

Topics in Error-Correcting Codes

Notes by Rohan Goyal

MIT — 6.S953

Lecture 10: Expander Codes

Today we study the Sipser–Spielman construction: binary linear codes of constant rate and linear distance with a linear-time algorithm for correcting a constant fraction of adversarial errors.

The Graph View of Parity Checks

Recall our convention that codewords are row vectors. A matrix $H \in \mathbb{F}_2^{n \times m}$ defines a linear code

$$\mathcal{C} = \{x \in \mathbb{F}_2^n : xH = 0\}.$$

The columns of H specify parity checks. If the transmitted word is $x \in \mathcal{C}$ and the received word is $y = x + e$, then

$$yH = eH.$$

Thus a nonzero entry of the *syndrome* yH identifies a check containing an odd number of errors. In particular, at least one of the positions participating in that check is incorrect.

A check involving only a few positions gives local information about the errors. This suggests using a sparse parity-check matrix, as in Gallager’s *low-density parity-check* (LDPC) codes.

Variables and constraints

View H as the adjacency matrix of a bipartite graph $G = (L, R, E)$:

- $L = \{v_1, \dots, v_n\}$ contains the *variable vertices*, one for each coordinate;
- $R = \{r_1, \dots, r_m\}$ contains the *check vertices*, one for each parity check;
- $(v_i, r_j) \in E$ exactly when $H_{i,j} = 1$.

An assignment $x \in \mathbb{F}_2^n$ is a codeword precisely when every check sees an even number of 1’s among its neighbors.

Gallager’s *ensemble* [1] samples sparse parity checks at random, with each bit participating in a fixed number of checks. With high probability, this gives good binary codes whose rate and distance approach the Gilbert–Varshamov bound arbitrarily closely as the constant degrees are chosen larger. Gallager also proposed simple iterative decoders. Sipser–Spielman’s contribution is an expansion-based analysis guaranteeing efficient correction of worst-case errors.

An MIT connection. This construction grew out of work at MIT by Michael Sipser and his then Ph.D. student Daniel Spielman. Expander codes formed part of [Spielman’s 1995 MIT thesis](#).

For example, consider

$$H = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \begin{aligned} r_1 : & \quad x_1 + x_2 + x_4 = 0, \\ r_2 : & \quad x_2 + x_3 + x_5 = 0, \\ r_3 : & \quad x_3 + x_4 + x_6 = 0. \end{aligned}$$

The word $x = (1, 1, 0, 0, 1, 0)$ satisfies all three checks. Flipping its second bit gives $y = (1, 0, 0, 0, 1, 0)$, for which $yH = (1, 1, 0)$.




Figure 1. Circles contain variable values; squares contain the resulting check parities. Changing the second bit toggles exactly its two neighboring checks. This small graph illustrates the representation; the expansion assumptions will be imposed below.

A brief connection to Tanner codes

Tanner’s viewpoint allows a check vertex of degree d to require that its ordered list of neighboring symbols belong to a fixed short linear code $\mathcal{C}_0 \subseteq \mathbb{F}_2^d$. For example, seven neighboring bits could be required to form a word of the $[7, 4, 3]$ Hamming code. Taking

$$\mathcal{C}_0 = \{z \in \mathbb{F}_2^d : z_1 + \cdots + z_d = 0\}$$

recovers a single parity check. We will focus on these parity-check constraints and their Sipser–Spielman decoder.

Vertex Expansion and Unique Neighbors

Throughout, graphs are simple and bipartite. Let $|L| = n$ and $|R| = m$. A graph is (c, d) -regular if every left vertex has degree c and every right vertex has degree d . Counting edges gives

$$cn = dm, \quad \frac{m}{n} = \frac{c}{d}.$$

We will use fixed degrees $c < d$, so there are fewer checks than variables.

Definition 1 (Vertex expansion). For $S \subseteq L$, its neighborhood is

$$\Gamma(S) = \{r \in R : r \text{ has a neighbor in } S\}.$$

Lossless expansion: another MIT connection. Since $|\Gamma(S)| \leq c|S|$, a near-optimal expansion guarantee is

$$|\Gamma(S)| \geq (1 - \epsilon)c|S|.$$

We require this for every $|S| \leq \alpha n$. For any fixed $\epsilon > 0$, we seek constant degrees and some constant $\alpha > 0$ with this property, called *lossless expansion*. Such graphs have long been known to exist probabilistically, and explicit one-sided constructions were already known. In 2025, Hsieh, Lubotzky, Mohanty, Reiner, and Zhang gave the first [explicit constant-degree two-sided lossless expanders](#), with near-maximal expansion for small sets on both sides. The work took place primarily at MIT: Jun-Ting (Tim) Hsieh, our guest lecturer in

Let $0 < \alpha < 1$ and $\gamma > 0$. The graph is an (α, γ) -vertex expander on the left if

$$|\Gamma(S)| \geq \gamma|S| \quad \text{whenever } S \subseteq L, \quad |S| \leq \alpha n.$$

There are exactly $c|S|$ edges leaving S , but several may meet the same right vertex. Expansion asks for many *distinct* neighbors. The requirement applies to every small set S , not just to individual vertices. We take graphs with the stated properties as given.




Figure 2. Both sets send out six edges. On the right, those edges reach more distinct vertices. Green vertices have one neighbor in S ; orange vertices have two. Only edges leaving S are shown. These pictures compare the neighborhoods of one set, rather than certify expansion for the whole graph.

Definition 2 (Unique neighbors). A right vertex is a *unique neighbor of S* if it has exactly one neighbor in S . Write

$$\Gamma_1(S) = \{r \in R : |\Gamma(r) \cap S| = 1\}.$$

A unique neighbor may have other neighbors outside S . Uniqueness is always relative to the set S .

Lemma 3 (Vertex expansion gives unique neighbors). *Suppose every left vertex has degree c . For every $S \subseteq L$,*

$$|\Gamma_1(S)| \geq 2|\Gamma(S)| - c|S|.$$

In particular, if $|S| \leq \alpha n$ in an (α, γ) -vertex expander, then

$$|\Gamma_1(S)| \geq (2\gamma - c)|S|.$$

Proof. Let $U = \Gamma_1(S)$ and $T = \Gamma(S) \setminus U$. Each vertex of U receives one edge from S , and each vertex of T receives at least two. Hence

$$c|S| \geq |U| + 2|T| = 2|\Gamma(S)| - |U|.$$

Rearranging proves the first bound; substituting $|\Gamma(S)| \geq \gamma|S|$ proves the second. $\square$

We will use this lemma twice. When $\gamma > c/2$, every nonempty small set has a unique neighbor. When $\gamma > 3c/4$, it has more than $c|S|/2$ unique neighbors. The first statement gives distance; the second will give an efficient decoder.

The Sipser–Spielman Code: Rate and Distance

Given G , define its binary parity-check code by

$$\mathcal{C}(G) = \left\{ x \in \mathbb{F}_2^n : \sum_{v_i \in \Gamma(r)} x_i = 0 \text{ for every } r \in R \right\}.$$

All sums are over $\mathbb{F}_2$. The parity-check matrix has one column per right vertex, as in Section 1.

Theorem 4 (Rate and distance). *Let G be a (c, d) -regular bipartite graph with $c < d$ and n left vertices. Suppose G is an (α, γ) -vertex expander with $\gamma > c/2$. Then*

$$\text{rate}(\mathcal{C}(G)) \geq 1 - \frac{c}{d}, \quad d_{\min}(\mathcal{C}(G)) > \alpha n.$$

Proof. There are m homogeneous linear constraints, so

$$\dim \mathcal{C}(G) = n - \text{rank}(H) \geq n - m.$$

Divide by n and use $m/n = c/d$ for the rate bound. The checks need not be linearly independent.

For distance, suppose a nonzero codeword x has weight at most αn , and let $S = \{v_i : x_i = 1\}$. By Lemma 3,

$$|\Gamma_1(S)| \geq (2\gamma - c)|S| > 0.$$

A check in $\Gamma_1(S)$ sees exactly one 1, so its parity is 1, contradicting $x \in \mathcal{C}(G)$. Thus every nonzero codeword has weight greater than αn . Since the code is linear, this is also a lower bound on its minimum distance. $\square$

For a family of such graphs with fixed c, d, α , these are asymptotically good binary codes. We next show that a stronger expansion condition supplies a very simple decoding algorithm.

Sipser–Spielman Bit-Flipping Decoding

Let y be the received word. A check is *unsatisfied* if the parity of its current neighboring bits is 1. If a variable has more unsatisfied than satisfied neighboring checks, flipping that variable reduces the total number of unsatisfied checks.

Sequential bit flipping

1. Set $z \leftarrow y$.
2. While some variable has more than $c/2$ unsatisfied neighboring checks, flip its bit in z .
3. If all checks are satisfied, return z ; otherwise report failure.

The choice of a flippable variable can be arbitrary.




Figure 3. Flipping a variable toggles all of its neighboring checks. Here $c = 5$ and three checks fail before the flip. The other variables adjacent to the checks are omitted. All checks outside this neighborhood are unchanged.

We use expansion to show that this process recovers the transmitted codeword.

Theorem 5 (Sipser–Spielman decoding). *Let G be a (c, d) -regular (α, γ) -vertex expander, where c, d are constants and $\gamma > 3c/4$. Set $s = \lfloor \alpha n \rfloor \geq 1$. Given a word y differing from $x \in \mathcal{C}(G)$ in at most $\lfloor s/2 \rfloor$ positions, sequential bit flipping recovers x in $O(n)$ time, given adjacency lists for G .*

Proof. Every flip makes more failed checks satisfied than satisfied checks fail. Thus the number of failed checks decreases by at least one at every step, and the algorithm stops after at most m flips.

Let S be the set of positions where the current word differs from the transmitted codeword x . Whenever $0 < |S| \leq s$, Lemma 3 gives

$$|\Gamma_1(S)| \geq (2\gamma - c)|S| > \frac{c}{2}|S|.$$

Every unique neighbor of S is a failed check: it contains exactly one error. By averaging, some variable in S has more than $c/2$ failed neighboring checks and can be flipped. Hence the algorithm cannot stop with between 1 and s errors.

A flip might change a correct bit, so we must also check that the number of errors stays small. Initially there are at most $\lfloor s/2 \rfloor$ errors, and therefore at most $c\lfloor s/2 \rfloor \leq cs/2$ failed checks. Each flip changes the number of errors by one. If this number ever reached s , the unique-neighbor bound would give more than $cs/2$ failed checks, contradicting the fact that their number only decreases. Thus there are always fewer than s errors. At termination there can be no errors left, so the output is x .

For fixed degrees, the algorithm can be implemented with $O(n)$ initialization and $O(1)$ amortized time per flip. Since $m = cn/d = O(n)$, the total running time is $O(n)$. $\square$

References and Further Reading

- [1] Robert G. Gallager. [Low-density parity-check codes](#). *IRE Transactions on Information Theory*, 8(1):21–28, 1962.
- [2] Michael Sipser and Daniel A. Spielman. [Expander codes](#). *IEEE Transactions on Information Theory*, 42(6):1710–1722, 1996.

- [3] Anup Rao. [Lecture 5: Expander codes](#). Lecture notes for *CSE 599: Modern Coding Theory*, University of Washington, October 9, 2019.
- [4] Jun-Ting Hsieh, Alexander Lubotzky, Sidhanth Mohanty, Assaf Reiner, and Rachel Yun Zhang. [Explicit lossless vertex expanders](#). In *Proceedings of the 66th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 894–911. IEEE Computer Society, 2025. [arXiv:2504.15087](#).
- [5] Madhu Sudan. [Lecture 17: Graph-based codes](#). Lecture notes for *6.440: Essential Coding Theory*, Massachusetts Institute of Technology, April 9, 2008.
- [6] Madhu Sudan. [Lecture 13: Concluding folded RS codes and graph codes](#). Lecture notes for *CS 229r: Essential Coding Theory*, Harvard University, October 23, 2023. Scribed by Yeongwoo Hwang.