Optimized Degree Distributions for Binary and Non-Binary LDPC Codes in Flash Memory

ISITA 2014

Kasra Vakilinia
Contributors: Dariush Divsalar, and Richard D. Wesel

Communication Systems Laboratory
University of California, Los Angeles

October 27, 2014

Outline

Outline

Outline

Outline

Outline

NAND Flash Memory

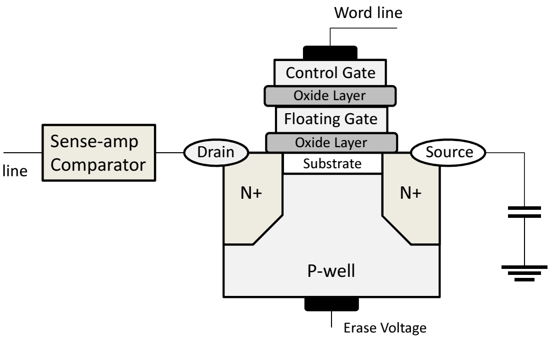

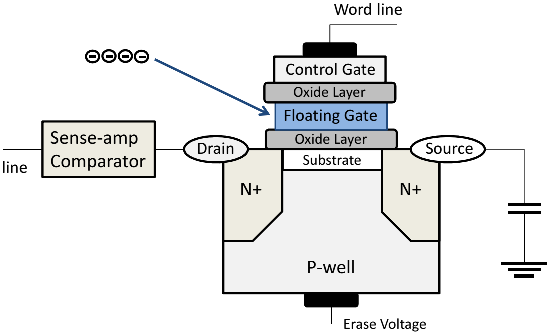

To store information, add a specified amount of charge to the floating gate.

NAND Flash Memory

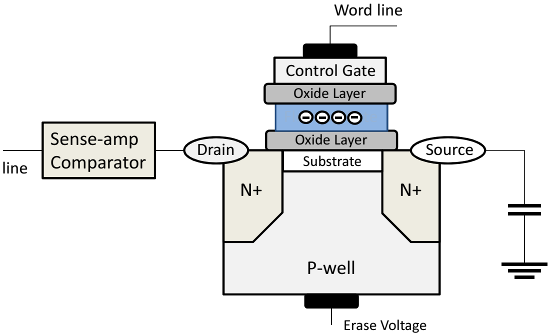

To store information, add a specified amount of charge to the floating gate.

NAND Flash Memory

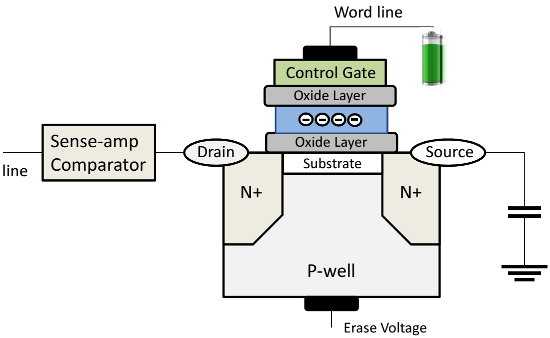

To read information, apply a specified word-line voltage to the control gate.

NAND Flash Memory

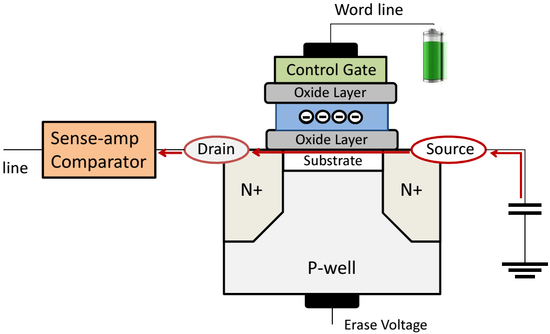

To read information, apply a specified word-line voltage to the control gate.
The sense-amp comparator then provides a single bit of information (whether the transistor is "on", i.e. the drain current is above a specified threshold).
The threshold voltage is the lowest voltage at which the transistor turns on.

Variations in Threshold Voltage

- The threshold voltage is proportional to the amount of charge in the floating gate.

- The threshold voltage is proportional to the amount of charge in the floating gate.
- The actual charge level written to the floating gate can vary with:
  Overcharge in the write operation
  Leakage in the retention period
  Interference from nearby cells

Variations in Threshold Voltage

- The threshold voltage is proportional to the amount of charge in the floating gate.
- The actual charge level written to the floating gate can vary with:
  Overcharge in the write operation
  Leakage in the retention period
  Interference from nearby cells
- So, there is a *distribution* associated with the threshold voltage.

One Read

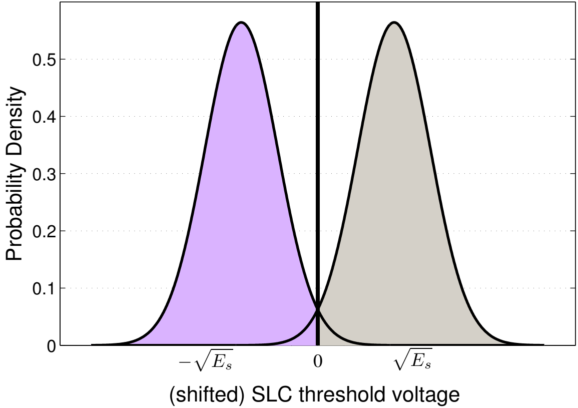

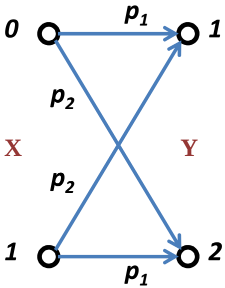

(a) One read

One read system is equivalent to binary symmetric channel (BSC).

Two Reads

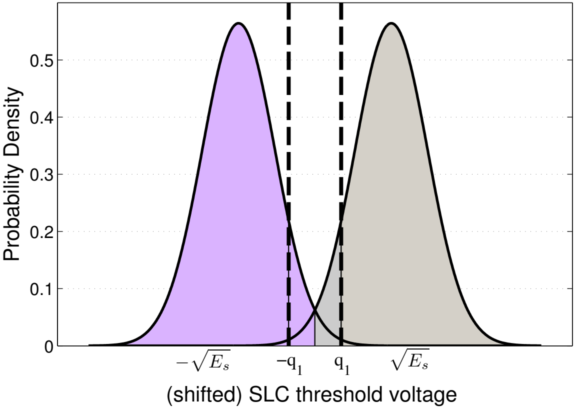

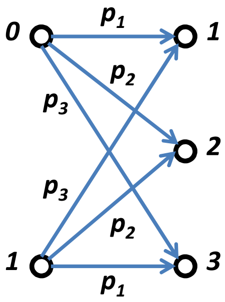

(b) Two reads

With two symmetric reads at q_1 and $-q_1$, the equivalent channel has an erasure region.

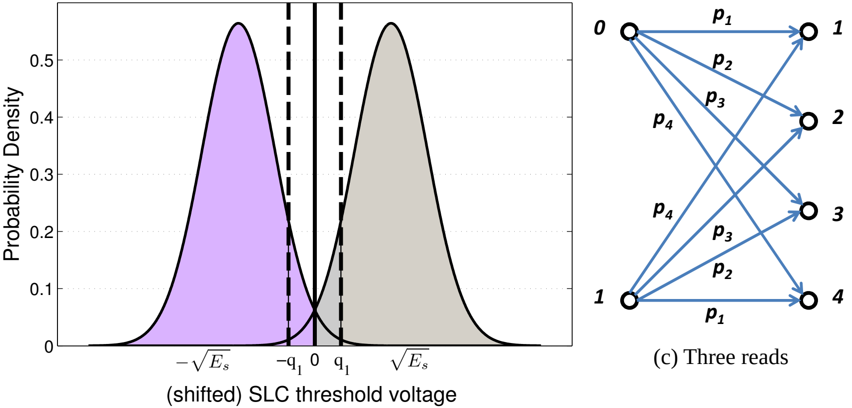

(c) Three reads

Additional reads result in additional outputs and more complicated equivalent communication channel.

Five Reads

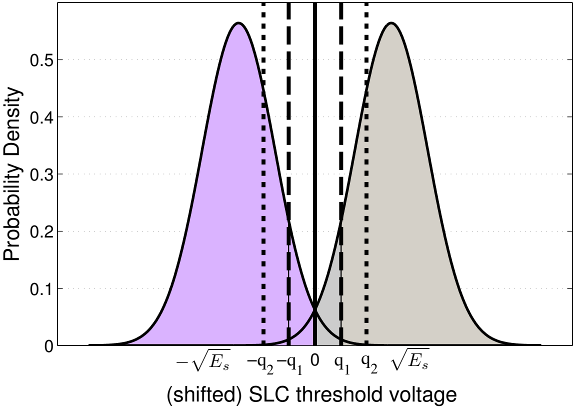

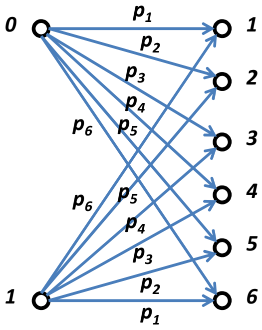

(d) Five reads

For n reads, the channel has n+1 outputs.

Mutual Information Maximization for One Read

Find the read voltages to maximize the mutual information.

$$\textit{maximize } I(X;Y) = H(Y) - H(Y|X)$$

For the 1-read system, the optimized read is at zero.

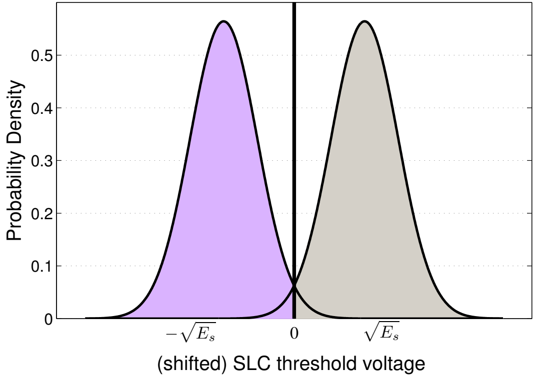

Find the threshold voltages to maximize the mutual information.

$$maximize\ I(X;Y) = H(Y) - H(Y|X)$$

For 2,3 reads, q_1 is the solution to $\frac{dI}{dq_1} = 0$.

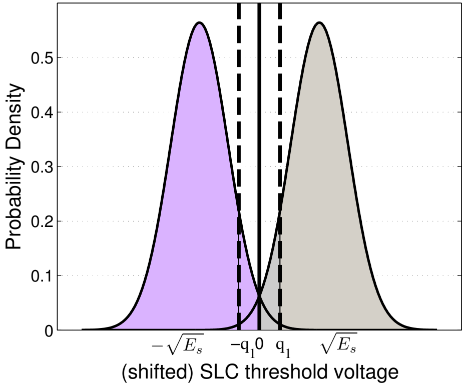

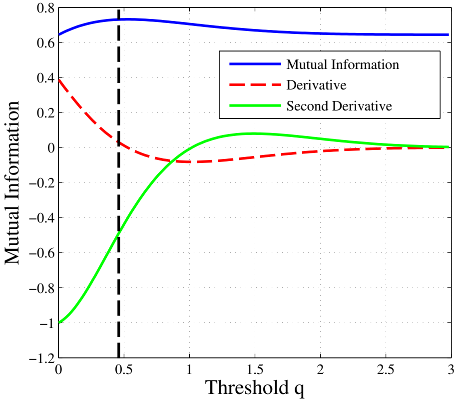

Mutual Information Maximization for Five Reads

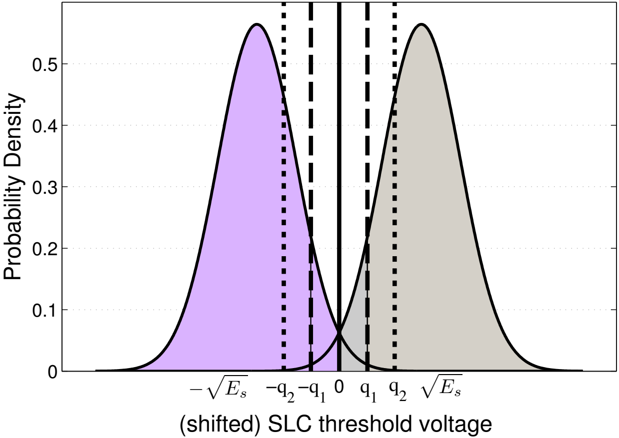

For the 5 reads, the Max MI is the solution of $\frac{dI}{dq_1} = 0$ and $\frac{dI}{dq_2} = 0$.

Mutual Information Maximization for Five Reads

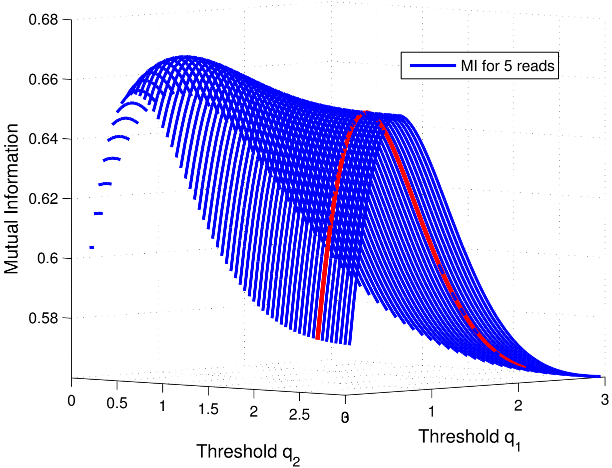

For the 5 reads, the Max MI is the solution of $\frac{dI}{dq_1} = 0$ and $\frac{dI}{dq_2} = 0$.

If $q_2 \geqslant q_1$, MI is quasi-concave in q_1 for a fixed value of q_2 and vice-versa.

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at three reads to optimize the MI with five reads.

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at three reads to optimize the MI with five reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at three reads to optimize the MI with five reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Progressive 3-read channel [x -0.4 0 0.4 x]: $MI = 0.643$

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at three reads to optimize the MI with five reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Progressive 3-read channel [x -0.4 0 0.4 x]: $MI = 0.643$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at three reads to optimize the MI with five reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Progressive 3-read channel [x -0.4 0 0.4 x]: $MI = 0.643$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at five reads to optimize the MI at three reads.

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at three reads to optimize the MI with five reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Progressive 3-read channel [x -0.4 0 0.4 x]: $MI = 0.643$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at five reads to optimize the MI at three reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at three reads to optimize the MI with five reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Progressive 3-read channel [x -0.4 0 0.4 x]: $MI = 0.643$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at five reads to optimize the MI at three reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$

Read Voltage Optimization

- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at three reads to optimize the MI with five reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Progressive 3-read channel [x -0.4 0 0.4 x]: $MI = 0.643$
- Fully optimized for 5-read channel [-0.9 -0.4 0 0.4 0.9]: $MI = 0.668$

- We can sacrifice some MI at five reads to optimize the MI at three reads.
- Fully optimized single-read channel [x x 0 x x]: $MI = 0.560$
- Fully optimized for 3-read channel [x -0.61 0 0.61 x]: $MI = 0.652$
- Progressive 5-read channel [-0.61 -0.28 0 0.28 0.61]: $MI = 0.663$

Concluding Remarks on MMI

Optimized single read	$MI = 0.560$	[x x 0 x x]
Optimized three reads	$MI = 0.652$	[x -0.61 0 0.61 x]
Optimized five reads	$MI = 0.668$	[-0.9 -0.4 0 0.4 0.9]
Progressive three reads	$MI = 0.643$	[x -0.4 0 0.4 x]
Progressive five reads	$MI = 0.663$	[-0.61 -0.28 0 0.28 0.61]

The MI differences are small, but large enough to have noticeable effects on frame error rate (according to [1]).

Optimizing q_1 for three-read performance maximizes the probability of decoding successfully after three reads.

Optimizing q_1 for the five-read scenario will produce the lowest probability of losing a page.

[1] Wang, J., et al. "Enhanced Precision Through Multiple Reads for LDPC Decoding in Flash Memories," IEEE JSAC , May 2014

LDPC Code Design and Threshold Calculation

In order to design LDPC codes that are matched to the Flash read channel, we use EXIT functions with the Gaussian approximation and Reciprocal Channel Approximation (faster alternative to Density Evolution) to compute LDPC decoding thresholds for different number of reads.

Two Read Scenario

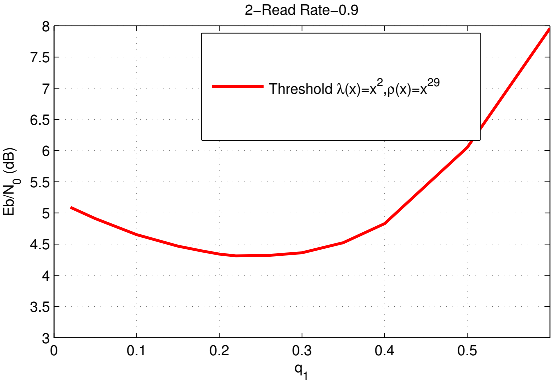

For a fixed degree distribution and q_1, RCA-EXIT analysis can determine the word-line voltages that minimize E_b/N_0 threshold for multiple-read Flash channels.

Two Read Scenario

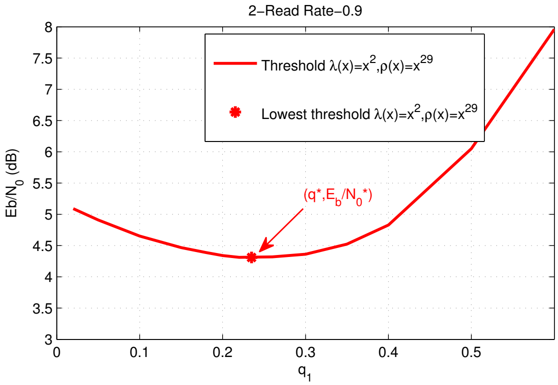

For a fixed degree distribution, there is a q_1 that has the absolute E_b/N_0 threshold for multiple-read Flash channels.

Eb/N_0 vs q_{MMI}

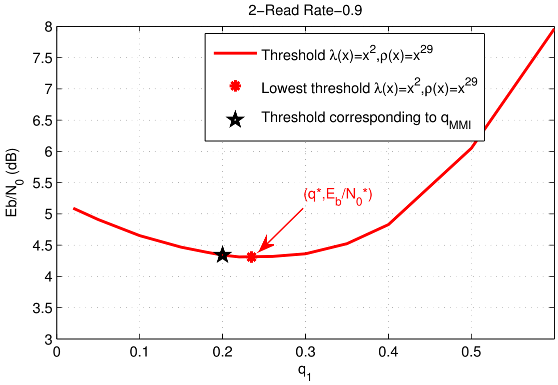

For each Eb/N_0 there is a q_1 that maximizes the MI.

Eb/N_0 vs q_{MMI}

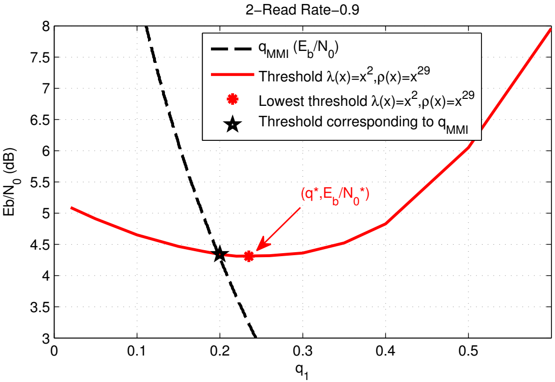

The threshold Eb/N_0 achieved with q_1 chosen to maximize MI is only 1% away from the best possible Eb/N_0^* achieved with a density-evolution optimized q_1.

Optimized Degree Distributions

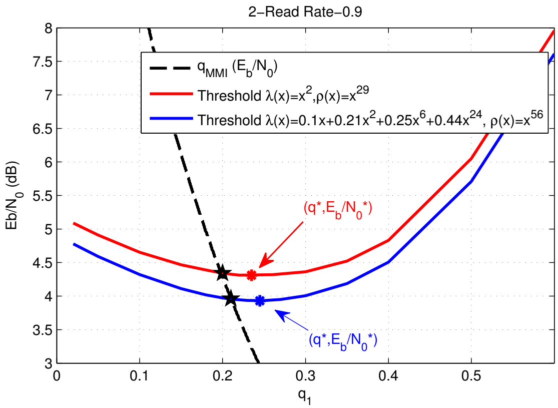

The optimized degree distribution is found by calculating the threshold at q_{MMI} and further optimizing the q_1.

Different Optimized Codes for Different Number of Reads

Table: Thresholds for optimized degree distributions on SLC flash with 1, 3, and 5 reads as well as soft information.

Target Threshold	1 read Th.	3 reads Th.	5 reads Th.	Soft Th.
Optimized for 1 read	**4.752**	3.728	3.542	3.398
Optimized for 3 reads	4.923	**3.640**	3.441	3.295
Optimized for 5 reads	4.926	3.649	**3.437**	3.288
Optimized for Soft	4.926	3.662	3.443	**3.275**
Shannon-Limit	4.400	3.495	3.328	3.198

Different codes have different thresholds for different levels of quantization.
A code designed for one level of quantization is not the best for another level.

Table: Gap of thresholds for optimized degree distributions on SLC flash with 1, 3, and 5 reads as well as soft information.

Target threshold gap	1 read Th. gap	3 reads Th. gap	5 reads Th. gap	Soft Th. gap
Optimized for 1 read	**0**	0.088	0.105	0.123
Optimized for 3 reads	0.171	**0**	0.004	0.020
Optimized for 5 reads	0.174	0.009	**0**	0.013
Optimized for Soft	0.1740	0.022	0.006	**0**

The code optimized for a single read performs well across multiple reads. There is reason to consider using the code designed for a single read.

Regular LDPC Code Baselines and Bounds

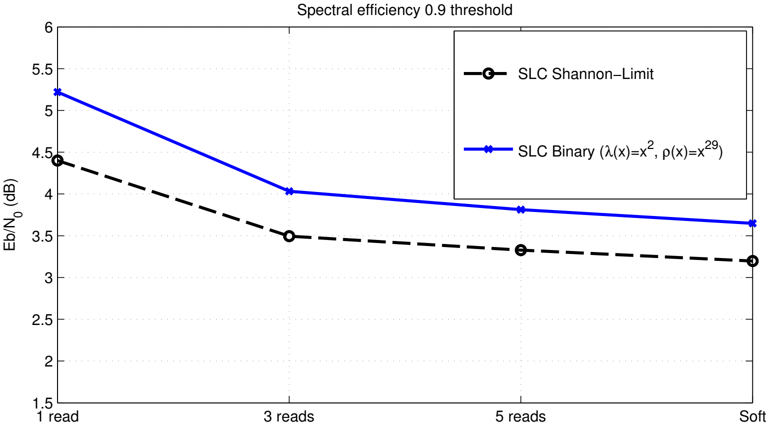

The threshold of two regular binary LDPC codes vs. Shannon-limit for different number of reads.

Non-Binary Code for SLC

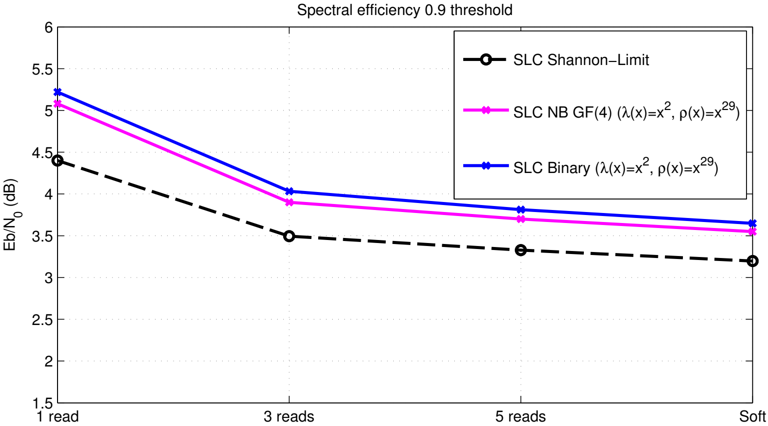

The threshold of regular NB-LDPC code for different number of reads.

Optimized Binary LDPC Code

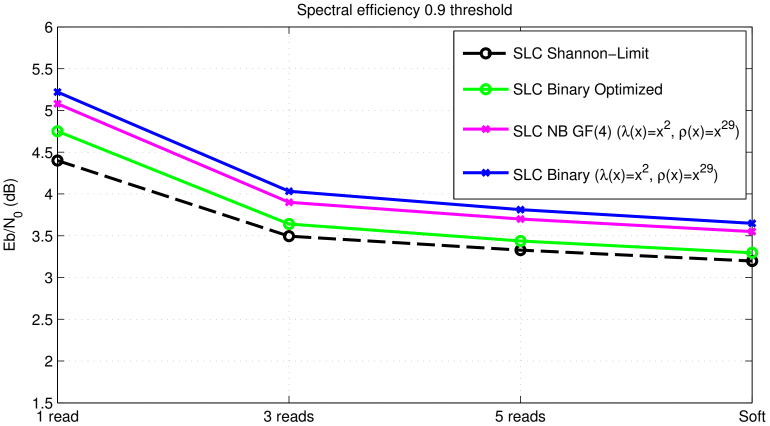

The threshold of an optimized binary LDPC code for different number of reads.

Non-Binary code on MLC

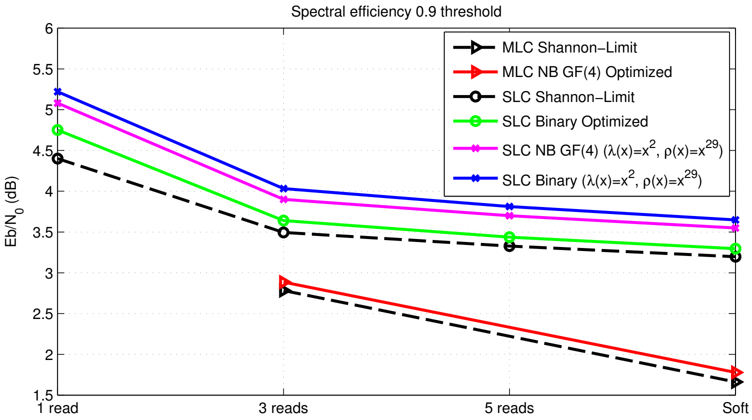

Huge gain can be achieved by using a rate-0.45 NB-LDPC code on MLC with the same number of reads and for the same spectral efficiency of 0.9.

Non-Binary code on MLC

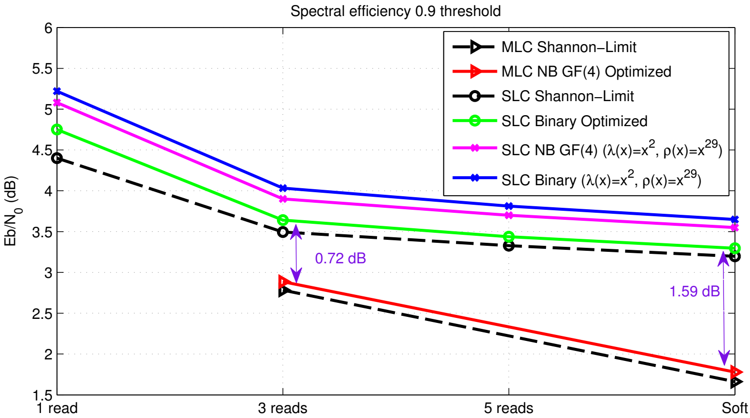

Huge gain can be achieved by using a rate-0.45 NB-LDPC code on MLC with the same number of reads and for the same spectral efficiency of 0.9.

Table of Degree Distributions

# Reads	Coefficients					
SLC 1 read	λ_2 0.07	λ_3 0.25	λ_7 0.11	λ_8 0.13	λ_{27} 0.44	ρ_{61} 1
SLC 3 reads	λ_2 0.1	λ_3 0.21	λ_6 0.11	λ_7 0.12	λ_{26} 0.46	ρ_{56} 1
SLC 5 reads	λ_2 0.11	λ_3 0.21	λ_5 0.09	λ_8 0.14	λ_{25} 0.45	ρ_{56} 1
SLC soft	λ_2 0.11	λ_3 0.21	λ_6 0.21	λ_8 0.14	λ_{26} 0.47	ρ_{56} 1
MLC	λ_2 0.16	λ_3 0.31	λ_4 0.1	λ_7 0.18	λ_8 0.1	λ_{11} 0.15
3 reads	ρ_5 0.45	ρ_6 0.16	ρ_8 0.1	ρ_{22} 0.29		
MLC	λ_2 0.15	λ_3 0.3	λ_4 0.1	λ_7 0.09	λ_9 0.22	λ_{11} 0.14
soft	ρ_5 0.44	ρ_6 0.07	ρ_8 0.1	ρ_9 0.11	ρ_{22} 0.27	

Conclusions

- q_{MMI} is the solution to $\frac{dI}{dq_1} = 0$ (and $\frac{dI}{dq_2} = 0$) and can be computed quickly.

Conclusions

- q_{MMI} is the solution to $\frac{dI}{dq_1} = 0$ (and $\frac{dI}{dq_2} = 0$) and can be computed quickly.
- Optimized Degree Distributions can be found by calculating the threshold at q_{MMI} and further optimizing the q to get the lowest E_b/N_0 thresholds.

Conclusions

- q_{MMI} is the solution to $\frac{dI}{dq_1} = 0$ (and $\frac{dI}{dq_2} = 0$) and can be computed quickly.
- Optimized Degree Distributions can be found by calculating the threshold at q_{MMI} and further optimizing the q to get the lowest E_b/N_0 thresholds.
- Even though a code designed for one level of quantization is not the best for another level of quantization, the code optimized for a single read performs well across multiple reads.

Conclusions

- q_{MMI} is the solution to $\frac{dI}{dq_1} = 0$ (and $\frac{dI}{dq_2} = 0$) and can be computed quickly.
- Optimized Degree Distributions can be found by calculating the threshold at q_{MMI} and further optimizing the q to get the lowest E_b/N_0 thresholds.
- Even though a code designed for one level of quantization is not the best for another level of quantization, the code optimized for a single read performs well across multiple reads.
- Using a low-rate LDPC code on MLC for the same spectral efficiency as a high-rate code on SLC can provide huge gain in terms of the thresholds.