Practical Approach to Determining SSD Reliability

Todd Marquart, Ph.D, Fellow-Micron Technology

Some Initial Definitions and Acronyms

Total Bytes Written (TBW):

- This is the amount of data transferred byt the host. Typically base 10.

Logical or User Density:

- This is the total amount of data a user can store on a drive at any given moment. Typically base 10.

Physical Density:

- The total amount of storage physically on the drive, typically base 2.

Write Amplification (WA):

- A measure of how much data is actually written to the SSD compared to what was requested by the host.

Mean Time To Failure (MTTF):

- The position parameter of the exponential distribution based on the time to first failure.
- The exponential distribution applies to failures that are truly random.
- MTBF (Mean Time Between Failure): Only valid for a repairable system.
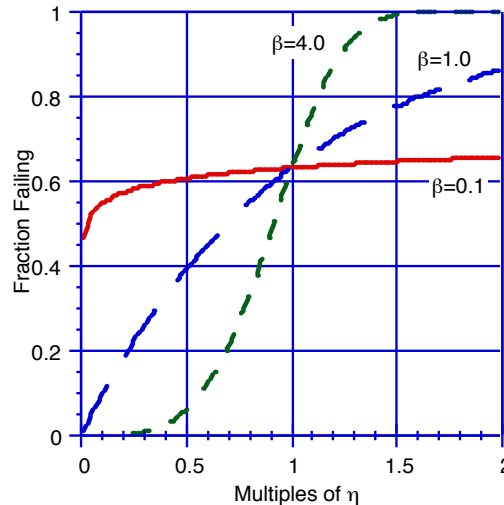
Annualized Failure Rate (AFR):

- The failure probability in 1 year, constant for an exponential distribution.

Reliability Statistics

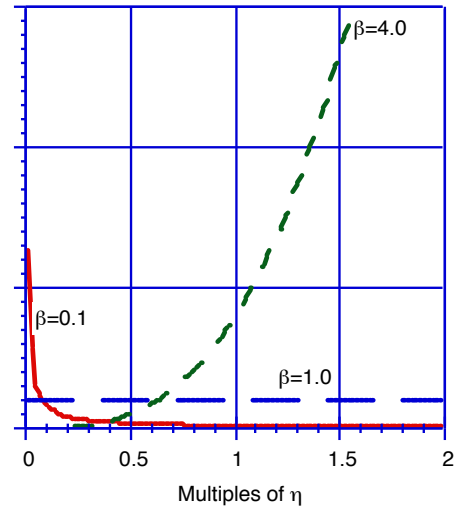F(t): Cumulative Distribution Function (CDF).

$$F(t) = 1 - S(t)$$

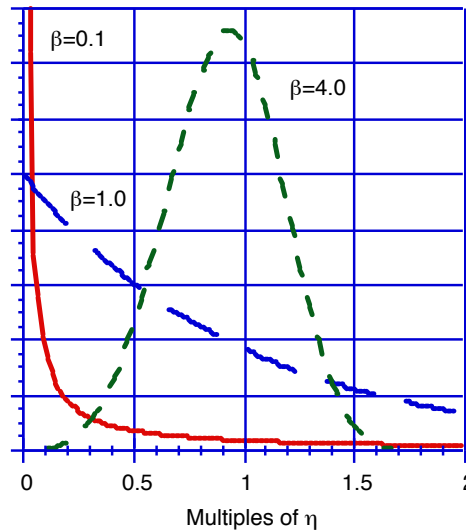$$\text{Where}: S(t) = \text{Survivor Function}$$

H(t): Hazard Function.

$$H(t) = \frac{f(t)}{1 - F(t)} = \frac{f(t)}{S(t)}$$

f(t): Probability Distribution Function (PDF)

$$f(t) = \frac{dF}{dt}$$

Reliability Statistics

Weibull Distribution

What is it?

- Large family of distributions that cover the entire life of many products.

What is it good for?

- Many failure modes can be described with or approximated by a Weibull distribution. The lognormal is not covered by the Weibull family. To distinguish Weibull from lognormal requires at least 20 points of GOOD data.

- Works well with very small sample sizes. Analysis is possible with 1 or even 0 failures.

- Good starting point for most analyses. If the data is actually lognormal, the Weibull distribution will give more conservative extrapolated lifetimes making it a safer starting point.

- Weakest link distribution.

What are the parameters?

- ß is the shape or slope parameter. It is the slope of the line when plotted on Weibull paper.

  - ß=1 gives the exponential distribution.

- η is the characteristic life or scale parameter. It is the time to 63.2% failure.

$$F(t) = 1 - \exp\left[-\left(\frac{t}{\eta}\right)^{\beta}\right]$$

$$f(t) = \left(\frac{t}{\eta}\right)^{\beta}\frac{\beta}{t}\exp\left[-\left(\frac{t}{\eta}\right)^{\beta}\right]$$

$$H(t) = \left(\frac{t}{\eta}\right)^{\beta}\frac{\beta}{t}$$

Reliability Statistics

In general it is advantageous to plot distributions on axes that generate straight lines in order to judge the fit of the data, examine signs of curvature etc…

Weibull CDF:
$$F(t) = 1 - \exp\left[-\left(\frac{t}{\eta}\right)^{\beta}\right]$$

$$\ln\ln\left(\frac{1}{1-F(t)}\right) = \beta\ln(t) - \beta\ln(\eta) = \ln(-\ln(1-F(t)))$$

$$Y \qquad = m\,X \;+\; b$$

Reliability Statistics

A plot of the Weibull CDF gives a good , intuitive view of the product lifetime.

The same reliability mechanisms will tend to have the same Weibull slope even when changing from one technode to another.
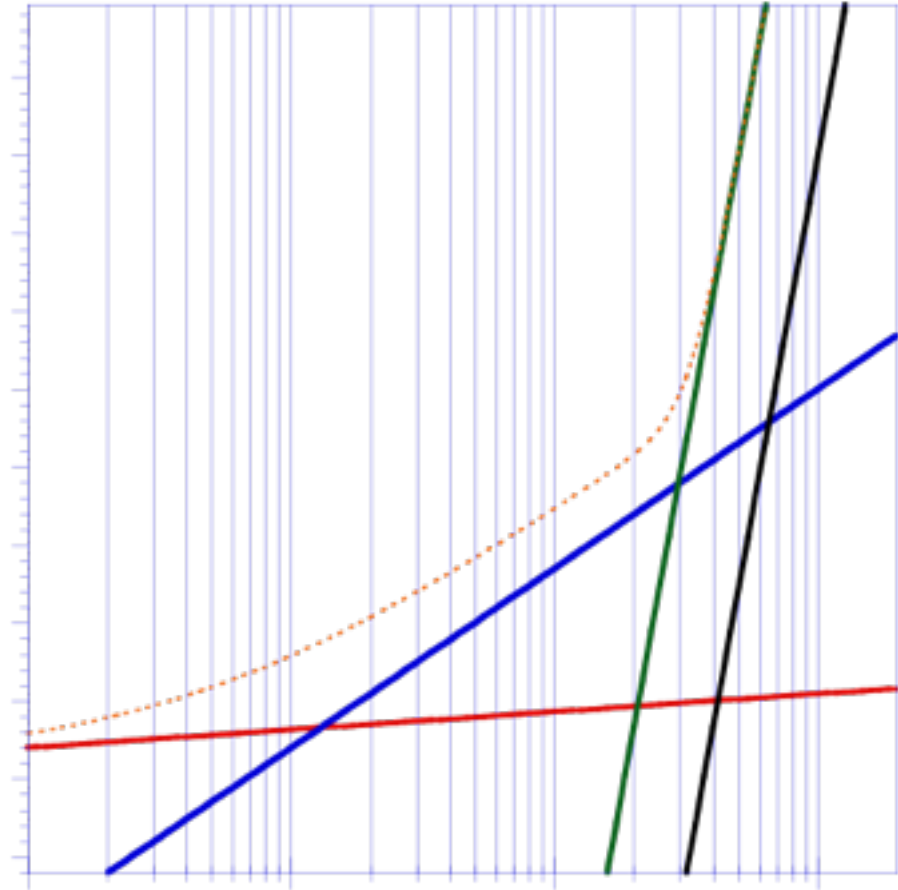
Reliability Statistics

Historical view of overall reliability lifetime.

Not my favorite view.

Early Life

Useful Life
Early Wear out

End of Life
Wear out

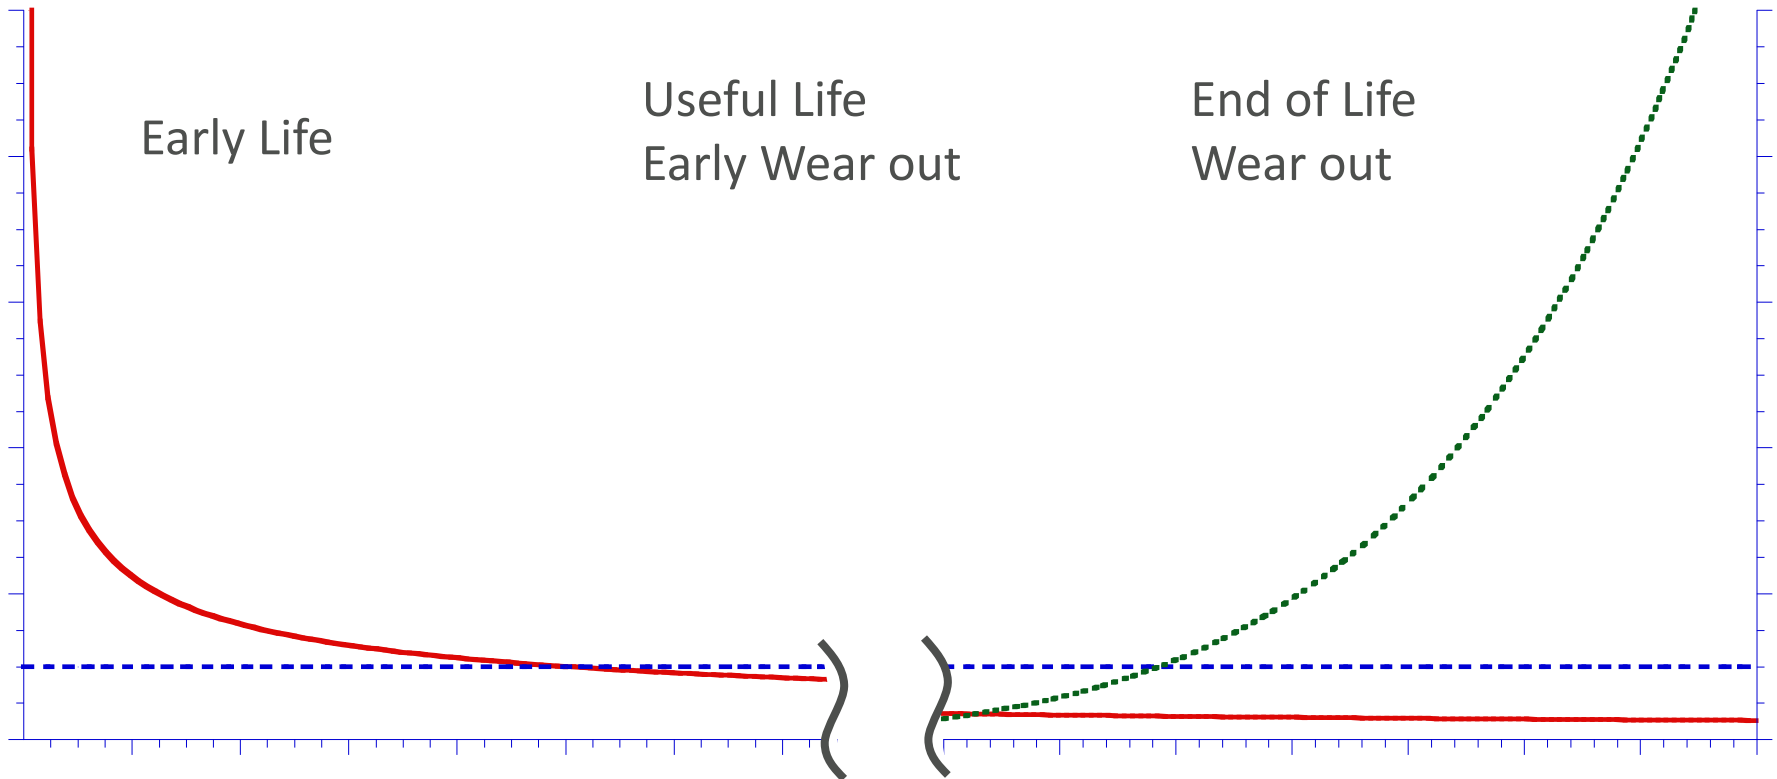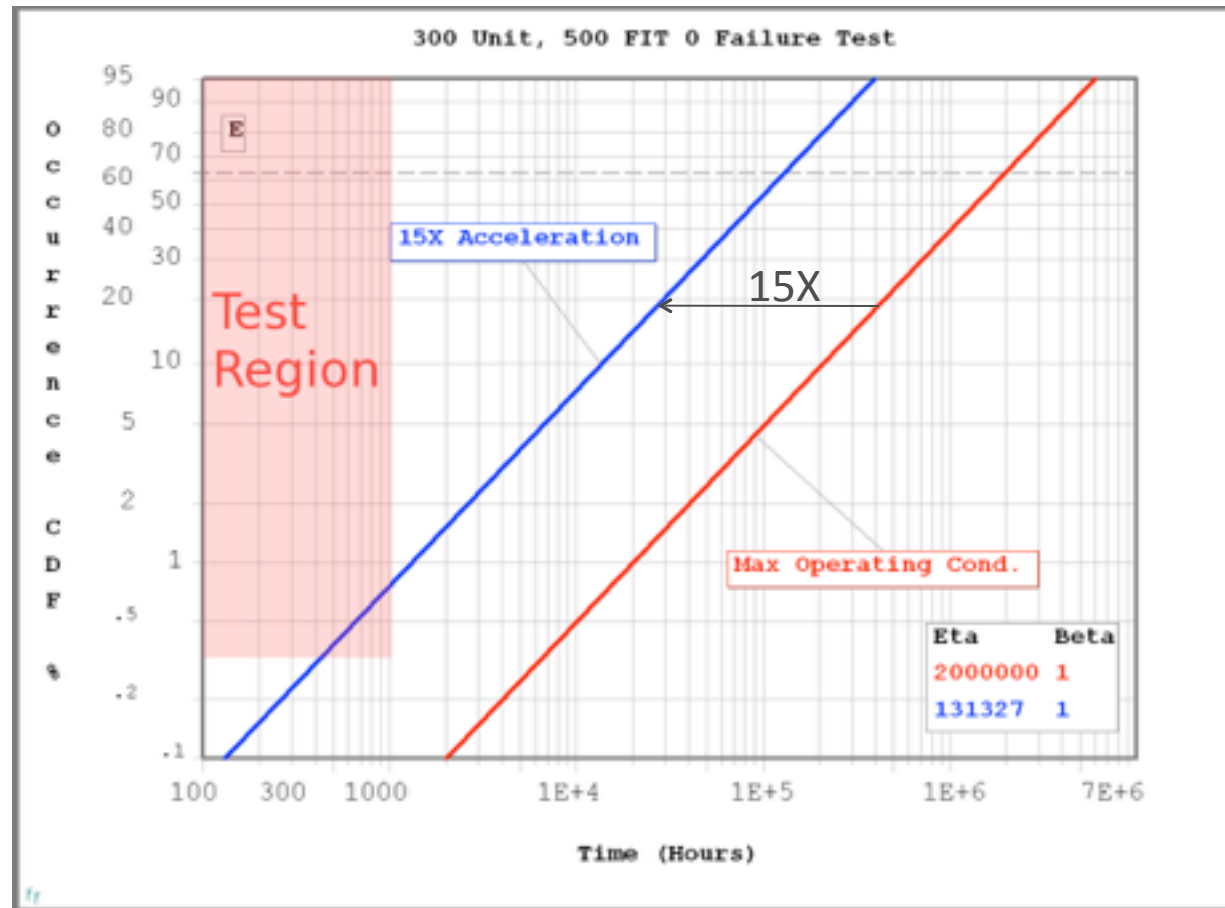# Acceleration Factors

Acceleration factors take the failure distribution and move it parallel to the original distribution.

- A change is slope typically indicates a shift in mechanism.

Typically it is used to move failure distributions into the region (stress and sample size) of interest.

300 Unit, 500 FIT 0 Failure Test

©2014 Micron Technology, Inc.

JEDEC JESD218

From JEDEC JESD218, Section 1:

"This standard defines JEDEC requirements for solid state drives. For each defined class of solid state drive, the standard defines the conditions of use and the corresponding endurance verification requirements. Although endurance is to be rated based upon the standard conditions of use for the class, the standard also sets out requirements for possible additional use conditions as agreed to between manufacturer and purchaser.

Qualification of a solid state drive involves many factors beyond endurance and retention, so such qualification is beyond the scope of this standard, but this standard is sufficient for the endurance and retention part of a drive qualification. This standard applies to individual products and also to qualification families as defined in this standard.

The scope of this standard includes solid state drives based on solid-state non-volatile memory (NVM). NAND Flash memory is the most common form on memory used in solid state drives at the time of this writing, and this standard emphasizes certain features of NAND. The standard is also intended to apply to other forms of NVM. "

JESD218-Testing

JESD218, is really about validating that the drive can meet its endurance specification, including post TBW data retention.
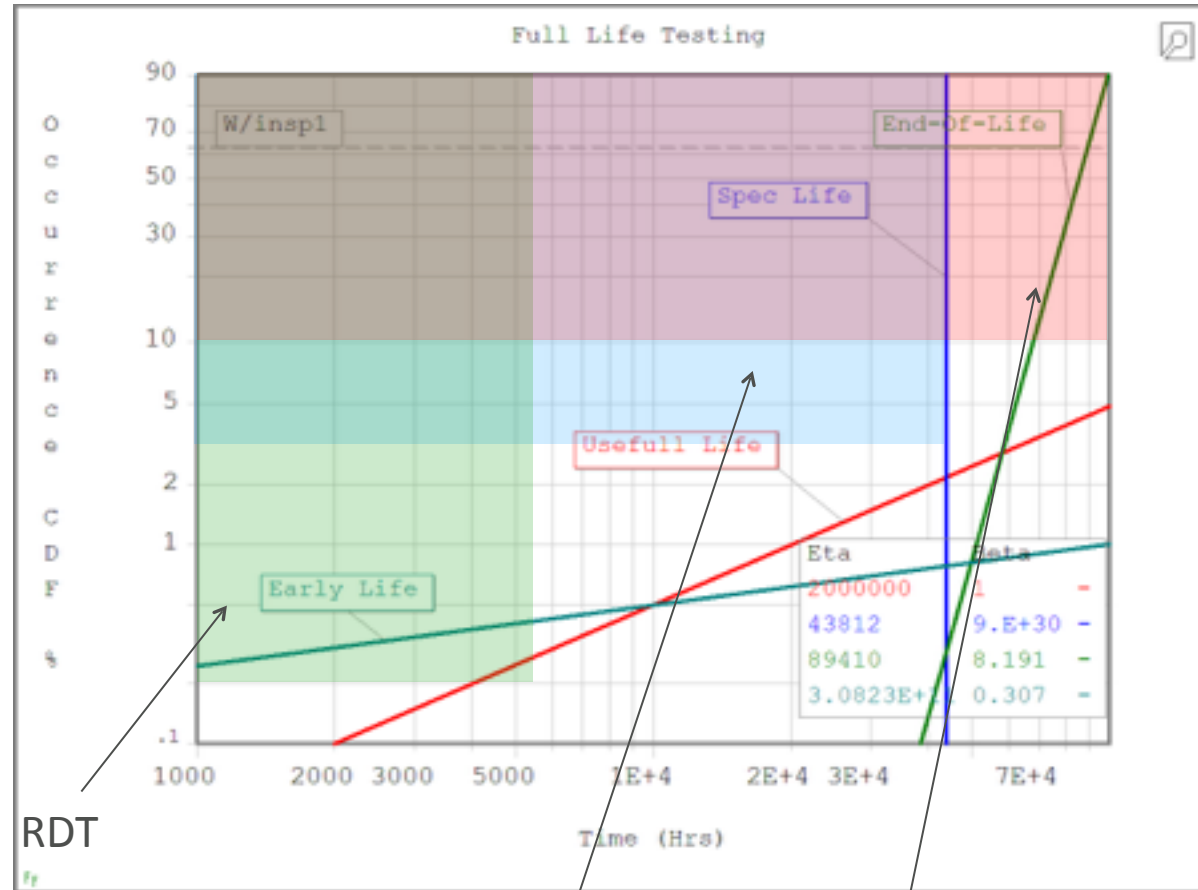
RDT

TBW Qual (JESD218)

Cycle-to-Death Testing

JESD218 Requirements

The general requirements for an SSD are shown here, as per JESD218.

- 2 Classes are defined, Client and Enterprise.

Table 1 — SSD Classes and Requirements

Application Class	Workload (see JESD219)	Active Use (power on)	Retention Use (power off)	Functional Failure Requirement (FFR)	UBER Requirement
Client	Client	40 °C 8 hrs/day	30 °C 1 year	≤3%	$\leq 10^{-15}$
Enterprise	Enterprise	55 °C 24hrs/day	40 °C 3 months	≤3%	$\leq 10^{-16}$

Functional Failure Rate

FFR is dominated by defectivity related issues.

Single pages, wordlines, blocks can fail as can entire die.

- In some cases failures can be graceful, not resulting in data loss, others are catastrophic.
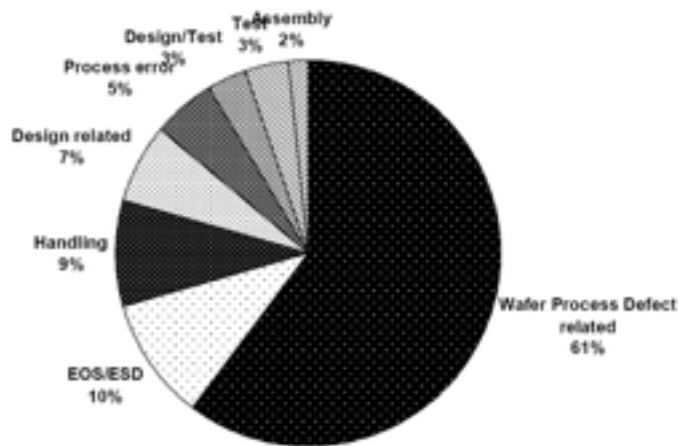

FIGURE 1. DISTRIBUTION OF THE ANALYZED FIELD
FAILURE CASES CATEGORIZED BY THE ROOT CAUSE
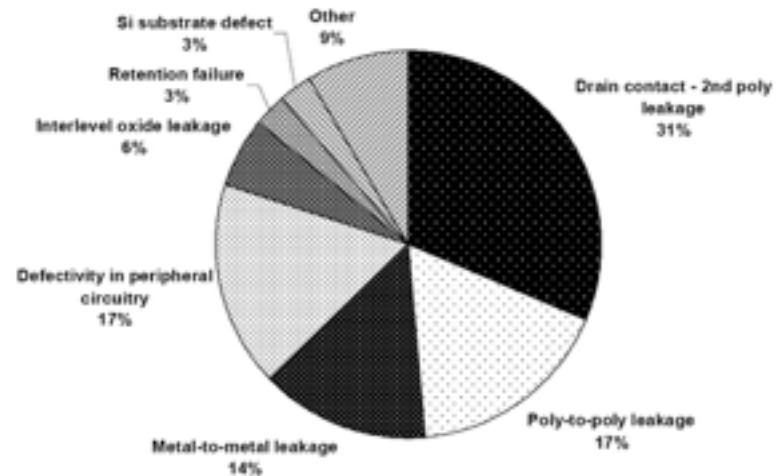
FIGURE 2. SUB-CLASSIFICATION OF THE WAFER
PROCESS RELATED DEFECTS.

P. Muroke, Proc IRPS, 2006

|

Functional Failure Rate

These are typically the defect related failures.

- Defective NAND, controller failures etc…

This is related to the MTTF/ AFR of the drive.

- Most drives specify the AFR/ MTTF of a drive that may not align with the JEDEC specification.

In most cases it is the purpose of the Reliability Demonstration Test (RDT) to demonstrate the MTTF.

Lifetime Fallout as a Function of MTTF

Legend:
- 1 Yr Lifetime
- 2 Yr Lifetime
- 3 Yr Lifetime
- 4 Yr Lifetime
- 5 Yr Lifetime

Y-axis: Lifetime Fallout (100%, 10%, 1%, 0.1%)
X-axis: MTTF (MHrs) (0.1, 1, 10)

Micron®
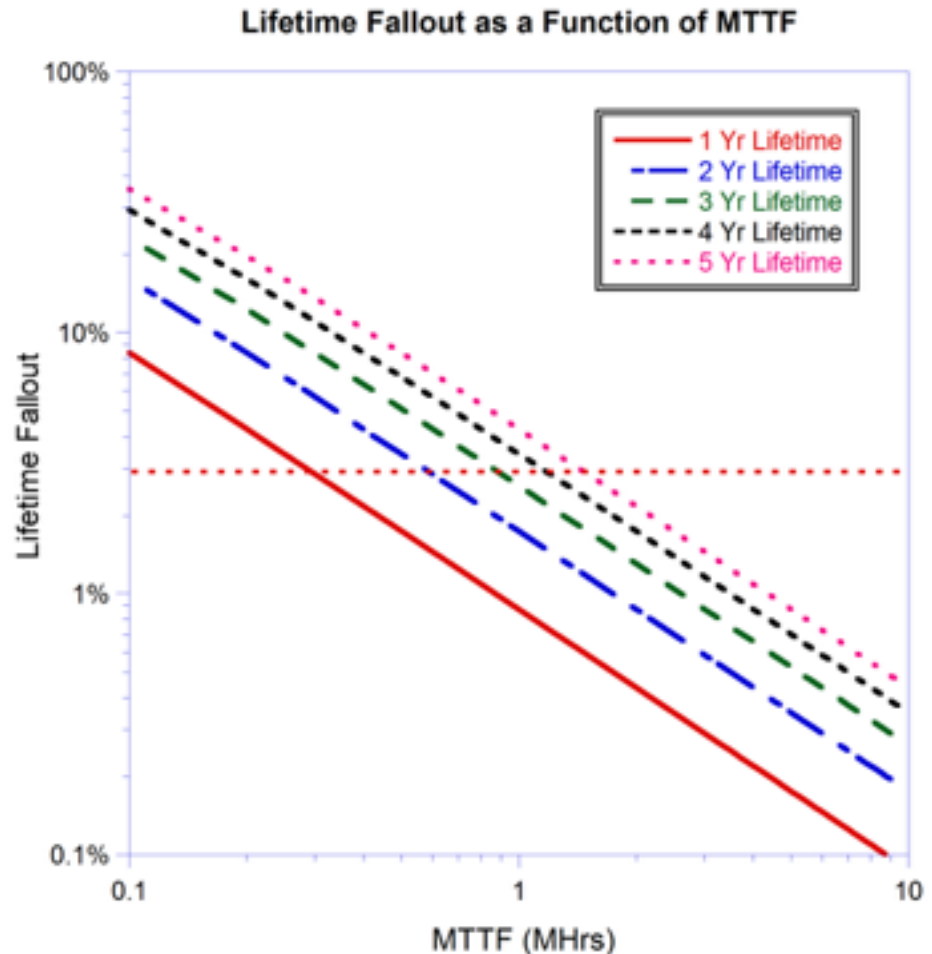
Functional Failure Rate-Sample Size

- The overall test sample size is calculated using:
  - $UCL(functional\_failures) \leq FFR \cdot SS$
  - UCL is the upper confidence limit based on the χ^2 distribution and 60% confidence.
  - functional_failures is the allowed number of failures
  - FFR is the functional failure rate
  - SS is the sample size.
- For example, if the number of allowed failures was 1 this would be:
  - $UCL(1) \leq 3\% \cdot SS$, Solve for SS
  - $SS \geq UCL(1)/3\%$
    - $UCL(1) = 0.5 \cdot \chi^2 (60\%, 2 \cdot functional\_failures + 2) = 0.5 \cdot \chi^2 (60\%, 2 \cdot 1 + 2) = 2.02$
    - $SS \geq 2.02/3\% = 68\ Drives$

Functional Failure Rate and RDT

The FFR at the max TBW gives an estimate of the average AFR over the product lifetime, but may not represent the fallout in the first year.

- For a purely exponential distribution, the numbers will be the same, however, in the non-exponential case it will not align.
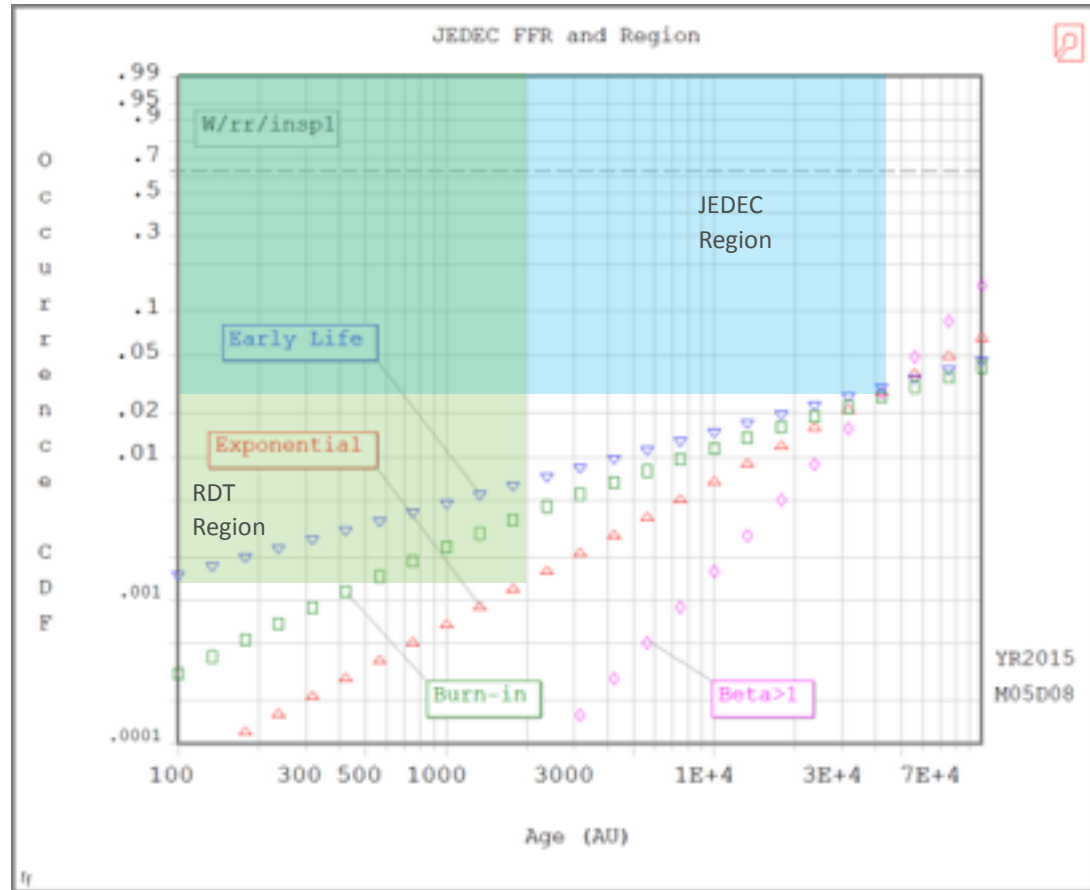
UBER-Uncorrectable Bit Error Rate

From JEDEC JESD218

3.22 Uncorrectable Bit Error Rate, or ratio (UBER)

A metric for the rate of occurrence of data errors, equal to the number of data errors per bits read. Mathematically,

$$UBER = \frac{number\ of\ data\ errors}{number\ of\ bits\ read} \qquad (1)$$

NOTE   Although the UBER concept is in widespread use in the industry, there is considerable variation in interpretation. In this standard, the UBER values for SSDs are to be *lifetime* values for the *entire population*. The numerator is the total count of data errors detected over the full TBW rating for the population of SSDs, or the sample of SSDs in the endurance verification. A sector containing corrupted data is to be counted as one data error, even if it is read multiple times and each time fails to return correct data. The denominator is the number of bits written at the TBW rating limit, which aligns to the usual definition of errors per bit read when the read:write ratio is unity. See 7.1.1 for a further discussion of UBER calculation.

UBER and Fail Probability

UBER is related to the failure probability of a codeword.

- The third piece of the equation below is essentially a usage model.

$$UBER(DR_{SSD}) = \frac{N_{failing\_sectors}}{N_{bits\_in\_sample} \cdot \left(\dfrac{N_{cycles}}{WA} + 1 \right)}$$

$$= \frac{F(t)_{codeword} \cdot N_{drives} \cdot N_{codewords\_per\_drive} \cdot N_{sectors\_per\_cw}}{N_{drives} \cdot N_{codewords\_per\_drives} \cdot N_{bits\_per\_cw} \cdot \left(\dfrac{N_{cycles}}{WA} + 1 \right)}$$

Instantaneous UBER

$$= \frac{F(t)_{codeword} \cdot \cancel{N_{drives}} \cdot \cancel{N_{codewords\_per\_drive}} \cdot N_{sectors\_per\_cw}}{\cancel{N_{drives}} \cdot \cancel{N_{codewords\_per\_drives}} \cdot N_{bits\_per\_cw} \cdot \left(\dfrac{N_{cycles}}{WA} + 1 \right)}$$

Usage Model

$$= \frac{F(t)_{codeword} \cdot N_{sectors\_per\_cw}}{N_{bits\_per\_cw} \cdot \left(\dfrac{N_{cycles}}{WA} + 1 \right)} = \frac{F(t)_{codeword}}{N_{bits\_per\_cw}} \cdot N_{sectors\_per\_cw} \cdot \frac{1}{\dfrac{N_{cycles}}{WA} + 1}$$

Flash Memory Reliability-UBER

Sometimes the function $F(t){\downarrow}cw\,/N{\downarrow}b\_p\_cw$ is thought of as the probability of a bit to fail ECC.

- If true, this would then allow scaling to different CW sizes.

- This assumption is based on a series system relationship and is <u>incorrect</u> since ECC is really a "k-out-of-s" system.

$$UBER(DR_{SSD}) = \frac{F(t)_{codeword}}{N_{bits\_per\_cw}} \cdot N_{sectors\_per\_cw} \cdot \frac{1}{\frac{N_{cycles}}{WA} + 1}$$

CW Fail Probability Scaling by Codeword
15b/512B, 5x10⁻⁴ RBER
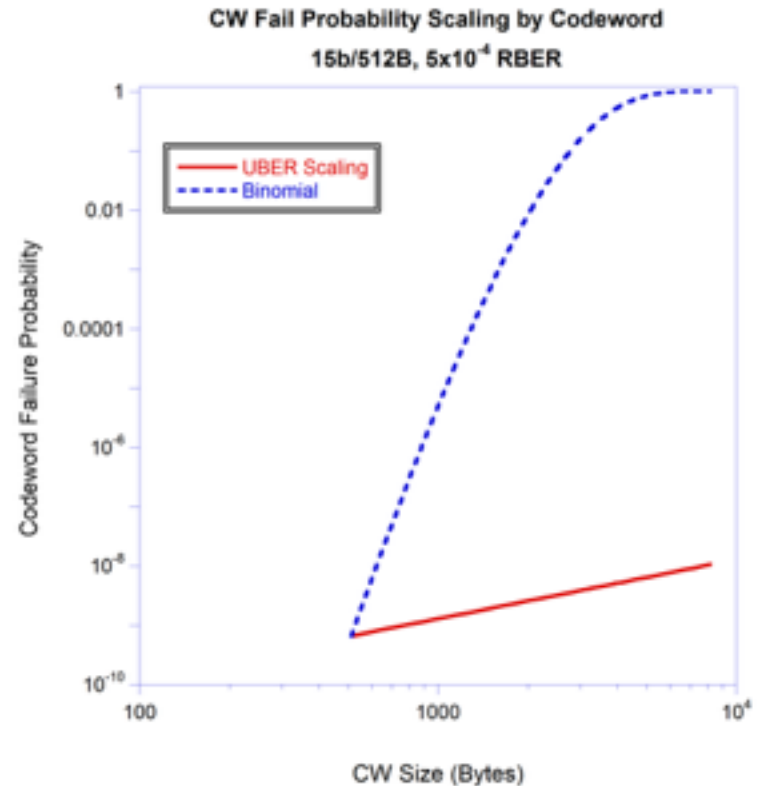
UBER at the Component vs Drive

The drive and component UBERs may not be aligned for a variety of reasons.

- The usage models that feed the UBER calculations may not be aligned.

- Additional error-management may be present at the drive level.

  - RAID

  - LDPC vs BCH

  - Different ECC levels

BCH Type Codes

LDPC Type Codes

Correction Probability

Number of errors in a Codeword
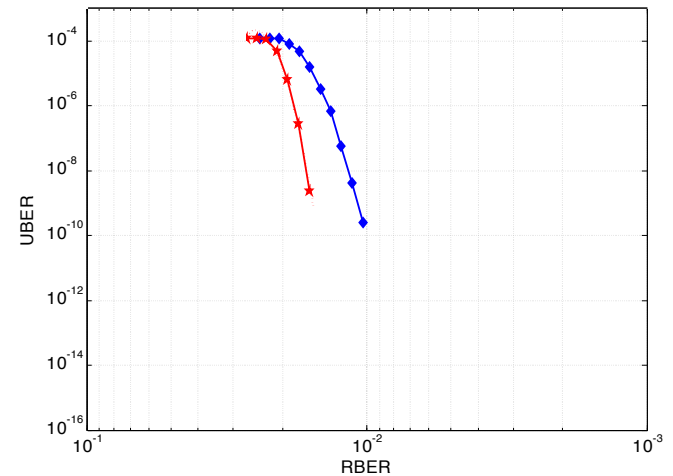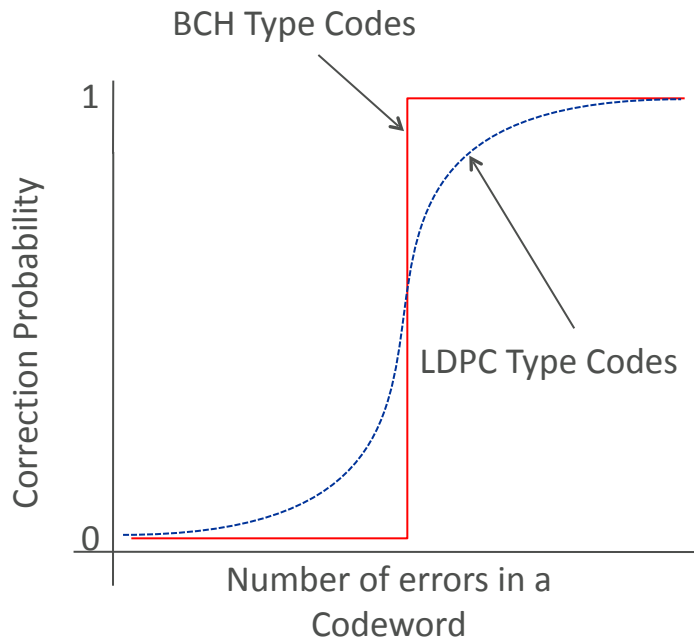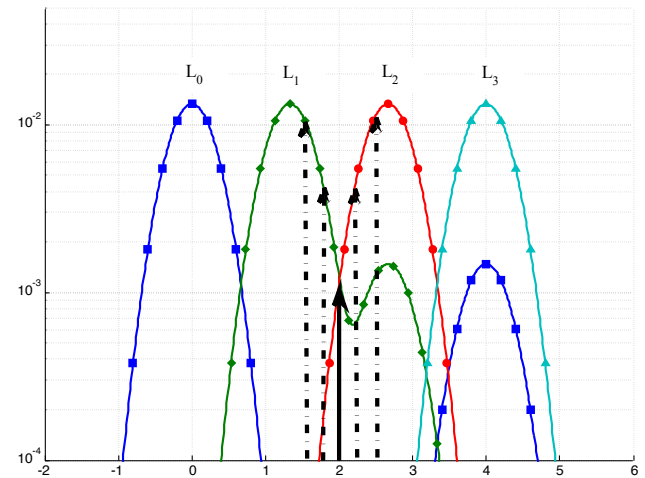
Khayat, P. R. et. al., IMC, 2015

Sample Size and UBER

- The overall test sample size is calculated using:
  - $UCL(data\_errors) \leq \min(TBW,TBR) \cdot 8 \cdot 10\uparrow12 \cdot UBER \cdot SS$
  - UCL is the upper confidence limit based on the χ^2 distribution and 60% confidence.
  - Data_errors is the allowed number of data errors
  - Min(TBW, TBR) is the minimum of the total bytes written (TBW) or total bytes read (TBR) in base 10 terabytes.
    - This is to prevent reduction of the apparent UBER by reading the same data multiple times.
  - SS is the sample size.
- For example, if the number of allowed failures was 1 for a client drive (UBER = 10^{-15}) with TBW=32TB and TBR=48TB would be:
  - $UCL(1) \leq \min(TBW,TBR) \cdot 8 \cdot 10\uparrow12 \cdot SS$, Solve for SS
  - $SS \geq UCL(1) / \min(TBW,TBR) \cdot 8 \cdot 10\uparrow12 \cdot UBER$
    - $UCL(1) = 0.5 \cdot \chi\uparrow2 \ (60\%, 2 \cdot functional\downarrow failures + 2) = 0.5 \cdot \chi\uparrow2 \ (60\%, 2 \cdot 1 + 2) = 2.02$
    - $SS \geq 2.02 / 32 \cdot 8 \cdot 10\uparrow12 \cdot 10\uparrow-15 \ = 8 \ Drives$

High Temperature Data Retention

Shifts

- Detrapping of electrons.

- Causes the entire distribution to shift and widen.

- The shift is of the entire population, the shift by bit is highly variable.
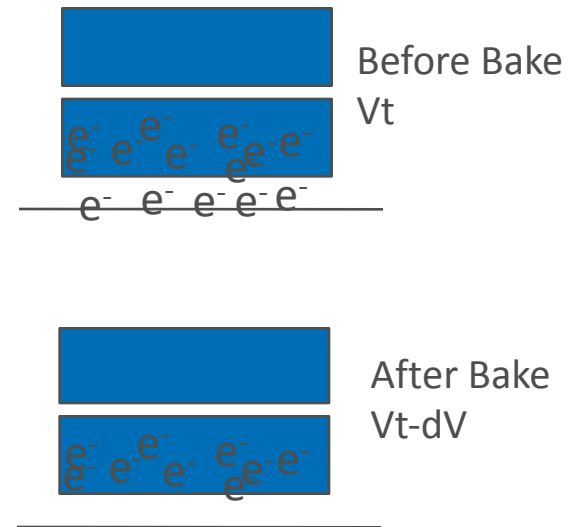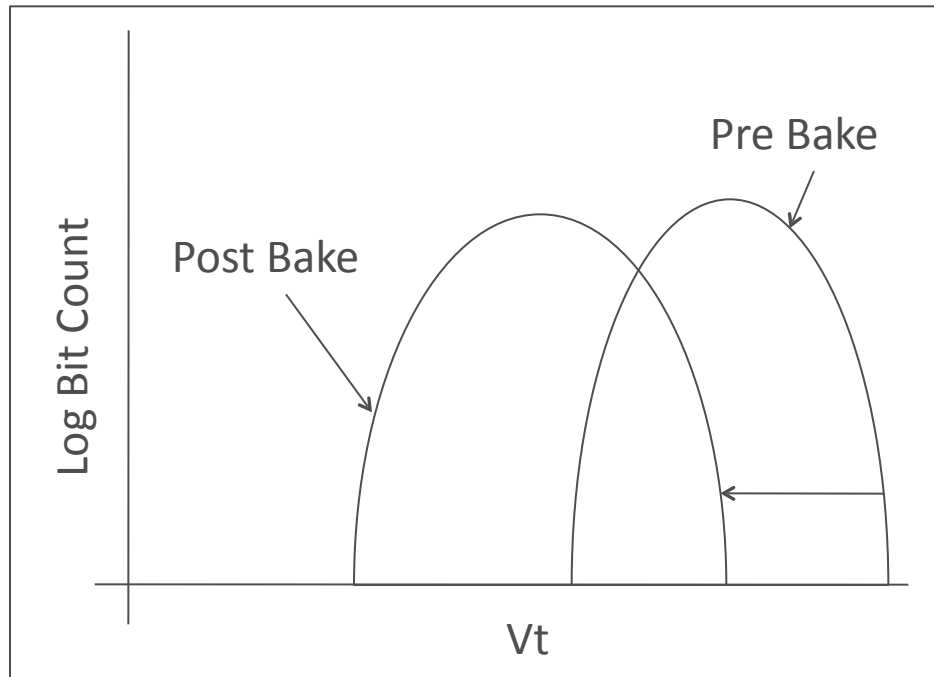
Flash Memory Reliability-High Temperature Data Retention

Highly temperature accelerated.

- Ea of ~1.1eV

Very sensitive to cycling speed.

- Higher speed allows for less detrapping between cycles and makes HTDRB worse.

Very sensitive to cycling temperature.

- Higher temperature allows more detrapping between cycles so HTDR gets better with higher cycling temp.

HTDR is somewhat field dependent.

- This implies L3 will move more than L2 or L1, but the impact is not nearly as strong as seen in room temperature data retention.
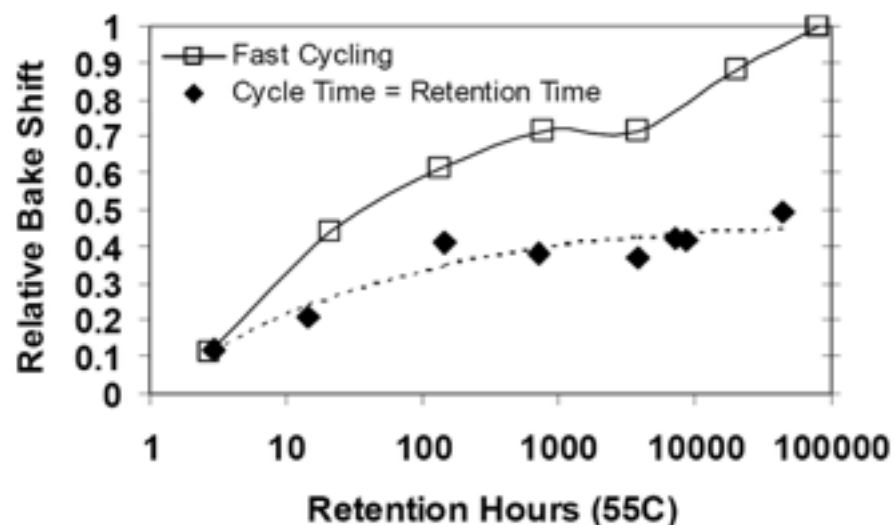
$$UDM \approx cycles^{0.5} \ln\left(1 + \frac{1}{A} \cdot \frac{t_{bake,n}}{t_{cycle,n}} \cdot \frac{e^{\frac{1.1eV}{KT_{use}} - \frac{1.1eV}{KT_{bake}}}}{e^{\frac{1.1eV}{KT_{use}} - \frac{1.1eV}{KT_{cycle}}}}\right)$$

$$= cycles^{0.5} \ln\left(1 + \frac{1}{A} \cdot \frac{t_{bake,n}}{t_{cycle,n}} \cdot e^{\frac{1.1eV}{KT_{cycle}} - \frac{1.1eV}{KT_{bake}}}\right)$$

N. Mielke et. al., Proc. 44[th] Annual IRPS, 2006

Data Retention Cyc Temp/Speed

The physical effect goes back to the NAND high-temperature device physics.

- The Vt shift is related to detrapping of electrons from the tunnel oxide.

If the number of traps generated is dependent primarily on electron fluence, then we need to only control trap occupation density and ensure it is equivalent to the assumed field usage model.

Cycle too fast, or too cold

JEDEC Field Conditions

Accelerated JEDEC conditions

Cycle too slow, or too hot
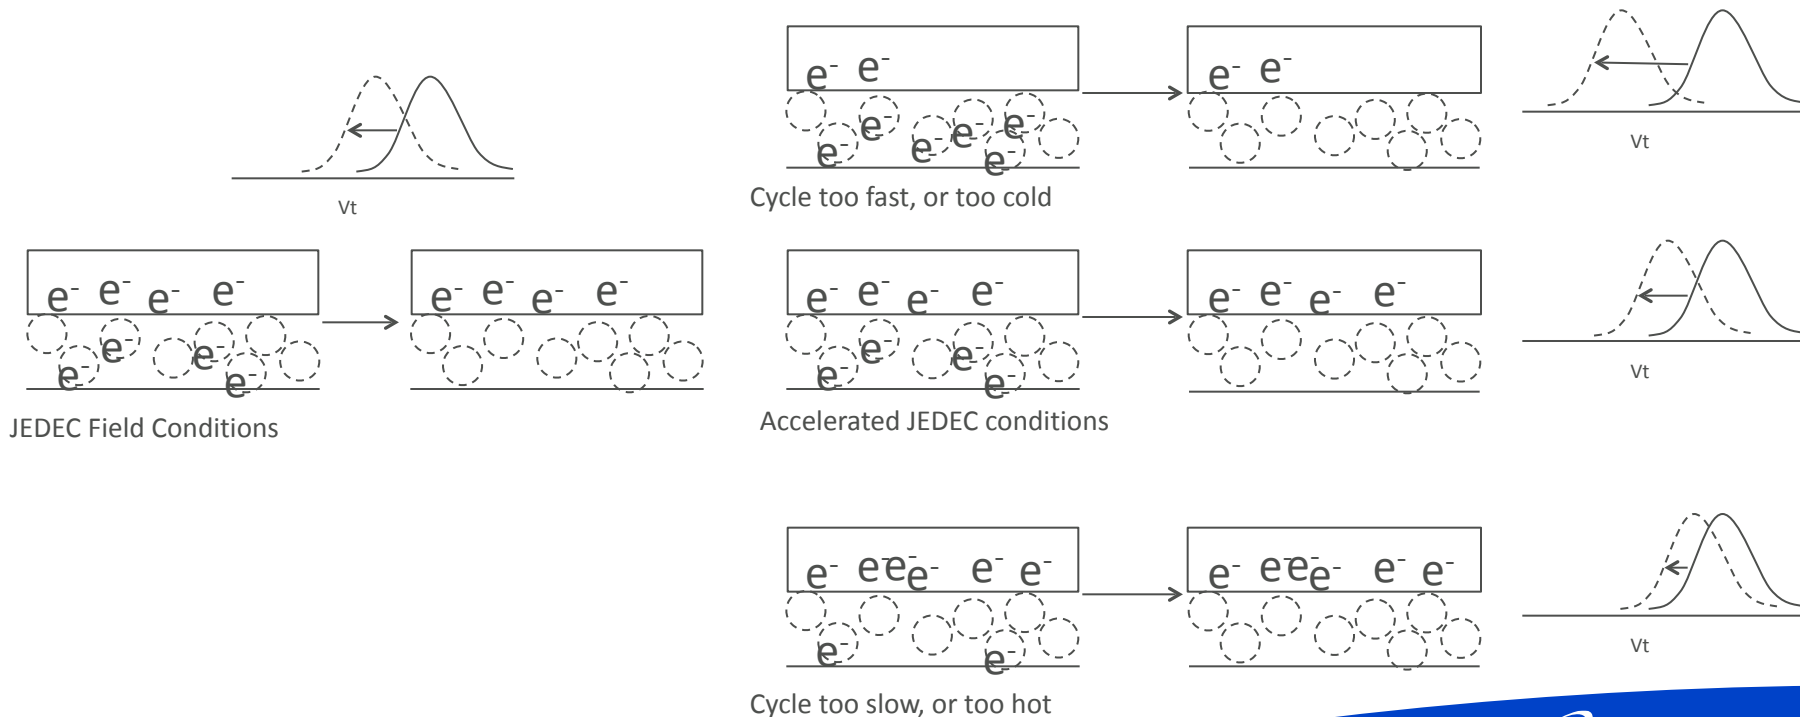
Flash Memory Reliability-High Temperature Data Retention

The cycling time/temperature relationship assumes that trap generation is essentially temperature independent, while occupancy is not.

In that case, the total time at temperature during cycling determines the occupancy.

- The goal is to get the same occupancy under accelerated conditions as under usage conditions.

- The trap generation is the same since both target the same number of P/E cycles (same fluence through the tox).

For example, assume parts are to cycle in the field to 5K cycles at 55C 24hrs/day over a 1 year period.

If the accelerated test will occur over 500hrs, the acceleration factor is simply the ratio of the times (365*24/500=17.52).

Using the Arrhenius equation, it is possible to calculate the stress temperature required to get this level of detrapping acceleration.

$$A(\text{detrap}) = \exp\left[\frac{Ea}{k}\left(\frac{1}{T_{use}} - \frac{1}{T_{stress}}\right)\right]$$

$$\frac{365 \cdot 24}{500} = \exp\left[\frac{1.1}{8.617 \cdot 10^{-5}}\left(\frac{1}{55 + 273.15} - \frac{1}{T_{stress} + 273.15}\right)\right]$$

Solve for Tstress Gives :

$$T_{stress} = 81°C$$

JESD218 Cyling Temperature/Time Requirements
1.1eV Ea, Split Approach
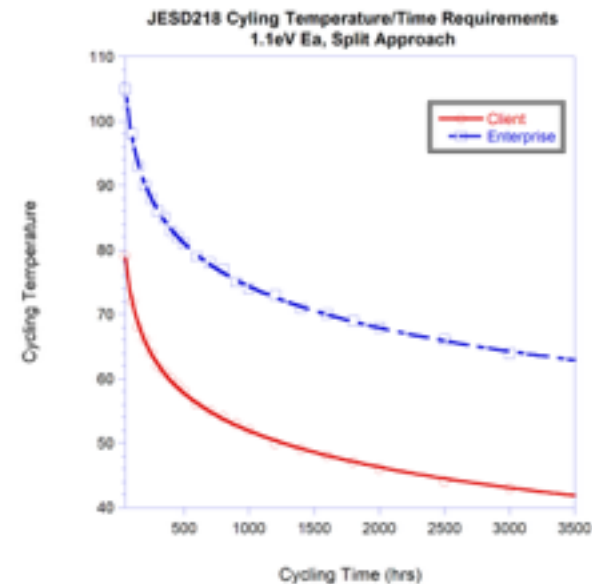
SSD Reliability-High Temperature Data Retention

For highly accelerated testing (for example high cycling enterprise drives), testing can be difficult.

- For example, for an enterprise drive to be cycled in 500 hours or less requires the temperature to be >80C.

- This can be an issue due to component temperature limitations, requiring lower temperature cycling.

Also note the extreme temperature sensitivity

- The difference between cycling over 500hours and cycling over 400hours is only 2 degrees.

Table 4 — Maximum endurance high temperature (T_{max}) vs. endurance stress times

Actual endurance stress hours	Split		Ramped	
	Client	Enterprise	Client	Enterprise
50	79	105	86	113
100	72	98	79	105
150	68	93	75	101
200	66	90	72	98
250	64	88	70	95
300	62	86	68	93
350	61	85	67	92
400	60	83	66	90
450	59	82	65	89
500	58	81	64	88
600	56	79	62	86
700	55	78	61	85
800	54	77	60	83
900	53	75	59	82
1000	52	74	58	81
1200	50	73	56	79
1400	49	71	55	78
1600	48	70	54	77
1800	47	69	53	75
2000	46	68	52	74
2500	44	66	50	72
3000	43	64	48	71

JEDEC JESD 218A

Flash Memory Reliability-High Temperature Data Retention

It is important to note that the widely quoted 1.1eV does have some uncertainty associated with it is pointed out in JEDEC JESD218:

"Although an SSD would be expected to reach its TBW rating over a lifetime of several years, for the specific purpose of calculating Tmax the full TBW is assumed to occur within a single year. This is a conservative assumption because a shorter time allows less relaxation between writes. This assumption is made to add margin
against possible inaccuracies in the 1.1eV acceleration model for high temperature data retention."

JEDEC JESD218

Flash Memory Reliability-Dealing With HTDR

Besides controlling usage conditions, in some cases other techniques may be used to minimize the impact of HTDR.

The case shown here suggests compensating for the shift by moving the Vt distributions back up to their initial positions.

(a) Threshold voltage distribution vs number of electrons in flash cell

(b) Retention errors cause threshold voltage shift to the left

(c) ISPP shifts threshold voltage to the right and can fix retention errors
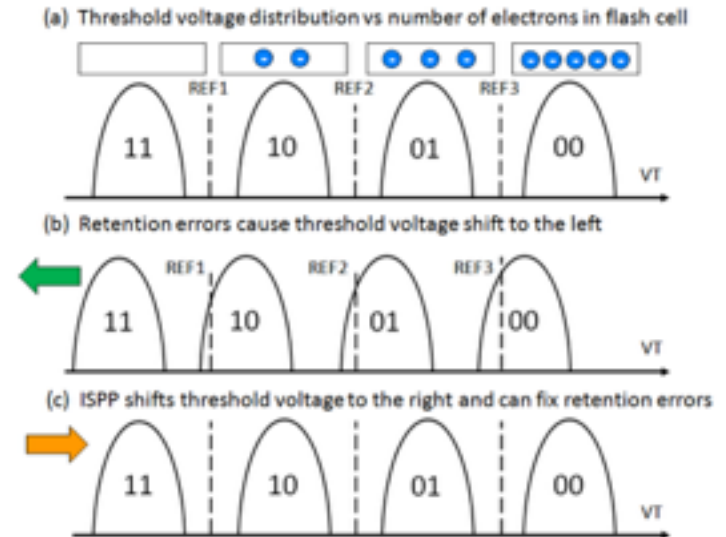
Y. Cai et.al., ICCD, 2012

Flash Memory Reliability-Low Temperature Data Retention

Leakage of electrons from the FG through trap-assisted-tunneling.

- Referred to as SILC or Stress-Induced-Leakage-Currents.

During data retention the cells will tend to move towards their lowest energy state.

- The exact same mechanism is responsible for read disturb, but tends to go the opposite direction (charge gain vs charge loss).
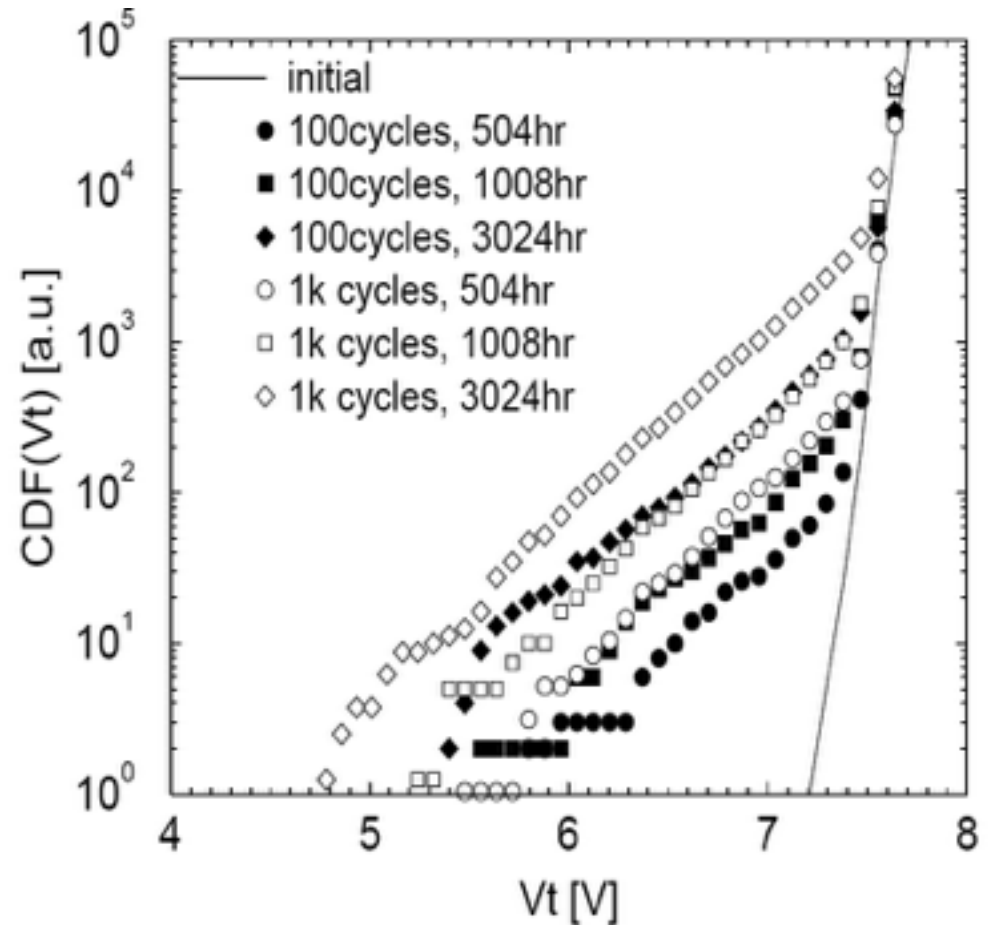
P. Cappelletti et.al., 2004 IEDM

Flash Memory Reliability-Low Temperature Data Retention

Vt Tail Formation

Main distribution remains relatively unchanged.

A. Hoefler et.al., Proc. 2002 IRPS, 2002

Flash Memory Reliability-Low Temperature Data Retention

Best to test to failure, but may be impractical.

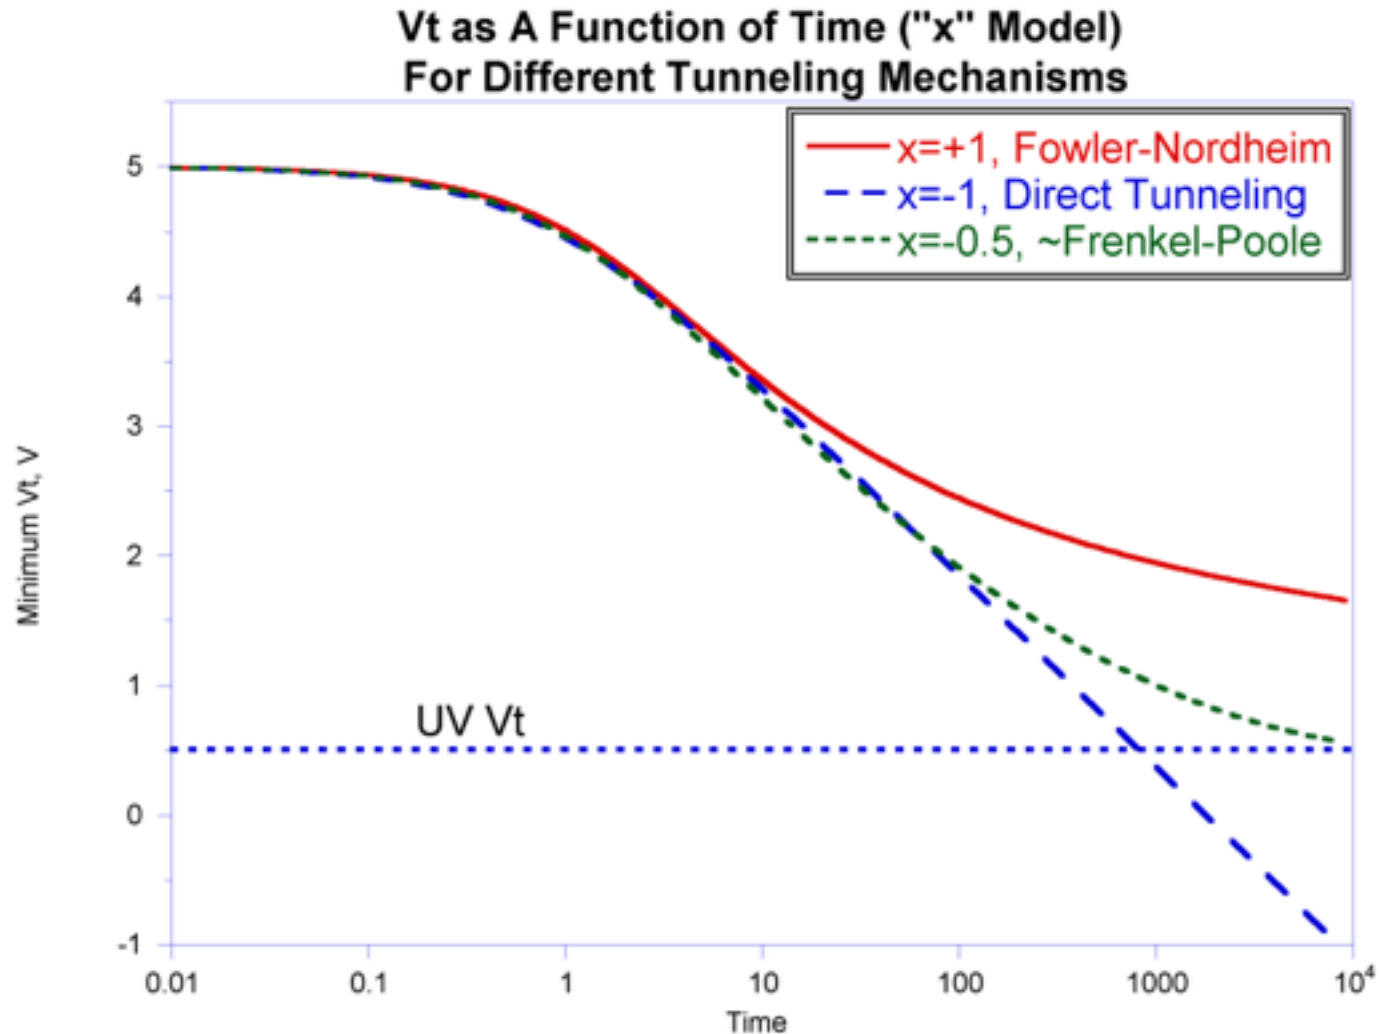May have to extrapolate to failure.

Several models.

- A. Hoefler et.al., "Statistical Modeling of the Program/Erase Cycling Acceleration of Low Temperature Data Retention in Floating Gate Nonvolatile Memories", Proc. IRPS, 2002

$$TTF = \frac{C_T t_{OX}}{AB^X sX} \left\{ \exp\left[s\left(\frac{B}{E(V_{tf})} \right)^X \right] - \exp\left[s\left(\frac{B}{E(V_{ti})} \right)^X \right] \right\}$$

- H. Belgal et. al., "A New Reliability Model for Post-Cycling Charge Retention of Flash Memories", Proc. IRPS, 2002

$$V_{tadj} = \left(-\frac{1}{b} \right) \cdot ln[t] + \left(-\frac{1}{b} \right) \cdot ln[I_0] + \left(-\frac{1}{b} \right) \cdot ln\left[\frac{b}{C_{tot} \cdot \alpha_g} \right] + V_{tna}$$

Vt as A Function of Time ("x" Model) For Different Tunneling Mechanisms

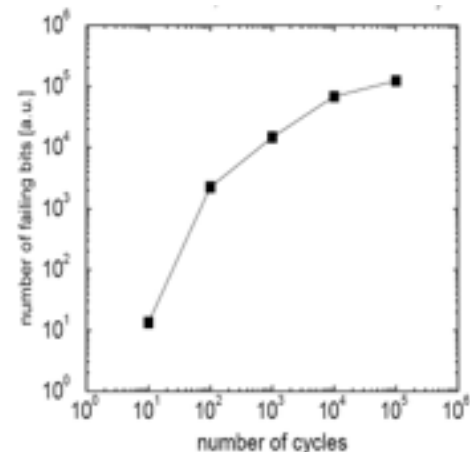# Flash Memory Reliability-Low Temperature Data Retention

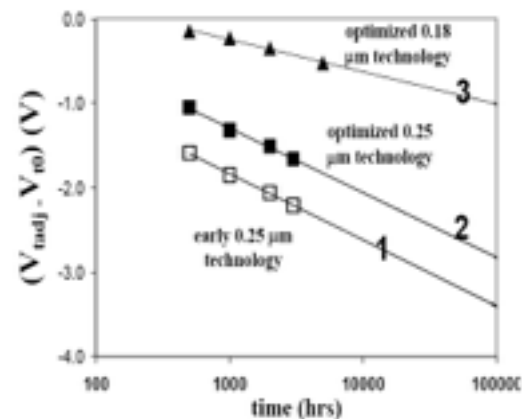LTDR is weakly temperature dependent (low Ea)

Strongly field dependent.

- Very sensitive to level with higher Vt's moving down faster.

In general extrapolation is required to get a good estimate of the lifetime.

Similar to HTDR, cycling will degrade LTDR.

A. Hoefler et.al., Proc. IRPS, 2002

H. Belgal et. al., Proc. IRPS, 2002

LTDR-Extrapolation

Because of the low activation energy, room temperature data retention typically requires extrapolation.

The example shown here uses the RBER to extrapolate to a target value.

- The target value should account for any non-idealities in the part otherwise the estimates will be unrealistic.
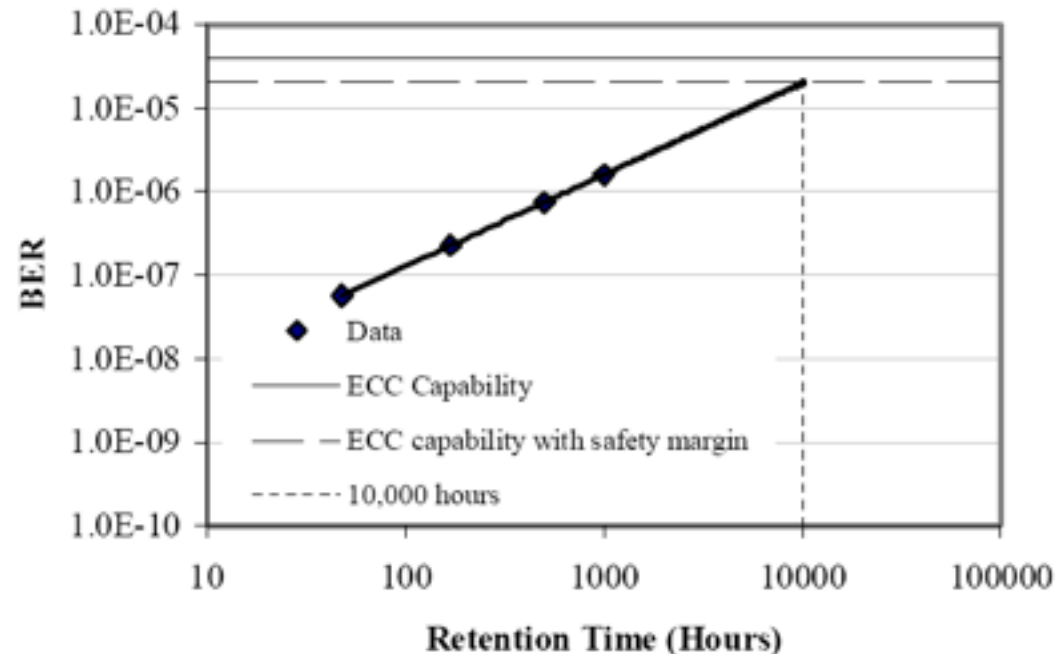
JEDEC JESD218

UBER Events vs FFR Events

Fails are not counted twice.

- A data loss event could be considered a UBER event or an FFR event.

7.1.2 Categorization of failures

Failing SSDs are to be divided into three categories: non-endurance failures, endurance functional failures, and endurance data errors. Non-endurance failures are to be excluded from consideration in the endurance verification but must of course be considered if relevant to other parts of drive qualification. The number of functional failures is to be held against the FFR acceptance criterion (equation 2), and the number of data errors is to be held against the UBER acceptance criterion (equation 3).

Failures are to be categorized as non-endurance failures only if compelling evidence exists that they were not caused by the act of writing the drive to its endurance limit, or by the subsequent retention stress. Failures that are not in the circuit path of the written data (for example, failures isolated to power supplies and capacitors) may be considered non-endurance failures. Failures in the circuit path of the written data, particularly the controller and the nonvolatile memory, will more often be considered endurance failures, but there may be exceptions. Failures that are in the circuit path of the

JEDEC JESD218

What About Everything Else?

JESD218 provides a good framework for demonstrating that an SSD meets its data-retention requirements at the maximum TBW rating.

- This is really just a test of intrinsic capability of the media.

A comprehensive qualification includes many other tests, all of which contribute to the understanding of the total drive reliability.

- An enterprise drive may have over 1000 components that could cause a drive to fail in the field.

Endurance (JESD 218)

Mechanical

RDT

Comprehensive Reliability Plan

4-Corners Test

Other

Manufact. Tests
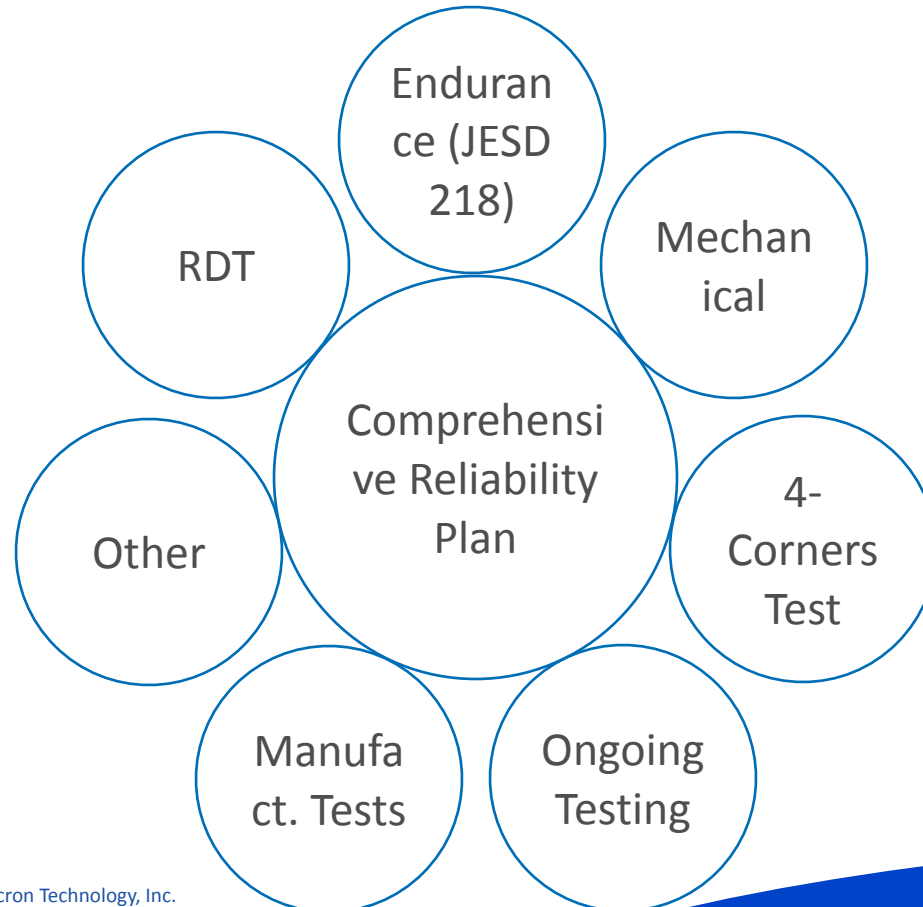
Ongoing Testing

Comments on Historical HDD Testing

Testing for SSDs need to be focused on SSD related concerns.

Tests that are historical, but were not developed for SSDs should be <u>reviewed</u> and determined if they apply.

- They may need to be modified or removed in favor of other tests that will help truly evaluate the quality and reliability of the product.

- For example, historical "altitude" testing for HDDs involve placing the HDD in a low pressure chamber due to concerns of the head not being able to "float" above the spinning media.

  - This is a good test for these types of mechanical systems, but have no value on SSDs.

Resources and time better spent to invest in developing new test and invest in equipment/capability that is meaningful.

All SSD testing/qualification procedures should be under constant review to ensure the procedures evolve as required.
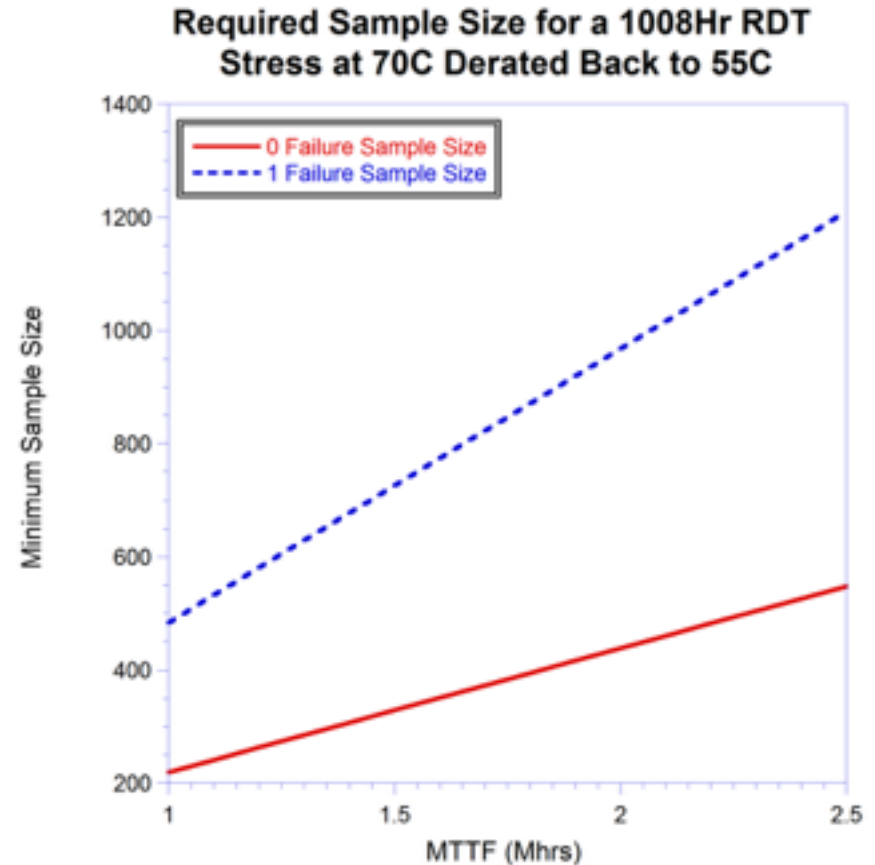
Micron®

Reliability Demonstration Test

The purpose of RDT is to demonstrate the reliability of the SSD.

- Typically specified in MTTF.

Fixed size/duration sample plans fail to comprehend the difference in MTTF in many cases.

Drive families should be defined in a way that makes sense.

Required Sample Size for a 1008Hr RDT Stress at 70C Derated Back to 55C

Legend:
— 0 Failure Sample Size
‑ ‑ ‑ 1 Failure Sample Size

Y-axis: Minimum Sample Size (200–1400)
X-axis: MTTF (Mhrs) (1–2.5)

Acceleration Factors

In order to split out mechanisms, the distributions need to plotted on a common axis.

This requires application of the appropriate acceleration factors.

Acceleration factors are applied based on the physics of failure.

Starting point, general rule-of-thumb

- NAND failures are treated as "TBW Accelerated"

- Hardware failures are treated as "Thermally Accelerated"

- Firmware failures are treated as un-accelerated.
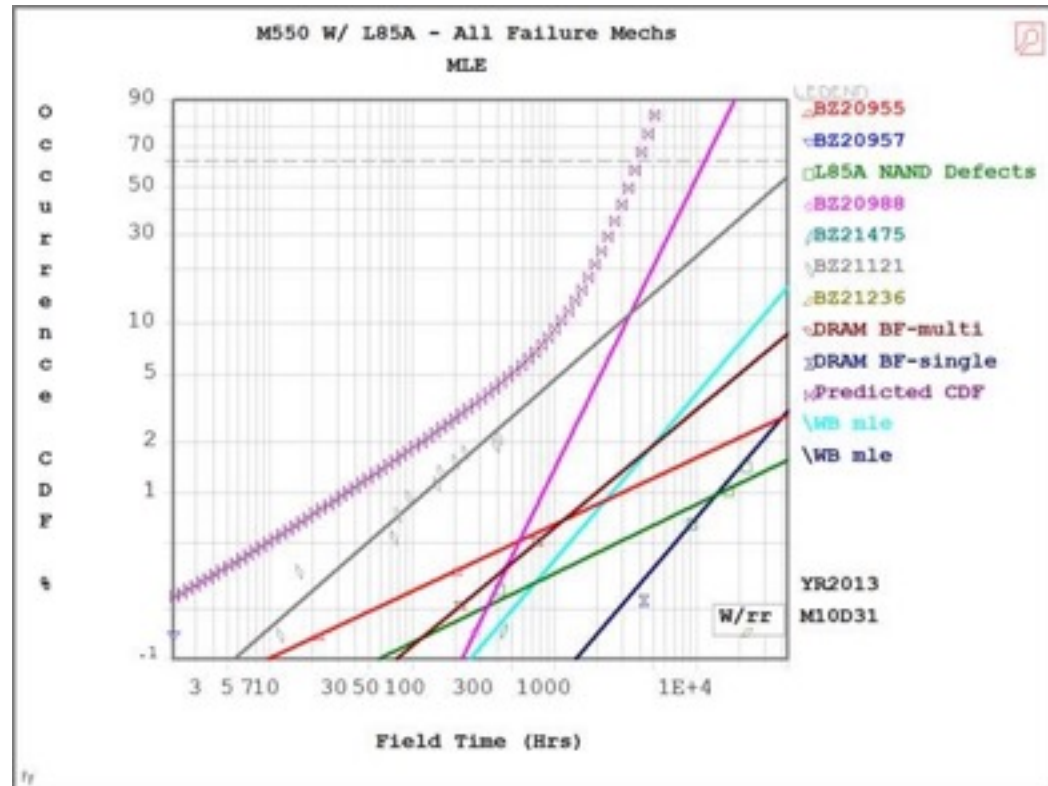
TBW Acceleration

If NAND cycling is expected to accelerate a failure mechanism, Micron uses a TBW acceleration factor when estimating drive lifetime.

- The acceleration factor is based on the data transfer rate in test versus what is expected in the field.

- Although not explicitly called out, this is embedded in the JESD218 methods

In some cases we will go beyond the specified drive TBW in order to get the best failure distributions.

- The tests are designed to generate failures (even if the drive meets spec), this gives better estimates of the AFR/MTTF as well as visibility into failure mechanisms as well as reliability margins.

$$A(TBW) = \frac{NAND\_Cyc/Hr\_Stress}{NAND\_Cyc/Hr\_Field}$$

$$= \left[\frac{10^{12} \cdot TBW_{stress} \cdot WA_{stress}}{2^{30} \cdot D_{physical} \cdot t_{stress}} \right] \cdot \left[\frac{2^{30} \cdot D_{physical} \cdot t_{field}}{10^{12} \cdot TBW_{field} \cdot WA_{field}} \right] = \frac{TBW_{stress} \cdot WA_{stress} \cdot t_{field}}{TBW_{field} \cdot WA_{field} \cdot t_{stress}}$$

For example, for a 100G drive transferring ~130TB in 1008 Hrs with a spec of 175TB over 5yrs, the acceleration factor would be:

$$A(TBW) = \frac{TBW_{stress} \cdot WA_{stress} \cdot t_{field}}{TBW_{field} \cdot WA_{field} \cdot t_{stress}} = \frac{130TB \cdot 1.2 \cdot 43800hrs}{175TB \cdot 1.2 \cdot 1008hrs} = 32$$

Thermal Acceleration

Thermal acceleration is typically based on an activation energy of 0.7eV.

The appropriate activation energy depends on the actual failure mechanism, and directly measured values will give the best estimates.

The literature and various specifications can also provide guidance.

$$\frac{ttf_{T_1}}{ttf_{T_2}} = \frac{A \bullet \exp\left(\frac{E_a}{kT_1}\right)}{A \bullet \exp\left(\frac{E_a}{kT_2}\right)} = \exp\left[\frac{E_a}{k}\left(\frac{1}{T_1} - \frac{1}{T_2}\right)\right] = Acceleration\_Factor$$

| Sect.	Failure Mode	Failure Mechanism	Activation Energy E_{aa} (eV)
5.8	Short, leakage	TDDB; traps & percolation	0.75
5.8	Short, leakage	Cu ion drift	1.0
5.9	Open	Al EM; vacancy transport	0.8
5.9	Open	Al EM; grain-boundary diffusion	0.68
5.9	Open	Al EM; interfacial diffusion	0.95
5.10	Open	Cu EM; vacancy transport	0.9
5.11	Open	Al corrosion (Chloride)	0.75
5.11	Open	Al corrosion (Chloride)	0.75
5.11	Open	Al corrosion (Phos Acid)	0.3
5.11	Open	Al corrosion (Chloride)	0.75
5.11	Leakage	Diffusion thru passivation cracks	0.79
5.11	Open	Ion transport thru Polyimide	1.15
5.11	I_{oo} quiescent	Water diffusion	0.73
5.11	Leakage	Ionic conductivity – lead frame coplanarity tape 1	0.74
5.11	Leakage	Ionic conductivity – lead frame coplanarity tape 2	0.77
5.12	Open	Al stress migration – vacancy diffusion & drift – voids coalesce	0.6 gb 1.0 bamboo
5.12	Open	Al stress migration – vacancy diffusion & drift – voids coalesce	0.6 gb 1.3 intra-grain
5.13	Open	Cu stress migration – vacancy diffusion & drift – voids coalesce	0.74-1.2 for Cu-cap interface - strong fcn of interface prep
5.14	Open	Crack propagation - stress concentration in PbSn solder fatigue if >30C	N/A
5.14	Open	Crack propagation - stress concentration in PbSn solder fatigue if <30C	N/A
5.14	Open	Crack Propagation – Fatigue SnAg solder	N/A
5.14	Open	Al wire	N/A
5.14	Open	Au$_4$Al IMC	N/A

JEDEC JEP122G

Other Acceleration Factors

Acceleration factors for other tests (such as temp cycle, HAST etc…) should be reviewed to ensure that the tests can be related to field usage conditions in order to allow for risk analysis as well as ensuring the tests are aligned with expectations.
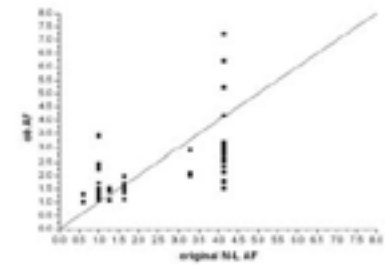
(a) Original N-L model
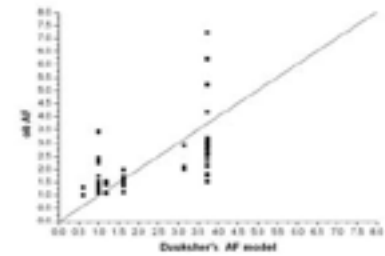
(b) Dauksher's model
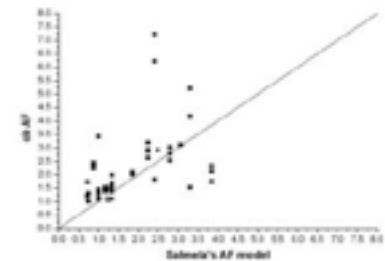
(c) Salmela's model

Fig. 5 The comparison of each N-L model based on our experimental data.
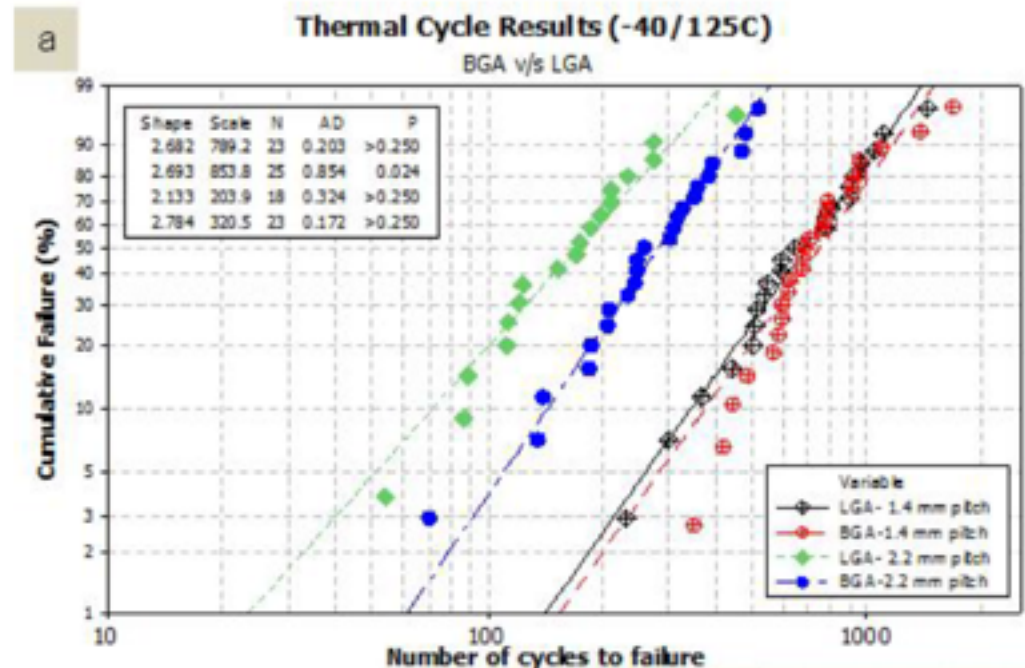
Kuan-Jung Chung et. Al., IMPACT, 201

Mechanical Testing

Temperature Cycling

- Used to evaluate solder-joint reliability.

HAST (Highly Accelerated Stress Test) and THP (Temperature Humidity Bias)

- Designed to evaluate the impact of moisture and temperature on components.
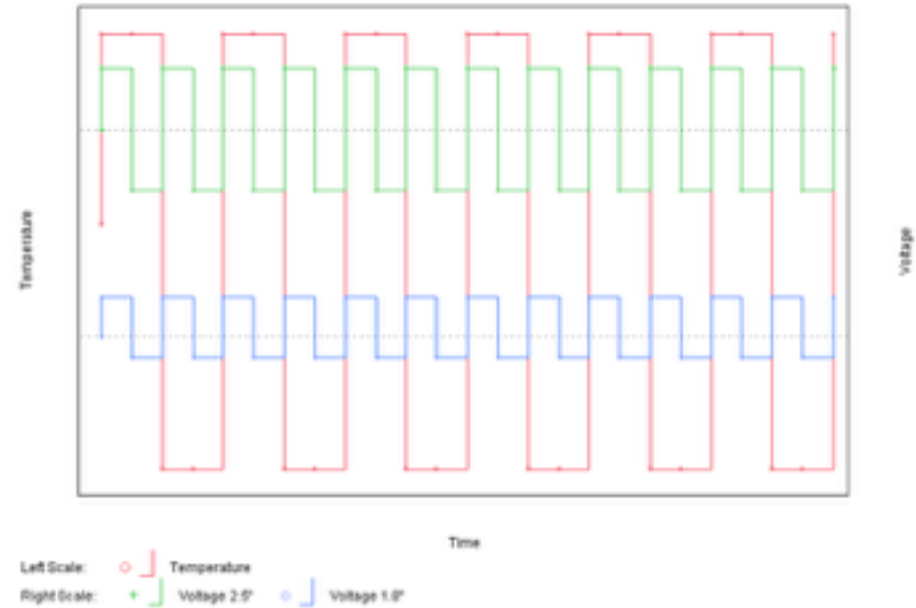

Babak Arfaei et.al., Electronic Comp. & Tech. Conf., 2014

4-Corners Testing

4-Corners testing measures the drive capability across temperature a voltage corners.

- Should include cross-temperature reads.

- Normally there are no operations during the temperature transitions.

| Left Scale: | ○ | Temperature |
| Right Scale: | + | Voltage 2.5° | | ○ | Voltage 1.8° |

Other Tests

There may be additional tests specific to particular products or markets.

Determining margin to failure, rather than pure substantiation testing.

August 25, 2015 | ©2014 Micron Technology, Inc.
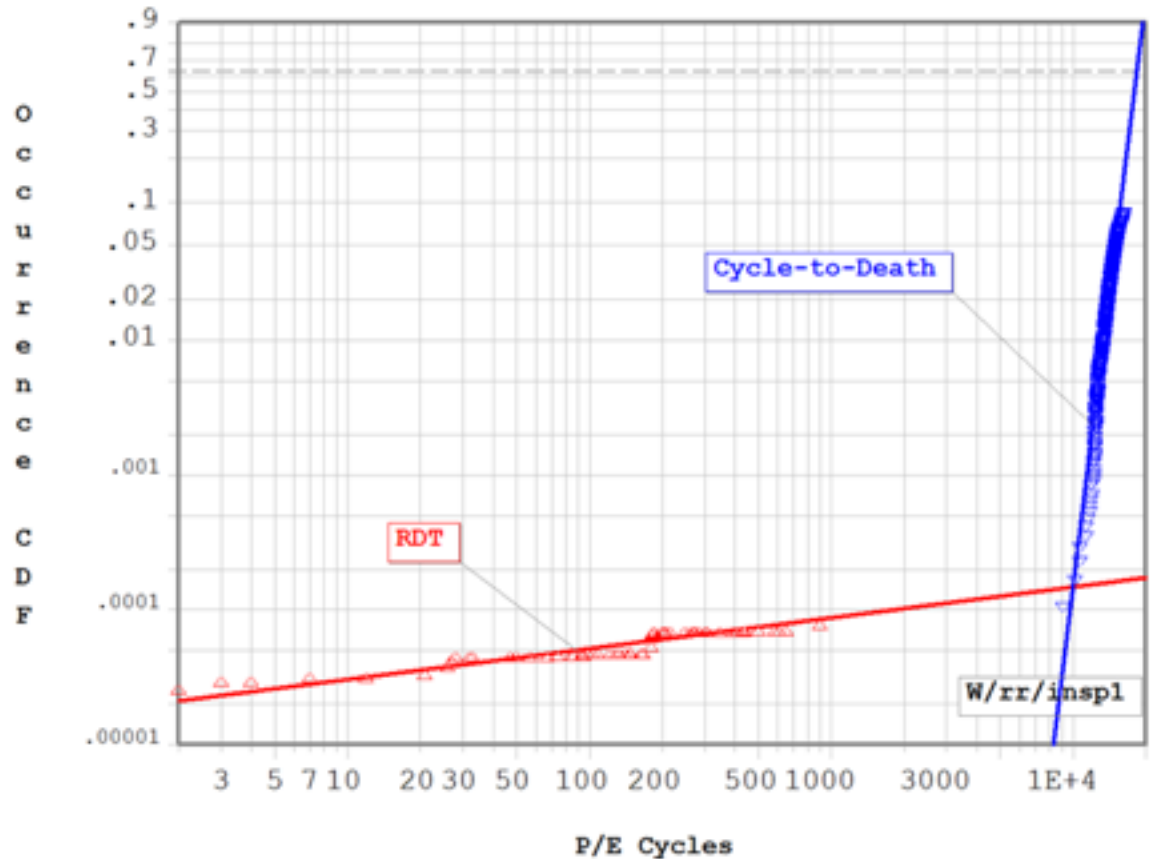
Additional Tests

Manufacturing Tests

- Examination of manufacturing tests of material going into qualification tests (and other testing), is important to pick up lower level issues and defectivity that the lower sample sizes of reliability testing may miss.

- Burn-in (if it exists) can provide important information about underlying failure distributions.

Ongoing Tests

- Post qualification it is important to monitor the reliability of the systems in order to ensure that it is not shifting down over the manufacturing lifetime.

Comments/Conclusions

JESD218 provides a method for evaluating endurance under specific conditions in a uniform manner across the industry.

- This is only 1 portion of a robust qualification strategy.

SSD qualification should include clear, statistically meaningful tests that cover the lifetime of the product under realistic conditions.

- Specs based on HDDs may no longer be valid due to what they were meant to detect and should be reviewed with a critical eye.

Models around testing conditions need to be constantly evaluated as the industry seeks to design SSDs closer to the media capability.

References

JEDEC Standard JESD218A, "Solid-State Drive (SSD) Requirements and Endurance Test Method", September 2010.

Muroke, P. , "Flash Memory Field Failure Mechanisms ", Reliability Physics Symposium Proceedings, 2006.

Khayat, P. R., Kaynak, M. N., Parthasarathy, S., Tehrani, S. S., "Performance Characterization of LDPC Codes for Large-Volume NAND Flash Data", IEEE IMC, 2015.

Mielke, N.; Belgal, H.P.; Fazio, A.; Meng, Q.; Righos, N., "Recovery Effects in the Distributed Cycling of Flash Memories", Reliability Physics Symposium Proceedings, 2006.

Cai, Yu ; Yalcin, G.; Mutlu, O.; Haratsch, E.F.; Cristal, A.; Unsal, O.S.; Ken Mai, "Flash correct-and-refresh: Retention-aware error management for increased flash memory lifetime", IEEE 30th International Conference on Computer Design (ICCD), 2012.

Cappelletti, P.; Bez, R.; Modelli, A.; Visconti, A., "What we have learned on flash memory reliability in the last ten years ", IEDM Technical Digest, 2004.

Hoefler, A.; Higman, J.M.; Harp, T.; Kuhn, P.J. , "Statistical modeling of the program/erase cycling acceleration of low temperature data retention in floating gate nonvolatile memories", Reliability Physics Symposium Proceedings, 2002.

Belgal, H.P.; Righos, N.; Kalastirsky, I.; Peterson, J.J.; Shiner, R.; Mielke, N., "A new reliability model for post-cycling charge retention of flash memories", Reliability Physics Symposium Proceedings, 2002.

JEDEC Standard JEP122G, "Failure Mechanisms and Models for Semiconductor Devices", October 2011

Kuan-Jung Chung; LiYu Yang; Bing-Yu Wang; Chia-Che Wu, "The Investigation of Nodified Norris-Landzberg Acceleration Models for Reliability Assessment of Ball Grid Array Packages", 5[th] International Microsystems Packaging Assembly and Circuits Technology Conference (IMPACT), 2010.

Babak Arfaei; Francis Mutuku; Keith Sweatman; Ning-Cheng Lee; Eric Cotts; Richard Coyle, "Dependence of Solder Joint Reliability on Solder Volumn, Composition and Printed Circuit Board Surface Finish", Electronic Components & Technology Conference, 2014.

Micron®

Focused on Memory | Engineered for ᴧ Innovation
                                          your