New Results on Locally Repairable Codes and Matroids

Camilla Hollanti
Department of Mathematics and Systems Analysis
Aalto University School of Science, Finland
camilla.hollanti@aalto.fi

Coding for Emerging Memories and Storage Technologies
Technion, May 3 2015

Joint work with T. Westerbäck, R. Freij, and T. Ernvall.

Outline

- Perspectives
- Almost affine codes and matroids
- Locally repairable codes and matroid invariants
- Singleton bounds and matroid operations
- Intuition for our construction

Perspectives: A _lot_ of data!

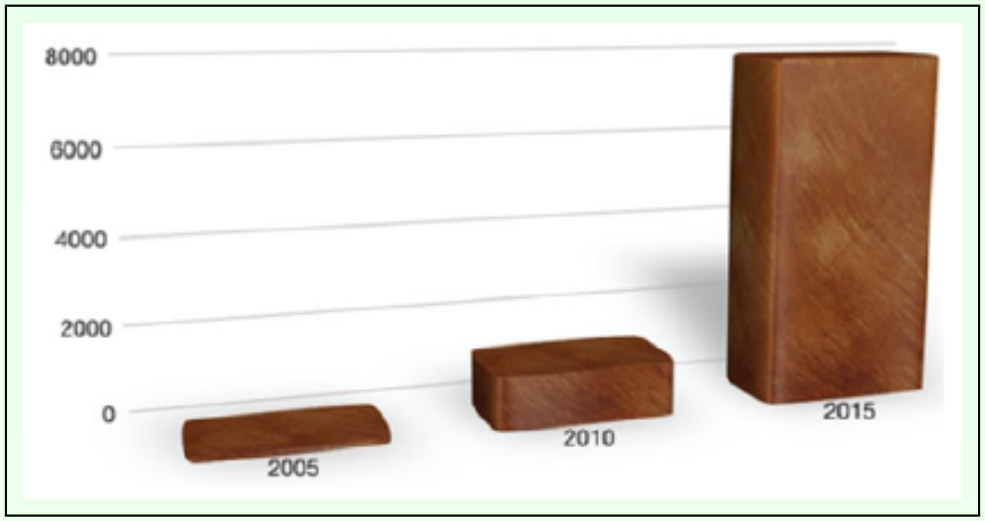

- An order of $5 \cdot 10^{21}$ bytes (zettabytes) of data stored worldwide, doubled every two years.
- Challenges come from physical storage space, energy consumption, bandwidth, quality of service, security...

Perspectives: A _lot_ of data!

- Facebook handles a million pictures a second at peak.

- NSA data centers use six million litres of water daily to cool their servers.
- Google used more than a million servers already in 2008[1].
- Data centers use about 2–5% of all electricity *worldwide*
  $\Rightarrow$ Effective data storage affects the environment on a global scale.[2]

[1] http://www.datacenterknowledge.com.

[2] Greenpeace report: How clean is your cloud?

Perspectives: A _lot_ of data!

- Data centers worldwide experience about 3 million hours of outage yearly resulting in huge financial losses (billions of dollars).

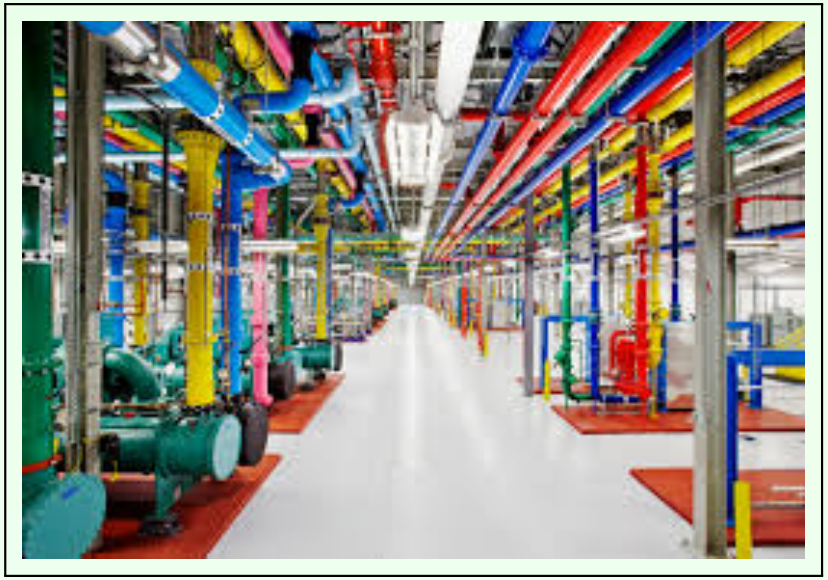

- How do we prevent data from becoming unavailable during these outages, without wasting valuable storage space?

Distributed systems (DSS)

- In a DSS, a file is divided into k packets, and distributed over $n \geq k$ nodes in a network. Assume this is done by using a code with minimum distance d.
- If the content of no more than $d - 1$ nodes are erased, their content can be reconstructed.

A toy example via replication

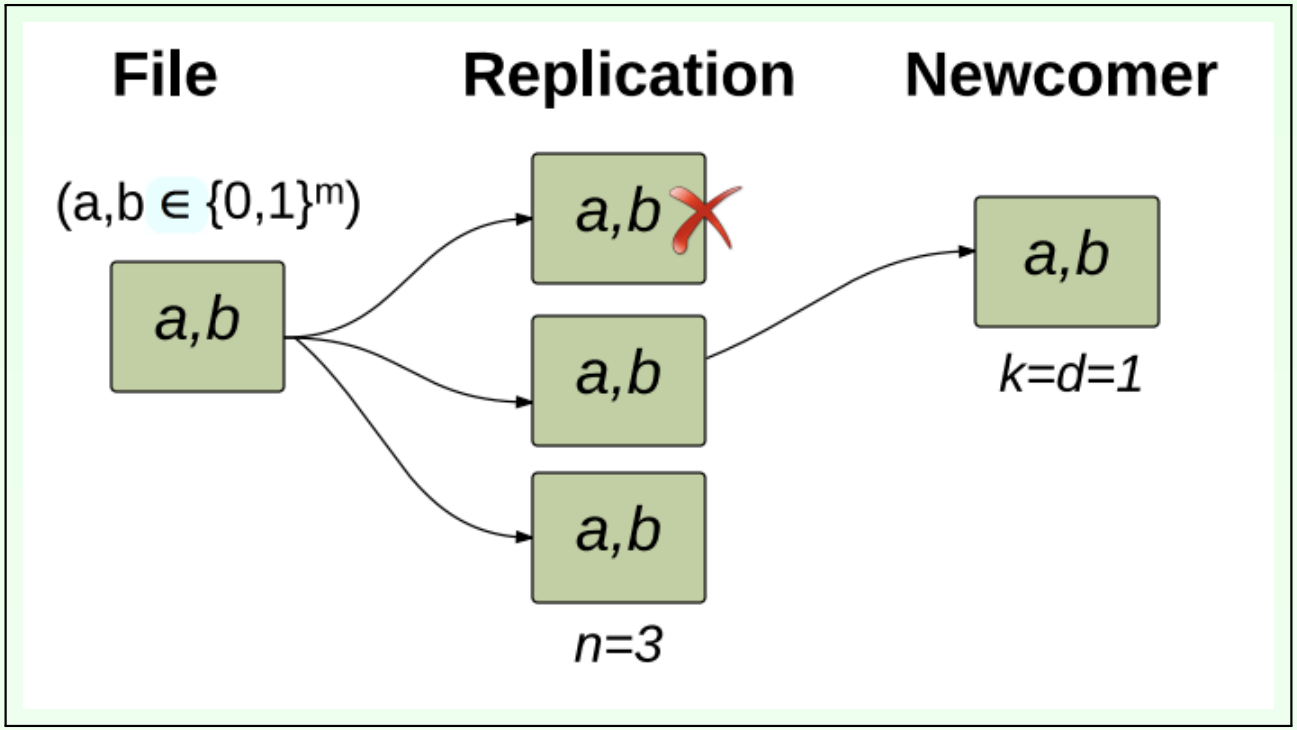

A toy example via network coding

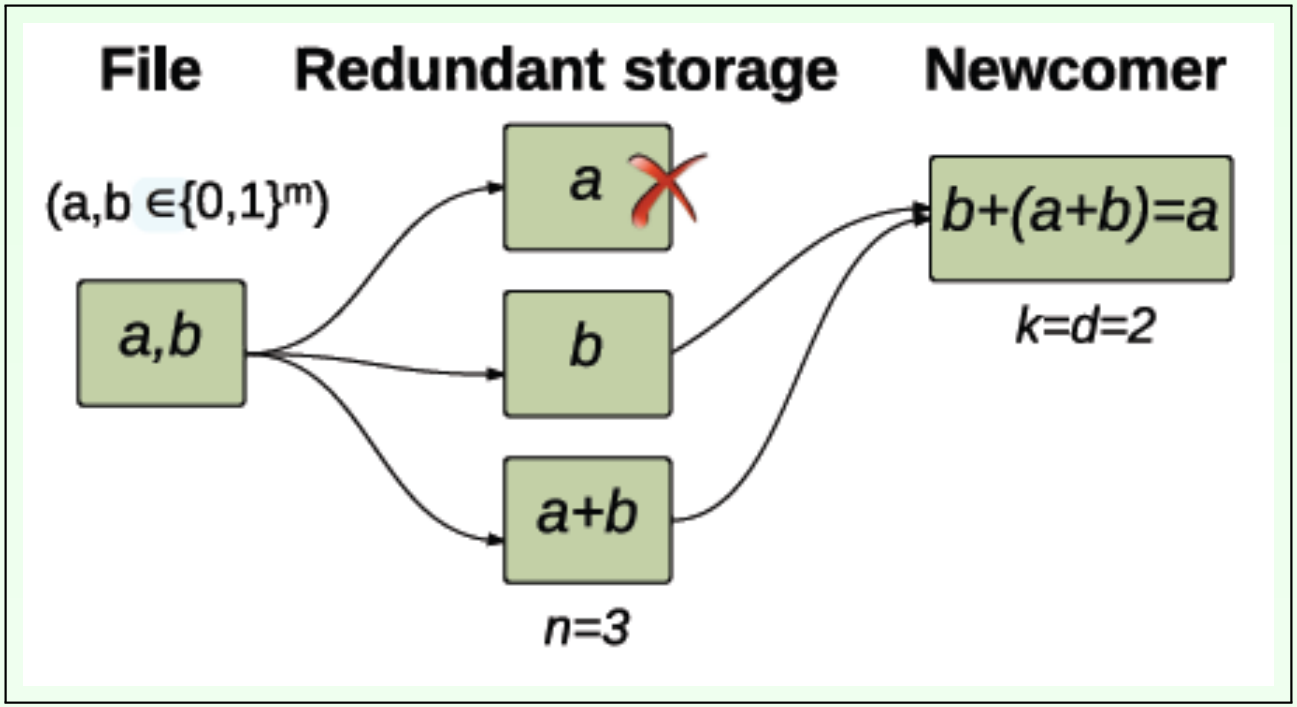

Contributions

- New links between matroids and linear[3] (n, k, d, r, δ)-LRCs, showing that the parameters are matroidal properties of the LRCs.
- New constructions of linear-LRCs with a wide range of different parameters (n, k, d, r, δ)-LRCs.
- Extending the existence and non-existence results on linear LRCs meeting the generalized Singleton bound.

[3] More generally almost affine codes.

About the constructions

- Our work here gives a new useful tool to use when working with LRCs.
- Thanks to this, the constructions are quite simple (though extremely technical) with good control on the parameters and repair groups.
- It is a combinatorial construction in the sense that it only uses the intersections of the repair groups, weighted graphs, and/or intersection schemes.

Matroids

- A matroid is a combinatorial structure that captures and generalises notions of **independence** (for example linear independence, algebraic independence, or acyclicity in graphs).
- Applications in geometry, topology, combinatorial optimization, network theory and coding theory.

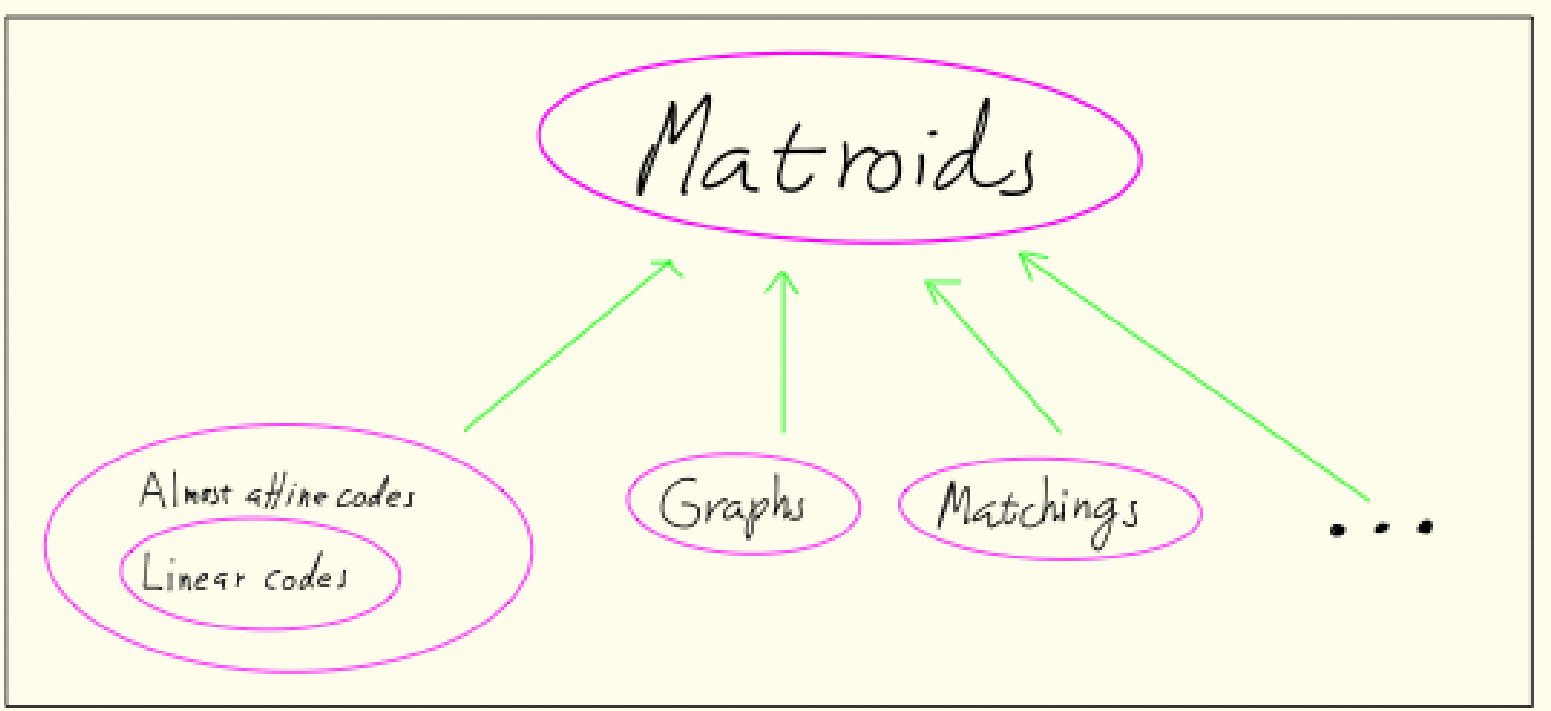

- For a finite set E, let 2^E denote the set of subsets of E.

Definition

$M = (\rho, E)$ is a *matroid* with a *rank function* $\rho : 2^E \to \mathbb{Z}$, if ρ has the following properties:

$(R1)$ $0 \leq \rho(X) \leq |X|$ for all $X \in 2^E$,
$(R2)$ If $X \subseteq Y \in 2^E$ then $\rho(X) \leq \rho(Y)$,
$(R3)$ If $X, Y \in 2^E$ then $\rho(X) + \rho(Y) \geq \rho(X \cup Y) + \rho(X \cap Y)$.

- A set $X \in 2^E$ is *independent* in M if $\rho(X) = |X|$, otherwise it is *dependent*.

Proposition (Alternative definition for topologists)

$\mathcal{I} \subset E$ is the collection of independent sets of a matroid if and only if I is a pure simplicial complex, all of whose induced subcomplexes are pure. The rank function is defined by $\rho : 2^E \to \mathbb{Z}$,

$$\rho(X) = \max_{Y \subseteq X, Y \in \mathcal{I}} |Y|.$$

- A third way to define matroids is via their set of bases.

- A dependent set X is a *circuit* if all proper subsets of X are independent.
- The *dual* of a matroid $M = (\rho, E)$ is a matroid $M^* = (\rho^*, E)$, where ρ^* is defined by:

$$\rho^*(X) = \rho(E \setminus X) + |X| - \rho(E), \text{ for all } X \in 2^E.$$

MDS (minimum distance separable) codes

Theorem (Singleton)

For any code of length n, dimension k and minimum distance d, over an arbitrary alphabet $\mathbb{A}$, the inequality

$$d \leq n - k + 1$$

holds.

MDS (minimum distance separable) codes

Theorem (Singleton)

For any code of length n, dimension k and minimum distance d, over an arbitrary alphabet $\mathbb{A}$, the inequality

$$d \leq n - k + 1$$

holds.

- A code achieving equality in the Singleton bound is an MDS-code.
- Explicit (linear) constructions of MDS-codes exist over all alphabets $\mathbb{A} = \mathbb{F}_q$ where $|\mathbb{A}| = q \geq n$ is a prime power.

MDS (minimum distance separable) codes

- A generic $n \times k$ matrix has every $k \times k$-minor non-degenerate, so it is the generator matrix of a code where any k nodes can reconstruct the information. This implies that the code is MDS.
- The associated matroid $M_{\mathcal{C}}$ is the uniform matroid U_n^k (circuits have exactly k+1 elements).
- Existence of MDS codes becomes a question of whether generic matrices exist over your favourite field.

Cooperative locally repairable codes

Gopalan *et al.*, Oggier *et al.*, and Papailiopoulos *et al.*

- $\mathcal{C}$ a code of length n, dimension k, rate k/n, minimum distance d.

Definition

An (r, δ)-*cloud* is a set F of nodes such that for every $(\delta - 1)$-tuple $x_1, \ldots, x_{\delta-1} \in F$, there are $y_1, \cdots y_m \in F \setminus \{x_i\}$, $m \leq r$, such that $x_i = f(y_{j_1}, \ldots, y_{j_m})$ for some function f.

Definition

$\mathcal{C}$ is a *locally repairable code (LRC)* with parameters (n, k, d, r, δ), if each of the n nodes is contained in an (r, δ)-cloud.

Example: A $(12, 6, 4, 3, 3)$-LRC

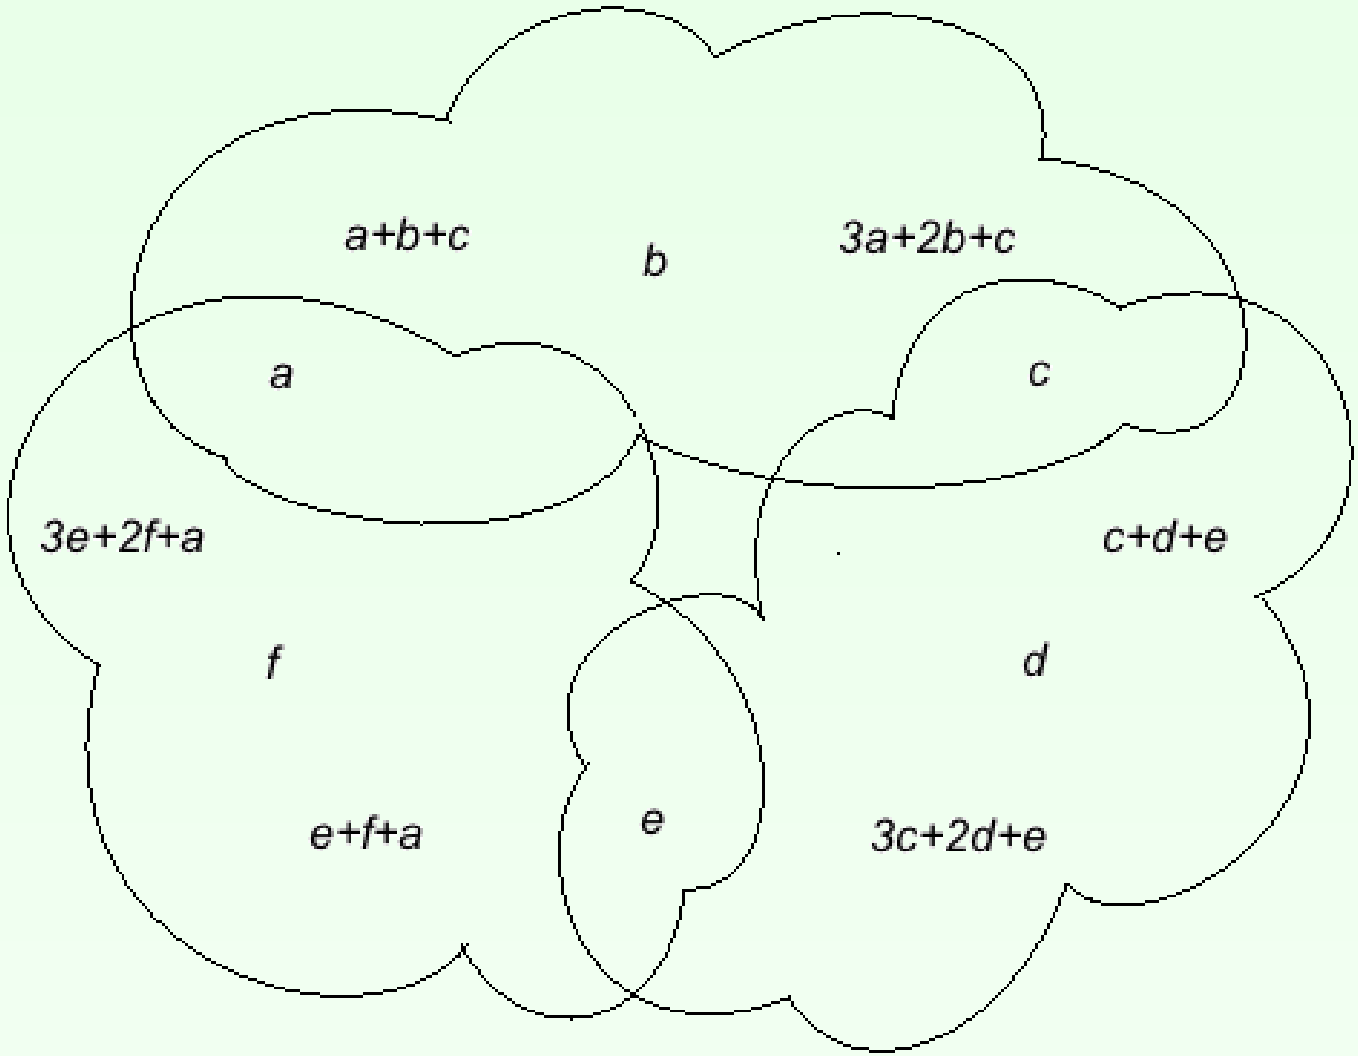

Example: A $(12, 6, 4, 3, 3)$-LRC

$$G = \begin{array}{c} \\ a \\ b \\ c \\ d \\ e \\ f \end{array} \begin{array}{c} \begin{array}{cccccccccccc} \mathbf{1} & \mathbf{2} & \mathbf{3} & \mathbf{4} & \mathbf{5} & \mathbf{6} & \mathbf{7} & \mathbf{8} & \mathbf{9} & \mathbf{10} & \mathbf{11} & \mathbf{12} \end{array} \\ \begin{pmatrix} 1 & & & & & & 1 & & 1 & 3 & & 1 \\ & 1 & & & & & 1 & & & 2 & & \\ & & 1 & & & & 1 & 1 & & 1 & 3 & \\ & & & 1 & & & & 1 & & & 2 & \\ & & & & 1 & & & 1 & 1 & & 1 & 3 \\ & & & & & 1 & & & 1 & & & 2 \end{pmatrix} \end{array}$$

Singleton bound for LRC

- The Singleton bound can be sharpened for locally repairable codes that are linear / almost affine (Prakash/Westerbäck *et al.*, 2012/2014)

$$d_{min}(\mathcal{C}) \leq n - k + 1 - (\delta - 1)\left(\left\lceil \frac{k}{r} \right\rceil - 1\right).$$

- We can also bound the rate

$$\text{rate}(\mathcal{C}) = \frac{k}{n} \leq \frac{r}{r + \delta - 1}.$$

- How do we construct LRC with equality? Using matroids!

Translation from LRC to matroids

Tamo *et al.* (2013), Westerbäck *et al.* (2014)

- Let $\mathcal{C}$ be *almost affine*, meaning

$$\mathcal{C}|_I = |\mathbb{A}|^{\rho(I)}$$

  for an integer $\rho(I)$, for every $I \subseteq [n]$.
- Then $(\rho, [n])$ is a (representable) matroid.
- The parameters (n, k, d, r, δ) can easily be generalized to arbitrary finite matroids.

Translation from LRC to matroids

Our contributions

- Let $\mathcal{C}$ be an almost affine code.
- The parameters (n, k, d, r, δ) can be read off from the associate matroid $M_{\mathcal{C}} = (\rho_{\mathcal{C}}, [n])$ as follows:
- $k = \rho_{\mathcal{C}}([n])$.
- $d = \min\{\, |X| \,:\, X \text{ cocircuit, } i.e., \text{ a circuit of } M^*_{\mathcal{C}}\}$.
- F is an (r, δ)-cloud if and only if F is a minimal cyclic flat of rank $\leq r$ and corank $\geq \delta$.

Translation from LRC to matroids

Our contributions

- Any matroid is uniquely determined by the *lattice of cyclic flats* (which is a lattice), and the rank function restricted to the cyclic flats.
- An extremal[4] (n, k, d, r, δ)-matroid has its lattice of cyclic flats generated by sets F_i corresponding to the clouds, with
  - $|F_i| - \rho(F_i) \geq \delta - 1$
  - $\rho(F_i) \geq r$
  - $|\cup_i F_i| = k + \sum_i (|F_i| - \rho(F_i))$
  - If
    $$\rho(\cup_{i \in I} F_i) < k, \rho(\cup_{j \in J} F_j) < k, \rho(\cup_{i \in I \cup J} F_i) = k,$$
    then
    $$|\cup_{i \in I \cup J} F_i| + \sum_{i \in I \cup J} (|F_i| - \rho(F_i)) \geq k.$$
- Determining whether such set systems exist, is a ~~boring~~ ~~tedious~~ simple exercise in hypergraph theory.

[4] *I.e.*, satisfies the generalized Singleton bound with equality.

- The inequality

$$d(\mathcal{C}) \leq n - k + 1 - (\delta - 1)\left(\left\lceil \frac{k}{r} \right\rceil - 1\right).$$

  now holds for matroids in general.
- For all parameters (n, k, r, δ), there is a matroid that satisfies

$$d(\mathcal{C}) = n - k - (\delta - 1)\left(\left\lceil \frac{k}{r} \right\rceil - 1\right).$$

- This is obtained as a direct sum of uniform matroids $U^r_{r+\delta-1}$, augmented with $d - \delta$ additional elements.

- Remember Singleton:

$$d(\mathcal{C}) \leq n - k + 1 - (\delta - 1)\left(\left\lceil \frac{k}{r} \right\rceil - 1\right).$$

- We can characterize (using graphs of overlapping clouds) for exactly which values we can improve (upon the direct sum construction), *i.e.*, satisfy the Singleton bound with equality.
- As a result, we settle the existence and construction of LRCs for a wider parameter set than earlier works.

Construction

- I will try to give some intuition as to how these matroids can be constructed, and in fact be realized as codes. For ugly technical details, consult our papers or Thomas and Ragnar.

Construction

Step 1 Identify a storage system with its generator matrix, the information symbols with the rows, and the storage nodes with the columns.

Key idea: first consider only linear dependencies among the columns (ignoring the rows).

Only in the end will we use some vintage algebraic combinatorics (gammoids[5]) to show that one can actually fill in the matrix to obtain exactly the desired dependencies.

[5] A matroid describing sets of vertices that can be reached by vertex-disjoint paths in a directed graph.

Construction

Step 2 These dependencies can be explained as starting from a set of n independent nodes, and imposing a whole lot of restrictions of the kind "the set S should have rank r, but otherwise be as independent as possible".

These restrictions can be captured in the form of a graph (in a not at all obvious way, Thomas's clever trick), and the matroid is isomorphic to the gammoid associated to this graph.

Construction

Step 3 *Gammoids are representable* by a theorem in algebraic combinatorics.

If we assign r-spaces to all local clouds, and make sure they intersect each other in the right dimensions (this can be done, which is nontrivial), then if we fill the columns in generically, we will get a LRC with the right parameters.

Here, "generically" just means that we need to avoid making any of a finite number of polynomials zero. Over a large enough field, this happens asymptotically almost surely (which is the geometric perspective on the random matrix result in our other paper).

References (non-exahaustive)

- A. Dimakis, P. Godfrey, Y. Wu, M. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE Trans. Inf. Theory*, 56(9), pp. 4539–4551, September 2010.
- A. Dimakis, slides of invited talk in *ACN Bordeaux*: http://users.ece.utexas.edu/d̃imakis/Bordaux_ACN_Talk.pdf, 2014.
- I. Tamo, D. Papailiopoulos, and A. Dimakis, "Optimal Locally Repairable Codes and Connections to Matroid Theory", *arxiv.1301.7693*, 2013.
- N. Silberstein, A.S. Rawat, O. Koyluoglu, and S. Vishwanath, "Optimal Locally Repairable Codes via Rank-Metric Codes", *arxiv.1301.6331*, 2013.
- A.S. Rawat, O. Koyluoglu, N. Silberstein and S. Vishwanath, "Optimal Locally Repairable and Secure Codes for Distributed Storage Systems", *arxiv.1210.6954*, 2012.
- A. Singh Rawat, A. Mazumdar, and S. Vishwanath, "Cooperative Local Repair in Distributed Storage", *arxiv.1409.3900*, 2014.

References (non-exhaustive)

- W. Song, S. H. Dau, C. Yuen, T. J. Li, "Optimal Locally Repairable Linear Codes", *IEEE Sel. Areas Commun.*, 32(5), 2014.
- T. Ernvall, T. Westerbäck, C. Hollanti, R. Freij, "Constructions and Properties of Linear Locally Repairable Codes", *arXiv.1410.6339*, 2014.
- T. Westerbäck, R. Freij, T. Ernvall, C. Hollanti, "Almost Affine Locally Repairable Codes from Matroids", *IEEE Information Theory Workshop (ITW)*, Hobart, Nov. 2014.
- T. Westerbäck, R. Freij, T. Ernvall, C. Hollanti, "On the Combinatorics of Locally Repairable Codes via Matroid Theory", *arXiv.1501.00153*, 2014.
- N. Silberstein and A. Zeh, "Optimal Binary Locally Repairable Codes via Anticodes" *arxiv.1501.07114*, 2015.

See also references within the above references!