon}{e + 1}.$$

PROOF. By Lemma 5.1, at least $e + 1$ of the channels modified by Eve must pass the integrity check performed by Bob. To achieve this, it is necessary that the hash value of the modified s'_i matches at least one verification tag v_{ij} sent across an honest channel.

The ε-almost universality of $\mathcal{H}$ implies that $\Pr_{h \leftarrow \mathcal{H}}[h(s_i) = h(s'_i)] \leq \varepsilon$ whenever $s_i \neq s'_i$. Hence, ε is an upper bound on the probability that a single corrupt channel agrees with a single honest channel. Since there are $t + 1$ honest channels, the probability for a modified channel to be consistent with at least one honest can be bounded above by $(t + 1)\varepsilon$.

Let X be the random variable counting the number of modified but un-caught channels. Since the hash keys k_{ij}, $k_{i'j'}$ are independent whenever $(i, j) \neq (i', j')$, the integrity checks of the modified channels can be considered as t independent Bernoulli trials, each with a success probability of at most $(t + 1)\varepsilon$. Thus, X follows a binomial distribution, and has expected value $\mathbb{E}[X] \leq t(t + 1)\varepsilon$. The Markov inequality now gives

$$\Pr[X \geq e + 1] \leq \frac{\mathbb{E}[X]}{e + 1} \leq \frac{t(t + 1)\varepsilon}{e + 1},$$

and the result follows. □

7

5.2. Number of bits transmitted

When the proposed protocol is used to transmit a message, the total number of $\mathbb{F}$-symbols transmitted is $n(a + n|\mathcal{V}| + n|\mathcal{K}|)$, where $|\mathcal{V}|$ and $|\mathcal{K}|$ denote the number of field symbols necessary to represent v_{ij} and k_{ij}, respectively.

5.3. Using polynomial evaluation

For concreteness, we analyse the reliability when $\mathcal{H}_{\mathsf{PEval}}^{\eta}$ is applied with $\mathcal{K} = \mathbb{F}$. Here, both the keys and the verification tags consist of η field elements. Hence, the total number of transmitted bits is $2\eta n^2 + an$. Depending on the message size, this can give various transmission rates, but under the assumption that η is some constant value, the optimal transmission rate of $\Theta(1)$ is obtained when both a and b are $\Theta(n)$. That is, when the message is of size $\Theta(n^2)$.

Since the hash family is $\frac{a^{\eta}}{|\mathbb{F}|^{\eta}}$-almost universal, it follows from Proposition 5.2 that we must require

$$\frac{t(t+1)a^{\eta}}{(e+1)|\mathbb{F}|^{\eta}} \leq \delta \quad \Longrightarrow \quad |\mathbb{F}| \geq a\left(\frac{t(t+1)}{(e+1)\delta}\right)^{\frac{1}{\eta}}.$$

in order to obtain reliability δ. In particular, we note that for $\eta = 1$, the original protocol by [3] requires $|\mathbb{F}| \geq n^3/\delta$. In the proposed protocol, we can set both b and e to be $\Theta(n)$ and obtain the requirement $|\mathbb{F}| \geq \Theta(n^2/\delta)$. In other words, by reducing the second dimension of the message, the required field size is reduced by a factor of n asymptotically. Furthermore, introducing the parameter η highlights the trade-off between the number of $\mathbb{F}$-symbols transmitted and the required field size.

6. Comparison with existing protocols

In order to compare the RMT-protocols proposed in Sections 4 and 5 to those already in the literature, we will restrict ourselves to the hash family $\mathcal{H}_{\mathsf{PEval}}^{\eta}$ from Definition 2.2 with $\mathcal{K} = \mathbb{F}$ and $\eta = 1$.

For five protocols, Table 1 gives an overview of the required field size given δ; the message size in $\mathbb{F}$-symbols; whether the protocol attains the optimal transmission rate; and whether it is computationally efficient. Here, *efficient* means polynomial in the number of available channels. We use the Θ-notation to keep the presentation as clear and self-contained as possible.

For the protocol of Section 5, we remark that $a = \Theta(n)$ was chosen even though it is in principle possible to use any value smaller than $|\mathbb{F}|$. Choosing greater values, however, also increases the required field size. We

Protocol	Field size	Message size	Optimal	Computational efficiency
[10, Sec. 4.1]	$\Theta(n^2/\delta)$	$\lfloor n/8 \rfloor + 1$	✓	✓
This work, Sec. 4	$\Theta(n^2/\delta)$	$\lfloor n/5 \rfloor + 1$	✓	✓
[3, Sec. 4]	n^3/δ	$\Theta(n^2)$	✓	✓
[4, Sec. 3.1]	$\Theta(n^4)$	$\Theta(n^2)$	✓	✗
This work, Sec. 5	$\Theta(n^2/\delta)$	$\Theta(n^2)$	✓	✓

Table 1: Comparison of one-round RMT-protocols. The second column shows the minimal field size given a desired reliability parameter δ. The third column gives the message size (in terms of $\mathbb{F}$-elements) that leads to an optimal transmission rate, and the fourth indicates whether such an optimal transmission rate is achievable. The final column states whether the computational cost is at most polynomial in the number of channels.

shall refrain from doing such analysis here since Table 1 already shows the desired improvement.

As the table indicates, the first two protocols are better suited for small message sizes. Although both have the same asymptotic performance, the modification suggested in Section 4 allows a larger message size. The remaining three protocols all have $\Theta(n^2)$ as the optimal message size, which suggests that they should fare better when transmitting larger messages. It may be noted that the protocol proposed in Section 5 achieves this while reducing the required field size by a factor of n asymptotically.

Even though Table 1 gives an overview of the general properties of each protocol, it does not reveal how they will perform in concrete problem instances. If the message size and the number of channels have already been fixed, a separate analysis is needed to determine the protocol that will perform the best.

7. Acknowledgements

[1] D. Dolev, C. Dwork, O. Waarts, M. Yung, Perfectly secure message transmission, J. ACM 40 (1) (1993) 17–47. doi:10.1145/138027.138036.

[2] M. Franklin, R. N. Wright, Secure communication in minimal connectivity models, J. Cryptol. 13 (1) (2000) 9–30. doi:10.1007/s001459910002.

9

[3] A. Patra, A. Choudhury, C. P. Rangan, K. Srinathan, Uncondition-
    ally reliable and secure message transmission in undirected synchronous
    networks: Possibility, feasibility and optimality, Int. J. Appl. Cryptogr.
    2 (2) (2010) 159–197. `doi:10.1504/IJACT.2010.038309`.

[4] R. Safavi-Naini, M. A. A. Tuhin, P. Wang, A general construc-
    tion for 1-round δ-RMT and $(0, \delta)$-SMT, in: F. Bao, P. Samarati,
    J. Zhou (Eds.), ACNS 2012, Springer, Heidelberg, 2012, pp. 344–362.
    `doi:10.1007/978-3-642-31284-7_21`.

[5] S. Agarwal, R. Cramer, R. de Haan, Asymptotically optimal
    two-round perfectly secure message transmission, in: C. Dwork
    (Ed.), CRYPTO 2006, Springer, Heidelberg, 2006, pp. 394–408.
    `doi:10.1007/11818175_24`.

[6] K. Kurosawa, K. Suzuki, Truly efficient 2-round perfectly secure message
    transmission scheme, IEEE Trans. Inf. Theory 55 (11) (2009) 5223–5232.
    `doi:10.1109/TIT.2009.2030434`.

[7] G. Spini, G. Zémor, Perfectly secure message transmis-
    sion in two rounds, in: TCC 2016-B, 2016, pp. 286–304.
    `doi:10.1007/978-3-662-53641-4_12`.

[8] Y. Wang, Y. Desmedt, Secure communication in multicast channels:
    The answer to Franklin and Wright's question, J. Cryptol. 14 (2) (2001)
    121–135. `doi:10.1007/s00145-001-0002-y`.

[9] D. Stinson, Universal hashing and authentication codes, Des. Codes
    Cryptogr. 4 (3) (1994) 369–380. `doi:10.1007/BF01388651`.

[10] A. Bishop, V. Pastro, R. Rajaraman, D. Wichs, Essentially optimal ro-
     bust secret sharing with maximal corruptions, in: EUROCRYPT 2016,
     2016, pp. 58–86. `doi:10.1007/978-3-662-49890-3_3`.

[11] M. Sudan, Decoding of reed solomon codes beyond the error-
     correction bound, J. Complexity 13 (1) (1997) 180 – 193.
     `doi:10.1006/jcom.1997.0439`.

[12] V. Guruswami, M. Sudan, Improved decoding of Reed-Solomon
     and algebraic-geometric codes, in: FOCS 1998, 1998, pp. 28–37.
     `doi:10.1109/SFCS.1998.743426`.