Flash Memory Summit

Error Correction Concepts for Modern Multi-Level Flash

MICROCHIP

A Leading Provider of Smart, Connected and Secure Embedded Control Solutions

Peter Graumann, Technical Fellow

August 8, 2023

SMART | CONNECTED | SECURE

Agenda

- Background
- LDPC FER vs. BER
- LDPC FER vs. Err
- NAND Flash Error Profile
- Performance Analysis
- Conclusions

Agenda

- **Background**
- LDPC FER vs. BER
- LDPC FER vs. Err
- NAND Flash Error Profile
- Performance Analysis
- Conclusions

Motivation

- **NAND Flash technology continues to advance with increasing layer counts and increasing numbers of bits per cell**

- **Bit-error rates of these NAND devices continue to increase requiring improved performance from error correction solutions**

- **Error correction area and power consumption must be reduced to keep system cost down**

- **How can error correction solutions be optimized to meet the needs of modern NAND devices?**
  - BER vs. Err error modeling
  - 4k or 16k byte block size

MICROCHIP

Definitions

- **FEC Block**
  - A block of data containing both user data and LDPC parity
  - Code Rate (CR) is the ratio of the data_size/(data_size + parity_size)
- **Bit-Error Rate (BER)**
  - Error insertion such that every bit in a block has an equal chance to be in error
- **Error Rate (Err)**
  - Error insertion such that every block has an equal number of errored bits
- **Frame-Error Rate (FER)**
  - The probability that an FEC block is in error after LDPC decoding

Agenda

- Background
- **LDPC FER vs. BER**
- LDPC FER vs. Err
- NAND Flash Error Profile
- Performance Analysis
- Conclusions

MICROCHIP

Test LDPC Codes

- **Test LDPC Codewords**
  - 4k byte block
    - 4,352 bytes of data, 4,736 bytes of data and parity (CR=0.919)
    - 4,352 bytes of data, 4,992 bytes of data and parity (CR=0.871)
  - 16k byte blocks
    - Exactly 4x the size of the 4k byte blocks
    - 17,408 bytes of data, 18,952 bytes of data and parity (CR=0.919)
    - 17,408 bytes of data, 19,968 bytes of data and parity (CR=0.871)

- **Experiments conducted with similar codes and using a generic min-sum decoding algorithm**

LDPC FER vs. BER

- Red curves are for CR 0.919 codes and blue curves are for CR 0.871 codes
- Dashed curves are 4k byte blocks and solid curves are 16k byte blocks
- 16k curves have much steeper waterfall curves

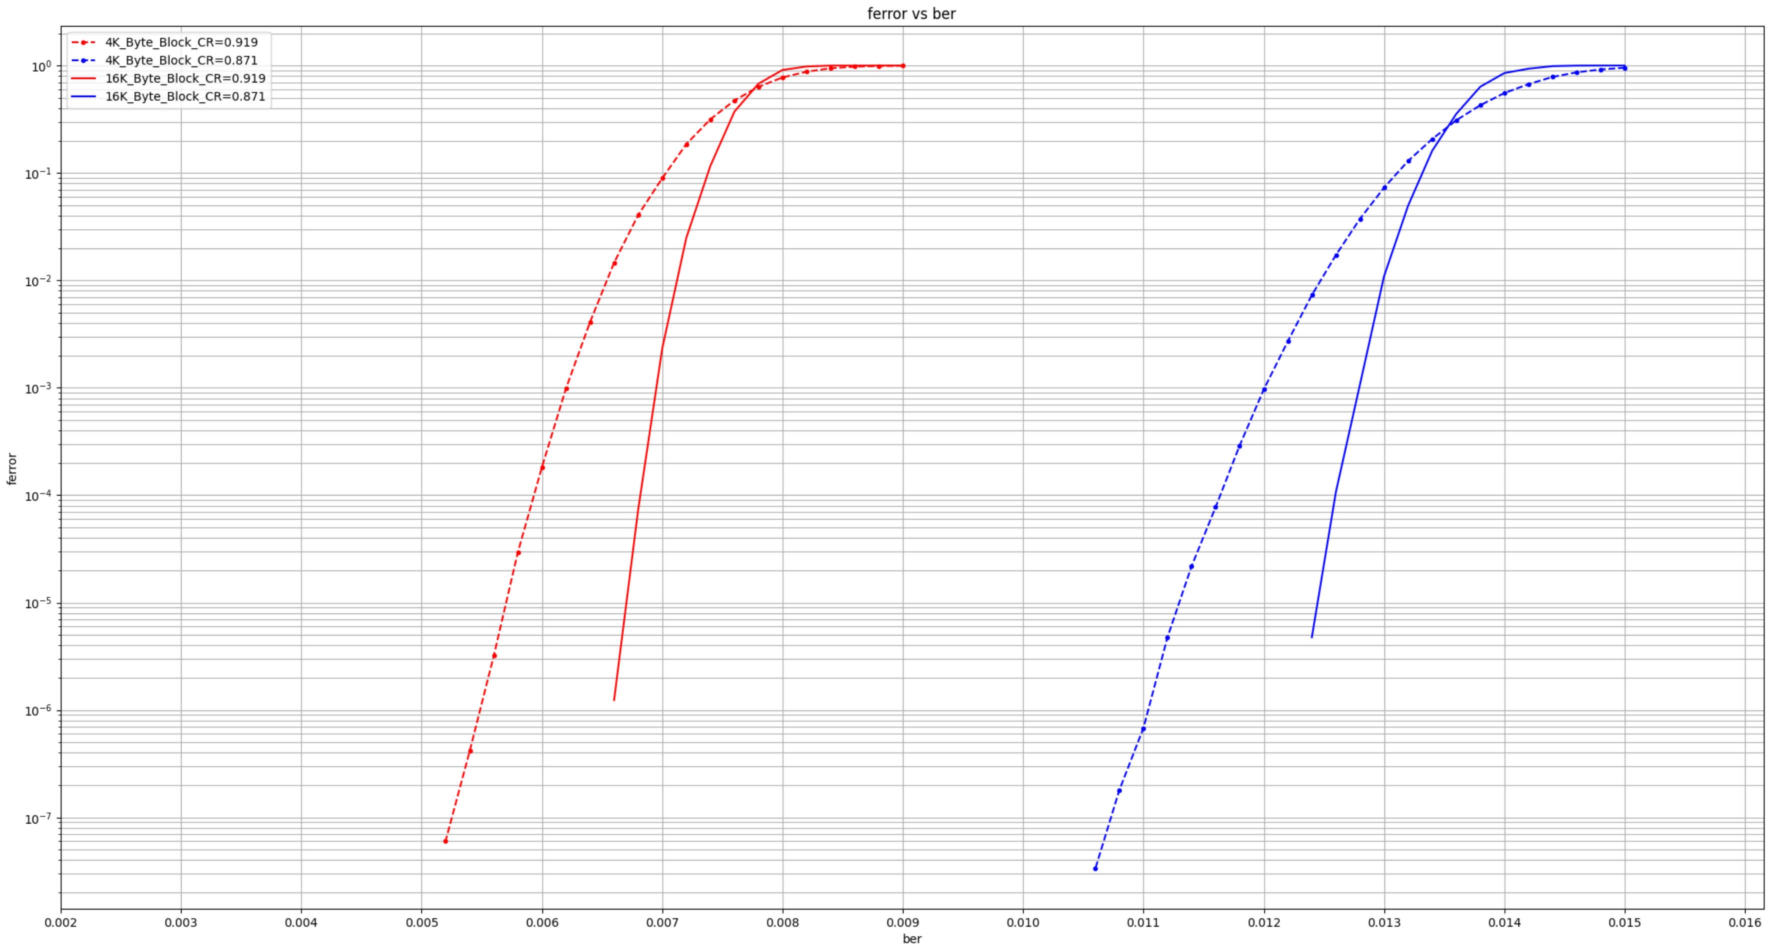

MICROCHIP

LDPC FER vs. BER (16k Equivalent Blocks)

- Fair FER comparison must be done on the same block size
- Scale the 4k curve to match 16k
- 16 keq = 1-(1-FER)^4
- ~15% Improvement in BER @ FER=1e-4
- Future 4k FER plots will show 16k equivalent

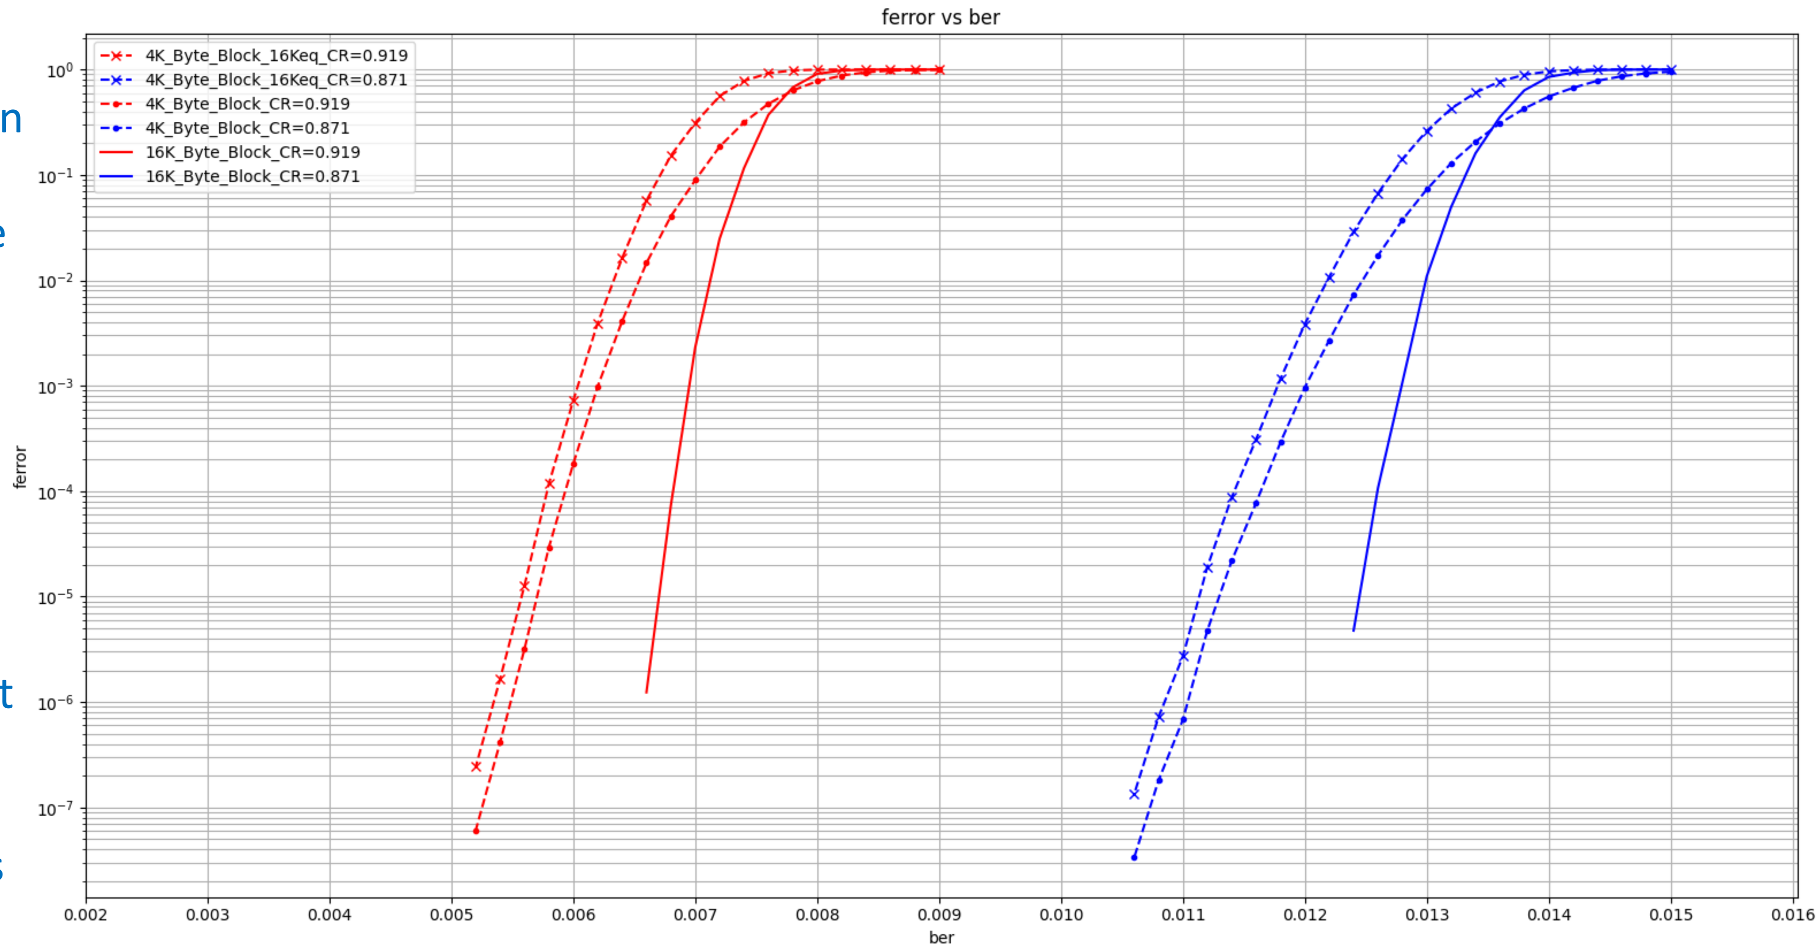

MICROCHIP

BER 4k Byte Error PDF

BER error distribution

- BER error model leads to Gaussian distribution of errors per block
- Plots show the number of errors per 4k byte block
- Higher error rate gives a wider distribution
- 16k blocks have tighter distributions
- Infinitely large blocks would have single-entry distribution

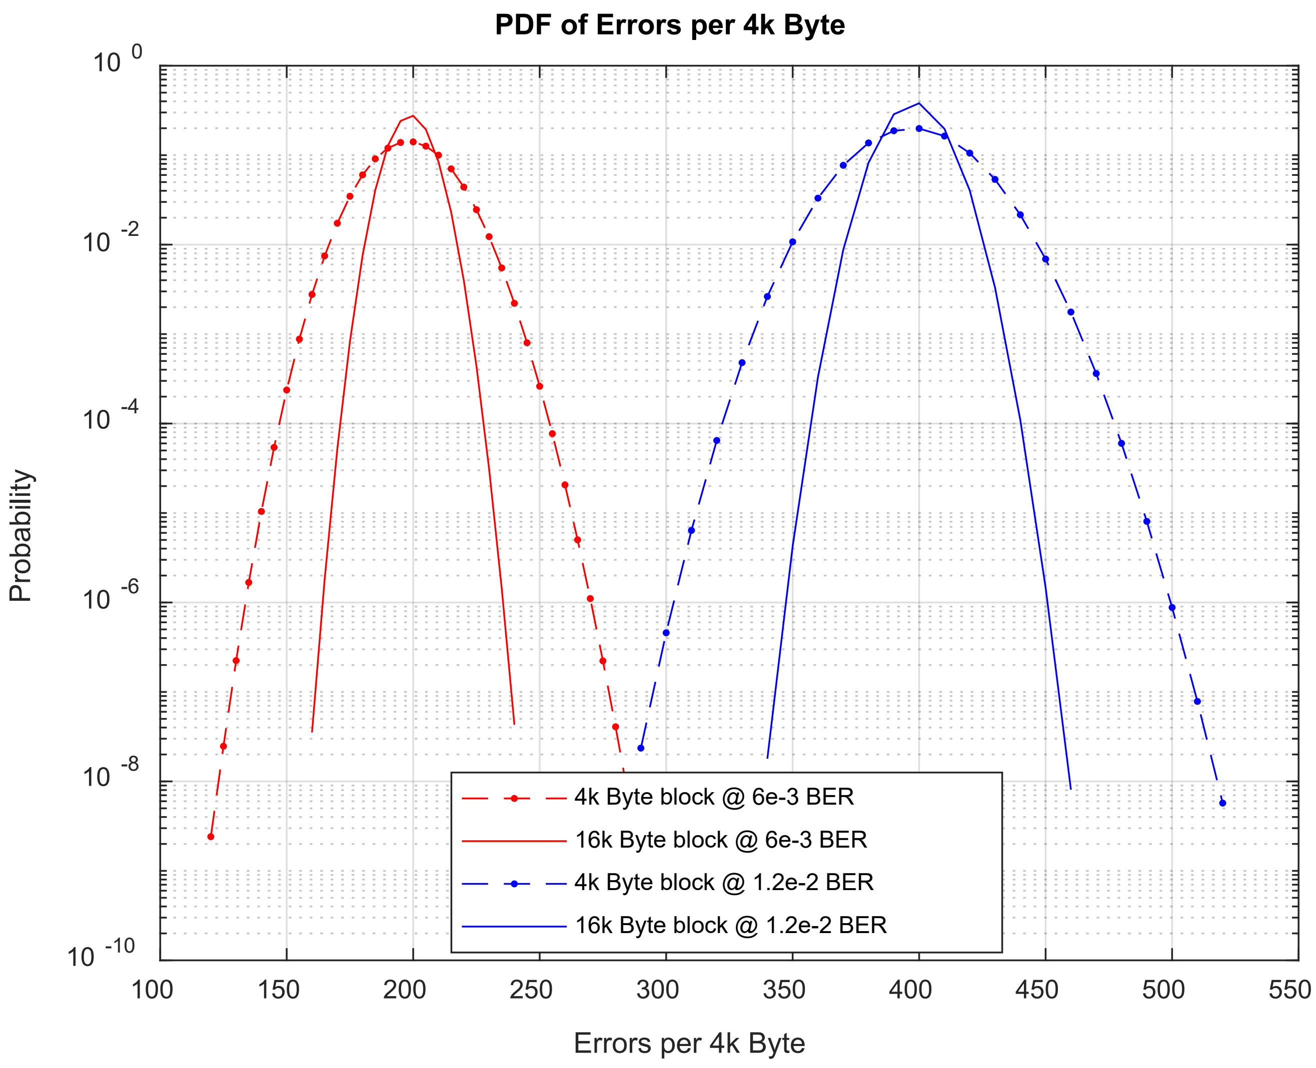

LDPC Average Iterations vs. BER

- We use average iterations as a proxy for throughput
- Throughput is inversely proportional to average iterations
- Some advantage to 4k at high BERs
- Some advantage to 16k blocks at lower BERs

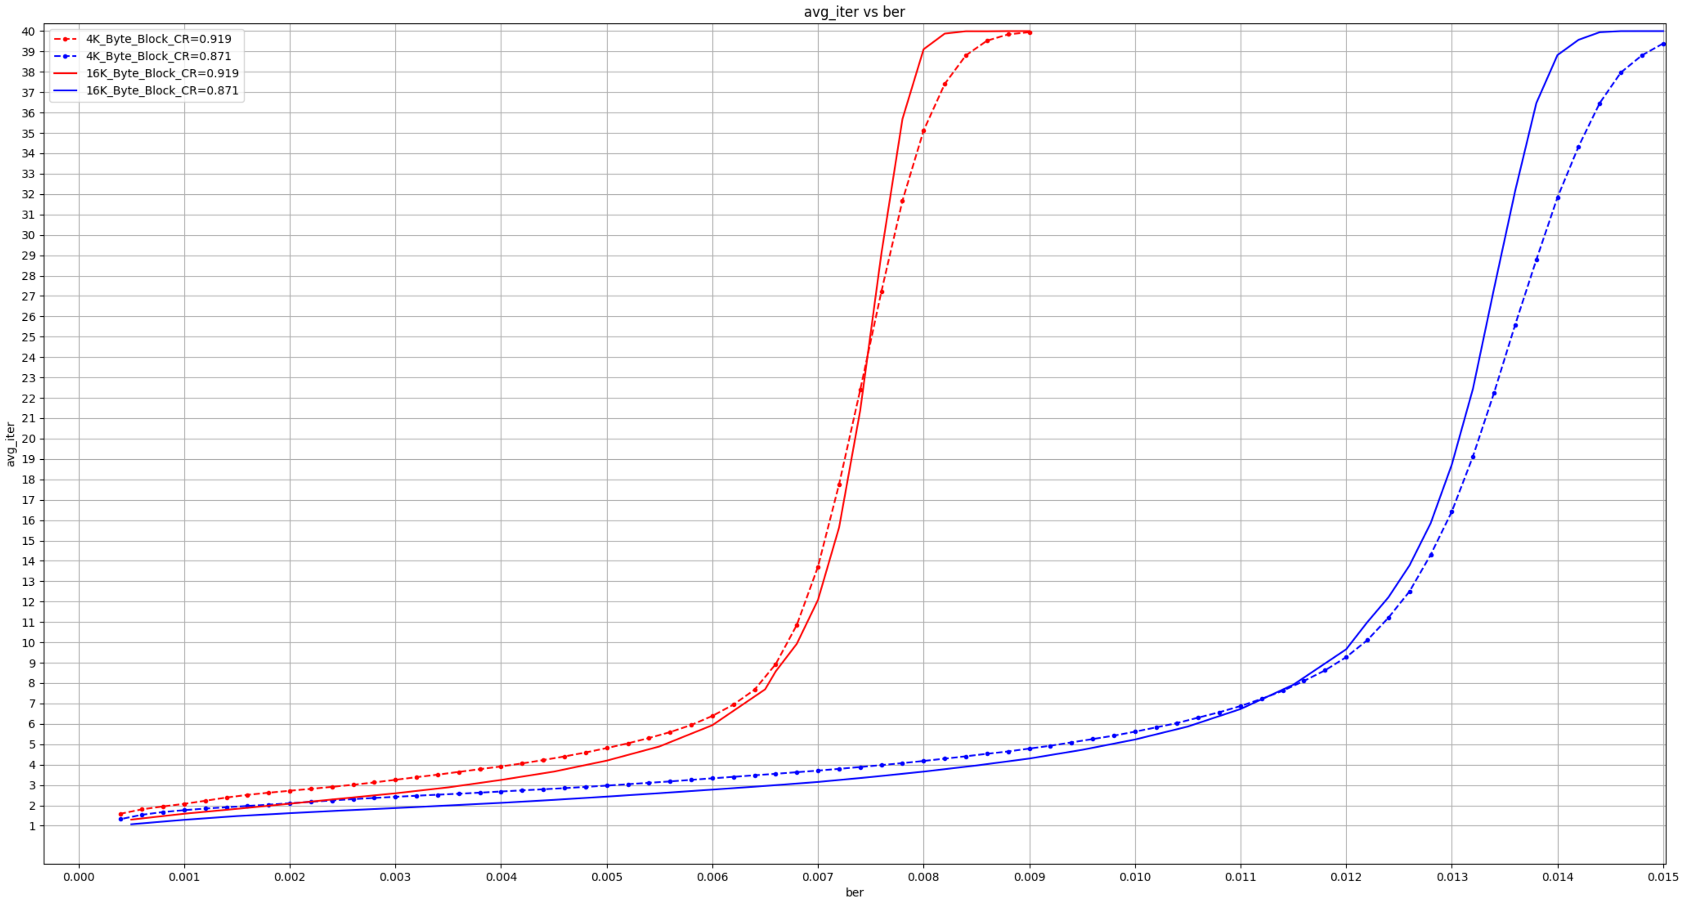

Agenda

- Background
- LDPC FER vs. BER
- **LDPC FER vs. Err**
- NAND Flash Error Profile
- Performance Analysis
- Conclusions

LDPC FER vs. Err (Scaled)

- 16k blocks are solid curves, 4k are dashed
- Red and blue are BER, green and purple are Err
- BER is better at the top of the curves, Err is much better otherwise
- The difference between 4k and 16k performance is reduced with Err model

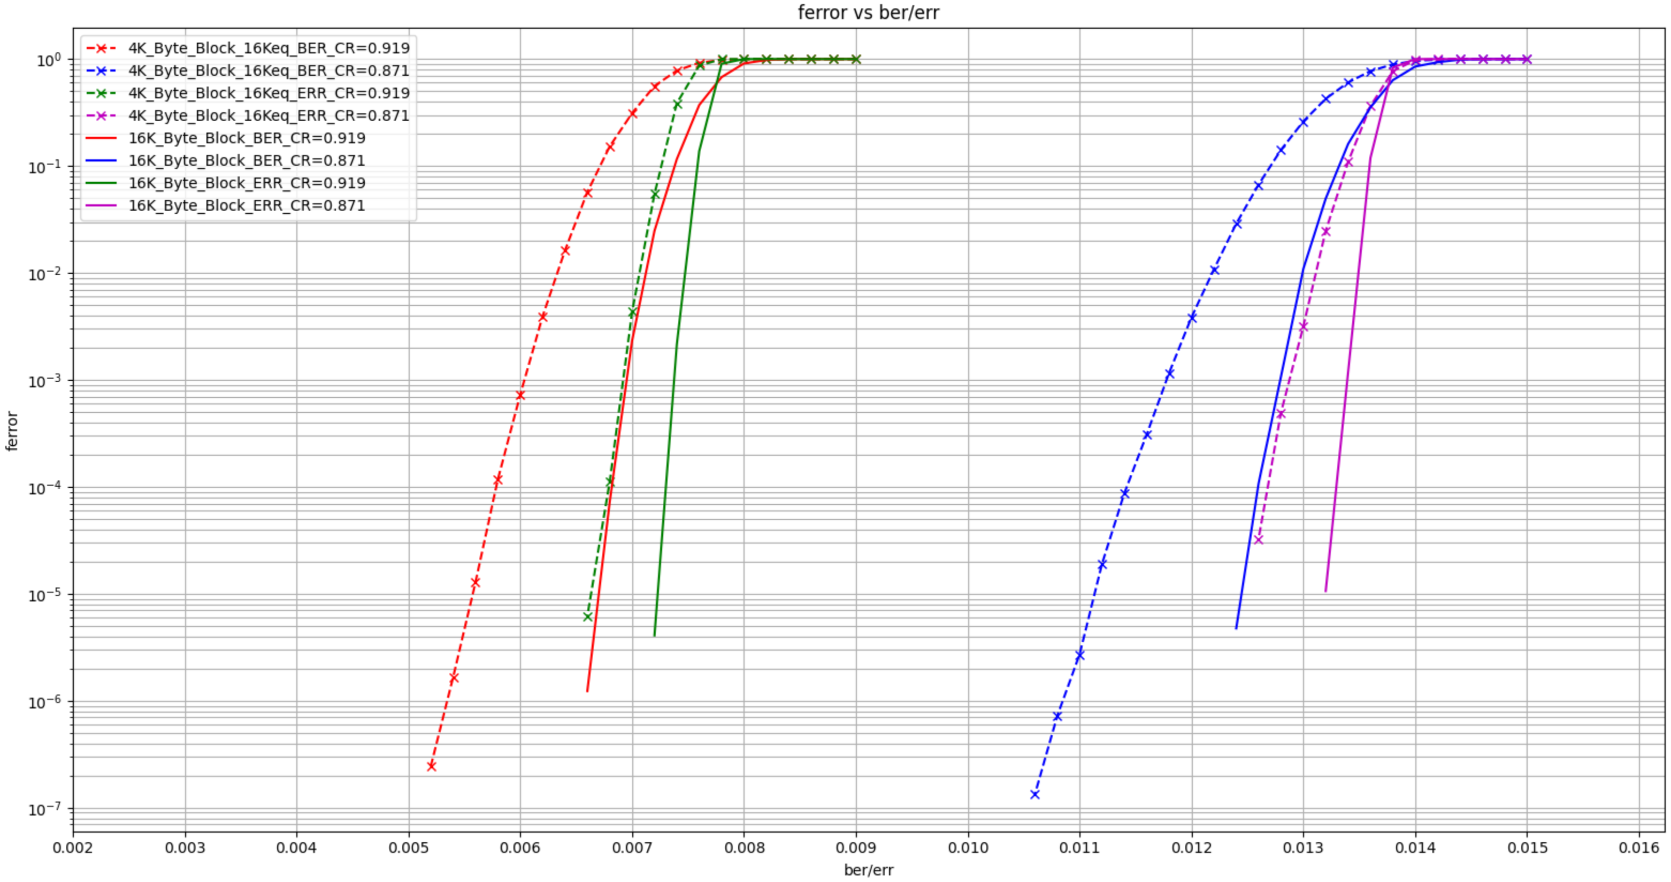

BER 4k Byte Error PDF

Err Error Distribution

- Single number of errors at every Err value instead of a distribution of errors
- Number of errors = Err * block size

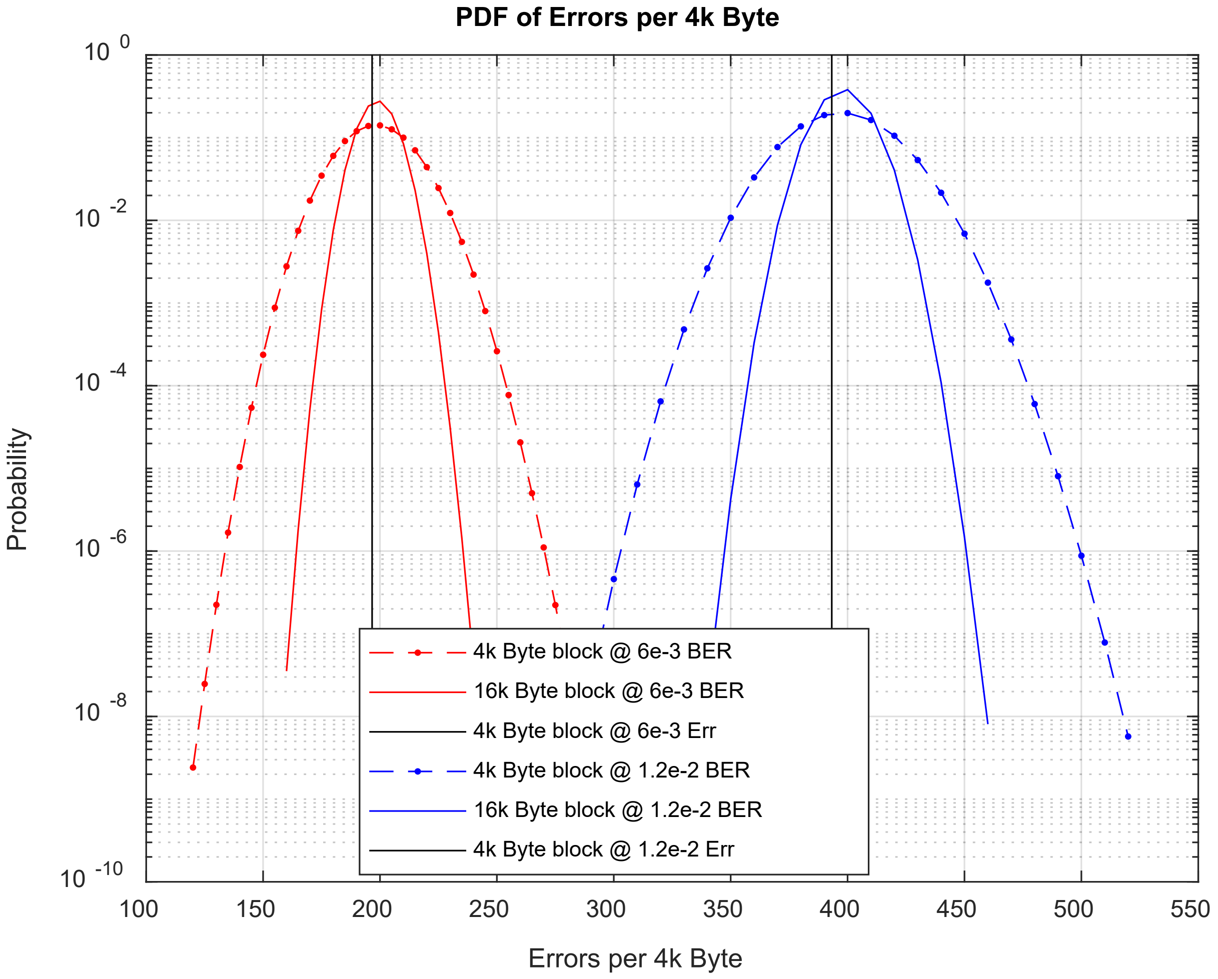

LDPC Average Iterations vs. Err

- Similar average iteration plots for Err compared to BER
- Average iteration measurements combine the full distribution of blocks

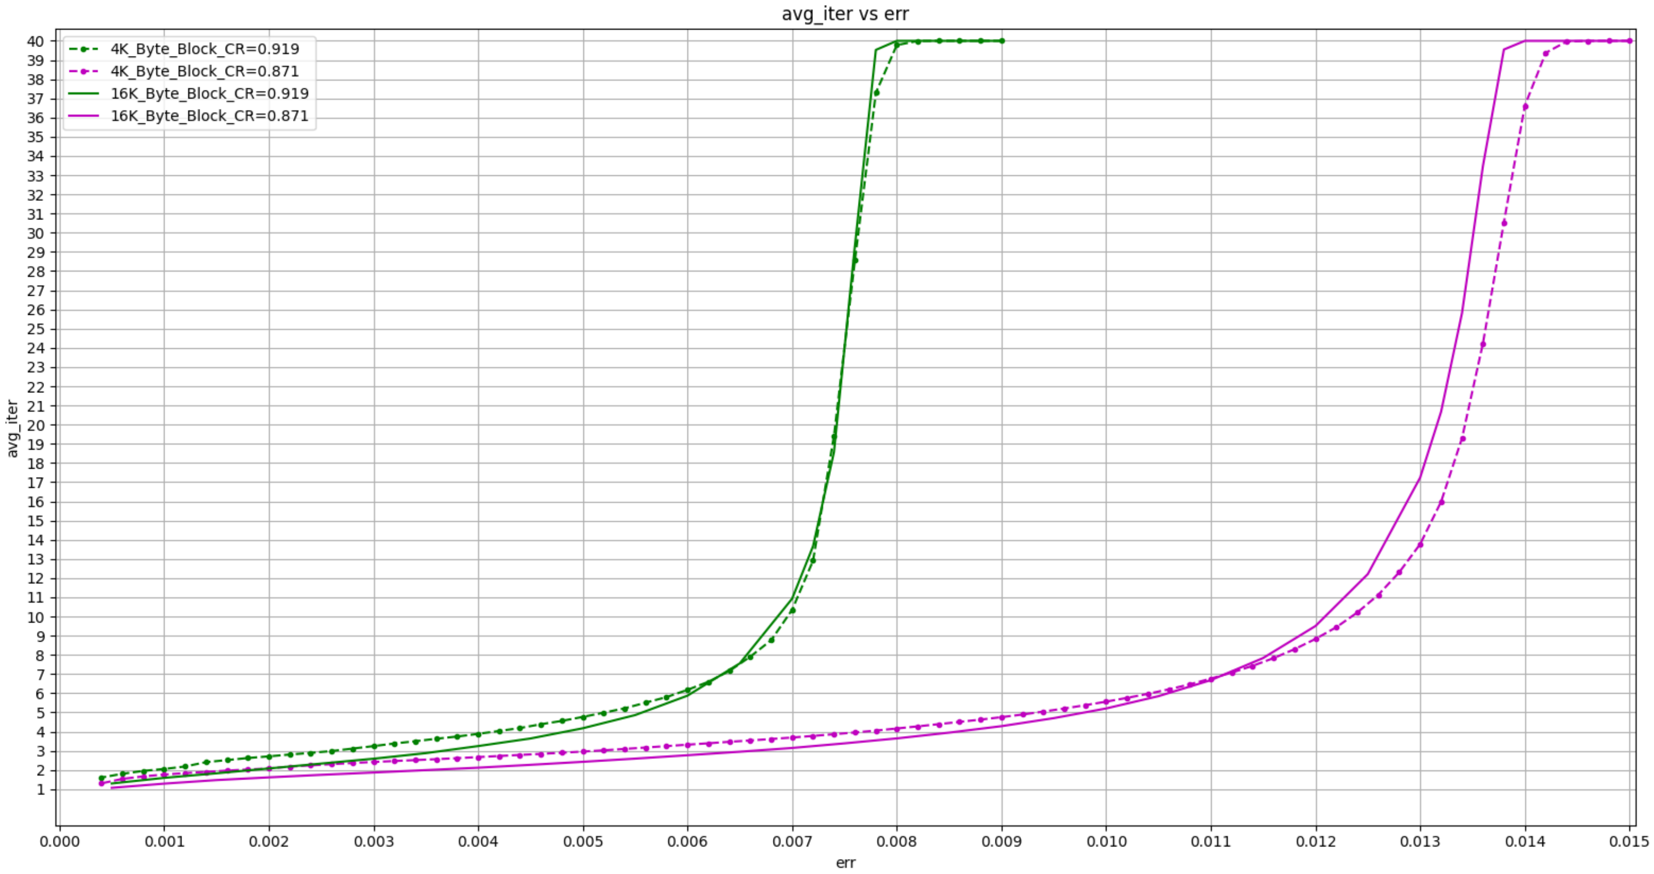

Agenda

- Background
- LDPC FER vs. BER
- LDPC FER vs. Err
- **NAND Flash Error Profile**
- Performance Analysis
- Conclusions

Sample NAND Flash Error Distribution

Sample Err PDFs from real Flash devices are shown

- Y-axis is the probability of occurrence
- X-axis is Err rate
  - Number of errors in any block = Err * block_size

- Most blocks have low Err but there is a long tail of blocks with high Err

- BER model does not provide a good match to this distribution

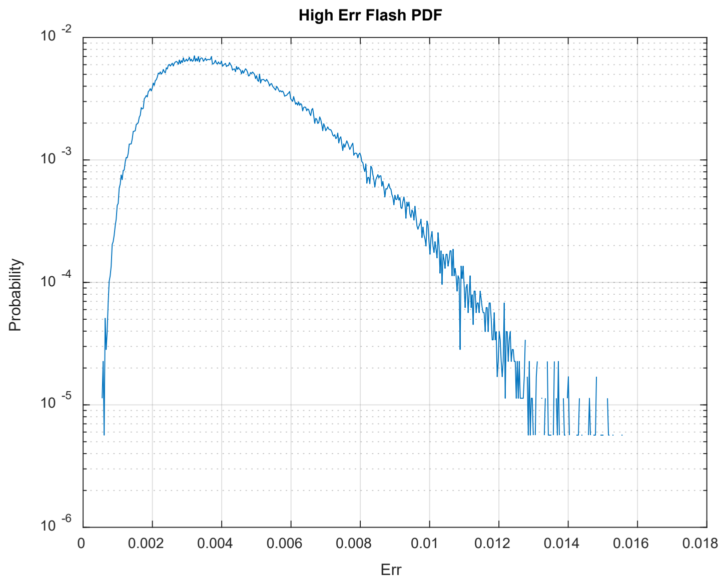

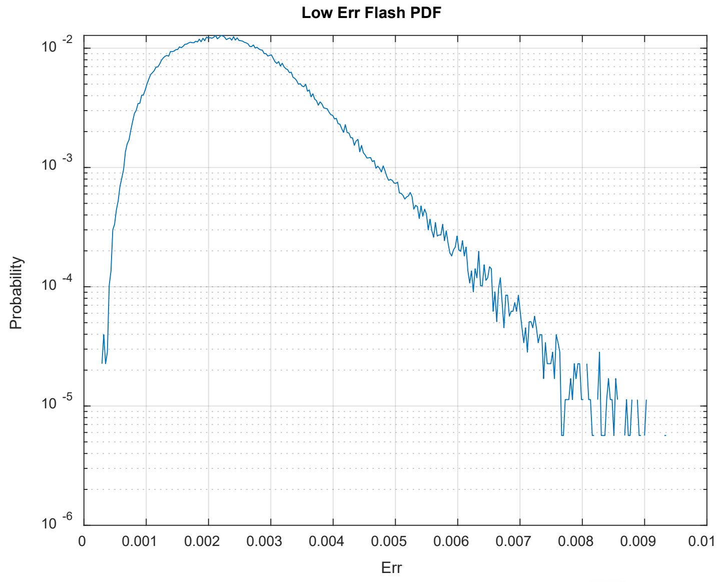

Error Modeling (PDF)

Log-Normal Distribution

- A random variable (x) whose logarithm (y) is normally distributed
  - y = ln(x) has a normal distribution
- We can select the mean (mu) and the standard deviation (sigma) of the LN distribution
- Two examples shown
  - (red) mu = -5.5, sigma = 0.25 (LN1)
  - (red) Average Err = 4.2e-3
  - (blue) mu = -5.0, sigma = 0.25 (LN2)
  - (blue) Average Err = 6.7e-3

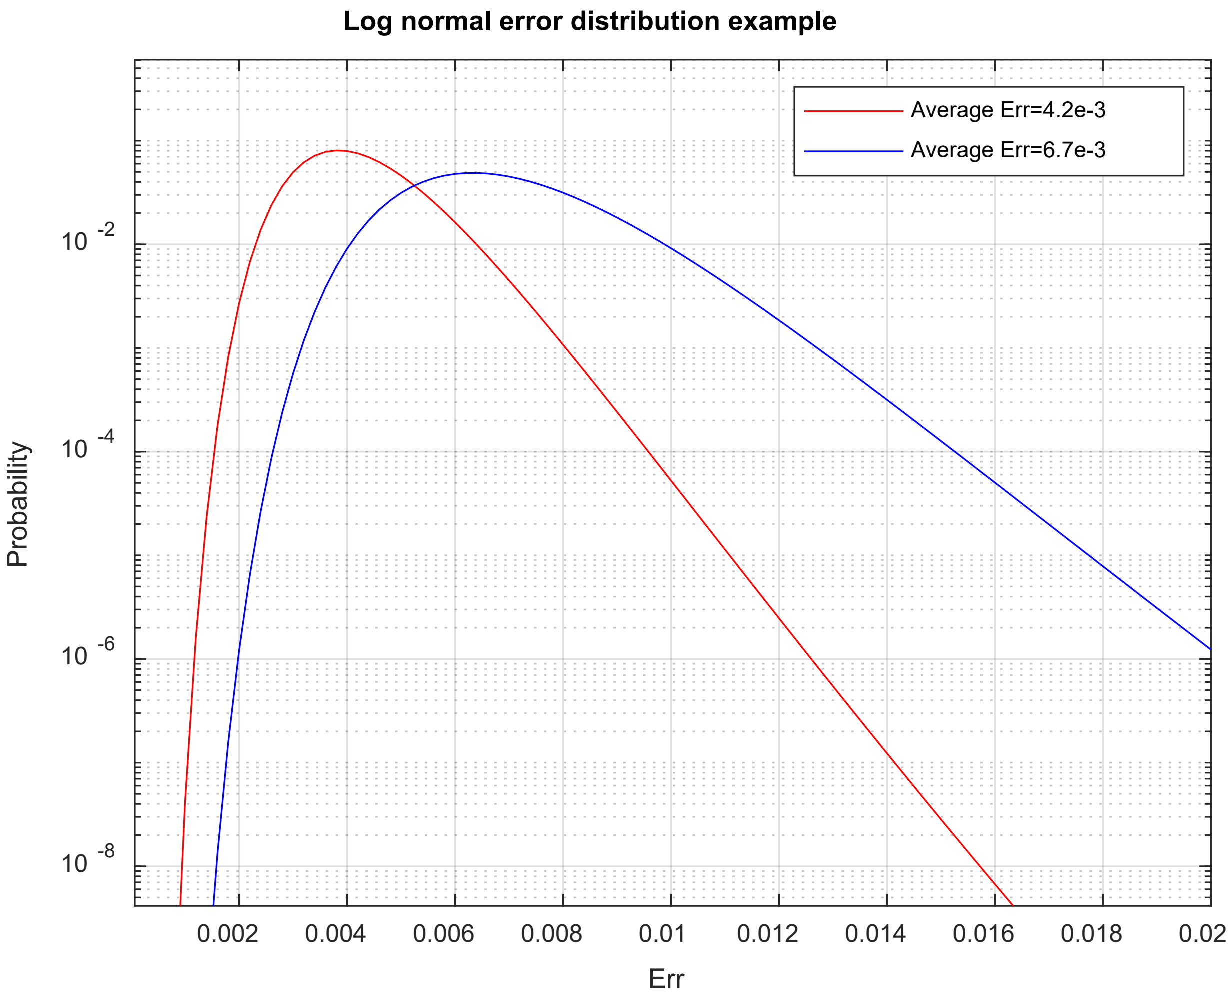

Error Modeling (CDF)

CDF

- Low average errors
  - (red) 99% < 7.4e-3
  - (blue) 99% < 1.2e-3
- Long tail
  - (red) 99.99% < 1.04e-2
  - (blue) 99.99% < 1.7e-2
- Throughput determined by the low average errors
- Soft trigger rate determined by the long tail

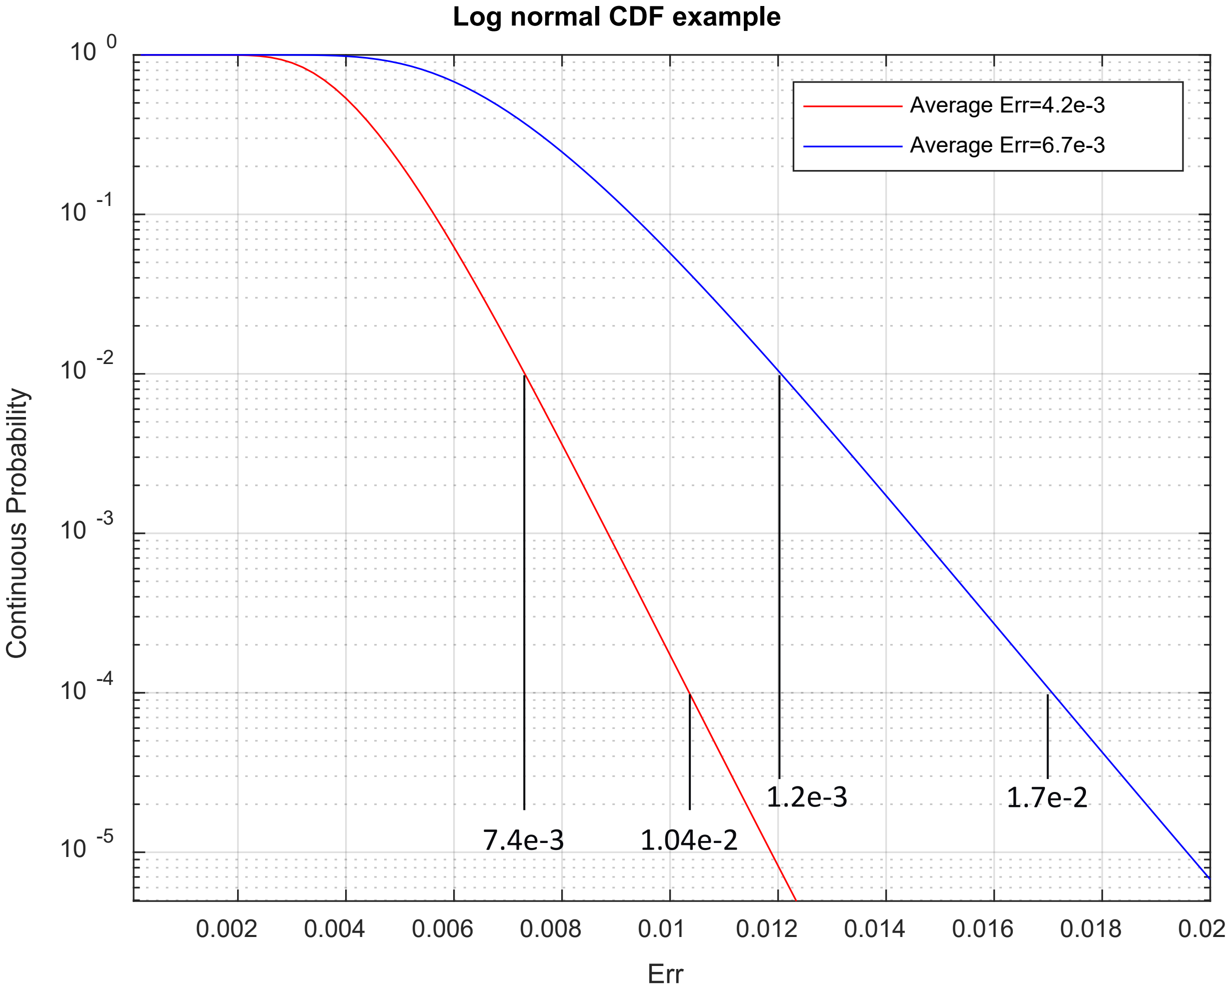

Agenda

- Background
- LDPC FER vs. BER
- LDPC FER vs. Err
- NAND Flash Error Profile
- **Performance Analysis**
- Conclusions

Performance

- **Testing error correction performance on LN error distributions**
  - Generate LN*FER and LN*Average Iteration plots

- **Must meet throughput requirements**
  - average iteration as a proxy
  - Which solution(s) have the lowest average iterations?

- **Must meet soft-trigger rate requirements**
  - Typically set between FER = 1e-2 and FER = 1e-6
  - Which solutions have the lowest FER?

LN1 FER Performance

Performance (CR=0.919, LN1)

- Top figure shows FER * LN1 for 4k and 16k Blocks
  - Summation of each curve gives total FER
  - Total FER (4k) = 7.6e-3
  - Total FER (16k) = 4.8e-3
  - Marginal FER Improvement for 16k blocks

- Bottom figure shows average iterations*LN1
  - Summation of each curve gives overall average iterations
  - Avg (4k) = 4.47
  - Avg (16k) = 3.90
  - Some average iteration improvement for 16k blocks

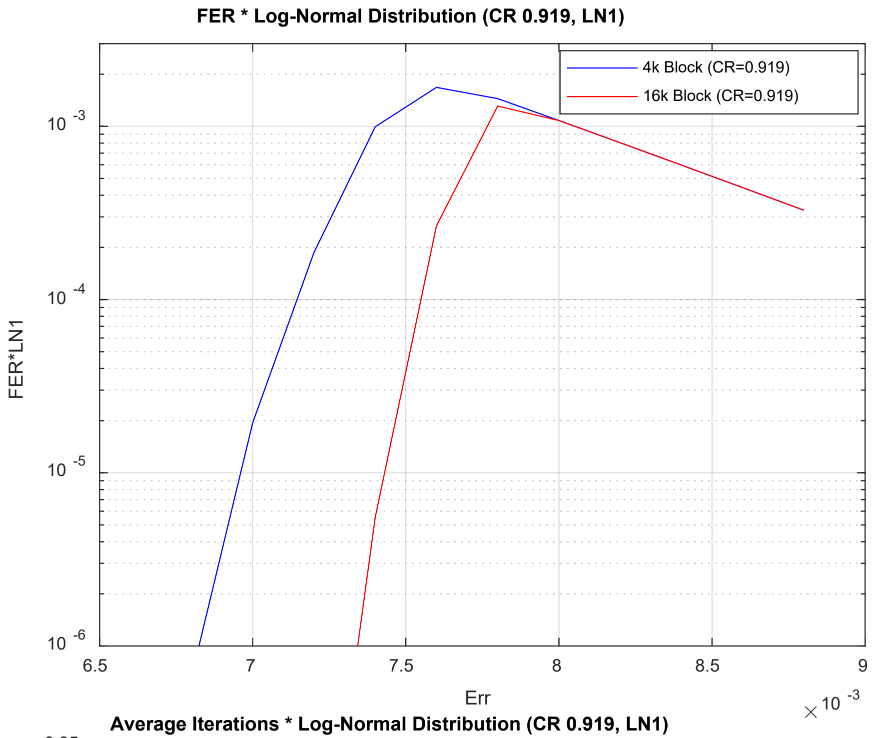

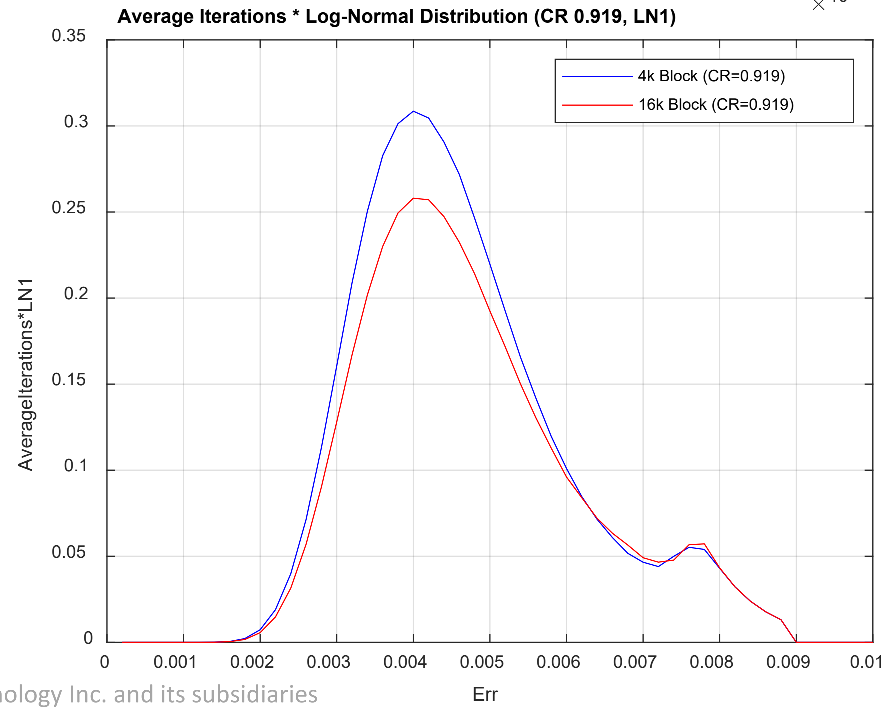

LN2 FER Performance

Performance (CR=0.871, LN2)

- Top figure shows FER*LN2 for 4k and 16k Blocks
  - Summation of each curve gives total FER
  - Total FER (4k) = 2.4e-3
  - Total FER (16k) = 2.2e-3
  - Marginal FER improvement for 16k blocks

- Bottom figure shows average iterations*LN2
  - Summation of each curve gives overall average iterations
  - Avg (4k) = 3.93
  - Avg (16k) = 3.46
  - Some average iteration improvement for 16k blocks (10-15%)

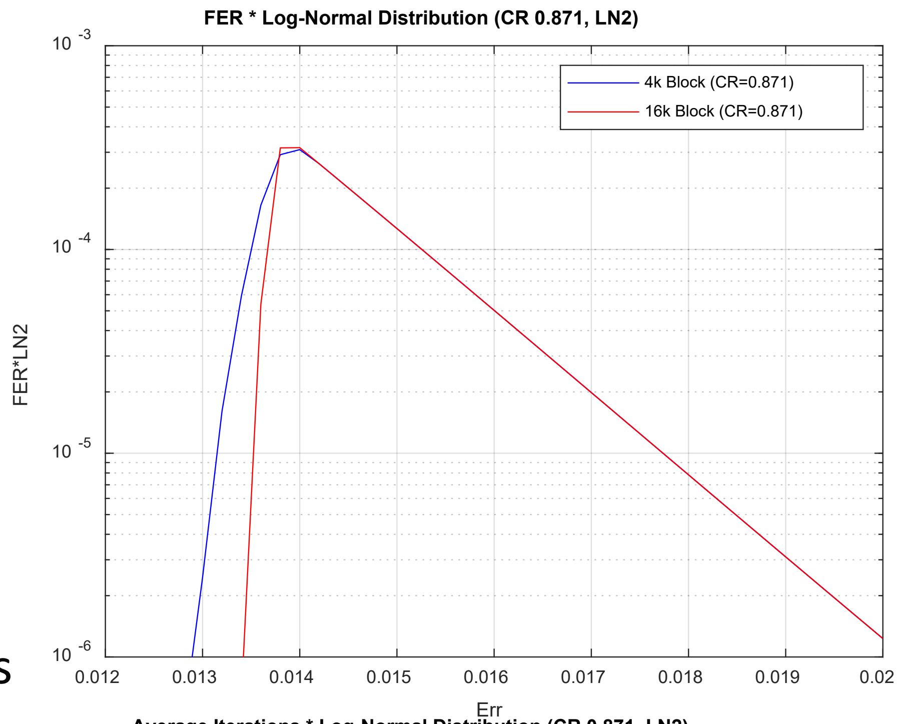

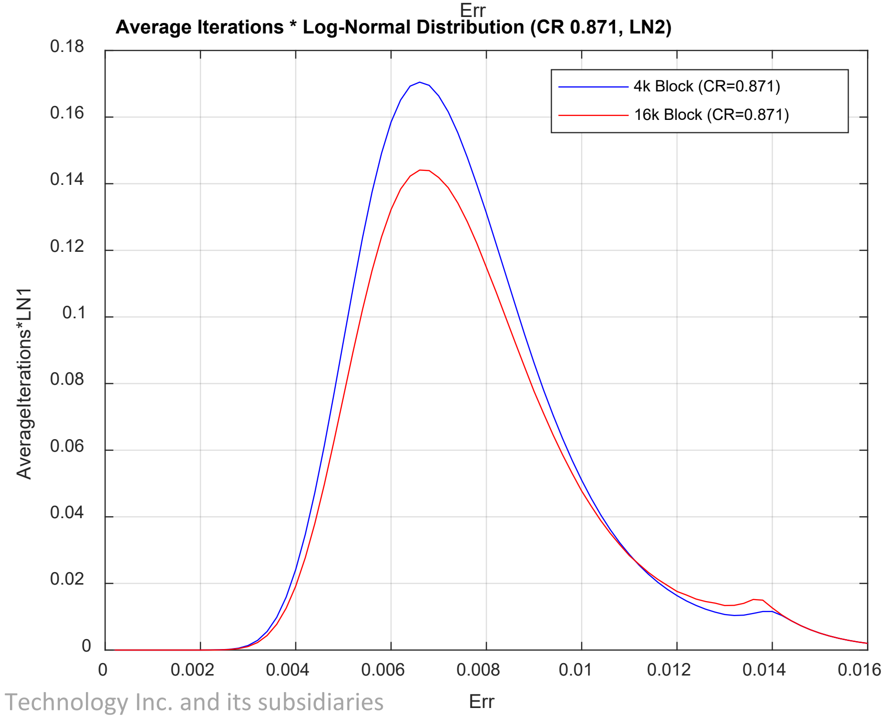

Conclusions

- 16k byte FEC blocks give ~15% improvement in BER at FER = 1e-4
- The BER error model does not match well with NAND Flash error distributions (Log-Normal)
- The Err model can be combined with Log-Normal error distributions to compare throughput FER performance
- Limited benefit for 16k byte FEC blocks when tested against Log-Normal error distributions

MICROCHIP

Thank you!
Visit Microchip at Booth 419.