



Flash Memory Summit



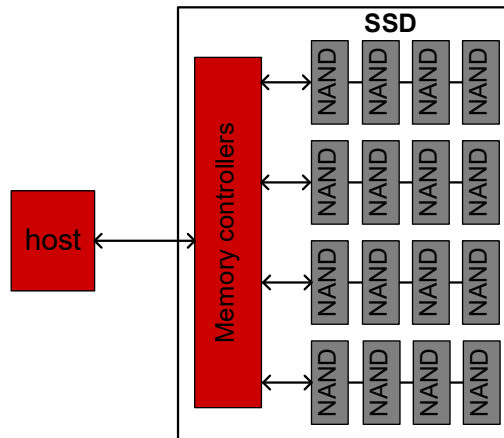
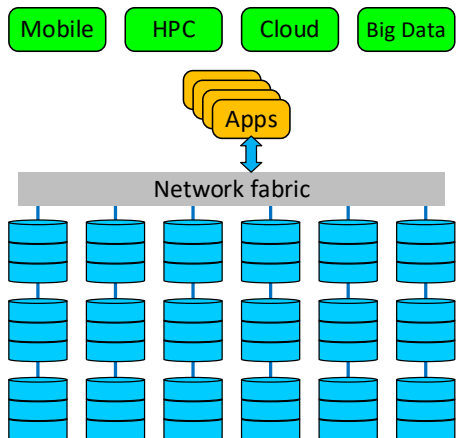
Improving the Locality of Generalized Integrated Interleaved Code

Xinmiao Zhang

The Ohio State University



Application of Generalized Integrated Interleaved (GII) Codes



- Redundancy is needed to tolerate and recover from failures in distributed storage
- GII codes are a type of locally recoverable erasure codes that substantially reduce the latency and overhead of failure recovery
- Enable short sub-codewords to be adopted for error correction without substantial redundancy increase
- Achieve >40GBytes/sec throughput with excellent correction capability



Outline


Flash Memory Summit

- Generalized integrated interleaved (GII) codes
- Modified two-layer GII codes with improved locality
- A generalization of three-layer integrated interleaved codes
- Conclusions



Generalized Integrated Interleaved (GII) Codes

Flash Memory Summit



$$\begin{bmatrix} \tilde{c}_0 \\ \tilde{c}_1 \\ \vdots \\ \tilde{c}_{v-1} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \alpha & \alpha^2 & \cdots & \alpha^{m-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha^{(v-1)} & \alpha^{2(v-1)} & \cdots & \alpha^{(v-1)(m-1)} \end{bmatrix} \begin{bmatrix} c_0 \\ c_1 \\ \vdots \\ c_{m-1} \end{bmatrix}$$

Nesting matrix G

Correction capability

$$t_v \geq t_{v-1} \geq \cdots \geq t_1 \geq t_0$$


$$\mathcal{C}_v \subseteq \mathcal{C}_{v-1} \subseteq \cdots \subseteq \mathcal{C}_1 \subseteq \mathcal{C}_0$$

$$\begin{matrix} \Psi & \Psi & & \Psi & \Psi \\ \tilde{c}_0 & \tilde{c}_1 & & \tilde{c}_{v-1} & c_0, c_1, \dots, c_{m-1} \end{matrix}$$



Decoding of GII Codes

Flash Memory Summit



Interleave syndromes

$$\begin{bmatrix} S_j^{(l_1)} \\ S_j^{(l_2)} \\ \vdots \\ S_j^{(l_b)} \end{bmatrix}$$

=

$$\begin{bmatrix} 1 & 1 & \dots & 1 \\ \alpha^{l_1} & \alpha^{l_2} & \dots & \alpha^{l_b} \\ \vdots & \vdots & \ddots & \vdots \\ \alpha^{(b-1)l_1} & \alpha^{(b-1)l_2} & \dots & \alpha^{(b-1)l_b} \end{bmatrix}^{-1}$$

Syndrome conversion matrix

nested syndromes

$$\begin{bmatrix} \tilde{S}_j^{(0)} \\ \tilde{S}_j^{(1)} \\ \vdots \\ \tilde{S}_j^{(b-1)} \end{bmatrix}$$

l_1, l_2, \dots, l_b : indices of interleaves with more than t_0 erasures (exceptional interleaves)

- t syndromes are needed to correct t erasures
- Higher-order syndromes for the interleaves are generated from the nested syndromes
- Syndrome conversion matrix is always invertible
- Each nested syndrome is generated by utilizing all interleaves




Modified GII Codes


Flash Memory Summit

$$G' = \begin{bmatrix} 1 & 1 & \dots & 1 & 1 & \dots & 1 \\ 0 & \alpha & \dots & \alpha^{v-2} & \alpha^{v-1} & \dots & \alpha^{m-1} \\ 0 & 0 & \dots & \alpha^{2(v-2)} & \alpha^{2(v-1)} & \dots & \alpha^{2(m-1)} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & \alpha^{(v-1)^2} & \dots & \alpha^{(v-1)(m-1)} \end{bmatrix}$$

Previous GII



Modified GII

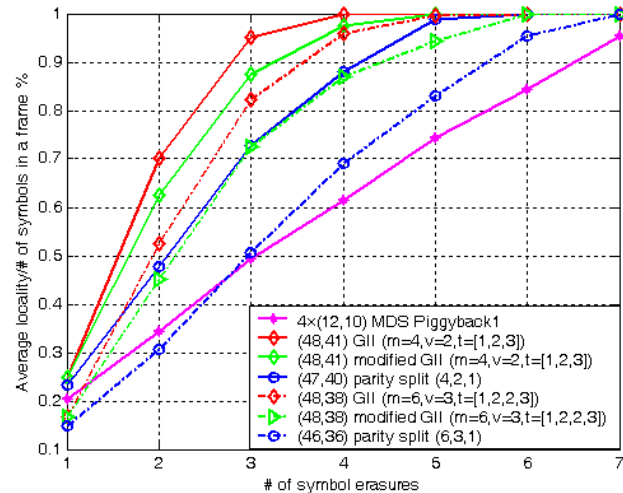
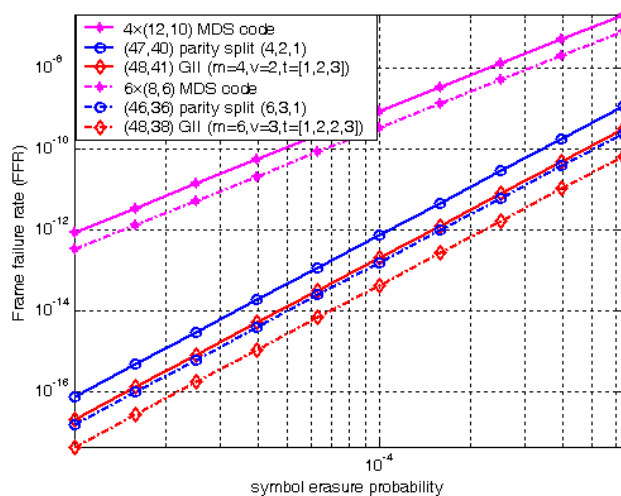


- Less powerful nestings involve fewer interleaves
- Form the syndrome conversion matrix using the bottom rows of G' as much as possible
 - The selected nestings should have sufficient correction capability
 - The selected nestings should cover the exceptional interleaves
 - Consecutive nestings are used to simplify the selection
- The syndrome conversion matrix is invertible if the number of interleaves does not exceed the values in the table

v	$GF(2^4)$	$GF(2^5)$	$GF(2^6)$	$GF(2^7)$	$GF(2^8)$	$GF(2^9)$
2	16	32	64	128	256	512
3	5	6	12	22	28	62
4	5	6	8	12	15	20



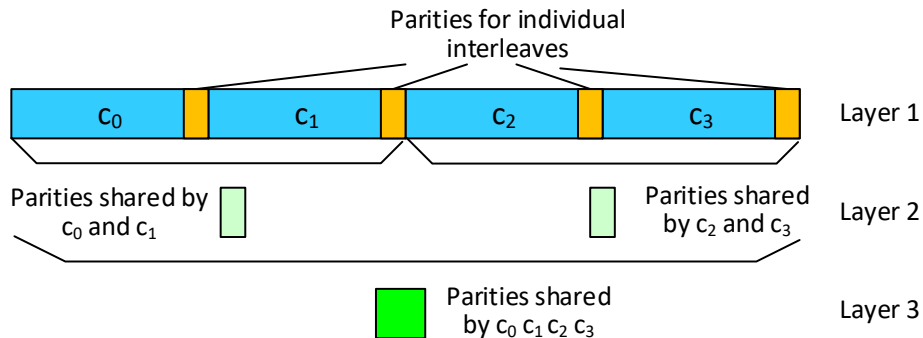
Correction Capability and Locality Comparisons



- Modified GII codes preserve the same correction capability as the original GII codes for most practical settings
- Modified GII codes require fewer interleaves to utilize the shared parities when there are fewer extra erasures to correct
- Have very small implementation overhead compared to the GII codes
- Achieve good tradeoff on locality and correction capability



Three-layer Integrated Interleaved Codes



- Layer 2 parities are shared by individual subgroups of interleaves
- Layer 3 parities are shared by all interleaves
- When there are fewer interleaves with fewer extra errors/erasures to correct, layer 2 parities and the other interleaves in the same sub-group are utilized



Nestings in Three-Layer Codes

Flash Memory Summit

G_1 : layer-2 nesting matrix

G_2 : layer-3 nesting matrix

$$\Gamma = \begin{bmatrix} G_1 & 0 & \cdots & 0 \\ 0 & G_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & G_1 \\ \hline & & & G_2 \end{bmatrix} = \begin{bmatrix} 11 \cdots 1 & 0 & \cdots & 0 \\ 0 & 11 \cdots 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 11 \cdots 1 \\ \hline 11 \cdots 1 & 11 \cdots 1 & 11 \cdots 1 & 11 \cdots 1 \end{bmatrix}$$

Joint nesting matrix

Previous 3-layer nesting

➤ Previous 3-layer integrated interleaved code

- One-level of nesting in each layer
- Only one exceptional interleave from each subgroup can be corrected
- Layer-3 parities only add to the correction of a single exceptional interleave in a subgroup



Generalized Three-layer Integrated Interleaved Codes

Flash Memory Summit

$$\Gamma = \left[\begin{array}{c|c|c|c} G_1 & 0 & \cdots & 0 \\ \hline 0 & G_1 & \cdots & 0 \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline 0 & 0 & \cdots & G_1 \\ \hline \underbrace{\hspace{10em}}_{G_2} \end{array} \right] = \left[\begin{array}{c|c|c|c} G_1 & 0 & \cdots & 0 \\ \hline 0 & G_1 & \cdots & 0 \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline 0 & 0 & \cdots & G_1 \\ \hline G_2^{r-1} & G_2^{r-1} & \cdots & G_2^{r-1} \\ \hline G_2^{r-2} & \alpha^{-1} G_2^{r-2} & \cdots & \alpha^{-(m_1-1)} G_2^{r-2} \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline G_2^0 & \alpha^{-(r-1)} G_2^0 & \cdots & \alpha^{-(m_1-1)(r-1)} G_2^0 \end{array} \right]$$

v_1 : # of levels of codes in layer 2
 m_1 : # of subgroups
 m_2 : # of interleaves in a subgroup
 r : # of subgroups with extra interleaves to be corrected
 s_i : # of extra interleaves in a subgroup to be corrected

$$G_1 = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \alpha & \alpha^2 & \cdots & \alpha^{m_2-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha^{(v_1-1)} & \alpha^{2(v_1-1)} & \cdots & \alpha^{(v_1-1)(m_2-1)} \end{bmatrix}$$

$$G_2^i = \begin{bmatrix} 1 & \alpha^{v_1} & \alpha^{2v_1} & \cdots & \alpha^{v_1(m_2-1)} \\ 1 & \alpha^{v_1+1} & \alpha^{2(v_1+1)} & \cdots & \alpha^{(v_1+1)(m_2-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha^{v_1+s_i-1} & \alpha^{2(v_1+s_i-1)} & \cdots & \alpha^{(v_1+s_i-1)(m_2-1)} \end{bmatrix}$$

- The syndrome conversion matrix formed by the columns corresponding to the interleaves with extra erasures and consecutive rows of $[G_1^T | (G_2^i)^T]^T$ in the joint nesting matrix is always invertible, if the numbers of groups and exceptional interleaves in the groups are not exceeded



Correctable Erasure Pattern Comparisons

Flash Memory Summit

Example 1: same redundancy
single-level nesting in 3-layer GII code

Layer-2: $t_1 = t_0 + 1$; layer-3: $t'_1 = t_0 + 2$

$$\Gamma_{II} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \quad \Gamma_{GII} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & \alpha & 1 & \alpha \end{bmatrix}$$

Pattern	GII	II	Pattern	GII	II
{1,0 0,0}	✓	✓	{1,0 1,0}	✓	✓
{1,1 0,0}	✓	✗	{1,1 1,0}	✓	✗
{2,0 0,0}	✓	✓	{2,0 1,0}	✓	✓
{2,1 0,0}	✓	✗	{2,1 1,0}	✓	✗

- Multiple code levels are allowed in layer 2 and 3
- Layer-3 codes can correct additional interleaves
- Layer-3 codes do not have to be stronger than layer-2 codes

Example 2: same redundancy
multi-level nesting in 3-layer GII code

II codes: $t_1 = t_0 + 2$; $t'_1 = t_0 + 3$

GII codes: $t_1 = t_2 = t_0 + 1$; layer-3: $t'_1 = t_0 + 1$; $t'_2 = t_0 + 2$

Extra erasure patterns correctable by the 3-layer II code but not the 3-layer GII code	{2,0,0 2,0,0 0,0,0}	{2,0,0 2,0,0 1,0,0}
	{2,0,0 2,0,0 2,0,0}	
	{3,0,0 0,0,0 0,0,0}	{3,0,0 1,0,0 0,0,0}
Extra erasure patterns correctable by the 3-layer GII code but not the 3-layer II code	{3,0,0 1,0,0 1,0,0}	{3,0,0 2,0,0 0,0,0}
	{3,0,0 2,0,0 1,0,0}	{3,0,0 2,0,0 2,0,0}
	{1,1,0 0,0,0 0,0,0}	{1,1,0 1,0,0 0,0,0}
	{1,1,0 1,0,0 1,0,0}	{1,1,0 1,1,0 0,0,0}
	{1,1,0 1,1,0 1,0,0}	{1,1,0 1,1,0 1,1,0}
	{1,1,1 0,0,0 0,0,0}	{1,1,1 1,0,0 0,0,0}
	{1,1,1 1,0,0 1,0,0}	{1,1,1 1,1,0 0,0,0}
	{1,1,1 1,1,0 1,0,0}	{1,1,1 1,1,0 1,1,0}
	{2,1,0 0,0,0 0,0,0}	{2,1,0 1,0,0 0,0,0}
	{2,1,0 1,0,0 1,0,0}	{2,1,0 1,1,0 0,0,0}
	{2,1,0 1,1,0 1,0,0}	{2,1,0 1,1,0 1,1,0}
	{2,1,1 0,0,0 0,0,0}	{2,1,1 1,0,0 0,0,0}
	{2,1,1 1,0,0 1,0,0}	{2,1,1 1,1,0 0,0,0}
	{2,1,1 1,1,0 1,0,0}	{2,1,1 1,1,0 1,1,0}
	{2,1,1 1,1,1 0,0,0}	{2,1,1 1,1,1 1,0,0}
	{2,1,1 1,1,1 1,1,0}	




Implementation Architectures for GII Decoders

➤ Two-layer GII codes based on Reed-Solomon codes over $GF(2^8)$

- 8 sub-codewords
- Length of each sub-codeword: 2040 bit
- Redundancy: 12.5%

➤ Hardware complexity

- <30% area increase compared to hard-decision Reed-Solomon decoder
- Clock frequency: >550Mhz on FPGA
- >40GByte/s throughput



* X. Zhang and Z. Xie, "Efficient Architectures for Generalized Integrated Interleaved Decoder," IEEE Trans. on Circuits and Systems-I, 2019.



Conclusions

Flash Memory Summit

- GII codes can achieve hyper throughput with excellent correction capability and low complexity
- Modified two-layer GII codes improve the locality of erasure correction without degradation on the correction capability for most practical settings
- Three-layer GII codes further improve the locality of erasure and error correction and the third-layer parities can be used in a flexible way
- Three-layer GII codes achieve better locality than two-layer GII codes at the cost of higher redundancy