Erasure Correction and Locality of Hypergraph Codes

Allison Beemer, Carolyn Mayer*, and Christine A. Kelley

University of Nebraska-Lincoln, Lincoln, NE, USA
cmayer@huskers.unl.edu

Abstract. We examine erasure correction and locality properties of regular and biregular hypergraph codes. We propose a construction of t-uniform t-partite biregular hypergraphs based on (c, d)-regular bipartite graphs. We show that for both the regular and biregular case, when the underlying hypergraph has expansion properties, the guaranteed erasure correcting capability for the resulting codes is improved. We provide bounds on the minimum stopping set size and cooperative locality of these hypergraph codes for application in distributed storage systems.

Keywords: Hypergraphs · Generalized LDPC codes · Locality · (r, ℓ)-cooperative locality · Availability · Stopping sets · Expander graphs

1 Introduction

The past decade has seen an increased interest in coding for distributed storage systems (DSS) due to the increasing amounts of data that need to be stored and accessed across many servers. A primary focus in this area is the design of codes with locality properties, where error correction of small sets of symbols may be performed efficiently without having to access all symbols or all information from accessed symbols, and where data may be protected by multiple repair groups. Coding techniques optimizing different aspects of the storage and retrieval problem yield different families of codes, such as batch codes, locally repairable codes, private information retrieval codes, and local regeneration codes, [1–3], with connections among them [4, 5].

Codes from regular hypergraphs with expansion-like properties were introduced and analyzed in [6] and [7]. As was shown for expander codes that have underlying expander graph representations, the authors show that better expansion (referred to as ϵ-homogeneity in the hypergraph case) implies improved minimum distance and error correction. Indeed, [8] gives a construction of one of the few explicit algebraic families of asymptotically good codes known to date, and there have since been several papers addressing the design and analysis of codes using expander graphs [9, 10]. In this paper, we consider codes based on regular hypergraphs (i.e. all vertices have the same degree) as well as biregular hypergraphs (i.e. vertices have one of two distinct degrees), and present bounds on their error-correction capabilities in the context of distributed storage, specifically the minimum stopping set size and cooperative locality of the codes.

This paper is organized as follows. Section 2 presents relevant background. We derive bounds on the minimum stopping set size and cooperative locality of regular hypergraph codes in Section 3. In Section 4, we show how to construct a t-uniform t-partite biregular hypergraph from an underlying (c, d)-regular bipartite graph and provide similar bounds. Section 5 concludes the paper.

2 Preliminaries

A *hypergraph* $\mathcal{H}$ is a set of vertices V along with a set of edges E, where E is a set of subsets of V. A hypergraph is *t-uniform* if every edge contains exactly t vertices, and it is *t-partite* if the vertex set V can be partitioned into t vertex sets $V_1, \dots, V_t$ such that no edge contains more than one vertex from any part. In this paper we will use $\mathcal{H} = (V_1, V_2, \dots, V_t; E)$ to denote a t-uniform t-partite hypergraph in which each edge contains exactly one vertex from each part. A hypergraph is *Δ-regular* if every vertex belongs to Δ edges. We assume that all hypergraphs considered in this paper have no parallel edges. That is, no two distinct edges are comprised of exactly the same set of vertices.

Such t-uniform t-partite Δ-regular hypergraphs were used in [6, 7] to design codes where the edges of the hypergraph represent the code symbols, and the vertices of the hypergraph represent the constraints. Specifically, when there are n vertices in each part, the block length of the corresponding *hypergraph code* is $n\Delta$ and the number of constraint vertices is nt. As in [11], each constraint node represents a linear block length Δ code (commonly referred to as a "subcode"), and is satisfied when the assignment on the edges incident to that constraint node form a codeword of the subcode. This is similar to generalized low density parity-check (LDPC) codes as defined in [11], in which the constraint nodes may be more sophisticated than simple parity-check codes.

Sipser and Spielman [8] show that the guaranteed error correction capabilites of a code may be improved when the underlying graph is a good expander. Let G be a (c, d)-regular bipartite graph with m left vertices of degree c and n right vertices of degree d, and let μ be the second largest eigenvalue (in absolute value) of the adjacency matrix of G. Then for subsets S and T of the left and right vertices, respectively, the number of edges in the subgraph induced by $S \cup T$ is at most $|E(S,T)| \leq \frac{d}{m}|S||T| + \frac{\mu}{2}(|S| + |T|)$ [12]. Loosely speaking, a graph is a good expander if small sets of vertices have large sets of neighbors; graphs with small second largest eigenvalue have this property. Bilu and Hoory introduced an analogous notion for hypergraph expansion as follows:

Definition 1. *[6] Let $\mathcal{H} = (V_1, V_2, \dots, V_t; E)$ be a t-uniform t-partite Δ-regular hypergraph with n vertices in each part. Then $\mathcal{H}$ is* ϵ-homogeneous *if for every choice of $A_1, A_2, \dots, A_t$ with $A_i \subseteq V_i$ and $|A_i| = \alpha_i n$,*

$$\frac{|E(A_1, A_2, \dots, A_t)|}{n\Delta} \leq \prod_{i=1}^{t} \alpha_i + \epsilon\sqrt{\alpha_{\sigma(1)}\alpha_{\sigma(2)}},$$

where $\sigma \in S_t$ is a permutation on $[t]$ such that $\alpha_{\sigma(i)} \leq \alpha_{\sigma(i+1)}$ for each $i \in [t-1]$, and $E(A_1, \ldots, A_t)$ denotes the set of edges which intersect all of the A_i's.

Let $[N, K, D]$ denote a binary linear code with block length N, dimension K, and minimum distance D. The following bounds on the rate and minimum distance of a code $\mathcal{Z}$ from an ϵ-homogeneous t-uniform t-partite Δ-regular hypergraph with n vertices in each part and a $[\Delta, \Delta R, \Delta\delta]$ subcode C at each constraint node are given in [6]:

$$\text{rate}(\mathcal{Z}) \geq tR - (t-1) \tag{1}$$

$$d_{\min}(\mathcal{Z}) \geq n\Delta \left(\delta^{\frac{t}{t-1}} - c(\epsilon, \delta, t)\right) \tag{2}$$

where $c(\epsilon, \delta, t) \to 0$ as $\epsilon \to 0$. Note that $n\Delta$ is the blocklength of $\mathcal{Z}$.

The *locality* of a code measures how many code symbols must be used to recover an erased code symbol. While there are a variety of locality notions relevant to coding for distributed storage, in this paper we focus on (r, ℓ)-cooperative locality and (r, τ)-availability [13, 14].

Definition 2. *A code C has* (r, ℓ)-cooperative locality *if for any $\mathbf{y} \in C$, any set of ℓ symbols in $\mathbf{y}$ are functions of at most r other symbols. Furthermore, C has* (r, τ)-availability *if any symbol in $\mathbf{y}$ can be recovered by using any of τ disjoint sets of symbols each of size at most r.*

3 Bounds on Regular Hypergraph Codes

In this section, we examine the erasure correction and cooperative locality of regular hypergraph codes. First, a *stopping set* in the Tanner graph representing a generalized LDPC code is a subset S of the variable nodes such that each neighbor of S has at least $d_{\min}(C)$ neighbors in S, where C is the subcode represented by the constraint vertices [15, 16]. It was initially shown in [17] that stopping sets characterize iterative decoder failure for simple LDPC codes (i.e. when the subcodes are simple parity check codes) on the binary erasure channel. We first define stopping sets for regular hypergraph codes and give a lower bound on the minimum stopping set size.

Definition 3. *Let $\mathcal{Z}$ be a code on a hypergraph $\mathcal{H} = (V_1, \ldots, V_t; E)$, with edges representing code symbols and vertices representing the constraints of a subcode C. Then a* stopping set *S is a subset of the edges of $\mathcal{H}$ such that every vertex contained in an element of S is contained in at least $d_{\min}(C)$ elements of S.*

Though the size of a minimum stopping set depends on both the hypergraph representation and the choice of subcode, we denote this size by $s_{\min}(\mathcal{H})$, and assume that the subcode is clear from context.

Theorem 1. *Let $\mathcal{H}$ be a t-uniform t-partite Δ-regular hypergraph. If the vertices of $\mathcal{H}$ represent constraints of a subcode C with minimum distance $d_{\min}(C)$ and block length Δ, then the size of the minimum stopping set, $s_{\min}(\mathcal{H})$, is bounded by*

$$s_{\min}(\mathcal{H}) \geq d_{\min}(C)^{t/(t-1)}. \tag{3}$$

Proof. Let $\mathcal{H}$ be as above, and let S be a minimum stopping set. Each edge in S contains exactly one constraint node from each of the t parts of $\mathcal{H}$, so each part of $\mathcal{H}$ has exactly $|S| = s_{\min}(\mathcal{H})$ incident edges belonging to S. Each constraint node contained in an edge in S must be contained in at least $d_{\min}(C)$ edges in S. By the pigeonhole principle, the number of vertices in any part of $\mathcal{H}$ that are contained in some edge in S is bounded above by $s_{\min}(\mathcal{H})/d_{\min}(C)$. Indeed, were there more than $s_{\min}(\mathcal{H})/d_{\min}(C)$ vertices incident to S in a single part, some vertex must have fewer than $s_{\min}(\mathcal{H})/(s_{\min}(\mathcal{H})/d_{\min}(C)) = d_{\min}(C)$ incident edges from S, a contradiction. Now consider the maximum size of S: this amounts to counting the number of edges possible, given that each edge is incident to exactly one vertex of (at most) $s_{\min}(\mathcal{H})/d_{\min}(C)$ vertices in each of the t parts of $\mathcal{H}$. That is, there are at most $(s_{\min}(\mathcal{H})/d_{\min}(C))^t$ edges in S. Thus,

$$\left(\frac{s_{\min}(\mathcal{H})}{d_{\min}(C)}\right)^t \geq s_{\min}(\mathcal{H}) \Rightarrow s_{\min}(\mathcal{H}) \geq d_{\min}(C)^{t/(t-1)}.$$

The bound of Theorem 1 is tight. For example, when $\mathcal{H}$ is a complete 3-uniform 3-partite hypergraph with at least two vertices in each part and constraint code C such that $d_{\min}(C) = 4$, it is easy to show that $s_{\min}(\mathcal{H}) = 8$.

Since the errors of particular relevance to DSS are erasures (such as a server going down), we can use the stopping set bound to characterize how many errors can be corrected. Theorem 1 guarantees that we may correct any $d_{\min}(C)^{t/(t-1)} - 1$ erasures using iterative decoding. If C is a code with locality r_1, at most $(s_{\min}(\mathcal{H})/d_{\min}(C)) \cdot r_1 \cdot t$ other codeword symbols are involved in the repair of the erasures in the decoding process. This yields the following:

Corollary 1. *If the subcodes C of the regular hypergraph code $\mathcal{Z}$ have r_1 locality, then $\mathcal{Z}$ has (r, ℓ)-cooperative locality where*

$$r = r_1 t s_{\min}(\mathcal{H})/d_{\min}(C) \tag{4}$$

$$s_{\min}(\mathcal{H}) - 1 \geq \ell \geq d_{\min}(C)^{t/(t-1)} - 1. \tag{5}$$

Observe that if the subcode C has (r, τ)-availability, then the hypergraph code $\mathcal{Z}$ has at least (r, τ)-availability.

We now extend the result to codes on hypergraphs with known ϵ-homogeneity.

Theorem 2. *Let $\mathcal{H} = (V_1, V_2, \ldots, V_t; E)$ be a t-uniform t-partite Δ-regular ϵ-homogeneous hypergraph where there are n vertices in each of the t parts. If the subcodes C have minimum distance $d_{\min}(C)$,*

$$s_{\min}(\mathcal{H}) \geq \left(\left(1 - \frac{\epsilon\Delta}{d_{\min}(C)}\right) \frac{n^{t-1} d_{\min}(C)^t}{\Delta} \right)^{1/(t-1)}. \tag{6}$$

For $\epsilon < \frac{d_{\min}(C)(n^{t-1}-\Delta)}{\Delta n^{t-1}}$, this gives an improvement on the bound in Theorem 1.

Proof. Let S be a minimum stopping set. By Theorem 1, $s_{\min}(\mathcal{H}) \geq d_{\min}(C)^{t/(t-1)}$. Now, let $A_i \subseteq V_i$ be the set of vertices in V_i, for $i \in [t]$, contained in an edge in S. By ϵ-homogeneity,

$$s_{\min}(\mathcal{H}) \;=\; |S| \leq |E(A_1, \ldots, A_t)| \leq n\Delta \left(\prod_{i=1}^{t} \alpha_i + \epsilon \sqrt{\alpha_{\sigma(1)}\alpha_{\sigma(2)}} \right).$$

Since $|A_i| \leq s_{\min}(\mathcal{H})/d_{\min}(C)$ for all i, $\alpha_i \leq s_{\min}(\mathcal{H})/n d_{\min}(C)$. Thus, the above inequality simplifies to obtain the result:

$$s_{\min}(\mathcal{H}) \leq n\Delta \left(\left(\frac{s_{\min}(\mathcal{H})}{n d_{\min}(C)} \right)^t + \epsilon \frac{s_{\min}(\mathcal{H})}{n d_{\min}(C)} \right).$$

Observe that we have shown in general that

$$s_{\min}(\mathcal{H}) \geq \left(\left(1 - \frac{\epsilon\Delta}{d_{\min}(C)}\right) \frac{n^{t-1}}{\Delta} \right)^{1/(t-1)} d_{\min}(C)^{t/(t-1)}.$$

Then, if $\left(\left(1 - \frac{\epsilon\Delta}{d_{\min}(C)}\right) \frac{n^{t-1}}{\Delta} \right)^{1/(t-1)} > 1$, this gives a better lower bound for $s_{\min}(\mathcal{H})$ than that found in Theorem 1. Simplifying, we have our condition on ϵ.

Corollary 2. *Using iterative decoding on a code $\mathcal{Z}$ based on a t-uniform t-partite Δ-regular ϵ-homogeneous hypergraph with vertices representing constraints of a subcode C, up to*

$$\left(\left(1 - \frac{\epsilon\Delta}{d_{\min}(C)}\right) \frac{n^{t-1} d_{\min}(C)^t}{\Delta} \right)^{1/(t-1)} - 1$$

erasures may be corrected.

In other words, if δ is the relative minimum distance of C, and N is the total number of edges in the hypergraph (that is, the block length of the code $\mathcal{Z}$), we may correct up to a $\delta(\delta - \epsilon)^{1/(t-1)} - \frac{1}{N}$ fraction of erasures.

Remark 1. We may correct up to a $\delta^{t/(t-1)} - \frac{1}{N} - c(\epsilon, \delta, t)$ fraction of erasures, where $c(\epsilon, \delta, t) \to 0$ as $\epsilon \to 0$. It can be shown that the bound in Corollary 2 improves the error correction capability of

$$\binom{t-1}{t/2}^{-2/t} \left(\frac{\delta}{2}\right)^{(t+2)/t} - c'(\epsilon, \delta, t)$$

in [6] for any $0 < \delta < 1$ and $t \geq 2$ (i.e. for all relevant cases) and large block length. Note that $c'(\epsilon, \delta, t) \neq c(\epsilon, \delta, t)$, but that both vanish as $\epsilon \to 0$. It is important to note that we are focusing solely on erasures, while [6] gives a decoding algorithm and correction capabilities for more general errors.

4 Bounds on Biregular Hypergraph Codes

A construction of t-uniform t-partite Δ-regular hypergraphs is presented in [6] based on an underlying regular bipartite expander graph. In this section we show how to obtain a t-uniform t-partite (Δ_1, Δ_2)-biregular hypergraph from a (c, d)-regular bipartite graph in a similar way. We provide bounds on the stopping set size, cooperative locality, rate, and minimum distance for the resulting hypergraph codes.

Definition 4. *We say that a t-uniform t-partite hypergraph $\mathcal{H} = (V_1, \ldots, V_t; E)$ is* (Δ_1, Δ_2)-biregular *if the parts can be labeled such that each vertex in an odd (resp., even) index part is contained in Δ_1 (resp., Δ_2) edges.*

Construction 1. *Let $G = V \cup W$ be a (c, d)-regular bipartite expander graph with $|V| \geq |W|$. For $t \in \mathbb{N}$, construct a t-uniform t-partite hypergraph $\mathcal{H}$ with parts $V_1, \ldots, V_t$ as follows. For odd (resp., even) i, let V_i be a copy of V (resp., W). Take $E(\mathcal{H})$ to be the set of edges corresponding to walks of length $t-1$ in G. That is $(v_1, \ldots, v_t)$ with $v_i \in V_i$ is in $E(\mathcal{H})$ if and only if $(v_1, \ldots, v_t)$ corresponds to a walk in G.*

Note that $\mathcal{H}$ is indeed t-uniform and t-partite, and has vertices of degree $\Delta_1 = c^{\lceil \frac{t}{2} \rceil} d^{\lfloor \frac{t}{2} \rfloor - 1}$ (resp., $\Delta_2 = c^{\lceil \frac{t}{2} \rceil - 1} d^{\lfloor \frac{t}{2} \rfloor}$) in odd (resp., even) index parts.

The definition of a stopping set may be extended to biregular hypergraph codes.

Definition 5. *Let $\mathcal{Z}$ be a code on a hypergraph $\mathcal{H} = (V_1, V_2, \ldots, V_t; E)$, with the edges representing the code symbols and the vertices representing the constraints of a subcode C_1 (resp., C_2) if the vertex is in an odd (resp., even) index part. Then a* stopping set *S is a subset of the edges of $\mathcal{H}$ such that every vertex contained in an element of S is contained in at least $d_{\min}(C_1)$ (resp., $d_{\min}(C_2)$) elements of S if the vertex is in an odd (resp., even) index part.*

We now give a bound on the minimum stopping set size and (r, ℓ)-cooperative locality of codes resulting from Construction 1. The proofs are similar to those in Section 3 and are thus omitted.

Theorem 3. *Let $\mathcal{H}$ be a t-uniform t-partite (Δ_1, Δ_2)-biregular hypergraph. If the vertices in an odd (resp., even) index part of $\mathcal{H}$ represent constraints of a subcode C_1 (resp., C_2) with block length Δ_1 (resp., Δ_2), then the size of the minimum stopping set, $s_{\min}(\mathcal{H})$, is bounded by*

$$s_{\min}(\mathcal{H}) \geq \left(d_{\min}(C_1)^{\lceil \frac{t}{2} \rceil} d_{\min}(C_2)^{\lfloor \frac{t}{2} \rfloor} \right)^{1/(t-1)}. \tag{7}$$

Corollary 3. *If the subcodes C_1 (resp., C_2) of the biregular hypergraph code $\mathcal{Z}$ have r_1 (resp., r_2) locality then $\mathcal{Z}$ has (r, ℓ)-cooperative locality where*

$$r = r_1 \lceil \frac{t}{2} \rceil \frac{s_{\min}(\mathcal{H})}{d_{\min}(C_1)} + r_2 \lfloor \frac{t}{2} \rfloor \frac{s_{\min}(\mathcal{H})}{d_{\min}(C_2)} \tag{8}$$

$$s_{\min}(\mathcal{H}) - 1 \geq \ell \geq \left(d_{\min}(C_1)^{\lceil \frac{t}{2} \rceil} d_{\min}(C_2)^{\lfloor \frac{t}{2} \rfloor} \right)^{1/(t-1)} - 1. \tag{9}$$

Observe that $\mathcal{Z}$ has at least the (r, τ)-availability of its subcodes.

We next extend the definition of ϵ-homogeneity to biregular hypergraphs and give an improved minimum stopping set bound for the corresponding codes.

Definition 6. *Let $\mathcal{H} = (V_1, V_2, \ldots, V_t; E)$ be a t-uniform t-partite (Δ_1, Δ_2)-biregular hypergraph with n_1 vertices in each odd index part and n_2 vertices in each even index part. We say that $\mathcal{H}$ is ϵ-*homogeneous *if for every choice of $A_1, A_2, \ldots, A_t$, with $A_i \subseteq V_i$,*

$$\frac{|E(A_1, A_2, \ldots A_t)|}{\Delta_1 n_1} \leq \prod_{i=1}^{t} \alpha_i + \epsilon \sqrt{\alpha_{\sigma(1)} \alpha_{\sigma(2)}},$$

where σ is a permutation on $[t]$ such that $\alpha_{\sigma(i)} \leq \alpha_{\sigma(i+1)}$ for each $i \in [t-1]$, and $|A_i| = \alpha_i n_1$ if i is odd and $|A_i| = \alpha_i n_2$ if i is even.

Theorem 4. *Let $\mathcal{H} = (V_1, \ldots, V_t; E)$ be a t-uniform t-partite (Δ_1, Δ_2)-regular ϵ-homogeneous hypergraph where there are n_1 (resp., n_2) vertices in each of the odd (resp., even) index parts. Let C_1 and C_2 be the subcodes of the odd and even index parts, respectively. Then $s_{\min}(\mathcal{H})$ is bounded below by*

$$\left(\frac{(n_1 d_{\min}(C_1))^{\lceil \frac{t}{2} \rceil} (n_2 d_{\min}(C_2))^{\lfloor \frac{t}{2} \rfloor}}{n_1 \Delta_1} \left(1 - \frac{\epsilon n_1 \Delta_1}{\min_{i=1,2}\{n_i d_{\min}(C_i)\}} \right) \right)^{\frac{1}{t-1}}. \tag{10}$$

For $\epsilon < \left(1 - \frac{n_1 \Delta_1}{n_1^{\lceil \frac{t}{2} \rceil} n_2^{\lfloor \frac{t}{2} \rfloor}} \right) \frac{\min_{i=1,2}\{n_i d_{\min}(C_i)\}}{n_1 \Delta_1}$, this improves the Theorem 3 bound.

We now give bounds on the rate and minimum distance of a length $n_1 \Delta_1$ code $\mathcal{Z}$ from an ϵ-homogeneous t-uniform t-partite (Δ_1, Δ_2)-regular hypergraph with n_1 (resp., n_2) vertices in each odd (resp., even) index part and $[\Delta_1, \Delta_1 R_1, \Delta_1 \delta_1]$ subcodes C_1 (resp., $[\Delta_2, \Delta_2 R_2, \Delta_2 \delta_2]$ subcodes C_2).

$$\text{rate}(\mathcal{Z}) \geq R_1 \lceil \frac{t}{2} \rceil + R_2 \lfloor \frac{t}{2} \rfloor - (t-1) \tag{11}$$

$$d_{\min}(\mathcal{Z}) \geq n_1 \Delta_1 \left((\delta_1^{\lceil \frac{t}{2} \rceil} \delta_2^{\lfloor \frac{t}{2} \rfloor})^{\frac{1}{t-1}} - c(\epsilon, \delta_1, \delta_2, t) \right) \tag{12}$$

where $c(\epsilon, \delta_1, \delta_2, t) \to 0$ as $\epsilon \to 0$. The proofs are similar to those for the regular hypergraph bounds given in [6] and are thus omitted.

Ramanujan graphs have the largest possible gap between their first and second largest eigenvalues (in absolute value), and are thus the best in terms of expansion. In [18], the authors give an explicit algebraic construction of regular Ramanujan graphs. The existence of infinite families of (c, d)-regular bipartite Ramanujan graphs was shown in [19]. Moreover, the regular t-uniform t-partite hypergraphs constructed in [6] from regular expander graphs with second largest eigenvalue λ were shown to be $2(t-1)\lambda$-homogeneous. We conjecture that when Construction 1 starts with a (c, d)-regular bipartite expander graph, the resulting hypergraph will be ϵ-homogeneous, where ϵ depends on the second largest eigenvalue of the underlying expander graph.

5 Conclusions

We examined the erasure correcting capability of regular hypergraph codes as well as a new code construction based on biregular hypergraphs. We provided lower bounds on the minimum stopping set size, cooperative locality, and availability of these codes, and gave improved parameters for ϵ-homogeneous hypergraph codes. Designing explicit families of hypergraphs that are optimal with respect to their corresponding codes' locality properties is a line for future work. An interesting open problem is whether the proposed construction of biregular hypergraphs can be proven to yield biregular ϵ-homogeneous hypergraphs when starting from biregular expander graphs.

Acknowledgements We thank the anonymous reviewers for their useful comments that improved the quality of the paper.

References

1. Y. Ishai, E. Kushilevitz, R. Ostrovsky and A. Sahai, Batch codes and their applications, *ACM Symp. on Thy. of Comp.*, pp. 262-271, June 2004, Chicago, IL.
2. A.G. Dimakis, K. Ramchandran, Y. Wu and C. Suh, A Survey on Network Codes for Distributed Storage,*Proc. of IEEE*, vol. 99, no. 3, Mar. 2011.
3. G.M. Kamath, N. Prakash, V. Lalitha and P.V. Kumar, Codes with local regeneration, *IEEE Info. Thy. and App. Workshop*, pp. 1-5, Feb. 2013, San Diego, CA.
4. T. Ernvall, T. Westerbäck, R. Freij-Hollanti, and C. Hollanti. A Connection Between Locally Repairable Codes and Exact Regenerating Codes. Proceedings of IEEE Intl. Symp. on Info. Thy. Barcelona, July 2016.
5. V. Skachek. Batch and PIR Codes and Their Connections to Locally Repairable Codes. axXiv:1611.09914v2, March 2017.
6. Y. Bilu and S. Hoory, On codes from hypergraphs, *European Journal of Combinatorics*, **25** (2004), 339–354.
7. A. Barg and G. Zémor, Codes on hypergraphs, *IEEE International Symposium on Information Theory*, pp. 156 - 160, July 2008, Toronto, ON.

8. M. Sipser and D. A. Spielman, Expander codes, *IEEE Transactions on Information Theory*, **42** (1996), 1710–1722.
9. H.L. Janwa and A. K. Lal, On expander graphs: parameters and applications, preprint, arXivcs/0406048.
10. G. Zémor, On expander codes, *IEEE Trans. on Info.Theory*, **47** (2001), no. 2, pp. 835–837.
11. R. M. Tanner, A recursive approach to low complexity codes, *IEEE Transactions on Information Theory*, **27** (1981), 533–547.
12. H. Janwa and A. K. Lal, On Tanner codes: minimum distance and decoding, *Proceedings of AAECC*, 13 (2003), pp. 335 - 347.
13. A.S. Rawat, A. Mazumdar and S. Vishwanath, Cooperative Local Repair in Distributed Storage, arXiv1409.3900.
14. A.S. Rawat, D.S. Papailiopoulos, A.G. Dimakis and S. Vishwanath, Locality and availability in distributed storage, IEEE Trans. on Info. Thy., Vol. 62, no. 8, pp. 4481-4493. Aug. 2016.
15. C.A. Kelley and D. Sridhara, Eigenvalue bounds on the pseudocodeword weight of expander codes, *Advances in Mathematics of Communications*, **1** (2007), 287 – 307.
16. N. Miladinovic and M. P. C. Fossorier, Generalized LDPC codes and generalized stopping sets, *IEEE Transactions on Communications*, **56** (2008), 201 – 212.
17. C. Di, D. Proietti, T. Richardson, E. Teletar and R. Urbanke, Finite-length analysis of low-density parity-check codes on the binary erasure channel, *IEEE Transactions on Information Theory*, **48** (2002), 1570 – 1579.
18. A. Lubotzky, R. Phillips and P. Sarnak, Ramanujan graphs, *Combinatorica*, **8** (1988), 261–277.
19. A. W. Marcus, D. A. Spielman and N. Srivastava, Interlacing families I: bipartite Ramanujan graphs of all degrees, *Annals of Mathematics*, **182** (2015), 307–325.