Locally repairable codes on multiple scales

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

Istanbul, 5.11.2015

Designs and Applications of Random Network Codes

LRC on multiple scales

Based on joint work with Camilla Hollanti, Thomas Westerbäck and Toni Ernvall

Distributed storage systems (DSS)

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

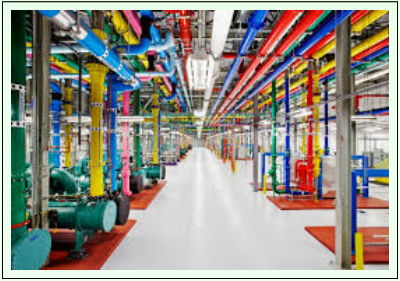

- How do we secure data from getting lost during these outages, without wasting valuable storage space?

Distributed storage systems (DSS)

- In a (linear) DSS with exact repair, a file is divided into k packets, each packet identified with an element in an alphabet (field) $\mathbb{A}$, and distributed over $n \geq k$ nodes in a network via a (linear) injection $\mathbb{A}^k \to \mathbb{A}^n$.
- If the content of no more than $d-1$ nodes are erased, their content can still be reconstructed.
- In a general DSS, there is no guarantee that an erased node can be recovered without restoring the entire file.
- In contrast, in a *locally repairable* system, few ($< \delta$) erasures can be repaired by few ($\leq r$) other nodes.

Cooperative locally repairable codes
Gopalan *et al.*, Oggier *et al.*, and Papailiopoulos *et al.*

- A code $\mathcal{C} \subseteq \mathbb{A}^n$ with $|\mathcal{C}| = |\mathbb{A}|^k$ is said to have *size* n and *rank* k.
- The minimum distance d is defined as

$$d = \min\{|X| : |\mathcal{C}_{|[n]\setminus X}| < |\mathcal{C}|\}.$$

- An (r, δ)-cloud of the code $\mathcal{C}$ is $Z \subseteq [n]$ such that the projection $\mathcal{C}_{|Z}$ has rank $\leq r$ and minimum distance $\geq \delta$.

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.

- Let $(\mathbf{n}, \mathbf{k}, \mathbf{d}) = ((n_i, k_i, d_i)_{i \geq 0})$ be a finite sequence of triples, coordinatewise decreasing.

Definition (Freij-Hollanti et al, 2015+)

An (n_0, k_0, d_0)-code $\mathcal{C}$ is said to have recoverability

$$(\mathbf{n}, \mathbf{k}, \mathbf{d}) = (n_i, k_i, d_i)_{i \geq 0}$$

if for every $x \in [n_0]$, there is $X \subseteq [n_0]$ with $x \in X$ such that $\mathcal{C}|_X$ has recoverability

$$(n_{i+1}, k_{i+1}, d_{i+1})_{i \geq 0}$$

- A matroid is a combinatorial structure that captures and generalises notions of **independence** (for example linear independence, algebraic independence, or acyclicity in graphs).

Definition

Let E be a finite set, and 2^E its power set. $M = (\rho, E)$ is a *matroid* with a *rank function* $\rho : 2^E \to \mathbb{Z}$, if ρ has the following properties:

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

Matroids can also be defined via their independent sets, which are the sets $X \subseteq E$ with $|X| = \rho(X)$.

- To a linear code (and more generally to any almost affine code) $\mathcal{C} \subseteq \mathbb{A}^n$ corresponds a code $M_{\mathcal{C}} = (\rho, [n])$, defined by

$$\rho(X) = \dim_{\mathbb{A}}(\mathcal{C}_{|X}).$$

- The parameters (n, k, d) are matroid invariants, where $k = \rho([n])$ and

$$d = \min\{\, |X| \;:\; Y \subsetneq E \text{ and } \rho(X \setminus Y) < \rho(X)\}$$

- An important invariant of linear codes is the *weight enumerating polynomial*

$$W(\mathcal{C}; x) = \sum_{c \in \mathcal{C}} x^{w(c)},$$

where $w(c)$ is the number of non-zero letters in the code word c.

- An important invariant of linear codes is the *weight enumerating polynomial*

$$W(\mathcal{C}; x) = \sum_{c \in \mathcal{C}} x^{w(c)},$$

where $w(c)$ is the number of non-zero letters in the code word c.

- Similarly, an important invariant of matroids is the Tutte polynomial

$$T(M; x, y) = \sum_{S \subseteq E} (x-1)^{\rho(E)-\rho(S)} (y-1)^{|S|-\rho(S)},$$

which is in a precise sense the most general polynomial that is recursively defined via deletion and contraction identities.

The weight enumerator $W(\mathcal{C}; x)$ of a code and the Tutte polynomial $T(M_{\mathcal{C}}; x, y)$ of the associated matroid are related via

Theorem (Greene, 1976)

$$W(\mathcal{C}; z) = z^{n-k}(1-z)^k T\left(M_{\mathcal{C}}; \frac{1+(q-1)z}{1-z}, \frac{1}{z}\right),$$

where $\mathcal{C}$ is a linear code over $\mathbb{F}_q$.

- A dependent set X is a *circuit* if all proper subsets of X are independent.
- A set X is a *flat* if $\rho(X \cup y) = \rho(X) + 1$ for all $y \in (E \setminus X)$. A flat is *cyclic* if it is a union of circuits.
- In the setting of codes, the cyclic flats Z are "repair sets", meaning that erasures inside Z can be repaired by other nodes in Z, but Z itself cannot repair any node outside Z.

- The collection of flats of a matroid is denoted by $\mathcal{L}(M)$, and is a geometric lattice ordered under inclusion.
  - A lattice is geometric if it is graded, atomic, and submodular, meaning that
  $$\rho(x) + \rho(y) \geq \rho(x \wedge y) + \rho(x \vee y)$$
  - Any geometric lattice is isomorphic to $\mathcal{L}(M)$ for some matroid M.
  - The lattice $\mathcal{L}(M)$ determines M up to isomorphism.

- The collection of cyclic flats of a matroid is denoted by $\mathcal{Z}(M)$, and is a lattice ordered under inclusion.
  - Any lattice is isomorphic to $\mathcal{Z}(M)$ for some matroid M.
  - The lattice $\mathcal{Z}(M)$ together with its rank function, determines M.
  - The elements of $\mathcal{Z}(M)$ must be thought of as sets, rather than abstract elements.

- The configuration $\mathcal{K}(M)$ of M is the triple $(K(M), \#, \rho)$, where $K(M)$ is the abstract lattice $\mathcal{Z}(M)$, and $(\#, \rho)$ are the cardinality and rank function on its nodes.
  - $\mathcal{K}(M)$ does not determine M.
  - However, $\mathcal{K}(M)$ does determine the Tutte polynomial $T(M; x, y)$. (Eberhardt, 2014)

THEOREM (EBERHARDT 2014, FREIJ-HOLLANTI ET AL 2015+)

Let M be a matroid, with configuration $(K(M), \#, \rho)$, and

$$\eta(S) = \#S - \rho(S).$$

Then

$$T(M; x, y) = \sum_{S \in K(M)} (x-1)^{k-\rho(S)} (y-1)^{\eta(S)} \cdot \left(1 + \sum_{R \lessdot S} \sum_{i=1}^{\rho(S)-\rho(R)-1} \binom{\#S - \#R}{i} (x-1)^i \right) \cdot \left(1 + \sum_{T \gtrdot S} \sum_{j=1}^{\eta(T)-\eta(S)} \binom{\#T - \#S}{j} (y-1)^{-j} \right).$$

Maximum Distance Separable (MDS) codes

Theorem (Singleton, 1964)

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

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

holds.

Maximum Distance Separable (MDS) codes

Theorem (Singleton, 1964)

For any linear 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.
- If $\mathcal{C}$ is an MDS-code, then the matroid $M_{\mathcal{C}}$ is the uniform matroid U_n^k, with $\mathcal{Z}(U_n^k) = \{\emptyset, [n]\}$ and $\rho([n]) = k$.
- 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.

- The parameters $(\mathbf{n}, \mathbf{k}, \mathbf{d}) = (n_i, k_i, d_i)_{i \geq 0}$ are matroid invariants.
- The Singleton bound can be generalised to matroids, and sharpened for codes and matroids with repairability $(\mathbf{n}, \mathbf{k}, \mathbf{d})$:

Theorem (Freij-Hollanti et al, 2015+)

Let M be a matroid with repairability $(\mathbf{n}, \mathbf{k}, \mathbf{d})$. Then

$$d_i(M) \leq n_i - k_i + 1 - (n_{i+1} - k_{i+1}) \left(\left\lceil \frac{k_i}{k_{i+1}} \right\rceil - 1 \right),$$

for every $i \geq 0$. Moreover, for every $i \geq 0$ we have

$$\frac{k_i}{n_i} \leq \frac{k_{i+1}}{n_{i+1}}.$$

- The parameters $(\mathbf{n}, \mathbf{k}, \mathbf{d}) = (n_i, k_i, d_i)_{i \geq 0}$ are invariants of the configuration $\mathcal{K}(M)$.
- Matroids (almost) achieving equality in the Singleton bounds have a nicely structured configuration:

Theorem (Freij-Hollanti et al, 2015+)

Let M be a matroid with repairability $(\mathbf{n}, \mathbf{k}, \mathbf{d})$, with

$$n_i - k_i + 1 - (n_{i+1} - k_{i+1}) \left\lceil \frac{k_i}{k_{i+1}} \right\rceil \leq d_i(M).$$

We call such a matroid locally nearly MDS. Then $\mathcal{K}(M)_{[k_{i+1}, k_i]}$ is a truncated Boolean lattice, generated by $\left\lceil \frac{n_i}{n_{i+1}} \right\rceil$ atoms of rank k_{i+1}, truncated at rank k_i.

- Let M be a locally nearly MDS matroid with repairability $(\mathbf{n}, \mathbf{k}, \mathbf{d})$.
- Then $T(M; x, y)$ can be written as a sum of polynomials

$$T(M; x, y) = \sum_{i \geq 0} \sum_{\substack{S \in K(M) \\ k_{i+1} \leq \rho(S) < k_i}} T_i(M; x, y).$$

- Each of the terms is a sum over a truncated Boolean lattice, and can be written out explicitly without reference to the lattice K.

- Let $\mathcal{C}$ be a locally nearly MDS code. Then

$$n_i - k_i + 1 - (n_{i+1} - k_{i+1}) \left\lceil \frac{k_i}{k_{i+1}} \right\rceil \leq d_i(M).$$

- Through the identity

$$W(\mathcal{C}; z) = z^{n-k}(1-z)^k T\left(M_{\mathcal{C}}; \frac{1+(q-1)z}{1-z}, \frac{1}{z}\right),$$

we get an explicit combinatorial formula for the weight enumerating polynomial, for any code with repairability $(\mathbf{n}, \mathbf{k}, \mathbf{d})$.

- For given parameters $(n_i, k_i, k_{i+1}, d_{i+1})$, satisfying

$$\begin{aligned}
& n_i - k_i + 1 - (n_{i+1} - k_{i+1}) \left\lceil \frac{k_i}{k_{i+1}} \right\rceil \\
\leq\, & d_i(M) \\
\leq\, & n_i - k_i + 1 - (n_{i+1} - k_{i+1}) \left(\left\lceil \frac{k_i}{k_{i+1}} \right\rceil - 1 \right),
\end{aligned}$$

  it is non-trivial to determine whether matroids with these parameters exist.
- Wang and Zhang (2015) and Westerbäck et al (2015+) give sufficient and necessary conditions, using linear programming and extremal graph theory respectively.
- When such matroids exist, they can be constructed explicitly via their lattice of cyclic flats.

- A *gammoid* is a matroid associated to a directed graph through flows, and are always representable as linear codes. (Oxley, 1961)
- The locally nearly MDS matroids can be constructed to be isomorphic to gammoids, associated to blow-ups of their configuration poset.
- Thus, we obtain linear codes with repairability $(\mathbf{n}, \mathbf{k}, \mathbf{d})$, when $(\mathbf{n}, \mathbf{k}, \mathbf{d})$ satisfy the generalised Singleton bounds, rate inequalities and certain congruences.
- If the generalised Singleton bounds are "almost" met with equality, we can also explicitly compute their weight enumeration in terms of $(\mathbf{n}, \mathbf{k}, \mathbf{d})$.

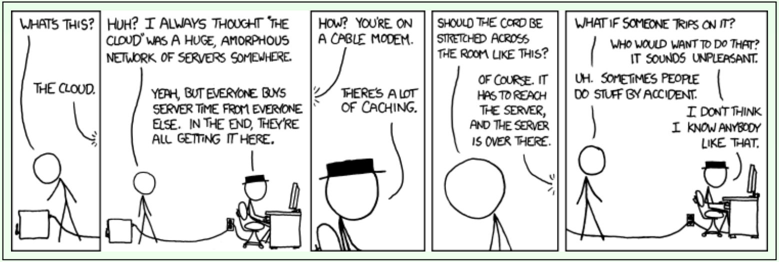
WHAT'S THIS?
THE CLOUD.
HUH? I ALWAYS THOUGHT "THE CLOUD" WAS A HUGE, AMORPHOUS NETWORK OF SERVERS SOMEWHERE.
YEAH, BUT EVERYONE BUYS SERVER TIME FROM EVERYONE ELSE. IN THE END, THEY'RE ALL GETTING IT HERE.
HOW? YOU'RE ON A CABLE MODEM.
THERE'S A LOT OF CACHING.
SHOULD THE CORD BE STRETCHED ACROSS THE 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.