



Chinese Remaindering with Errors

Oded Goldreich* Dana Ron† Madhu Sudan‡

November 22, 1998

Abstract

The Chinese Remainder Theorem states that a positive integer m is uniquely specified by its remainder modulo k relatively prime integers p_1, \dots, p_k , provided $m < \prod_{i=1}^k p_i$. Thus the residues of m modulo relatively prime integers $p_1 < p_2 < \dots < p_n$ form a redundant representation of m if $m < \prod_{i=1}^k p_i$ and $k < n$. This suggests a number-theoretic construction of an “error-correcting code” that has been implicitly considered often in the past. In this paper we provide a new algorithmic tool to go with this error-correcting code: namely, a polynomial-time algorithm for error-correction. Specifically, given n residues r_1, \dots, r_n and an agreement parameter t , we find a list of all integers $m < \prod_{i=1}^k p_i$ such that $(m \bmod p_i) = r_i$ for at least t values of $i \in \{1, \dots, n\}$, provided $t = \Omega(\sqrt{kn \frac{\log p_n}{\log p_1}})$. We also give a simpler algorithm to decode from a smaller number of errors, i.e., when $t > n - (n - k) \frac{\log p_1}{\log p_1 + \log p_n}$. In such a case there is a unique integer which has such agreement with the sequence of residues.

One consequence of our result is a strengthening of the relationship between average-case complexity of computing the permanent and its worst-case complexity. Specifically we show that if a polynomial time algorithm is able to guess the permanent of a random $n \times n$ matrix on $2n$ -bit integers modulo a random n -bit prime with inverse polynomial success rate, then $P^{\#P} = BPP$. Previous results of this nature typically worked over a fixed prime moduli or assumed success probability very close to one (as opposed to bounded away from zero).

*Department of Computer Science, Weizmann Institute of Science, Rehovot, ISRAEL. Email: oded@wisdom.weizmann.ac.il. Work done in part while visiting MIT, and partially supported by DARPA grant DABT63-96-C-0018.

†Department of Electrical Engineering – Systems, Tel Aviv University, Ramat Aviv, ISRAEL. Email: danar@eng.tau.ac.il. Work done in part while visiting MIT, supported by an ONR Science Scholar Fellowship of the Bunting Institute.

‡Department of Electrical Engineering and Computer Science, MIT, 545 Technology Square, Cambridge, MA 02139, USA. Email: madhu@mit.edu. Research supported in part by a Sloan Foundation Fellowship and a MIT-NEC Research Initiation Grant.

1 Introduction

The Chinese Remainder Theorem states that a positive integer m is uniquely specified by its remainder modulo k relatively prime integers p_1, \dots, p_k , provided $m < \prod_{i=1}^k p_i$. Thus if we pick $n > k$ relatively prime integers $p_1 < \dots < p_n$ such that $m < \prod_{i=1}^k p_i$, then the remainders of m modulo the p_i 's form a redundant encoding of m . Specifically, m can be recovered given any k of the n remainders. Thus this representation of integers yields a natural error-correcting code: given any two integers $m, m' < \prod_{i=1}^k p_i$, the sequences $\{(m \bmod p_1), \dots, (m \bmod p_n)\}$ and $\{(m' \bmod p_1), \dots, (m' \bmod p_n)\}$ differ in at least $n - k + 1$ coordinates.

This redundancy property of the Chinese remainder representation has been exploited often in theoretical computer science. The Karp-Rabin pattern matching algorithm is based on this redundancy [17]. This representation was used to show the strength of probabilistic communication over deterministic communication protocols (cf. [19, Exercise 3.6]). The representation allows for easy arithmetic — addition, multiplication, subtraction and division — on large integers and was even proposed as a potential representation for numbers in computers¹. The ability to reduce computation over large integers to that over small integers is also employed in complexity-theoretic settings, with a notable example being its use in showing the hardness of computing the permanent of 0/1 matrices [33]. In fact it is this application that motivates the main question studied here.

The redundancy of the Chinese remainder representation of integers and its similarity to error-correcting codes raises a natural algorithmic question:

Given a sequence of integers $\langle r_1, \dots, r_n \rangle$ that are obtained from taking residues of an integer $m < \prod_{i=1}^k p_i$ modulo relatively prime integers $p_1 < \dots < p_n$, where some of the residues are erroneous, can we find m ?

If the number of residues that are erroneous is less than $\frac{n-k}{2}$, then m is uniquely specified by the vector $\langle r_1, \dots, r_n \rangle$. However this fact is not algorithmic — it is not clear how to recover m in polynomial time (i.e., in time polynomial in n and $\log p_n$). Even in the case where the number of errors e is larger (but not larger than $n - \sqrt{nk}$), there exists a small list containing all integers whose Chinese remainder representations differ from the vector $\langle r_1, \dots, r_n \rangle$ in at most e coordinates [13]. Again it is not clear how to recover this list in polynomial time.

In this paper we present efficient algorithms for solving the above problems. Specifically we provide polynomial-time algorithms for the following two tasks:

1. Unique Decoding: Given n relatively prime integers $p_1 < \dots < p_n$; n residues r_1, \dots, r_n , with $0 \leq r_i < p_i$; and an integer k ; find an integer $m < \prod_{i=1}^k p_i$ satisfying $(m \bmod p_i) \neq r_i$ for at most $(n - k) \frac{\log p_1}{\log p_1 + \log p_n}$ values of $i \in \{1, \dots, n\}$, if such an integer exists. (Theorem 6.)
2. List Decoding (for large error): Given n relatively prime integers $p_1 < \dots < p_n$; n residues r_1, \dots, r_n , with $0 \leq r_i < p_i$; and an integer k ; construct a list of all integers m satisfying $m < \prod_{i=1}^k p_i$ and $(m \bmod p_i) = r_i$ for at least $\sqrt{2n(k+2)} \frac{\log p_n}{\log p_1} + \frac{k+3}{2} + 2 \log n = \Theta(\sqrt{nk} \frac{\log p_n}{\log p_1})$ values of $i \in \{1, \dots, n\}$. (Theorem 11.) (We comment that this list contains at most $\sqrt{2n/k}$ integers; cf., [13].)

¹Unfortunately, it does not allow for easy inequality comparisons — which is presumably why it was not employed.

In the context of coding theory, our algorithms add a new dimension to the family of codes that are efficiently correctable. The known examples of asymptotically good error-correcting codes with efficient algorithms can be classified in one of two categories:

1. **Algebraic codes:** These are codes defined using the properties of low-degree polynomials over finite fields and include a wide variety of codes such as Reed-Solomon codes, BCH codes, Alternant codes and algebraic-geometry codes. Such codes admit efficient error-correction algorithms; in fact all the algorithms (for unique-decoding) are similar in spirit and can be unified quite nicely [26, 18, 8].
2. **Combinatorial codes:** A second class of codes with efficient decoding algorithms evolve from combinatorial concepts such as expanders, super-concentrators etc. Examples of this family include the codes of Sipser and Spielman [29], and Spielman [30]. In both cases, the description of the code is captured by a graph; and the existence of a decoding algorithm is then related to combinatorial properties of the graph.

Our work provides the first example of a number theoretic code that is efficiently correctable. To the best of our knowledge - this is the only example which does not fall into one of the two classes above.

Our algorithms are obtained by abstracting from known paradigms for correcting algebraic codes: The first of our algorithms abstracts from a large collection of (unique) error-correcting algorithms for algebraic codes [27, 4, 25, 35]. In fact, an elegant unification of these results (see [26, 18, 8] or Appendix A) provides the inspiration for our algorithm. The second algorithm described above abstracts from the recent works on “list-decoding” algorithms [3, 31, 28, 15]. We stress however, that the translation of the above mentioned algorithms to our case is not immediate. In particular, the usual “interpolation” methods, that come in very handy in the algebraic case are not applicable here. In fact our code is not even linear in the usual sense and so linear algebra is not applicable in our case. Thus for solving analogies of “simple” problems in the algebraic case, we employ integer programming algorithms (in fixed dimensions) [21] for the Unique Decoding task, and the approximate basis reduction algorithm (in varying dimension) [20] for the List Decoding task. Our final algorithms achieve decoding capabilities comparable to those in algebraic cases and in particular, if $p_n = p_1^{O(1)}$ we can decode uniquely from a constant fraction of errors. We also get a list-decoding algorithm to recover from $n - o(n)$ errors, provided $k = o(n)$.

Permanent of random matrices One primary motivation for studying the Chinese remainder representation of integers was to study the “random self-reducibility” property of the permanent [22].

The standard presentation of this property *fixes* a prime $p > n + 1$, and consists of a randomized reduction of computing the permanent modulo p of a given $n \times n$ matrix to computing the permanent modulo p over uniformly distributed $n \times n$ matrices. Thus we are taking a two parameter problem (such as Quadratic Non-Residuosity and DLP) and the process of self-reduction fixes one parameter (here, the prime p) and randomizes over the second (here, the matrix). This is analogous to the results of [14, 6] but not to the recent result of Ajtai [1]. Thus, unlike Ajtai’s result, the above only relates the average and worst case complexities of computing the permanent modulo p for any fixed

p . What we want is a relation between the average and worst case complexities, when average-case complexity refers to all parts of the input.

Consider, for example, the product distribution on pairs (M, p) , parameterized by size n , where p is a uniformly distributed n -bit prime and M is a uniformly distributed n -by- n matrix with $2n$ -bit entries.

A naive analysis of the complexity of the permanent on such instances would work as follows. Suppose we have a heuristic to compute the permanent on instances from the above distribution. Then, given any pair (M, p) , pick at random many primes p_1, \dots, p_t , and then compute the permanent of M modulo p_i for every i . In each case use the random-self-reducibility of the permanent modulo p_i to reduce the computation of the permanent of M modulo p_i to $n + 1$ “random” (but not independent) instances of the permanent modulo p_i . If the heuristic does not make errors very often (say has error probability less than $\frac{1}{3(n+1)t}$) then with high probability (resp., probability at least $2/3$) all calls to the heuristic get answered correctly. Thus if t is large enough (e.g., $t = O(n)$ will do), then (applying the Chinese Remainder Theorem) we obtain the value of the permanent of M (over the integers), and can now reduce this modulo p to get the desired output.

However the reduction as described above is not very tolerant of errors. This problem has been addressed before in the case of one of the two parameters, namely in the choice of the matrix: The results of [11, 12, 31] imply that if for any prime p , the heuristic computes (M, p) on even a tiny but non-negligible fraction of the instances correctly then the permanent can be computed correctly on worst case instances of matrices, but over the same fixed prime p .

Our result complements the above, by allowing a similar treatment of the second parameter as well. Thus by combining the two results, we get the following natural statement:

If there exists a heuristic that computes the permanent of a random pair (M, p) , from the above distribution, with non-negligible probability (over the choice of (M, p)), then $\text{P}^{\#\text{P}} = \text{BPP}$.

Organization of this paper: In Section 2 we define the Chinese Remainder Code. In Sections 3 and 4 we give decoding algorithms for the Chinese Remainder Code, for small and large error, respectively. Section 5 gives the application to the permanent.

2 The Chinese Remainder Code

Notation: For positive integers M, N , Let \mathbb{Z}_M denote the set $\{0, \dots, M-1\}$, and let $[N]_M$ denote the remainder of N when divided by M . Note that $[N]_M \in \mathbb{Z}_M$.

Definition 1 (Chinese Remainder Code) Let $p_1 < \dots < p_n$ be relatively prime integers, and $k < n$ an integer. The Chinese Remainder Code with basis p_1, \dots, p_n and rate k is defined for message space \mathbb{Z}_K , where $K \stackrel{\text{def}}{=} \prod_{i=1}^k p_i$. The encoding of a message $m \in \mathbb{Z}_K$, denoted $E_{p_1, \dots, p_n}(m)$, is the n -tuple $\langle [m]_{p_1}, \dots, [m]_{p_n} \rangle$.

Thus the Chinese Remainder Code does not have a “fixed alphabet” (the alphabet depends on the coordinate position) and it is not linear in the usual sense (as the natural arithmetic here is done modulo p_i for the i 'th coordinate). Distance of a code can however be defined as usual; i.e., the distance between two “words” of block length n is the number of coordinates on which they differ; and the distance of a code is the minimum distance between any pair of distinct codewords. The distance properties of this code are very similar to those of Reed-Solomon and BCH codes; and follow immediately from the Chinese Remainder Theorem:

Theorem 2 (Chinese Remainder Theorem — CRT) *If q_1, \dots, q_ℓ are relatively prime positive integers and r_1, \dots, r_ℓ are integers such that $r_i \in \mathbb{Z}_{q_i}$, then there exists a unique integer $r \in \mathbb{Z}_{\prod_{i=1}^\ell q_i}$ such that $[r]_{q_i} = r_i$. Furthermore, $r = [\sum_{i=1}^n c_i \cdot Q_i \cdot r_i]_Q$, where $Q = \prod_{j=1}^\ell q_j$, $Q_i = Q/q_i$, and c_i is the multiplicative inverse modulo q_i of Q_i .*

Corollary 3 *For any n relatively prime integers p_1, \dots, p_n and any integer $k < n$, the Chinese Remainder Code with basis p_1, \dots, p_n and rate k has distance $n - k + 1$. That is, for any two messages m_1, m_2 , the code words $E_{p_1, \dots, p_n}(m_1)$ and $E_{p_1, \dots, p_n}(m_2)$ disagree on at least $n - k + 1$ coordinates.*

Thus if p_1, \dots, p_n are all $(1 + o(1)) \cdot \log n$ -bit primes, then the information rate and the distance of the Chinese Remainder Code are comparable with those of the Reed-Solomon code or the BCH code. For our purposes, it is more useful to consider a variant of the notions of block length, rate and distance as defined below.

Definition 4 (amplitude) *For a Chinese Remainder Code with basis p_1, \dots, p_n and rate k , the amplitude of the encoding is defined to be $N = \prod_{i=1}^n p_i$; the amplitude of the message space is defined to be $K = \prod_{i=1}^k p_i$. For vectors $\vec{v} = \langle v_1, \dots, v_n \rangle$ and $\vec{w} = \langle w_1, \dots, w_n \rangle \in \mathbb{Z}^n$ with $v_i, w_i \in \mathbb{Z}_{p_i}$, the amplitude of the distance between \vec{v} and \vec{w} is defined to be $\prod_{i:v_i \neq w_i} p_i$. The amplitude of agreement between \vec{v} and \vec{w} is defined to be $\prod_{i:v_i = w_i} p_i$. Notice that the product of the amplitudes of agreement and distance equals the amplitude of the encoding.*

It is easy to see that if the distance between \vec{v} and \vec{w} is d , and the amplitude of the distance between \vec{v} and \vec{w} is D ; then $d \log p_1 \leq \log D \leq d \log p_n$. In case of traditional codes that are defined over fixed alphabets, i.e., $p_1 = p_2 = \dots = p_n$, d is directly proportional to $\log D$ and hence there is no need to consider the latter separately. In our case, the latter parameter provides a more refined look at the performance of the algorithms. From the Chinese Remainder Theorem it follows immediately that the amplitude of distance between any two codewords is larger than N/K .

Our goal is to solve the following error-correction problems (for as large an error parameter as possible).

The Error-correction/List decoding Problem

Given: (1) n relatively prime integers $p_1 < \dots < p_n$ and rate parameter k specifying a Chinese Remainder Code; (2) n integers r_1, \dots, r_n , with $r_i \in \mathbb{Z}_{p_i}$ and an error-parameter e .

Task: Find (all) message(s) $x \in \mathbb{Z}_K$, where $K = \prod_{i=1}^k p_i$, s.t. $[x]_{p_i} \neq r_i$ for at most e values of i .

It follows from the distance of the Chinese Remainder Code that the answer is unique if $e < \frac{n-k}{2}$. In this case the problem corresponds to the traditional error-correction problem for error-correcting codes. If e is larger, then there may be more than one solution. We will expect the algorithm to return a list of all codewords x with at most e errors.

3 The Decoding Algorithm for Small Error

The first algorithm we present is a simple algorithm to recover from a small number of errors. The algorithm recovers from error of amplitude at most $\sqrt{N/K}$. Translating to classical measures this yields an error-correcting algorithm for $e \leq (n-k) \frac{\log p_1}{\log p_1 + \log p_n}$ (and in particular, if $p_n = p_1^{O(1)}$, then the algorithm can handle a constant fraction of errors).

The algorithm is described below formally. The inspiration for the algorithm comes from a general paradigm for decoding of many algebraic codes (see [26, 18, 8] or Appendix A). Given a received word $\langle r_1, \dots, r_n \rangle$ that is close to the encoding of (a unique) message m , the algorithm Unique-Decode tries to find two integers y and z such that $y \cdot m = z$. To this end it first reconstructs the integer $r \in \mathbb{Z}_N$ that corresponds to the received word $\langle r_1, \dots, r_n \rangle$ (i.e., $[r]_{p_i} = r_i$, for every i). It then searches for integers y and z such that $y \cdot r \equiv z \pmod{N}$ (where $N = \prod_{i=1}^n p_i$), and both y and z are of bounded sizes. In the analysis of the algorithm we show that the equality (modulo N) between $r \cdot y$ and z together with the restrictions on the sizes of y and z implies that $y \cdot m$ is equal to z (over the integers). Furthermore, as we show in Appendix B, y has an *error-detection* property. Namely, for every index i such that $r_i \neq [m]_{p_i}$, it holds that $[y]_{p_i} = 0$, and moreover, the message m can be reconstructed from the remaining r_i 's. Though we do not use this property explicitly in the algorithm described below (as well as in its analysis), it can be used to obtain a variant of the algorithm, described in Appendix B, which is more clearly related to the general decoding paradigm.

Unique-Decode($p_1, \dots, p_n, k, r_1, \dots, r_n$).

- Set $K = \prod_{i=1}^k p_i$, $N = \prod_{i=1}^n p_i$, and $F = (K-1)E$, with E to be determined later.
- Let $r \in \mathbb{Z}_N$ be s.t. $r_i = [r]_{p_i}$ (as defined by CRT).

1. Find integers y, z s.t.

$$\left. \begin{array}{l} 1 \leq y \leq E \\ 0 \leq z \leq F \\ y \cdot r \equiv z \pmod{N} \end{array} \right\} \quad (1)$$

2. Output z/y if it is an integer.

The above algorithm can be implemented in polynomial time in the bit sizes of p_1, \dots, p_n . The main realization is that Step 1 can be computed using an algorithm for integer programming in fixed number of variables, due to [21]. To see how to formulate our problem in this way, we let the final equality be expressed as $y \cdot r = z + j \cdot N$. Our task thus reduces to computing y and j s.t. $0 < y \leq E$ and $0 \leq y \cdot r - j \cdot N < F$. Step 2 is straightforward.

We now analyze the performance of this algorithm. We first describe it in terms of the amplitude of the distance between the message m and the received word r .

Lemma 5 *If r is such that for some $m \in \mathbb{Z}_K$ the amplitude of the distance between $\langle r_1, \dots, r_n \rangle$ and $\langle [m]_{p_1}, \dots, [m]_{p_n} \rangle$ is at most E , and $N > E^2 \cdot K$ then $\text{Unique-Decode}(p_1, \dots, p_n, k, r_1, \dots, r_n)$ returns m .*

We prove the lemma using the following two claims.

Claim 5.1 *For r as in the lemma, there exist y, z satisfying Eq. (1).*

Claim 5.2 *Let r and m be as stated in the lemma, and $N \geq (K - 1) \cdot E^2$. For any pair (y, z) satisfying Eq. (1) it holds that $y \cdot m = z$.*

We prove the two claim momentarily, and first show how Lemma 5 follows from the claims.

Proof of Lemma 5: By Claim 5.1, Step 1 of the algorithm always returns a pair (y, z) satisfying Eq. (1). By Claim 5.2, any pair (y, z) that may be the outcome of Step 1 satisfies $y \cdot m = z$. Thus $z/y = m$ is an integer and the output of the algorithm is m . \blacksquare

We now prove Claims 5.1 and 5.2.

Proof of Claim 5.1: Let $y = \prod_{\{i | r_i \neq [m]_{p_i}\}} p_i$ (so that y equals the amplitude of the distance between $\langle r_1, \dots, r_n \rangle$ and $\langle [m]_{p_1}, \dots, [m]_{p_n} \rangle$), and $z = y \cdot m$. Then notice that $y \neq 0$, and $y \leq E$ (as required by the first item of Eq. (1)). Since $m < K$, we have $z = m \cdot y \leq (K - 1) \cdot E$, and since $z \geq 0$, the second item of Eq. (1) also holds. Finally, by CRT, the condition $y \cdot r \equiv z \pmod{N}$ holds since the condition holds modulo every p_i : For any fixed $i \in \{1, \dots, n\}$, either $r_i = [m]_{p_i}$ or $[y]_{p_i} = 0$. In either case, we have $z = ym \equiv yr \pmod{p_i}$. \blacksquare

Proof of Claim 5.2: For every i s.t. $[m]_{p_i} = r_i$, we have

$$y \cdot m \equiv y \cdot [m]_{p_i} \equiv y \cdot r_i \equiv y \cdot r \equiv z \pmod{p_i}.$$

Thus, by CRT, $y \cdot m \equiv z \pmod{T}$ where $T = \prod_{\{i | [m]_{p_i} = r_i\}} p_i \geq N/E$ is the amplitude of the agreement between $\langle r_1, \dots, r_n \rangle$ and $\langle [m]_{p_1}, \dots, [m]_{p_n} \rangle$. But z and $m \cdot y$ are both at most $(K - 1)E < N/E$. Thus $z = m \cdot y$. \blacksquare

As an immediate consequence of Lemma 5, and the observation relating amplitudes of distance to classical distance, we get the following theorem.

Theorem 6 *$\text{Unique-Decode}(p_1, \dots, p_n, k, r_1, \dots, r_n)$ solves the error-correction problem in polynomial time for any value of the error parameter up to $(n - k) \frac{\log p_1}{\log p_1 + \log p_n}$, with the setting $E = \prod_{i=n-e+1}^n p_i$.*

Proof: Using $N = \prod_{i=1}^n p_i$, $K = \prod_{i=1}^k p_i$ and $E \leq \prod_{i=n-e+1}^n p_i$, Lemma 5 can be applied if $\prod_{i=1}^n p_i > (\prod_{i=1}^k p_i) \cdot (\prod_{i=n-e+1}^n p_i)^2$, which is equivalent to $\prod_{i=k+1}^{n-e} p_i > \prod_{i=n-e+1}^n p_i$. In turn this condition holds if $p_1^{n-(e+k)} > p_n^e$. The theorem follows by taking logarithms of both sides. \blacksquare

4 Decoding for Large Error

In this section we will describe an algorithm that recovers from possibly many more errors than described in the previous section. In particular, if we fix $k = \epsilon n$ and let $n \rightarrow \infty$, the fraction of errors that can be corrected goes to $1 - \sqrt{2\epsilon \frac{\log p_n}{\log p_1}}$. As $\epsilon \rightarrow 0$, this quantity approaches 1. This algorithm is inspired by the recent progress in list-decoding algorithms [3, 31, 28, 15]. Our algorithm and analysis follow the same paradigm, though each step is different. A closer comparison is included at the end of this section.

The algorithm List-Decode can be viewed as a generalization of Unique-Decode. In both algorithms, given the received word $\langle r_1, \dots, r_n \rangle$, the algorithm first finds, using CRT, an integer $r \in \mathbb{Z}_N$ corresponding to the received word (i.e., $[r]_{p_i} = r_i$ for every i). In Unique-Decode the algorithm then attempts to find integers y and z (restricted in size), such that $y \cdot r \equiv z \pmod{N}$, and outputs z/y . In other words, the algorithm searches for integers y, z satisfying $y \cdot r - z \equiv 0 \pmod{N}$, and outputs the (unique) root of the (degree-1) polynomial $y \cdot x - z$. In List-Decode, the algorithm instead searches for a *sequence* of integers c_0, \dots, c_ℓ (of certain bounded sizes), such that $\sum_i c_i r^i \equiv 0 \pmod{N}$ and outputs *all* roots of the polynomial $\sum_i c_i x^i$. As we show subsequently, the increase in the degree of the polynomial that the algorithm searches for (together with the particular restrictions on the sizes of its coefficients) allows us to decode for much larger error.

List-Decode($p_1, \dots, p_n, k, r_1, \dots, r_n$).

Set $N = \prod_{i=1}^n p_i$; $K = \prod_{i=1}^k p_i$; and $F = 2^{\frac{\ell+2}{2}} \cdot \sqrt{\ell+2} \cdot N^{\frac{1}{\ell+1}} \cdot K^{\frac{\ell+1}{2}}$, with ℓ to be determined shortly.

Let $r \in \mathbb{Z}_N$ s.t. $[r]_{p_i} = r_i$ for every i (as defined by CRT).

1. Find integers c_0, \dots, c_ℓ satisfying

$$\left. \begin{array}{l} \forall 0 \leq i \leq \ell \quad |c_i| \leq \frac{F}{K^i} \\ \text{s.t.} \quad \sum_{i=0}^{\ell} c_i r^i \equiv 0 \pmod{N} \\ \langle c_0, \dots, c_\ell \rangle \neq \vec{0} \end{array} \right\} \quad (2)$$

2. Output all roots of the integer polynomial $C(x) = \sum_{i=0}^{\ell} c_i x^i$.

The running time of Step 2 above is bounded by a polynomial in $n, \ell, \log N$ and $\log F$ (one can use LLL's algorithm for factoring polynomials over the integers if required, though faster algorithms exist for this simpler task of "root-finding"). We need to show how to implement Step 1. Mainly the idea is to set up a lattice whose short vectors correspond to small values of the coefficients c_i 's. We show first that very small vectors of this form exist; and then use the basis reduction algorithm of LLL to find short (but not shortest) vectors in this lattice; and this will suffice for Step 1.

Lemma 7 (Algorithm for Step 1.) c_i 's as required in Step 1 of List-Decode exist and can be found in polynomial time.

Proof: We set up an $\ell+2$ -dimensional integer lattice using basis vectors v_0, \dots, v_ℓ and w described next. Let M be a very large integer (to be determined later as a function of N and ℓ). For $j \in \{0, \dots, \ell+1\}$, the j th coordinate of the vector v_i , denoted $(v_i)_j$ is given by:

$$(v_i)_j = \begin{cases} K^i & \text{if } j = i \\ M \cdot r^i & \text{if } j = \ell + 1 \\ 0 & \text{otherwise.} \end{cases}$$

The vector w is zero everywhere except in the last coordinate where $(w)_{\ell+1} = M \cdot N$.

A generic vector in this lattice is of the form $u = \sum_{i=0}^{\ell} c_i v_i + dw$, for integers c_0, \dots, c_ℓ and d . Explicitly the j th coordinate of u is given by:

$$(u)_j = \begin{cases} c_j K^j & 0 \leq j \leq \ell \\ M \cdot (\sum_{i=0}^{\ell} c_i r^i + dN) & \text{if } j = \ell + 1. \end{cases}$$

We are interested in showing that this lattice contains “short” vectors whose last coordinate equals 0, and every other coordinate has absolute value at most F (thus satisfying Eq. (2)). Furthermore, we would like to show that such vectors can be found efficiently. To this end, we first prove the following technical lemma.

Lemma 8 For integers r, N if B_0, \dots, B_ℓ are positive integers such that $\prod_{i=0}^{\ell} B_i > N$, then there exist integers c_0, \dots, c_ℓ , such that $|c_i| < B_i$, $\langle c_0, \dots, c_\ell \rangle \neq \vec{0}$ and $\sum_{i=0}^{\ell} c_i r^i \equiv 0 \pmod{N}$.

Proof: Consider the function $f : \mathbb{Z}_{B_0} \times \dots \times \mathbb{Z}_{B_\ell} \rightarrow \mathbb{Z}_N$ given by $f(c_0, \dots, c_\ell) = [\sum_{i=0}^{\ell} c_i r^i]_N$. Since the domain has larger cardinality than the range, there exist different $\langle d_0, \dots, d_\ell \rangle$ and $\langle e_0, \dots, e_\ell \rangle$ s.t. $f(d_0, \dots, d_\ell) = f(e_0, \dots, e_\ell)$. Setting $c_i = d_i - e_i$, we get $|c_i| < B_i$, $\sum_i c_i r^i = 0$, and $\langle c_0, \dots, c_\ell \rangle \neq \vec{0}$ as required. ■

Using Lemma 8 with $B_i = N^{\frac{1}{\ell+1}} \cdot K^{\frac{\ell+1}{2}-i}$, we observe that the lattice defined above has a (short) non-zero vector (where the c_i 's are as guaranteed by the lemma and $d = -\sum_{i=0}^{\ell} c_i r^i / N$) with the last coordinate identically 0, and each other coordinate has absolute value at most $B_i \cdot K^i = N^{\frac{1}{\ell+1}} \cdot K^{\frac{\ell+1}{2}}$. Thus, the L_2 -norm of this vector is at most $\sqrt{\ell+2} \cdot N^{\frac{1}{\ell+1}} \cdot K^{\frac{\ell+1}{2}}$. By using the “approximate shortest vector” algorithm of [20], we find, in polynomial time, a vector of L_2 -norm at most $F = 2^{\frac{\ell+2}{2}} \cdot \sqrt{\ell+2} \cdot N^{\frac{1}{\ell+1}} \cdot K^{\frac{\ell+1}{2}}$. For sufficiently large M (any $M > F$ will do), all “short” vectors (i.e., with L_2 -norm at most F) have a last coordinate identical to 0, and thus yield a sequence of c_i 's satisfying $\sum_i c_i r^i \equiv 0 \pmod{N}$ and $|c_i \cdot K^i| \leq F$. This sequence is as required in Step 1. ■

Now we move on to Step 2 of List-Decode. We argue next that any solution to the list-decoding problem is a root of the polynomial whose coefficients are given by *any* solution to Step 1. Instead of performing the analysis in terms of the amount of error in the received word, we do so in terms of the amount of agreement with some message.

Lemma 9 *If r is such that for some $m \in \mathbb{Z}_K$ the amplitude of the agreement between (r_1, \dots, r_n) and $\langle [m]_{p_1}, \dots, [m]_{p_n} \rangle$ is greater than $2(\ell + 1)F$, and c_0, \dots, c_ℓ are integers satisfying Eq. (2), then $\sum_{j=0}^{\ell} c_j m^j = 0$ (i.e., m is a root of the polynomial $C(x)$).*

Proof: We first observe that since the c_j 's are small, $\sum_j c_j m^j$ is small in absolute value:

$$\begin{aligned} \left| \sum_{j=0}^{\ell} c_j m^j \right| &\leq (\ell + 1) \cdot \max_j \{|c_j m^j|\} \\ &\leq (\ell + 1) \cdot \max_j \{|c_j K^j|\} \\ &\leq (\ell + 1) \cdot F. \end{aligned}$$

Now we observe that for i such that $[m]_{p_i} = r_i$ it holds that

$$\sum_{j=0}^{\ell} c_j m^j \equiv \sum_{j=0}^{\ell} c_j [m]_{p_i}^j \equiv \sum_{j=0}^{\ell} c_j r_i^j \equiv \sum_{j=0}^{\ell} c_j r^j \equiv 0 \pmod{p_i}.$$

Define $P = \prod_{\{i | r_i = [m]_{p_i}\}} p_i$. By CRT, $\sum_{j=0}^{\ell} c_j m^j \equiv 0 \pmod{P}$. Since the sum $\sum_{j=0}^{\ell} c_j m^j$ has absolute value at most $(\ell + 1)F$, the hypothesis $P > 2 \cdot (\ell + 1)F$ implies that the sum is identically zero as required. \blacksquare

As an immediate consequence of the last two lemmas, we get a proof of the correctness of List-Decode. The following proposition describes the performance in terms of amplitude (for any choice of ℓ).

Proposition 10 *For any choice of the parameter ℓ , List-Decode($p_1, \dots, p_n, k, r_1, \dots, r_n$) produces a list of up to ℓ integers which includes all messages $m \in \mathbb{Z}_K$ such that the amplitude of agreement between $\langle [m]_{p_1}, \dots, [m]_{p_n} \rangle$ and \vec{r} is at least $2(\ell + 2)^{3/2} 2^{\frac{\ell+2}{2}} N^{\frac{1}{\ell+1}} K^{\frac{\ell+1}{2}}$.*

Proof: By Lemma 7, c_i 's satisfying Eq. (2) exist and are found in Step 1. By Lemma 9, any m as in the lemma is a root of the polynomial $\sum_j c_j x^j$, and thus is included in the output. \blacksquare

The following theorem is obtained by optimizing the choice of the parameter ℓ in the above proposition.

Theorem 11 *List-Decode($p_1, \dots, p_n, k, r_1, \dots, r_n$) with parameter $\ell = \left\lceil \sqrt{\frac{2n \log p_n}{k \log p_1}} - 1 \right\rceil$ solves the error-correction problem in polynomial time, for $e < n - \sqrt{2(k+3)n \frac{\log p_n}{\log p_1}} - \frac{k+6}{2}$.*

Remark: If $k/n = \epsilon$, then the above theorem indicates that approximately $1 - \sqrt{2 \cdot \left(\frac{\log p_n}{\log p_1}\right) \cdot \epsilon} - \epsilon/2$ fraction of errors can be corrected. In particular this fraction approaches 1 as $\epsilon \rightarrow 0$.

Proof: Suppose we want to find all codewords which agree with $\langle r_1, \dots, r_n \rangle$ on t coordinates. Setting $f_1 \stackrel{\text{def}}{=} (\ell + 2)^{3/2}$ and $f_2 \stackrel{\text{def}}{=} 2^{\frac{\ell+2}{2}}$, and applying Proposition 10, it suffices to show that

$$\prod_{i=1}^t p_i \geq f_1 \cdot f_2 \cdot \left(\prod_{i=1}^n p_i \right)^{\frac{1}{\ell+1}} \cdot \left(\prod_{i=1}^k p_i \right)^{\frac{\ell+1}{2}}$$

Setting $t = t_1 + t_2 + t_3$, we will find t_1, t_2, t_3 s.t.

$$p_1^{t_1} \geq f_1 \tag{3}$$

$$p_1^{t_2} \geq f_2 \tag{4}$$

$$\text{and } \prod_{i=1}^{t_3} p_i \geq \left(\prod_{i=1}^n p_i \right)^{\frac{1}{\ell+1}} \cdot \left(\prod_{i=1}^k p_i \right)^{\frac{\ell+1}{2}} \tag{5}$$

We start with an analysis of the last inequality. For this we need

$$\left(\prod_{i=1}^{t_3} p_i \right)^{1 - \frac{1}{\ell+1}} \geq \left(\prod_{i=t_3+1}^n p_i \right)^{\frac{1}{\ell+1}} \cdot \left(\prod_{i=1}^k p_i \right)^{\frac{\ell+1}{2}}$$

Let $q = (\prod_{i=1}^k p_i)^{1/k}$. Then $q \geq p_1$ and $(\prod_{i=1}^{t_3} p_i) \geq q^{t_3}$, provided $t_3 \geq k$. Thus it suffices to show

$$q^{t_3 \ell / (\ell+1)} \geq p_n^{(n-t_3)/(\ell+1)} \cdot q^{k \cdot \frac{\ell+1}{2}} \tag{6}$$

Fact: Eq. (6) holds if $t_3 \geq \frac{k(\ell+1)}{2} + \frac{n \log p_n}{(\ell+1) \log p_1}$ and $\ell \geq 1$.

Proof: By the hypothesis, $t_3 \geq k$ and so

$$\begin{aligned} & \left(t_3 - \frac{k(\ell+1)}{2} \right) \log p_1 \geq \frac{n}{\ell+1} \log p_n \\ \Rightarrow & \left(\frac{\ell}{\ell+1} t_3 - \frac{k(\ell+1)}{2} \right) \log q \geq \frac{n-t_3}{\ell+1} \log p_n \\ \Rightarrow & \frac{\ell t_3}{\ell+1} \log q \geq \frac{n-t_3}{\ell+1} \log p_n + \frac{k(\ell+1)}{2} \log q \end{aligned}$$

and Eq. (6) follows. \square

Setting $\ell + 1 = \left\lceil \sqrt{\frac{2n \log p_n}{k \log p_1}} \right\rceil$, shows that $t_3 = \sqrt{2kn \frac{\log p_n}{\log p_1}} + \frac{k}{2}$ suffices to achieve Eq. (5). This setting of ℓ also implies that to satisfy Eq. (4), which is equivalent to $t_2 \geq \frac{\ell+2}{2 \log p_1}$, it suffices to set $t_2 = \sqrt{\frac{2n \log p_n}{k \log p_1}} + \frac{3}{2}$. Finally, to satisfy Eq. (3), which is equivalent to $t_1 \geq \frac{3 \log(\ell+2)}{2 \log p_1}$, it suffices to set $t_1 = \frac{\log(2n \log p_n / k \log p_1)}{\log p_1}$, which is smaller than $2 \log t_2 / \log p_1 \ll t_2$.

Thus we find that it suffices to have

$$t = 2 \sqrt{\frac{2n \log p_n}{k \log p_1}} + 3 + \sqrt{2kn \frac{\log p_n}{\log p_1}} + \frac{k}{2}$$

$$\begin{aligned}
&= \left(1 + \frac{2}{k}\right) \cdot \sqrt{2kn \frac{\log p_n}{\log p_1}} + \frac{k+6}{2} \\
&< \sqrt{1 + 3 \cdot \frac{2}{k}} \cdot \sqrt{2kn \frac{\log p_n}{\log p_1}} + \frac{k+6}{2} \\
&= \sqrt{2(k+3)n \frac{\log p_n}{\log p_1}} + \frac{k+6}{2}
\end{aligned}$$

Setting $e < n - t$ yields the theorem. \blacksquare

Comparison with [3, 31] Our algorithm List-Decode is similar to those of [3, 31] in the basic steps. In their case also, they first find a polynomial “explaining” the corrupted word and then factor it to retrieve a list of messages. However the specifics are quite different: They look for a bivariate polynomial explanation; their criterion is to find a non-zero polynomial of low degree; they find it by solving a linear system; and then employ a bivariate factorization step. We look for a univariate polynomial explanation; our criterion is the size of the coefficients; we find it by (essentially) solving Diophantine systems; and finally employ univariate factorization. Similarly our analysis follows the same steps. The existence proof (Lemma 8) is similar to an analogous step in [31]; though our proof here appears to be more general than his proof. In particular, the pigeonhole argument could also be applied to his case achieving analogous results. Finally, Lemma 9 is also analogous in spirit to similar lemmas in [3, 31] - again our proofs are different since our criteria are different.

5 The Permanent of Random Matrices

In this section we show that computing the permanent of a random matrix modulo a random prime is very hard. The distribution of matrices and primes we consider is the following:

\mathcal{D} is an ensemble of distributions $\{\mathcal{D}_s\}$ where \mathcal{D}_s consists of pairs (T, p) where T is an $s \times s$ matrix whose entries are chosen uniformly and independently from $\mathbb{Z}_{2^{2s}}$, and p is a prime chosen uniformly from \mathbb{Z}_{2^s} .

The distributional problem we consider is: Given a randomly chosen pair (T, p) from \mathcal{D}_s , compute the permanent of T modulo p . We show that no polynomial time algorithm is likely to have inverse polynomial probability of solving this distributional problem.

Lemma 12 ([2] following [23]; cf., [7]) *Suppose there exists a probabilistic polynomial time algorithm A' and a polynomial $r : \mathbb{Z} \rightarrow \mathbb{Z}$ such that on input M , an $s \times s$ matrix of $2s$ -bit integer elements, $A'(M)$ outputs a list of $r(s)$ integers such that the permanent of M is included in this list (with probability at least, say, $\frac{1}{2}$ over the internal coin tosses of A'). Then $P^{\#P} = BPP$.*

We complement this lemma with an algorithm that utilizes a subroutine for computing the permanent on random instances, and uses it to compute a list of values of the permanent on worst-case instances.

Lemma 13 *Suppose there exists a polynomial time algorithm A and a function $\epsilon : \mathbb{Z} \rightarrow [0, 1]$ such that for every positive integer s ,*

$$\Pr_{(T,p) \in \mathcal{D}_s} [A(T, p) = [\text{perm}(T)]_p] \geq \epsilon(s).$$

Then there exists a randomized $\text{poly}(s/\epsilon(s))$ -time algorithm A' that on input an $s \times s$ matrix M with entries from $\mathbb{Z}_{2^{2s}}$, outputs a list of at most $O(1/\epsilon(s)^4)$ integers, which includes the permanent of M with high probability.

Proof: Assume, w.l.o.g, that when given a pair (T, p) , algorithm A first reduces each entry of T modulo p . Our algorithm for reconstructing the permanent of any s -by- s matrix, M , is given below:

Algorithm Perm(M).

- Parameters $n = \text{poly}(s/\epsilon(s))$, $n' = O(s/\epsilon(s)^2)$
- Uniformly select n random primes p_1, \dots, p_n in the interval $[2^{s/2}, 2^s]$.
- For $i = 1$ to n do /* try to obtain $[\text{perm}(M)]_{p_i}$ */
 - Subroutine Mod-Perm(M, p_i).
 - * Uniformly select an $s \times s$ random matrix R with entries from \mathbb{Z}_{p_i} .
 - * For $j = 1$ to n' do /* try to obtain $[\text{perm}(M + jR)]_{p_i}$ */
 - Let $v_j = A(M + j \cdot R, p_i)$;
 - * Reconstruct a list of all degree s univariate polynomials $\{f_1, \dots, f_{\ell'}\}$ that satisfy $f_h(j) = v_j$ for at least an $\epsilon(s)/16$ fraction of the v_j 's.
 - * Uniformly select a random $h \in \{1, \dots, \ell'\}$ and set $r_i = f_h(0)$.
 - /* with probability $\text{poly}(\epsilon(s))$ (taken over the choice of p_i and the internal coins of Mod-Perm), we will have $r_i = [\text{perm}(M)]_{p_i}$ */
- Reconstruct a list of all integers $x \leq s!2^{s^2}$ such that $[x]_{p_i} = r_i$ for at least $t = O(\epsilon(s)^4) \cdot n$ of the i 's, and output this list. Namely, apply List-Decode with parameters p_1, \dots, p_n , $k = 6s$ (as $K = s!2^{s^2} < 2^{3s^2}$ and $\forall i, p_i \geq 2^{s/2}$), and r_1, \dots, r_n .

The polynomial reconstruction step may be performed using the algorithm of [31], which requires $n' \geq 2s \cdot (\epsilon(s)/16)^{-2}$. (To recover polynomials of degree s from a list of values at n' places, the algorithm requires the agreement t' to satisfy $t' > \sqrt{2sn'}$.) The reconstruction of integers satisfying the Chinese Remainder Property uses Theorem 11 and works when $n = \Omega(s/\epsilon(s)^8)$. (Here to recover all sequences with agreement t out of n places, the algorithm requires $t = \Omega(\sqrt{kn}) = \Omega(\sqrt{sn})$.)

Let P_s denote the set of primes in the interval $[2^{s/2}, 2^s]$. Let \mathcal{D}'_s be the distribution over pairs (T', p') where p' is chosen uniformly in P_s (rather than among the primes in \mathbb{Z}_{2^s} , as defined by \mathcal{D}_s), and then T' is chosen uniformly from the set of $s \times s$ matrices with entries from $\mathbb{Z}_{p'}$ (rather than by reducing modulo p' a matrix with entries chosen independently and uniformly in $\mathbb{Z}_{2^{2s}}$). We notice that the statistical difference between the two distributions is at most $O\left(\frac{2^{s/2}/(s/2)}{2^{s/2}}\right) + s^2 \cdot \frac{2^s}{2^{2s}}$,

which is negligible (where the first term comes from the probability that in D_s a prime smaller than $2^{s/2}$ is selected, and the second from uneven wrap-around in the reduction modulo a prime). In particular this implies that

$$\Pr_{(T', p') \in \mathcal{D}'_s} [A(T', p') = [\text{perm}(T')]_{p'}] \geq \frac{\epsilon(s)}{2}.$$

Say that a prime p' (from P_s) is *good* if

$$\Pr_{T' \in \mathbb{Z}_{p'}^{s \times s}} [A(T', p') = [\text{perm}(T')]_{p'}] \geq \frac{\epsilon(s)}{4}.$$

A simple counting argument shows that at least $\epsilon(s)/4$ fraction of the primes in P_s are good.

For any fixed good prime p' , and for any $j \in \{1, \dots, n'\}$, we thus have that

$$\Pr_{R \in \mathbb{Z}_{p'}^{s \times s}} [A(M + jR, p') = [\text{perm}(M + jR)]_{p'}] \geq \frac{\epsilon(s)}{4}$$

(recall that we assume that when given a pair (T, p) , algorithm A first reduces each entry of T modulo p). Say that a matrix R is *compatible* with p' if

$$\Pr \left[|\{j : A(M + jR, p') = [\text{perm}(M + jR)]_{p'}\}| > \frac{\epsilon(s)}{16} n' \right] > \frac{\epsilon(s)}{16},$$

(where the probability here is taken only over the coin flips of A). It is not hard to verify that the probability that a random R is compatible with p' is at least $\epsilon(s)/8$. It follows that for any good p' ,

$$\Pr [\text{Mod-Perm}(M, p') = [\text{perm}(M)]_{p'}] \geq \frac{\epsilon(s)}{8} \cdot \frac{\epsilon(s)}{16} \cdot \frac{1}{\ell'}$$

where the first term ($\epsilon(s)/8$) is the probability that R is compatible with p' ; the second ($\epsilon(s)/16$) is the probability that A returns the correct output for at least $\epsilon(s)/16$ fraction of the j 's (so that the polynomial reconstruction can work), conditioned on R being compatible; and the third term ($1/\ell'$) is the probability of selecting the correct index h . As $\ell' \leq 2 \cdot (\epsilon(s)/16)^{-1}$ (cf., [31]), the above probability is $\Omega(\epsilon(s)^3)$.

Recall that the probability that each p_i (uniformly selected in P_s) is good is at least $\epsilon(s)/4$. Hence, the probability, taken over the choice of p_i and the random coin flips of Mod-Perm that $\text{Mod-Perm}(M, p_i) = [\text{perm}(M)]_{p_i}$, is $\Omega(\epsilon(s)^4)$. Finally, since the success events of the various i 's are independent, by applying a Chernoff bound, we get that with high probability, the number of p_i 's for which $r_i = [\text{perm}(M)]_{p_i}$ is at least $\Omega(\epsilon(s)^4) \cdot n$. In this case List-Decode will succeed in reconstructing a list that includes $\text{perm}(M)$. ■

By combining Lemma 12 and Lemma 13 we get

Theorem 14 *Suppose there exists a polynomial time algorithm A and a positive polynomial function $q : \mathbb{Z} \rightarrow \mathbb{Z}$ such that for every positive s ,*

$$\Pr_{(T, p) \in \mathcal{D}_s} [A(T, p) = [\text{perm}(T)]_p] \geq \frac{1}{q(s)}$$

Then $\text{P}^{\#\text{P}} = \text{BPP}$.

Acknowledgments

We would like to thank Venkatesan Guruswami for bringing the work of [8] to our attention. We would like to thank Valentine Kabanets for pointing out an error in the earlier version of this paper.

References

- [1] M. AJTAI. Generating hard instances of lattice problems (extended abstract). Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, pages 99-108, Philadelphia, Pennsylvania, 22-24 May 1996.
- [2] A. AMIR, R. BEIGEL, AND W. GASARCH. Cheatable, P-terse, and P-superterse sets, manuscript, Dec. 1989.
- [3] S. AR, R. LIPTON, R. RUBINFELD AND M. SUDAN. Reconstructing algebraic functions from mixed data. *SIAM Journal on Computing*, 28(2):488-511, 1999. Preliminary version in *FOCS*, 1992.
- [4] E. R. BERLEKAMP. *Algebraic Coding Theory*. McGraw Hill, New York, 1968.
- [5] E. R. BERLEKAMP. Bounded Distance +1 Soft-Decision Reed-Solomon Decoding. *IEEE Transactions on Information Theory*, 42(3):704-720, 1996.
- [6] M. BLUM AND S. MICALI. How to Generate Cryptographically Strong Sequences of Pseudo-Random Bits. *SICOMP*, Vol. 13, pages 850-864, 1984. Preliminary version in *23rd FOCS*, 1982.
- [7] J. CAI AND L. A. HEMACHANDRA. A note on enumerative counting. *Information Processing Letters*, 38(4):215-219, 31 May 1991.
- [8] I. M. DUURSMA. *Decoding codes from curves and cyclic codes*. Ph.D. Thesis, Eindhoven, 1993.
- [9] P. ELIAS. List decoding for noisy channels. Technical Report 335, Research Laboratory of Electronics, MIT, 1957.
- [10] P. ELIAS. Error-correcting codes for list decoding. *IEEE Transactions on Information Theory*, 37:5-12. 1991.
- [11] P. GEMMELL, R. LIPTON, R. RUBINFELD, M. SUDAN AND A. WIGDERSON. Self-testing/correcting for polynomials and for approximate functions. Proceedings of the Twenty Third Annual ACM Symposium on Theory of Computing, pages 32-42, New Orleans, Louisiana, 6-8 May 1991.
- [12] P. GEMMELL AND M. SUDAN. Highly resilient correctors for multivariate polynomials. *Information Processing Letters*, 43(4):169-174, 1992.
- [13] O. GOLDBREICH, R. RUBINFELD AND M. SUDAN. Learning polynomials with queries: The highly noisy case. *36th FOCS*, pages 294-303, 1995. Revised version available from *ECCC*, 1998.

- [14] S. GOLDWASSER AND S. MICALI. Probabilistic Encryption. *JCSS*, Vol. 28, No. 2, pages 270–299, 1984. Preliminary version in *14th STOC*, 1982.
- [15] V. GURUSWAMI AND M. SUDAN. Improved decoding for Reed-Solomon and algebraic-geometric codes. To appear, *FOCS* 1998.
- [16] E. KALTOFEN. Polynomial factorization 1987–1991. *LATIN '92*, I. Simon (Ed.) Springer LNCS, v. 583:294–313, 1992.
- [17] R. M. KARP AND M. O. RABIN. Efficient randomized pattern-matching algorithms. Technical report TR-31-81, Aiken Computation Laboratory, Harvard University, 1981.
- [18] R. KOTTER. A unified description of an error locating procedure for linear codes. *Proceedings of Algebraic and Combinatorial Coding Theory*, Voneshta Voda, Bulgaria, 1992.
- [19] E. KUSHILEVTITZ AND N. NISAN. *Communication Complexity*. Cambridge University Press, 1997.
- [20] A. K. LENSTRA, H. W. LENSTRA AND L. LOVASZ. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515–534, 1982.
- [21] H. W. LENSTRA. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8: 538–548, 1983.
- [22] R. J. LIPTON. New directions in testing. *Distributed Computing and Cryptography*, J. Feigenbaum and M. Merritt (ed.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, American Mathematics Society, 2:191–202, 1991.
- [23] C. LUND, L. FORTNOW, H. KARLOFF, AND N. NISAN. Algebraic Methods for Interactive Proof Systems. *JACM*, Vol. 39, No. 4, pages 859–868, 1992. Preliminary version in *31st FOCS*, 1990.
- [24] F. J. MACWILLIAMS AND N. J. A. SLOANE. *The Theory of Error-Correcting Codes*. North-Holland, Amsterdam, 1981.
- [25] J. L. MASSEY. Shift register synthesis and BCH decoding. *IEEE Transactions on Information Theory*, 15:122–127, 1969.
- [26] R. PELLIKAAN. On decoding linear codes by error correcting pairs. *Eindhoven University of Technology*, preprint, 1988.
- [27] W. W. PETERSON. Encoding and error-correction procedures for Bose-Chaudhuri codes. *IRE Transactions on Information Theory*, IT-60:459–470, 1960.
- [28] M. A. SHOKROLLAHI AND H. WASSERMAN. Decoding algebraic-geometric codes beyond the error-correction bound. *STOC*, 1998.
- [29] M. SIPSER AND D. A. SPIELMAN. Expander codes. *IEEE Transactions on Information Theory*, 42(6):1710–1722, 1996.

- [30] D. A. SPIELMAN. Linear-time encodable and decodable error-correcting codes. *IEEE Transactions on Information Theory*, 42(6):1723–1732, 1996.
- [31] M. SUDAN. Decoding of Reed-Solomon codes beyond the error-correction bound. *Journal of Complexity*, 13(1):180-193, 1997.
- [32] J. H. VAN LINT. *Introduction to Coding Theory*. Springer-Verlag, New York, 1982.
- [33] L. G. VALIANT. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189-201, April 1979.
- [34] A. VARDY. Algorithmic complexity in coding theory and the minimum distance problem. *STOC*, 1997.
- [35] L. WELCH AND E. R. BERLEKAMP. Error correction of algebraic block codes. *US Patent* Number 4,633,470, issued December 1986.

A The paradigm for decoding linear codes

Our algorithm in Section 3 was inspired by the following general method for decoding linear codes (which satisfy some special properties). The exposition below is reproduced from [8] where it is attributed to [26] and [18]. In Appendix B we present a variant of our algorithm that exhibits the relation to the general paradigm more clearly.

Definition 15 For integers n, k , and d and a field \mathbb{F} , a linear code \mathcal{C} over \mathbb{F} of block length n is a linear subspace of \mathbb{F}^n . The rate of the code \mathcal{C} , denoted $r(\mathcal{C})$, is the rank of \mathcal{C} . The distance of \mathcal{C} , denoted $d(\mathcal{C})$, is the largest integer d such that any two vectors in \mathcal{C} disagree on at least $d - 1$ coordinates.

In what follows, we will assume that a linear code \mathcal{C} over \mathbb{F} , of block length n and rate k , is specified by an $n \times k$ matrix $M_{\mathcal{C}} \in \mathbb{F}^{n \times k}$ such that $\mathcal{C} = \{M_{\mathcal{C}} \cdot x \mid x \in \mathbb{F}^k\}$.

For vectors $x, y \in \mathbb{F}^n$, let $x * y \in \mathbb{F}^n$ denote the vector given by $(x * y)_i = x_i \cdot y_i$. For sets $X, Y \subseteq \mathbb{F}^n$, let $X * Y \subseteq \mathbb{F}^n$ denote the set $\{x * y \mid x \in X, y \in Y\}$.

Definition 16 (t -error correcting pair) A t -error correcting pair for a linear code \mathcal{C} over \mathbb{F} of block length n is a pair of linear codes \mathcal{A}, \mathcal{B} satisfying:

$$\mathcal{A} * \mathcal{C} \subseteq \mathcal{B} \tag{7}$$

$$r(\mathcal{A}) > t \tag{8}$$

$$d(\mathcal{A}) > n - d(\mathcal{C}) \tag{9}$$

$$d(\mathcal{B}) > t \tag{10}$$

Theorem 17 For $t \leq n$, let $\mathcal{A}, \mathcal{B}, \mathcal{C}$ be linear codes over \mathbb{F} of block length n , such that $(\mathcal{A}, \mathcal{B})$ form a t -error correcting pair for \mathcal{C} . Then, given $\mathcal{A}, \mathcal{B}, \mathcal{C}$ and a “received word” $r \in \mathbb{F}^n$, a codeword $c \in \mathcal{C}$ that differs from r in at most t locations can be found in $O(n^3)$ time.

Proof: We start by describing the algorithm for finding the codeword c .

Linear-Decode($\mathcal{A}, \mathcal{B}, \mathcal{C}, r; \mathbb{F}, n$);

1. Find a, b satisfying:

$$a \in \mathcal{A} - \{0^n\}, b \in \mathcal{B} \text{ and } a * r = b \quad (11)$$

2. Let $I \stackrel{\text{def}}{=} \{i : [a]_i \neq 0\}$. Find $c \in \mathcal{C}$ s.t. $c_i = r_i$ for every $i \in I$ and output it.

Let $M_{\mathcal{A}}, M_{\mathcal{B}}$ and $M_{\mathcal{C}}$ be the generating matrices of \mathcal{A}, \mathcal{B} and \mathcal{C} respectively. Then Step 1 above amounts to finding vectors $x_A \in \mathbb{F}^{r(\mathcal{A})}$ and $x_B \in \mathbb{F}^{r(\mathcal{B})}$ such that $(M_{\mathcal{A}} \cdot x_A)_i r_i = (M_{\mathcal{B}} \cdot x_B)_i$. This amounts to solving a linear system with n equations and $r(\mathcal{A}) + r(\mathcal{B}) \leq 2n$ unknowns and can certainly be solved using $O(n^3)$ field operations. Similarly, Step 2 amounts to finding $x_C \in \mathbb{F}^{r(\mathcal{C})}$ such that $(M_{\mathcal{C}} \cdot x_C)_i = r_i$ for $i \in I$. Again these form at most n linear equations in $r_{\mathcal{C}} \leq n$ unknowns and can be solved in $O(n^3)$ time. This yields the claim about the running time.

We prove the correctness in the following three claims. The first of the claims shows that if r is in fact close to some codeword c , then a pair satisfying Eq. (11) must exist. The second claim shows that for any pair a, b satisfying Eq. (11), every codeword c that is close to r will satisfy the conditions required in Step 2. The third claim shows that for any pair a, b , there is at most one solution to Step 2, thus completing the proof that c is the unique answer output by our algorithm. (The properties in Eq. (7)-Eq. (10) will be used in the claims below. The property $\mathcal{A} * \mathcal{C} \subseteq \mathcal{B}$ is used everywhere. The property on the rate of \mathcal{A} is used in the first claim. The property on the distance of \mathcal{B} is used in the second and the property on the distance of \mathcal{A} in the third.)

Claim 17.1 *If there exists a codeword $c \in \mathcal{C}$ that differs from r in at most t coordinates, then there exist a, b satisfying Eq. (11).*

Proof: Let $E = \{i | r_i \neq c_i\}$. We first claim there exists an $a \in \mathcal{A} - \{0^n\}$ s.t. $a_i = 0$ for $i \in E$. This is true since \mathcal{A} has rank at least $t + 1$ and we have added at most t constraints of the form $a_i = 0$ (for $i \in E$). This yields a linear subspace of \mathbb{F}^n which has rank at least 1 and hence has a non-zero vector.

For such an a , we claim that the vector $b = a * c$ satisfies the conditions of Eq. (11). $b \in \mathcal{B}$, since $\mathcal{A} * \mathcal{C} \subseteq \mathcal{B}$; and $a * r = b$ since for every $i \in \{1, \dots, n\}$, exactly one of the following holds:

1. $r_i = c_i$ and hence $a_i \cdot r_i = a_i \cdot c_i = b_i$.
2. $r_i \neq c_i$ and hence $a_i = 0$ and hence $a_i * r_i = 0 = a_i * c_i = b_i$.

■

Claim 17.2 *For any solution a, b to Eq. (11), and any codeword $c \in \mathcal{C}$ such that c and r disagree on at most t coordinates, $a * c = a * r$ (and hence $c_i = r_i$ if $a_i \neq 0$).*

Proof: Consider the vector $b' \in \mathcal{B}$ given by $b' = a * c$. $(b')_i$ and b_i may differ only if $c_i \neq r_i$ (since $b_i = a_i \cdot r_i$); and this happens for at most t values of $i \in \{1, \dots, n\}$. Thus b_i and b'_i differ in at most t locations; but since both are members of a linear code of distance at least $t + 1$, they must be identical. The claim follows. \blacksquare

Thus we now know that a codeword $c \in \mathcal{C}$ satisfying the requirements in Step 2 does exist. Finally we need to show that any codeword returned in Step 2 is close to r . Notice that by Step 2, $r_i \neq c_i$ implies $i \notin I$.

Claim 17.3 *For any pair $a \in \mathcal{A}$, $b \in \mathcal{B}$, there exists at most one codeword $c \in \mathcal{C}$ satisfying $a * c = b$.*

Proof: Assume $c_1, c_2 \in \mathcal{C}$ satisfy $a * c_1 = a * c_2 = b$. Then $a * (c_1 - c_2) = 0^n$. Thus for every index i , it must be that $a_i = 0$ or $(c_1 - c_2)_i = 0$. But $a_i = 0$ for at most $n - d(\mathcal{A})$ values of $i \in \{1, \dots, n\}$ and $(c_1 - c_2)_i = 0$ for at most $n - d(\mathcal{C})$ values of $i \in \{1, \dots, n\}$. Thus we have that $n - d(\mathcal{A}) + n - d(\mathcal{C}) \geq n$ which contradicts the fact that $d(\mathcal{A}) > n - d(\mathcal{C})$. \blacksquare

This concludes the proof of the correctness of algorithm **Linear-decode**; and this yields the theorem. \blacksquare

The above paradigm for decoding unifies most of the known (unique) decoding algorithms. For example, the Welch-Berlekamp algorithm [35, 5] (as described in [12]) can be obtained as a special case of the above. Recall the definition of a Reed-Solomon code: For integers n and $k \leq n$ and field \mathbb{F} of cardinality $n + 1$, the Reed-Solomon code $\text{RS}_{n,k}$ over \mathbb{F} and block length n is given by

$$\text{RS}_{n,k} = \{(f(x))_{x \in \mathbb{F} - \{0\}} \mid f \text{ is a polynomial of degree } k - 1\}.$$

It is easy to verify that $\text{RS}_{n,k}$ has rate k and distance $n - k + 1$. Notice further that $\text{RS}_{n,k} * \text{RS}_{n,t} \subseteq \text{RS}_{n,k+t-1}$. Thus setting $t = \lfloor \frac{n-k+1}{2} \rfloor$ we notice that the pair $\text{RS}_{n,t}, \text{RS}_{n,k+t-1}$ form a t -error correcting pair for Reed-Solomon codes; and by Theorem 17 has an efficient algorithm for correcting up to t errors. Decoding algorithms for other algebraically specified codes such as the BCH codes, Alternant codes and algebraic-geometric codes follow by similar constructions of error-correcting pairs.

B A Variant of Unique-Decode

As mentioned previously, our algorithm Unique-Decode (presented in Section 3) was inspired by the paradigm described in Appendix A. To gain more intuition about the relation between our algorithm and the general paradigm, we present below a variant of our algorithm, which more directly exhibits this relation. Note that Step 1 of the alternative algorithm is the same as Step 1 of the original algorithm, but the integers y and z found in this step are used differently in the remaining steps. In view of this difference, the goal of Step 1 is interpreted differently. Details follow.

Given a received word $\langle r_1, \dots, r_n \rangle$ that is close to the encoding of (a unique) message m , similarly to what is done in Linear-Decode, the algorithm Unique-Decode' tries to detect the indices for which

$r_i \neq [m]_{p_i}$. It then reconstructs the message, using CRT, from those r_i 's that are assumed to be correct (or more precisely, from $[r]_{p'_i}$, where p'_i is a factor of p_i determined by the algorithm). The above detection is done by finding an integer y that satisfies $[y]_{p_i} = 0$ whenever $r_i \neq [m]_{p_i}$. By restricting y to be relatively small we ensure that $[y]_{p_i}$ does not equal 0 for many i satisfying $r_i = [m]_{p_i}$ (so that CRT can in fact be applied). To find this y , we need some way to (describe and) exploit the fact that there exists some small m s.t., for every i , $[y]_{p_i} = 0$ or $[m]_{p_i} = r_i$; or equivalently $[y]_{p_i} \cdot [m]_{p_i} \equiv [y]_{p_i} \cdot r_i \pmod{p_i}$. The final condition suggest that we may attempt to find $z \stackrel{\text{def}}{=} y \cdot m$ such that $[z]_{p_i} \equiv [y]_{p_i} \cdot r_i \pmod{p_i}$. While ideally we would like to specify further that z is a multiple of y , we relax this and simply use the fact that z is also small (since both y and m are small). This leads to the following algorithm:

Unique-Decode'($p_1, \dots, p_n, k, r_1, \dots, r_n$).

- Set $K = \prod_{i=1}^k p_i$, $N = \prod_{i=1}^n p_i$, and $F = (K - 1)E$, with E to be determined later.
- Let $r \in \mathbb{Z}_N$ be s.t. $r_i = [r]_{p_i}$ (as defined by CRT).

1. Find integers y, z s.t.

$$\left. \begin{array}{l} 1 \leq y \leq E \\ 0 \leq z \leq F \\ y \cdot r \equiv z \pmod{N} \end{array} \right\} \quad (12)$$

- 2. Let $I \stackrel{\text{def}}{=} \{i : [y]_{p_i} \neq 0\}$. For every $i \in I$ let $p'_i = p_i / \gcd(y, p_i)$, and set $x_i = [r]_{p'_i}$. (Note that if p_i is prime, $p'_i = p_i$ and so $x_i = r_i$.)
- 3. Find $x \in \mathbb{Z}_K$ s.t. $[x]_{p'_i} = x_i$ for every $i \in I$ (if such an x exists) and output it.

Similarly to the algorithm Unique-Decode, the above algorithm can be implemented in polynomial time in the bit sizes of p_1, \dots, p_n . Step 1 is as in Unique-Decode, Step 2 is straightforward, and Step 3 is just an application of the Chinese Remainder Theorem (i.e., find $x \in \mathbb{Z}_{\prod_{i \in I} p'_i}$ via CRT and check if it is smaller than K).

Theorem 18 Unique-Decode'($p_1, \dots, p_n, k, r_1, \dots, r_n$) solves the error-correction problem in polynomial time for any value of the error parameter up to $(n - k) \frac{\log p_1}{\log p_1 + \log p_n}$, with the setting $E = \prod_{i=n-e+1}^n p_i$.

The proof of Theorem 18 follows from the next lemma exactly as the proof of Theorem 6 follows from Lemma 5.

Lemma 19 If r is such that for some $m \in \mathbb{Z}_K$ the amplitude of the distance between (r_1, \dots, r_n) and $([m]_{p_1}, \dots, [m]_{p_n})$ is at most E , and $N > E^2 \cdot K$ then Unique-Decode'($p_1, \dots, p_n, k, r_1, \dots, r_n$) returns m .

Proof: We first observe that Claims 5.1 and 5.2, which were used to prove Lemma 5 hold here as well since Step 1 of Unique-Decode' is identical to Step 1 of Unique-Decode. By Claim 5.1, Step 1 of the algorithm always returns a pair (y, z) satisfying Eq. (1). By Claim 5.2, any pair (y, z) that may be the outcome of Step 1 satisfies $y \cdot m = z$. Since $y \cdot r \equiv z \pmod{N}$, it follows that for every i , $y \cdot r \equiv y \cdot m \pmod{p_i}$, and so for every $i \in I$ (where I is as defined in Step 2), we have $r \equiv m \pmod{p'_i}$ (since y has a multiplicative inverse modulo p'_i). Thus, $m \in \mathbb{Z}_K$ is a valid solution for the task in Step 3 (since $\forall i \in I$, $[m]_{p'_i} = [r]_{p'_i} = x_i$).

It remains to show that m is the only possible solution. For this, let $\bar{I} = \{1, \dots, n\} \setminus I$ (i.e., $\forall i \in \bar{I}$, $y \equiv 0 \pmod{p_i}$). Let $q_i \stackrel{\text{def}}{=} p_i/p'_i$, and observe (by definition of p'_i) that $y \equiv 0 \pmod{q_i}$ for every $i \in I$. (Note that for the special case in which the p_i 's are prime numbers (and not only relatively prime), $p'_i = p_i$ and so $q_i = 1$). By CRT, $y \equiv 0 \pmod{\prod_{i \in \bar{I}} p_i \prod_{i \in I} q_i}$, and since $y > 0$ it follows that $y \geq \prod_{i \in \bar{I}} p_i \prod_{i \in I} q_i$. Since $y \leq E$, we have $\prod_{i \in \bar{I}} p_i \prod_{i \in I} q_i \leq E$ and

$$\prod_{i \in I} p'_i = \frac{N}{\prod_{i \in \bar{I}} p_i \prod_{i \in I} q_i} \geq \frac{N}{E} > K$$

follows. Thus, again by CRT, the message m is the only solution in \mathbb{Z}_K to the system $\{x \equiv r \pmod{p'_i}\}_{i \in I}$. ■