Locally repairable codes through almost uniform matroids

Ragnar Freij
Aalto University, Finland
ragnar.freij@aalto.fi

ALCOMA 19.3.2015

Based on joint work with T. Ernvall, C. Hollanti och T. Westerbäck.

Outline

- Perspectives and a cute picture of a cat
- Almost affine codes and matroids.
- Locally repairable codes and matroid invariants.
- Singleton bounds and matroid operations.

Perspectives: A heck of a lot of data!

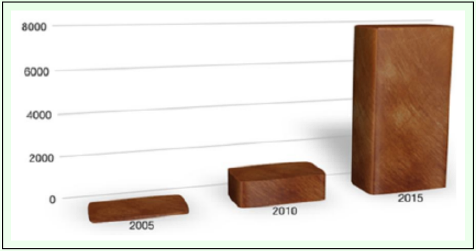

- EMC 2011: $1.8 * 10^{21}$ bytes (zettabytes?) of data stored world wide, doubled every two years.
- Challenges come from physical storage space, energy consumption, bandwidth, security...

Perspectives: A heck of 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 per cent of all electricity *world wide*. Effective date storage affects the environment on a global scale.[2]

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

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

Perspectives: A heck of a lot of data!

- Data centers worldwide experience about 3 million hours of outage yearly.

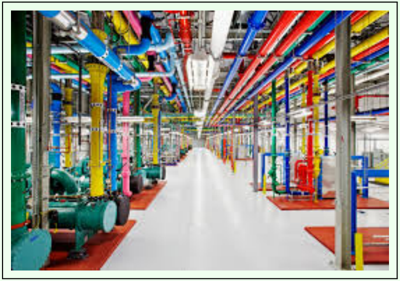

- Hur do we secure data from getting lost during these outages, without wasting valuable storage space?

Distributed systems (DSS)

- In a DSS, a file is divided into k packets, and distributed over $n \geq k$ nodes in a network.
- If the content of no more than $d-1$ nodes are erased, or no more than $\frac{d-1}{2}$ nodes are corrupted, their content can be reconstructed.
- This setting is known as *exact repair*. One is also interested in *functional repair*, where only reconstructability of some DSS with the same parameters is required.

A toy example via network coding

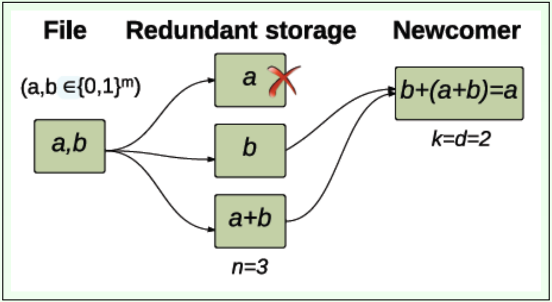

Matroids

- 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.

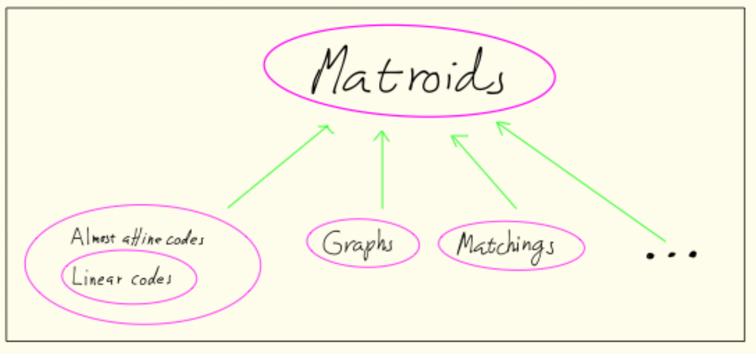

- For a finite set E, let 2^E denote the set of subsets of E.

Definition

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

$$\begin{array}{ll} (R1) & 0 \leq \rho(X) \leq |X| \text{ for all } X \in 2^E, \\ (R2) & \text{If } X \subseteq Y \in 2^E \text{ then } \rho(X) \leq \rho(Y), \\ (R3) & \text{If } X, Y \in 2^E \text{ then } \rho(X) + \rho(Y) \geq \rho(X \cup Y) + \rho(X \cap Y). \end{array}$$

- 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 sub complexes 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 is the generator matrix of a code where every k nodes can reconstruct the code word. This implies that the code is MDS.
- The matroid $M_{\mathcal{C}}$ is the uniform matroid U_n^k.
- 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 $x_1, \cdots x_{\delta-1} \in F$, there are $y_1, \cdots y_r \in F \setminus \{x_i\}$ such that $\{f(x_i)\}$ is a function of $\{f(y_i)\}$.

Definition

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

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

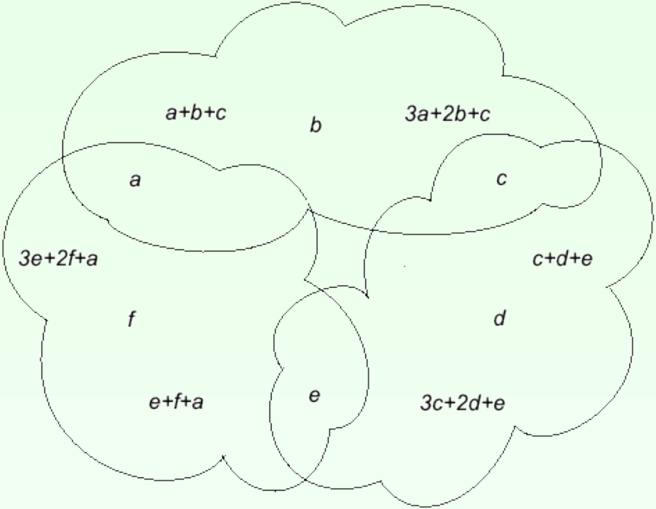

Example: A $(12, 6, 4, 3, 3)$-LRC

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

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

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

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

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

- 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 generalised to arbitrary finite matroids:

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

- 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 $(n, k, d_{min}, r, \delta)$-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.

- 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 disjoint union of copies of $U^r_{r+\delta-1}$, augmented with $d - \delta$ addictional elements.

- Remember Singleton:

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

- We can characterise (using graphs of overlapping clouds) for exactly which values this can be improved to satisfy the Singleton bound with equality.
- n and k has to satisfy certain congruences modulo r, $r+1$ δ and $\delta - 1$.
- Thomas will explain how these matroids can be constructed, and in fact be realized as codes.

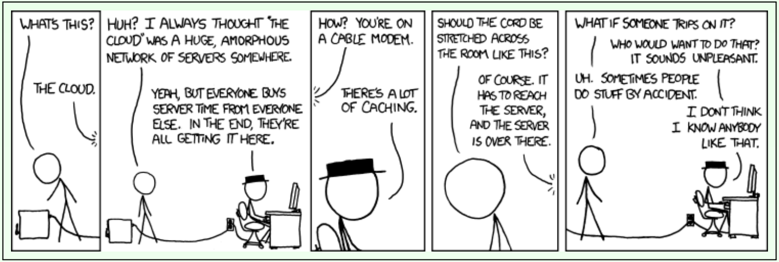
WHAT'S THIS?
THE CLOUD.
HUH? I ALWAYS THOUGHT "THE CLOUD" WAS A HUGE, AMORPHOUS NETWORK OF SERVERS SOMEWHERE.
YEAH, BUT EVERYONE BUYS SERVER TIME FROM EVERYONE ELSE. IN THE END, THEY'RE ALL GETTING IT HERE.
HOW? YOU'RE ON A CABLE MODEM.
THERE'S A LOT OF CACHING.
SHOULD THE CORD BE STRETCHED ACROSS THE ROOM LIKE THIS?
OF COURSE. IT HAS TO REACH THE SERVER, AND THE SERVER IS OVER THERE.
WHAT IF SOMEONE TRIPS ON IT?
WHO WOULD WANT TO DO THAT? IT SOUNDS UNPLEASANT.
UH. SOMETIMES PEOPLE DO STUFF BY ACCIDENT.
I DON'T THINK I KNOW ANYBODY LIKE THAT.