Eliminating controller-based reverse-engineering in NAND Flash chip-off data recovery

Robin England (Hardware R&D Team Lead, Ontrack)
Oliver Hambrey (Research Engineer, Siglead)

Ontrack® by KLDiscovery SIGLEAD

Data Recovery Techniques for SSD / NAND

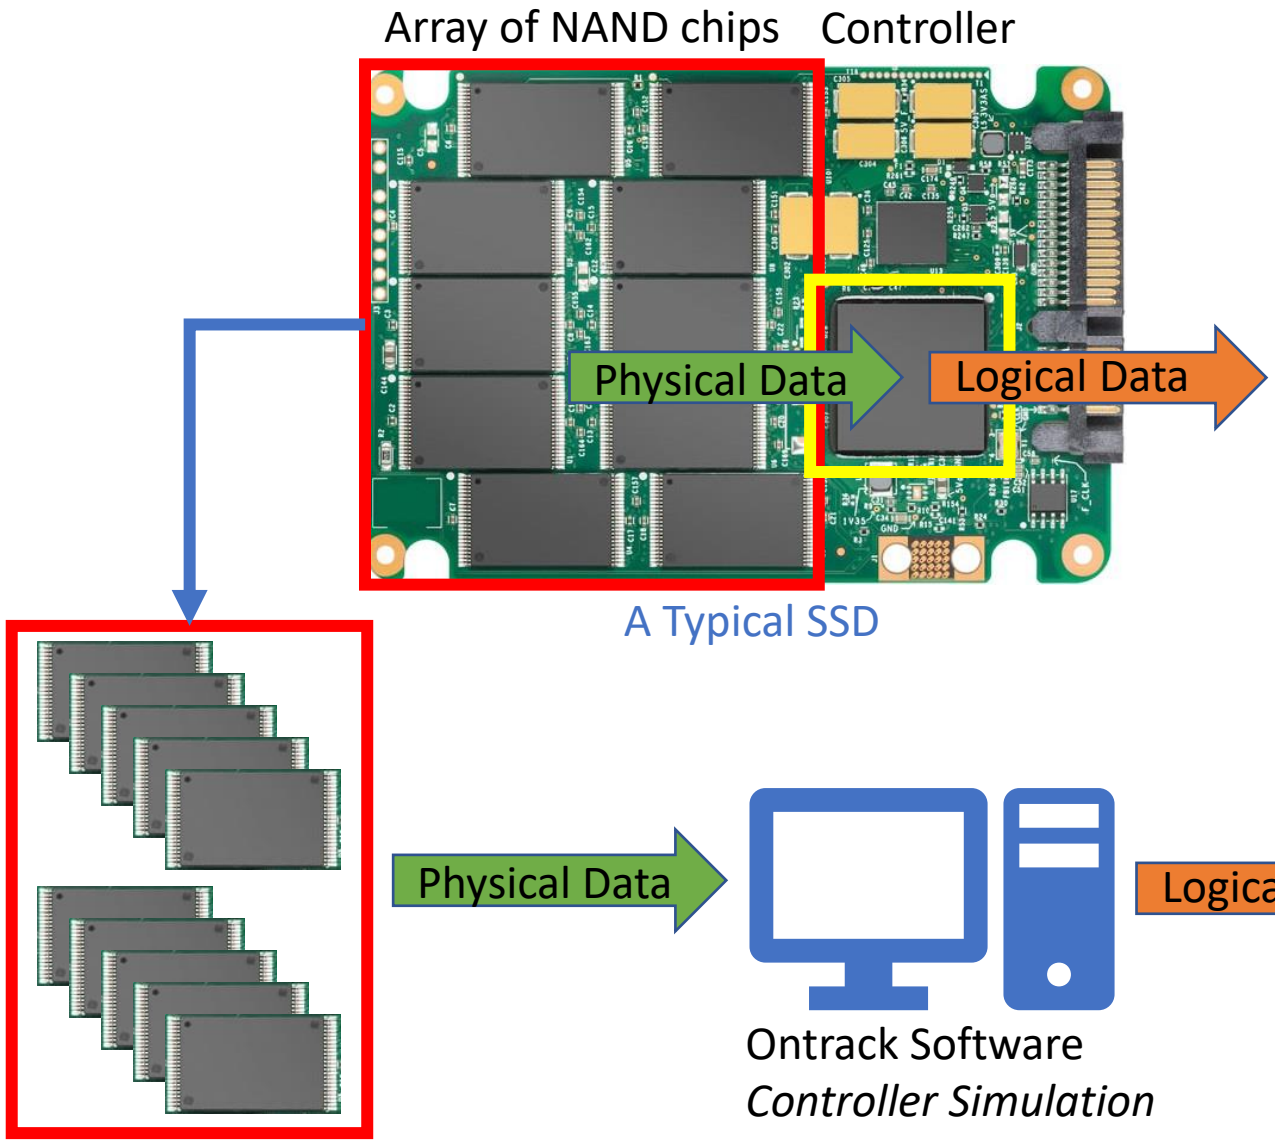

"Chip-On" Method – (Use Original Controller)

- ➢ Examination & diagnosis
- ➢ Component repair
- ➢ Transfer of critical components to custom hardware
- ➢ Firmware modification / System Area repair
- ➢ Further techniques applied as needed during recovery
- ❑ Simpler approach – *if drive design & failure mode allows*
- ❑ Risk of loss to existing data, may not be able to access all data, can take a very long time to read data

"Chip-Off" Method – (Process Physical Data)

- ➢ Remove NAND packages
- ➢ Read *best possible* copy of physical data from each NAND… *can we correct bit errors?*
- ➢ Process physical data to reconstruct logical data
- ❑ Read-only access to a snapshot of all physical data, limited further risk to data
- ❑ Complex – need to understand controller specifics & hardware encryption may preclude

NAND Flash & Bit Errors

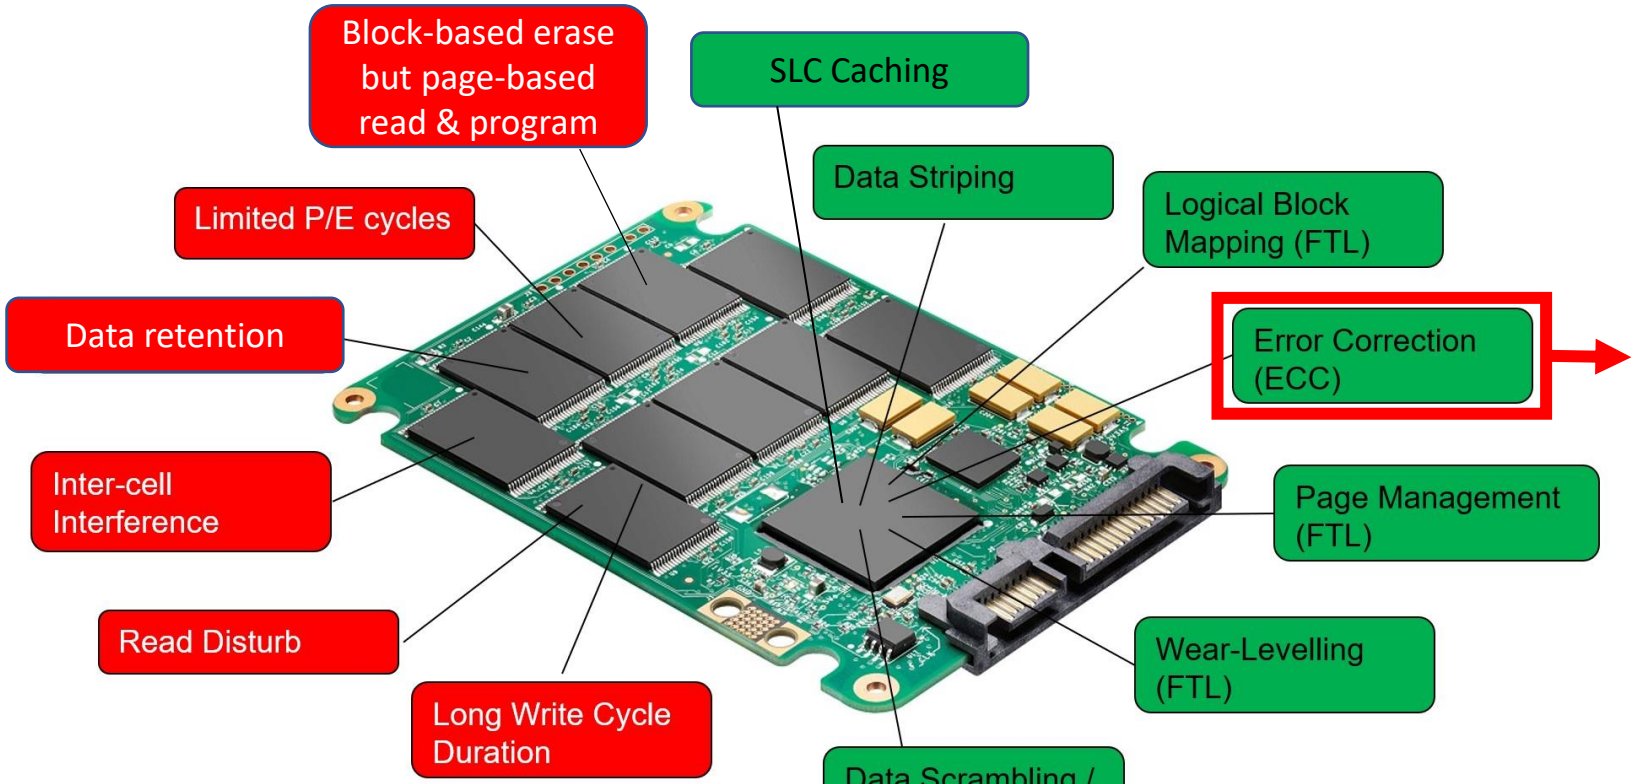

NAND's Challenges…
(A tendency toward bit errors)

Controller's Complexity…
(Fighting the effects of bit errors)

Causes of Data Loss

- Firmware / SA corruption
- Electronic component failure
- Physical damage
- **Un-correctable bit errors due to…**
  - Reduced endurance (P/E cycles)
  - Data Retention (electron de-trapping)
  - Read disturb (nearby cells)
  - Operating & storage temperatures

How do bit errors affect our data?

"The quick, brown fox jumps over the lazy dog!

Backed up by the world's largest R&D team in data recovery, as well as exceptional customer support, we make sure that your data recovery experience is first class. With Data Recovery services to suit customers ranging from home users to the largest businesses, Ontrack can help get your data back."

Source file:
`ASCII_Text.txt`
349 bytes (2792 bits)

1 random bit error (0.035%)

ASCII text with bit errors

```
The quick, brown fox jumps over
the lazy dog!

Backed up by the world's largest
R&D team in data recovery, as
well as exceptional customer
support, we make sure t`at your
data recovery experience is first
class. With Data Recovery
services to suit customers
ranging from home users to the
largest businesses, Ontrack can
help get your data back.
```

10 random bit errors (0.35%)

ASCII text with bit errors

```
The quick, brown fox!jumps over
the lazy dog!

Backed up by the world's largest
R&D team in data recovery, ac
well as exceptional custo-er
support, we make 3ure that your
data recovery experience is first
class. With Data Recovery
sertices to suit customers
ranging from jome usErq to the
largest businesses, Ontraci can
help get your data back.
```

100 random bit errors (3.5%)

ASCII text with bit errors

```
Tle quick, brown fïx jumðs"ovgr
thå la~ù dnf¡

BecKed up"by TLe gmrld's lapgust
R&  4eam in taJa rãc0veòY,
as`welm as"exceptional busuomer
surporo,0we iake's5re that you2
dat! rec_ve2y my`èvyejcu`is firs|
clcss. With Data
Rec/weryàservices po#suiô
ãustomers vanfyNo(from iome uóevs
t0vje lAr#ost busiæEsses,
Ontr!#k can hEmp get xiur Data
baco.
```

1000 random bit errors (35%)

ASCII text with bit errors

```
TÛå
qu{eO( iãÎ{r¢Go¨`êýip_°.qi(ð`B`]1
xhi`wï"/Po7àGi‹%†u¤

 ÆoÀblu* t-eiŒñdáoní Soìq
`àÅèä@ághu¸%`6ñ£mGaí*övbWfUds*m1r
yñ)Yïù $F9Á
rYú®>`ûY$aidmpkãNÃe¤°ñãvyj[°dçl%u
(åQbæ‘ßaekKe6›d:Bz:xaDU(Ôo{E¹0a{

)ie!ìroocÈ
\m5+sk>p@T© °L€Ècòç%7d
#whkMíste;œäuÐ¨ds n Igì]€o-4Hùs=2*

?
```

How do bit errors affect our data?

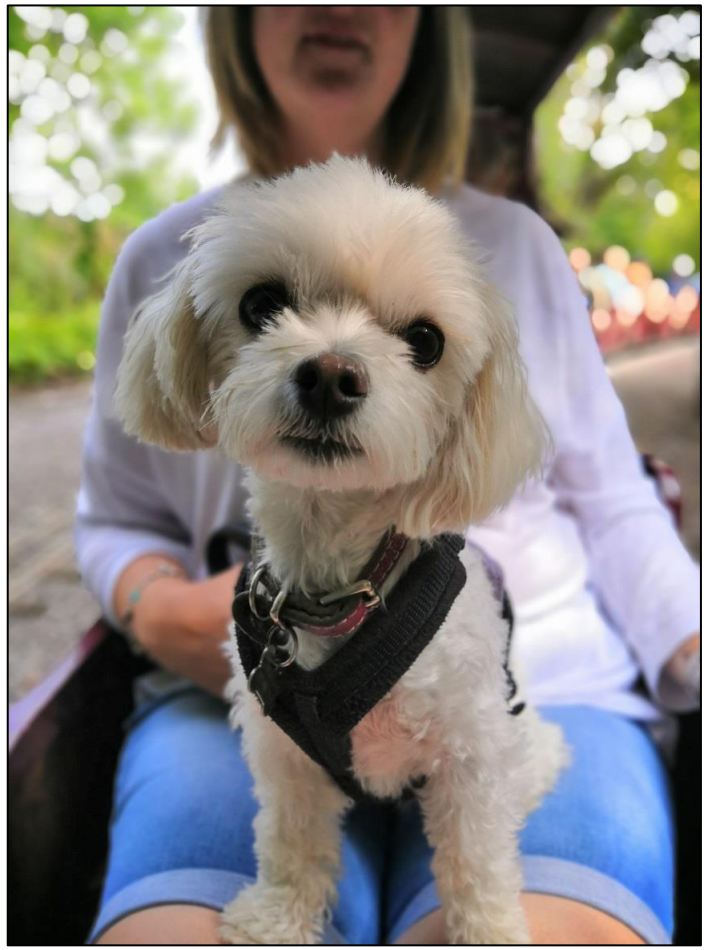

Source file: `Holly.jpg`
214,745 bytes (1,717,960 bits)

1 random bit error (0.000058%)

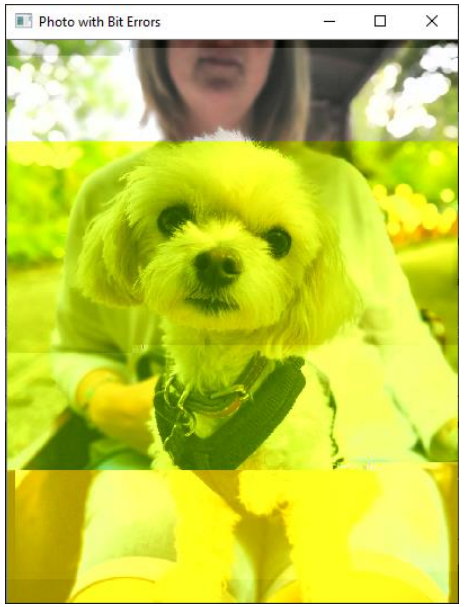

10 random bit errors (0.00058%)

100 random bit errors (0.0058%)

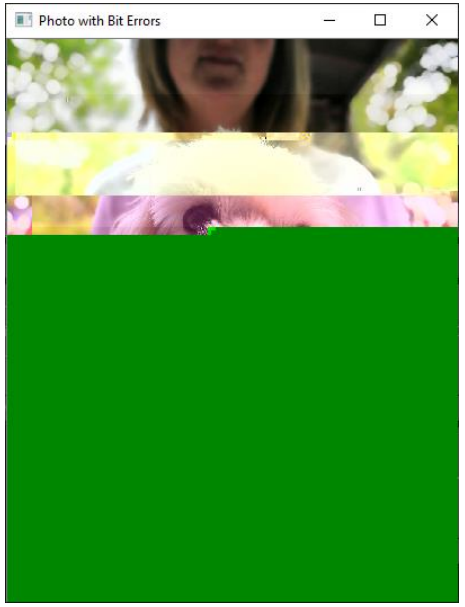

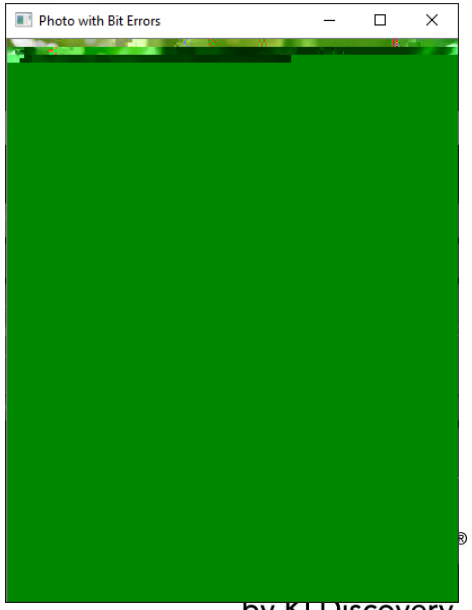

1000 random bit errors (0.058%)

Also affected by bit errors…

- Drive firmware
  Code & critical firmware data needed for data retrieval (e.g., Flash Translation Layer)
- Host Operating System
  Boot, executable and system files
- Filesystem / Volume
  File pointers, directory entries, filenames, cluster chains, file attributes
- Data security measures
  Full disk and file encryption, authentication, data validation (e.g., checksum / hash)
- Even data not *directly* affected by uncorrectable bit errors may become unavailable due to bit errors elsewhere!

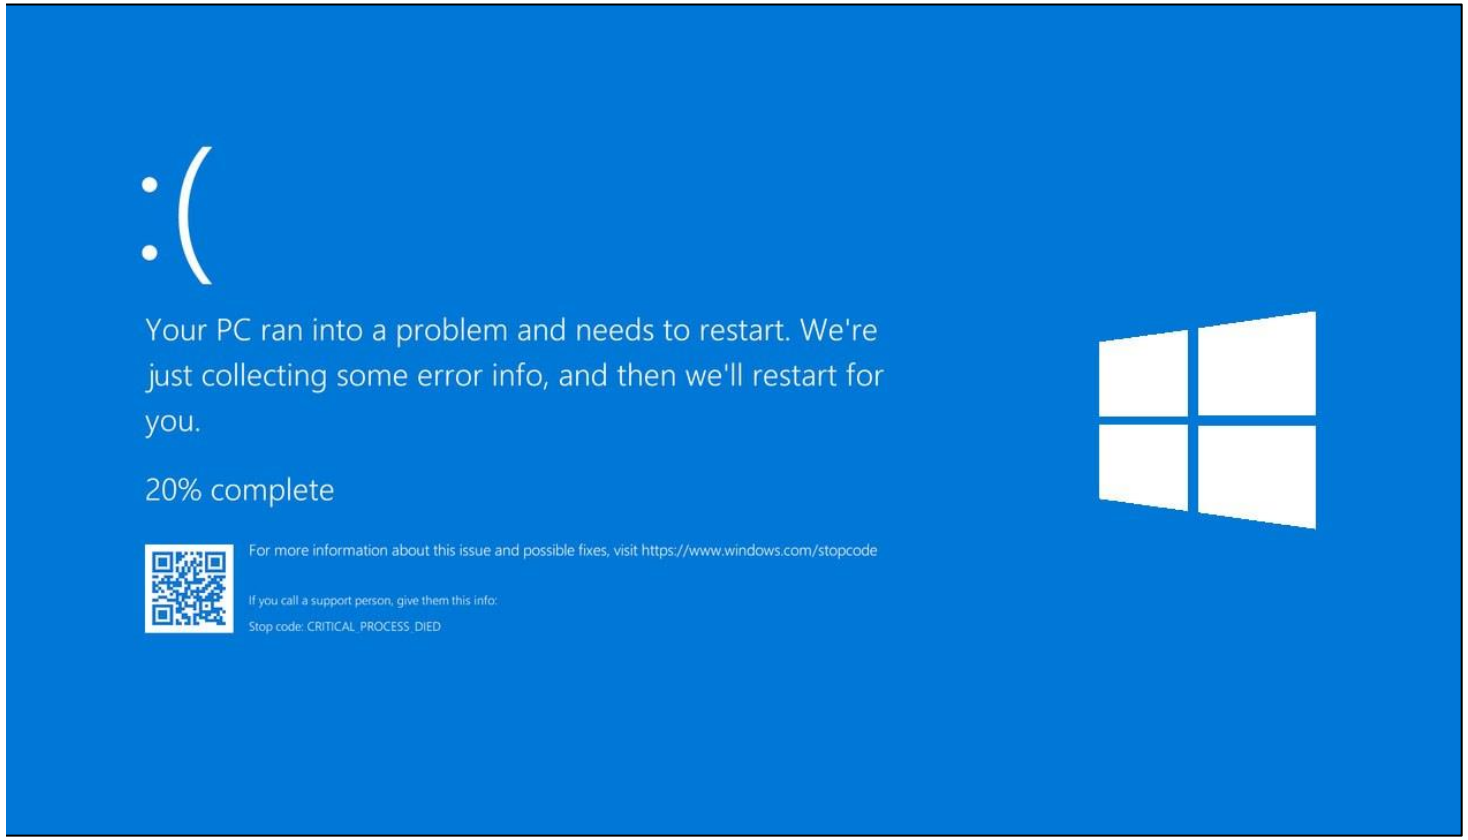

Cell Data Density, Bit Error Rate, ECC

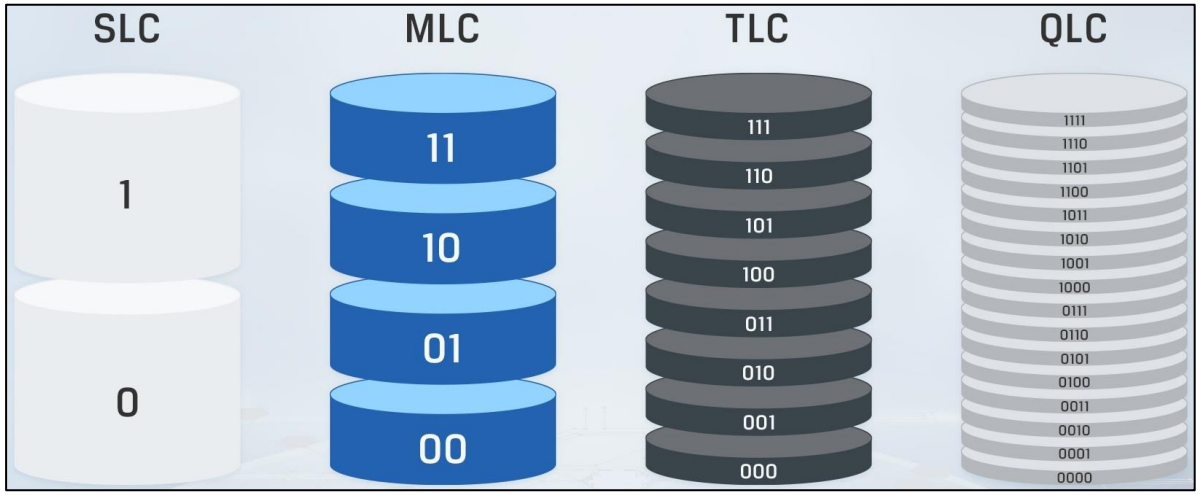

Single-Level Cell
1 data bit per cell
100k P/E cycles

Multi-Level Cell
2 data bits per cell
5k-10k P/E cycles

Triple-Level Cell
3 data bits per cell
Up to 3k P/E cycles

Quad-Level Cell
4 data bits per cell
Up to 1k P/E cycles

Bit Error Rate

NAND Cell Data Density

- As cell data density increases it becomes more difficult to reliably read back the correct data
- A much smaller change in cell charge level leads to a greater potential for bit errors in QLC when compared with SLC
- Bit errors arise when there is uncertainty that what we have read from the cell is what was written to the cell
- **Uncorrectable** bit errors occur when the number of cells in a codeword that have bit errors exceeds the error correction capability of the ECC used…**drive cannot return data ☹**

Read Retry, Read Offset, Auto Read Calibration

Read Retry

- Introduced in some of the last 2D NAND generations
- Provides a set of recommended preset read offset values that can be used for subsequent reads
- Coarse adjustment and finite set of values
- Relatively quick to try all recommended values, but lower chance of success

Read Offset

- Introduced in the last 2D NAND generation
- Mandatory for 3D TLC and QLC NAND
- Finer control and greater number of potential adjustment points to try for subsequent reads, but no guidance from NAND manufacturer on which to use
- Slower than Read Retry but slightly better chance of success; more possible “settings” available

Auto Read Calibration

- Introduced in 3D NAND
- Provides on-the-fly dynamically recommended values to be used with Read Offset
- Intended to be a “best of both worlds”; offers the finer calibration provided by Read Offset along with some assistance with selecting offsets
- High-latency

Ontrack® by KLDiscovery SIGLEAD

Existing Calibration Methods: Observations

- Effectiveness varies considerably depending on cell density, NAND age and usage
- No standard for implementation - open to interpretation by firmware engineers
- Calibration sometimes sacrificed in favour of drive performance… i.e., drive firmware limits efforts to correct data
- During each read attempt, the data must be always processed with the original ECC to measure success or failure
- Auto Read Calibration – industry's newest approach but in practice can underperform - most notable on new drives using QLC - as soon as bit error rates begin to increase

So… we need a better way to achieve optimal calibration…

Ontrack® by KLDiscovery SIGLEAD

What does Optimal Calibration look like?

- ☑ Must outperform currently available calibration methods
- ☑ Must be independent of ECC (BCH or LDPC variant)
- ☑ Must include feedback and verification of effectiveness
- ☑ Must be independent of NAND manufacturer; does not rely on vendor-unique commands
- ☑ Must use existing and standard features of NAND; no decapsulation or modification of the NAND package
- ☑ Must not require prior knowledge of the stored data
- ☑ Must be flash controller and firmware independent

Ontrack® by KLDiscovery SIGLEAD

Read Voltage Thresholds

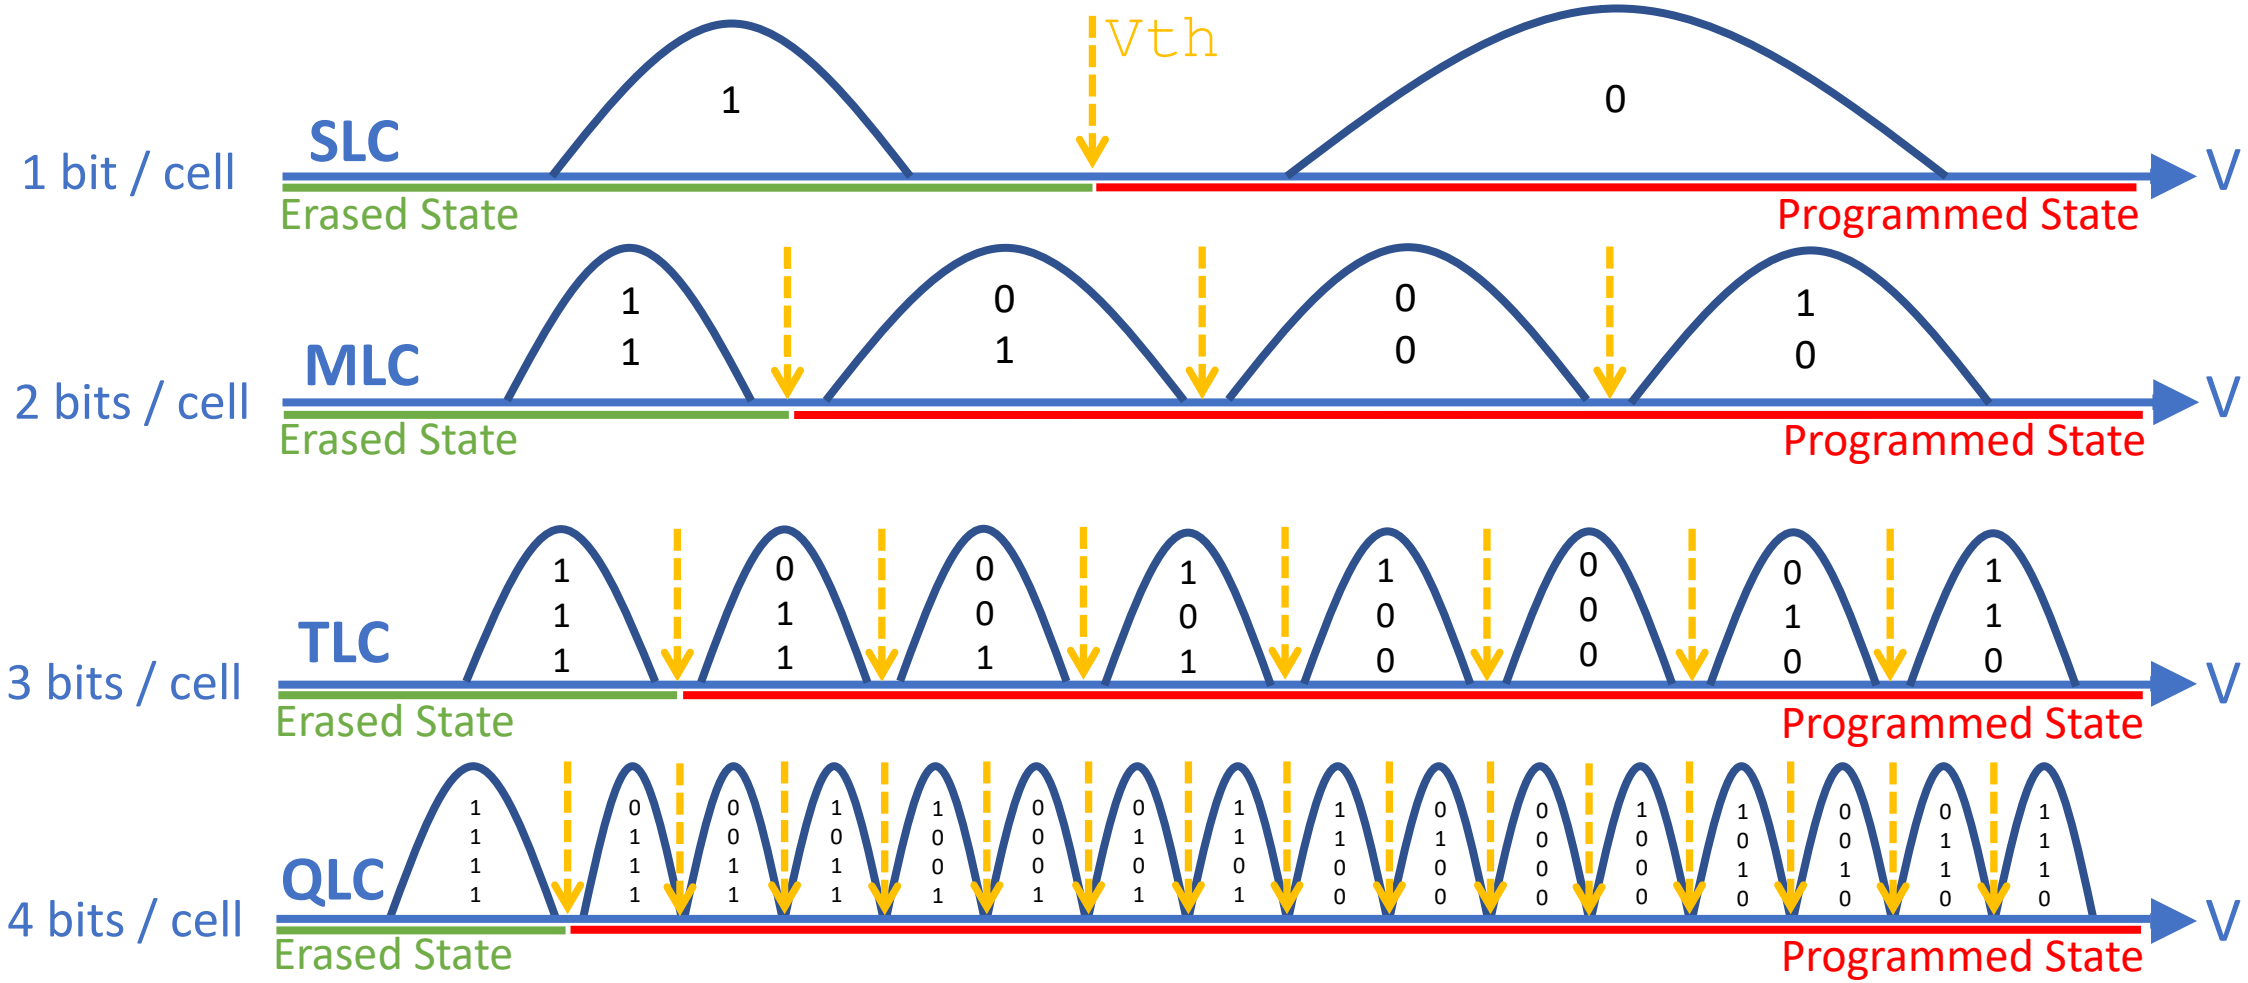

Gaussian Mixture Calibration – Voltage Distribution

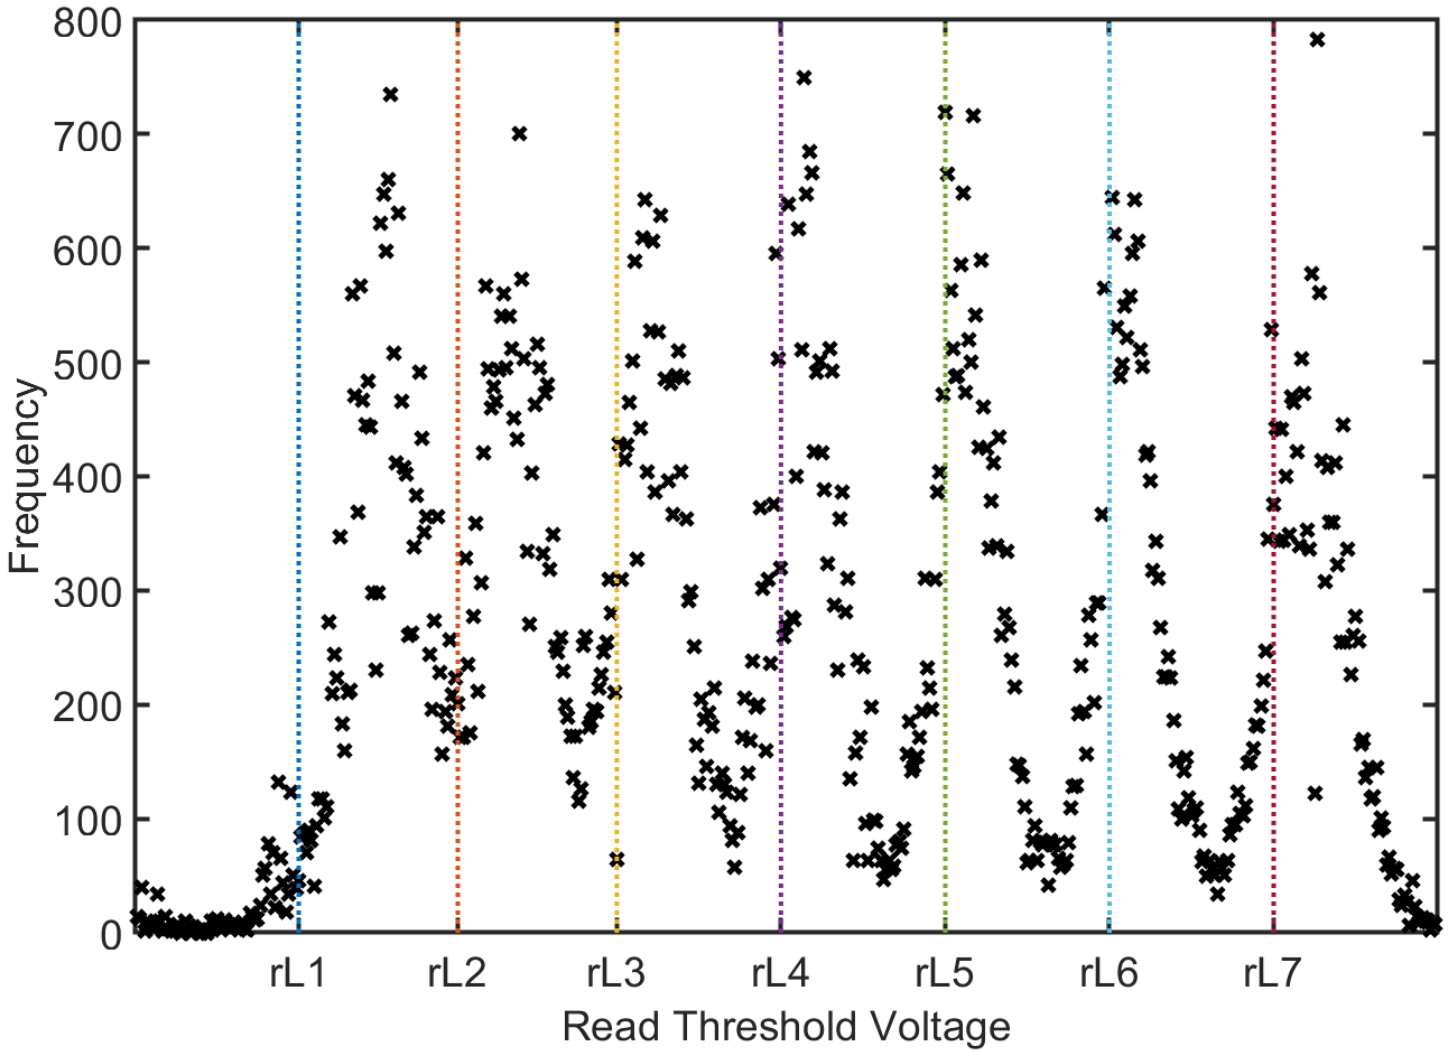

1. Make *"high-resolution"* image of the NAND

- Read NAND many times
- Each time, make a small increment in all read thresholds

2. Process high-resolution image to obtain the voltage distribution

- Processing of high-resolution image will precisely tell us the voltage level of each cell

- **Challenges:**

❑ Making high resolution image
- Significantly slower than reading once
- Requires control over NAND's read thresholds
- Read disturb affects data integrity

❑ Processing high-resolution image
- Offline computational power
- Dealing with transient noise

Gaussian Mixture Calibration – Data Fitting

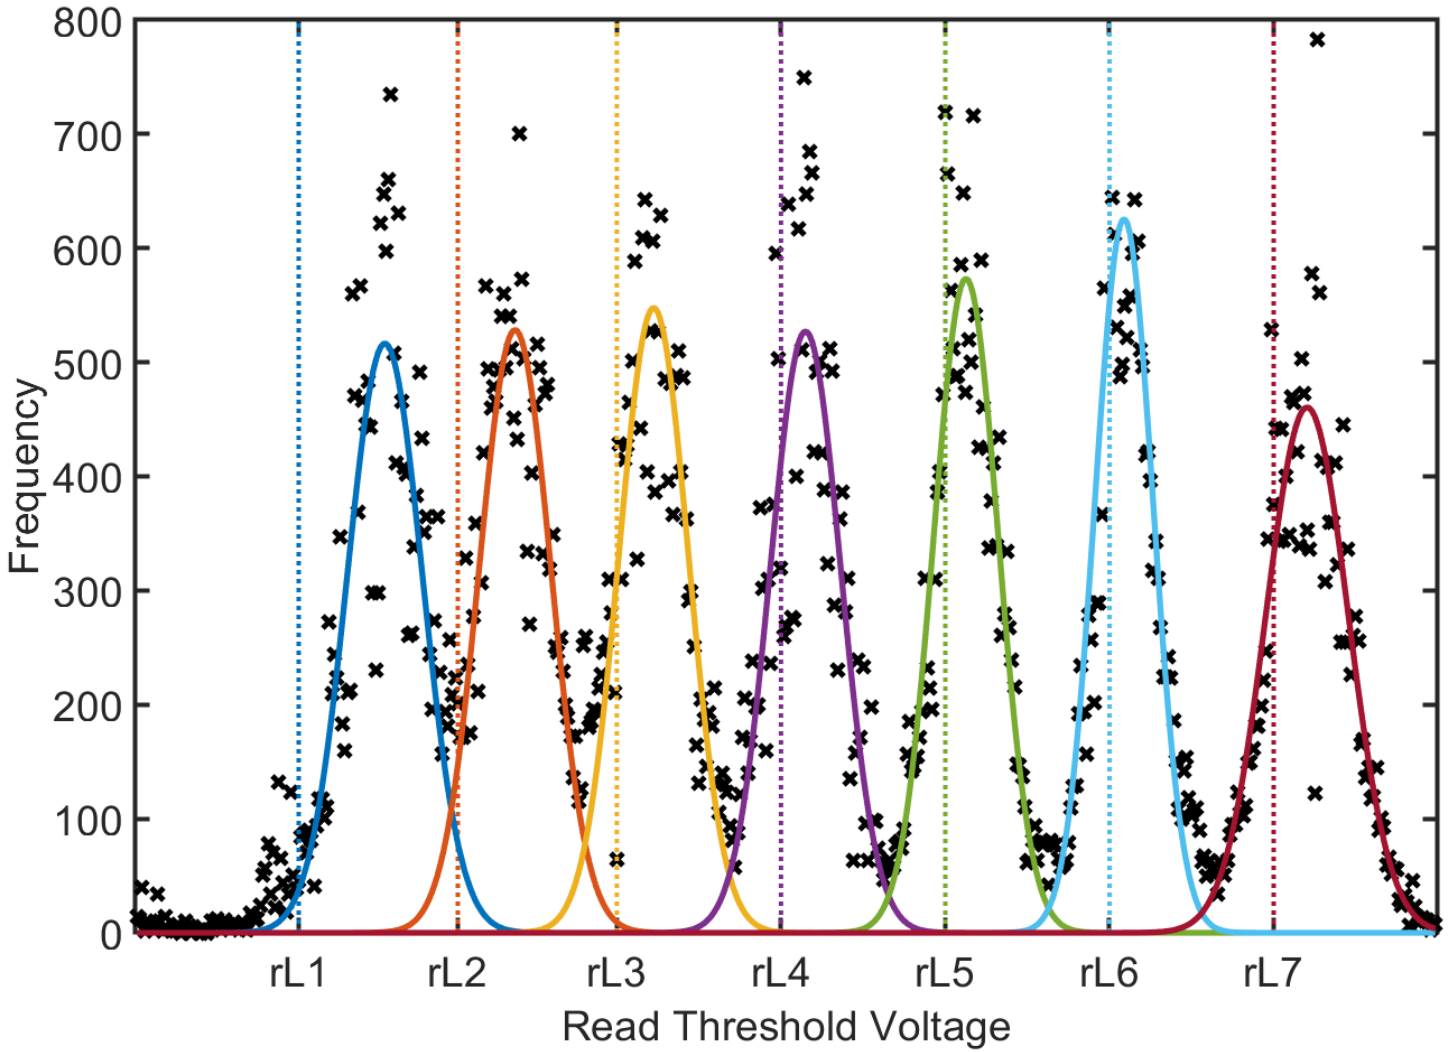

3. Perform data fitting to the voltage distribution

- We find a *"line of best fit"* using a mixture of Gaussian distributions (bell curves)
  - Reasonable approximation to data
  - Low dimensional parameter space
  - Algorithmically *"simple"* to achieve
- A smooth estimate of the voltage distribution resolves ambiguities arising from transient error effects
- Allows us to estimate and minimize data integrity error

- **Challenges:**
  - ❑ Suitability of model
  - ❑ Data fitting is typically iterative, sensitive to initial condition & convergence not guaranteed

Gaussian Mixture Calibration – Estimating Data Integrity Error

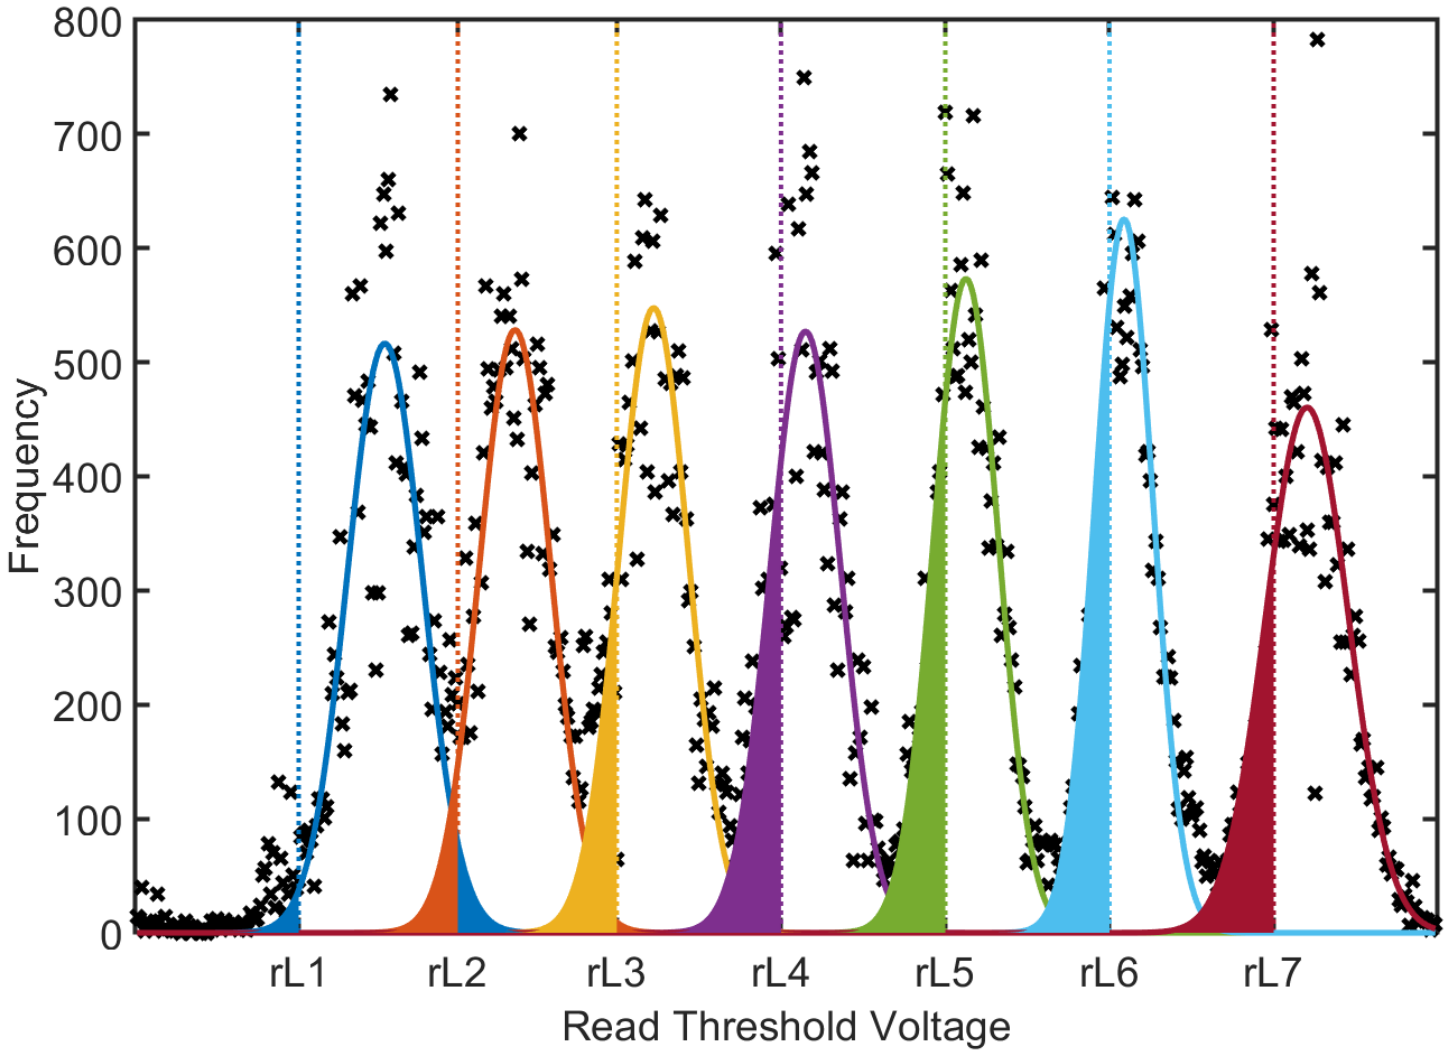

- **How do we estimate level of data integrity error?**
  - A cell contributes to data integrity error if its voltage level no longer lies between the lower and upper read thresholds of the state it was programmed to
  - Amount of data integrity error is estimated by the area under each bell curve sitting outside of its lower and upper read thresholds
- **How do we minimise data integrity error?**
  - Move the read thresholds to make the error area as small as possible

Gaussian Mixture Calibration – Optimal Read Thresholds

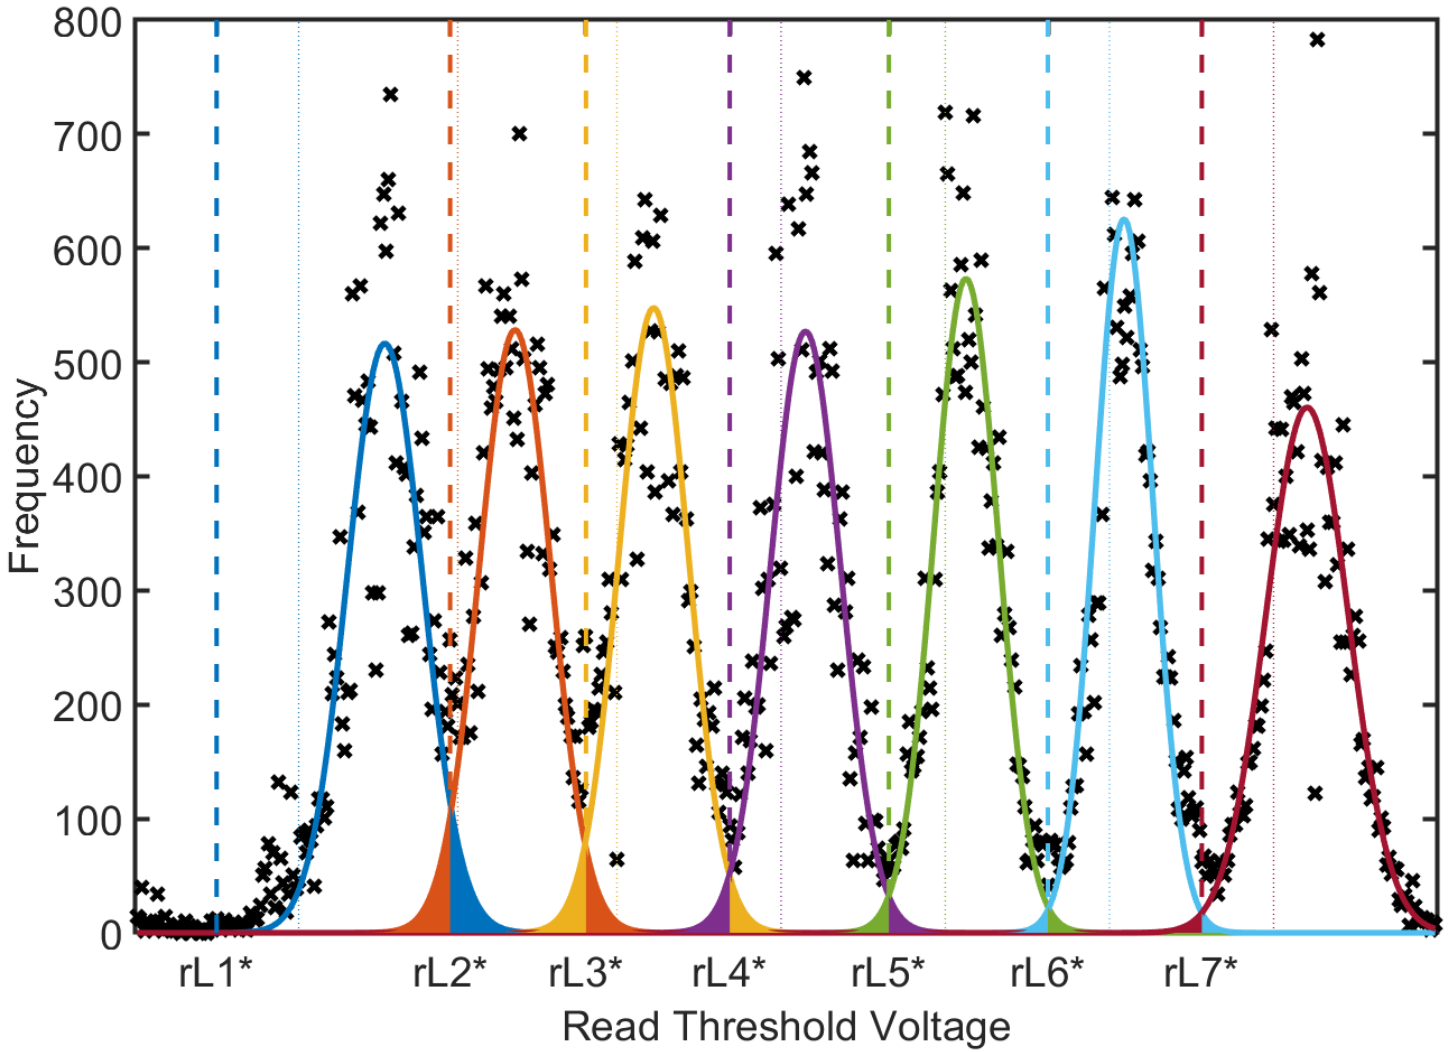

4. Determine optimal read thresholds

- The intersections or *"troughs"* of the bell curves provide an estimation of the optimal read thresholds
- When the read thresholds are positioned at the troughs, the number of cells lying outside of their programmed thresholds is (approximately) minimised

5. Construct optimal image

- We apply the optimal read thresholds to the high-resolution image
- We obtain an image containing (approximately) the lowest amount of data integrity error possible
- The controller's ECC and data processing schemes have a higher chance of correcting the data from the optimal image

Gaussian Mixture Calibration – Advantage over Gradient Methods

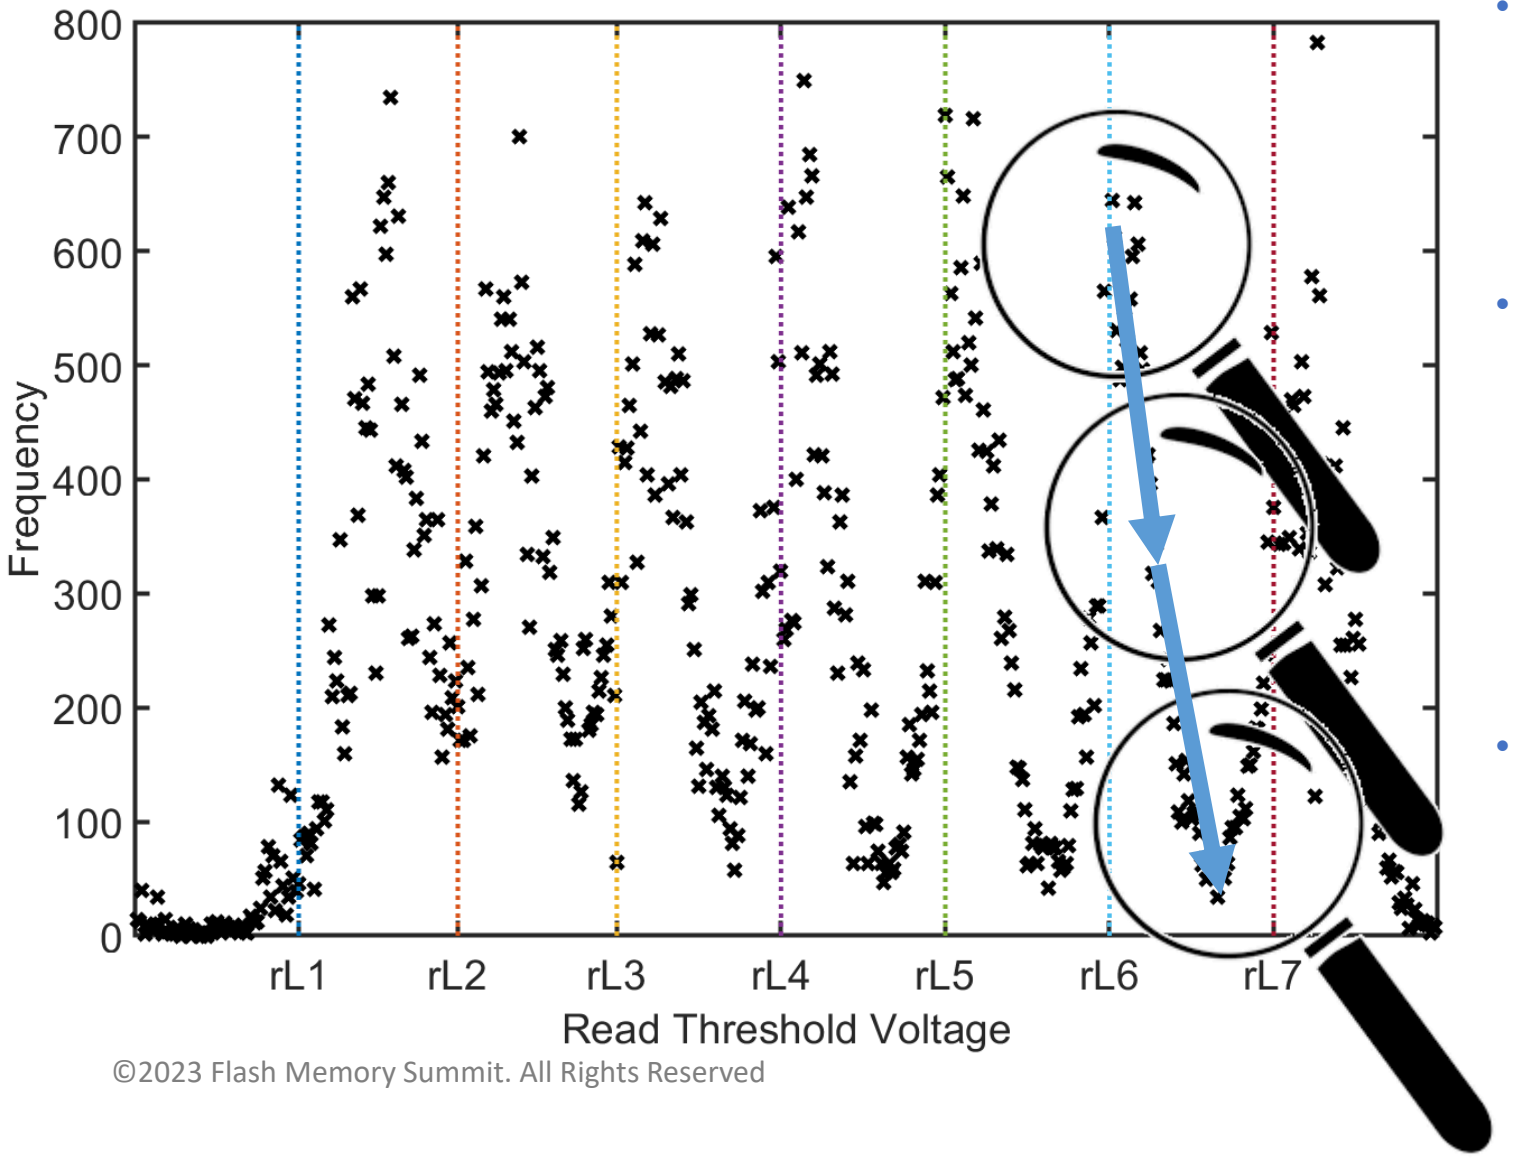

- **Optimization using data fitting (such as our Gaussian Mixture Calibration) is less vulnerable to finding the incorrect trough**
  - Uses the *"global"* voltage distribution to determine all optimal read thresholds simultaneously
- **Gradient methods (such as built-in auto read calibration) examine the *"local"* voltage distribution**
  - Build the voltage distribution around some initial threshold (default)
  - Proceed *"downhill"*, constructing the voltage distribution *"as you go"*
  - Determine when the *"valley"* (trough) is reached
- **In high data integrity error regimes, gradient methods are vulnerable to catastrophic data integrity error**
  - Incorrect positioning of the initial threshold may cause the wrong trough to be determined (as shown)

Gaussian Mixture Calibration – Advantage over Gradient Methods

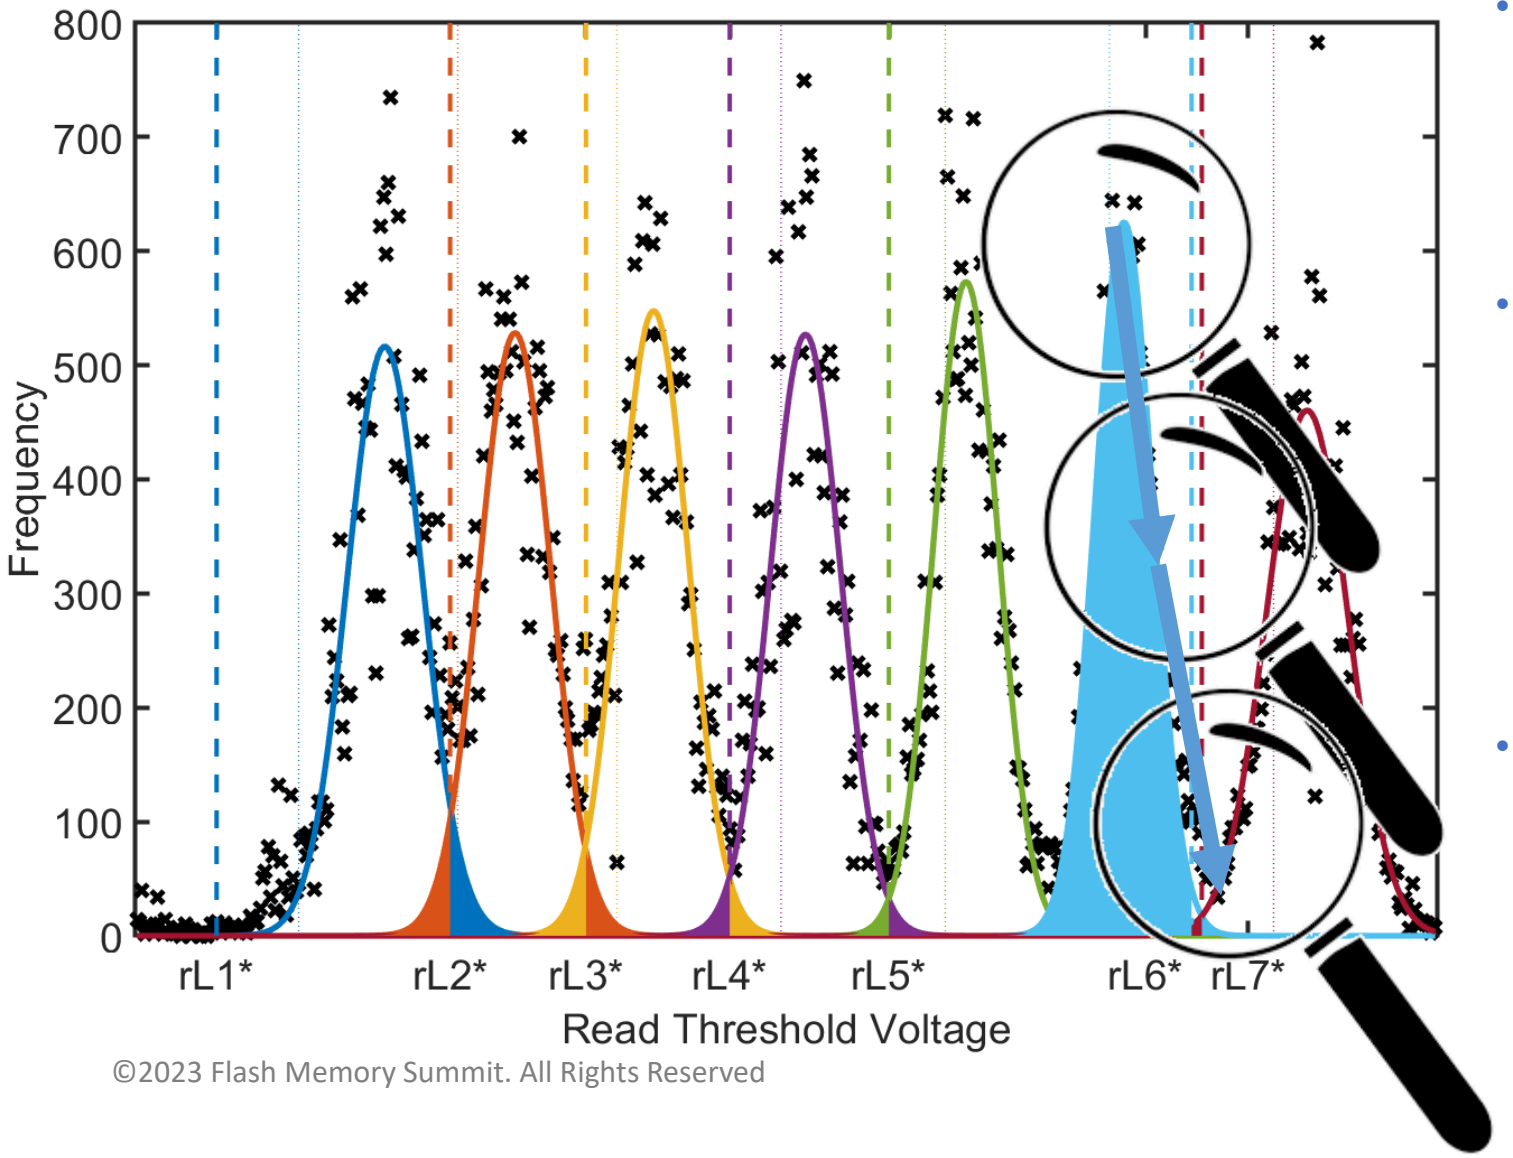

- **Optimization using data fitting (such as our Gaussian Mixture Calibration) is less vulnerable to finding the incorrect trough**
  - Uses the *"global"* voltage distribution to determine all optimal read thresholds simultaneously
- **Gradient methods (such as built-in auto read calibration) examine the *"local"* voltage distribution**
  - Build the voltage distribution around some initial threshold (default)
  - Proceed *"downhill"*, constructing the voltage distribution *"as you go"*
  - Determine when the *"valley"* (trough) is reached
- **In high data integrity error regimes, gradient methods are vulnerable to catastrophic data integrity error**
  - Incorrect positioning of the initial threshold may cause the wrong trough to be determined (as shown)

Offline Read Calibration – Data Recovery

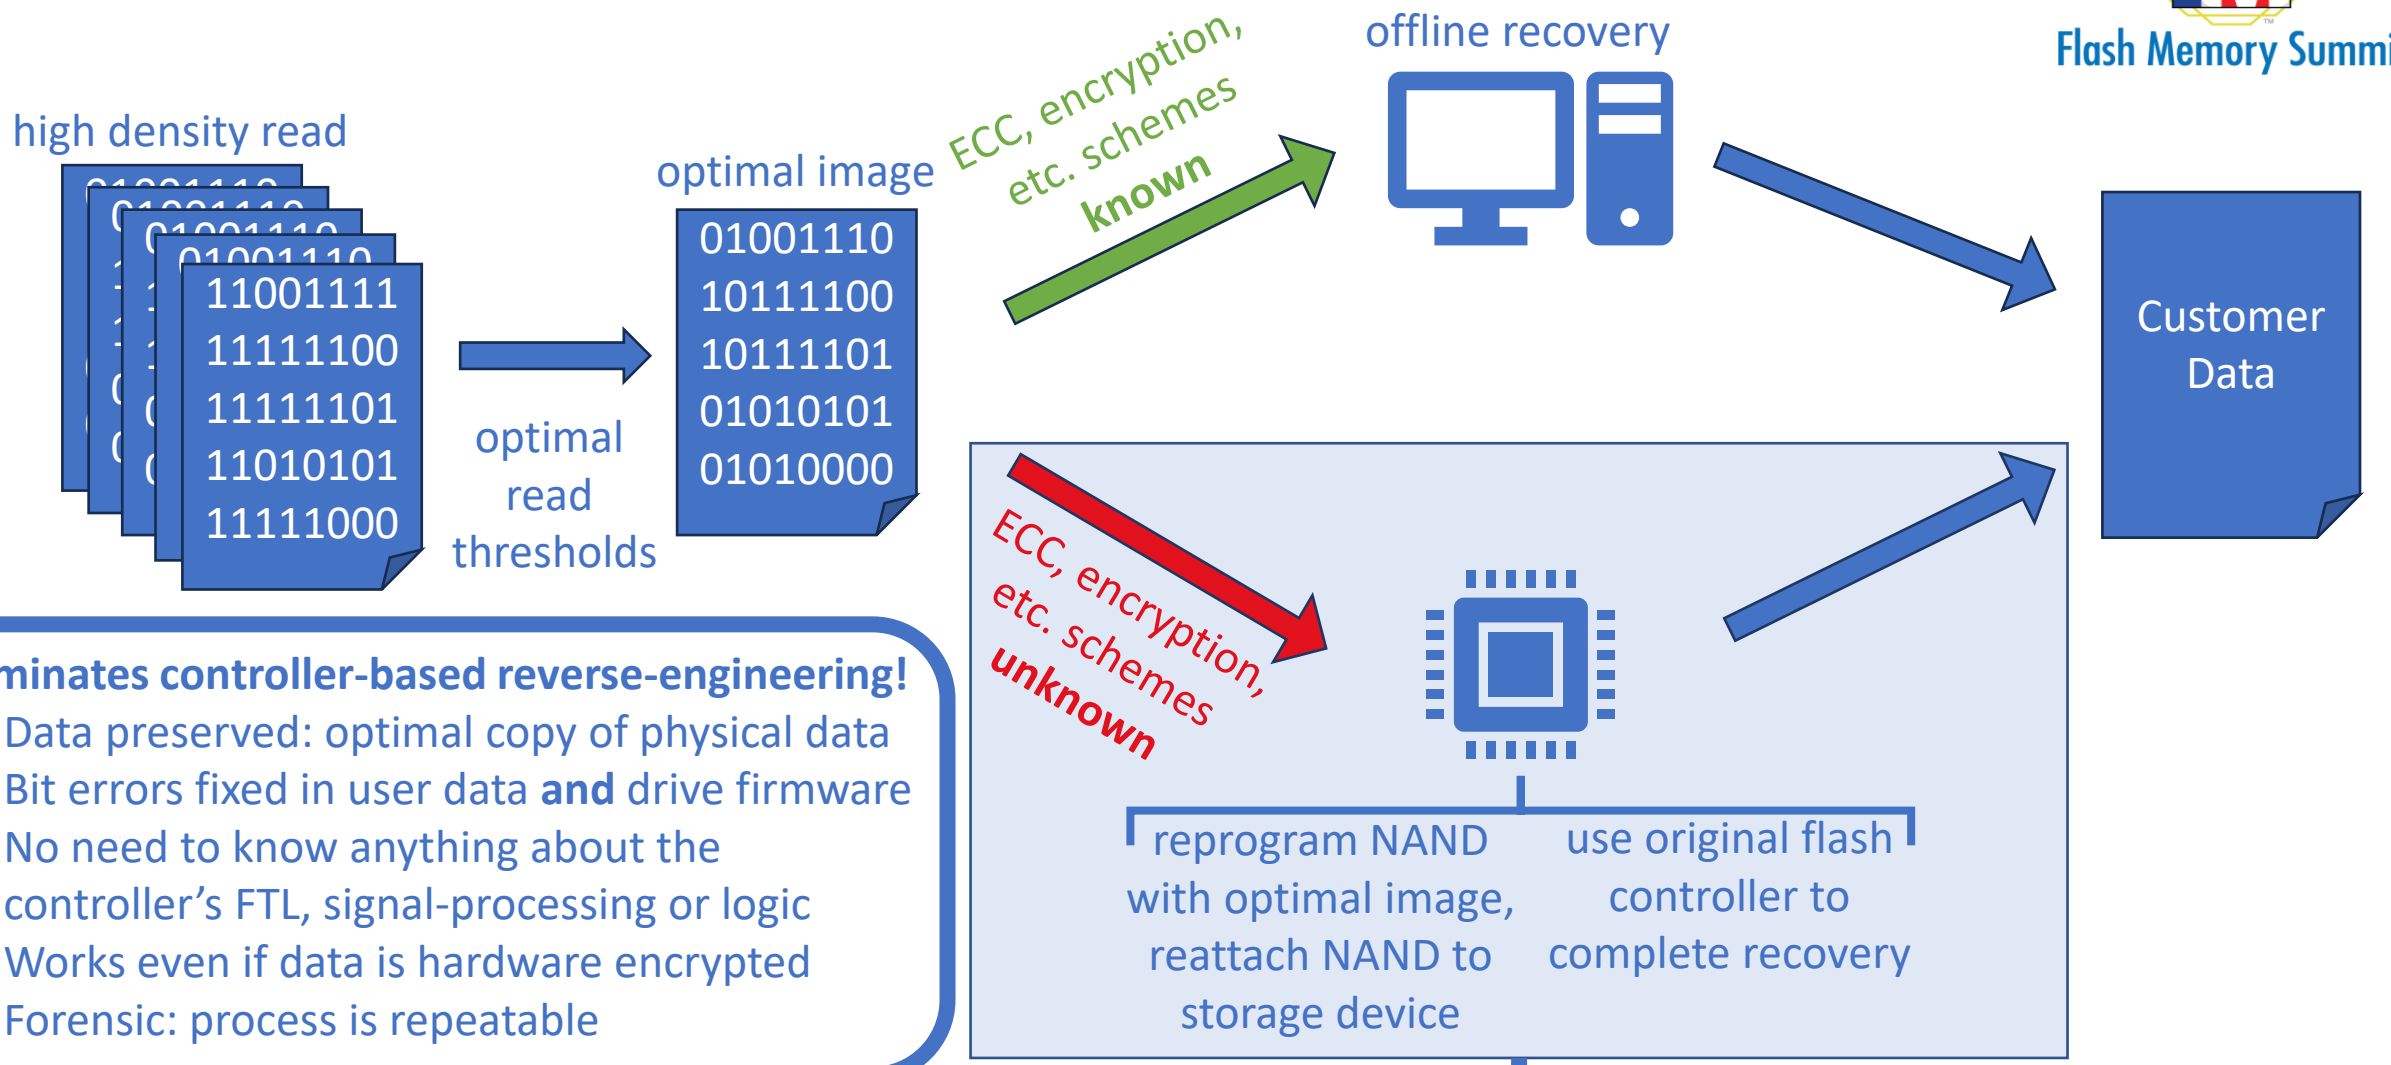

Eliminates controller-based reverse-engineering!

- ✓ Data preserved: optimal copy of physical data
- ✓ Bit errors fixed in user data **and** drive firmware
- ✓ No need to know anything about the controller's FTL, signal-processing or logic
- ✓ Works even if data is hardware encrypted
- ✓ Forensic: process is repeatable

Ontrack® by KLDiscovery SIGLEAD

Findings so far – TLC device

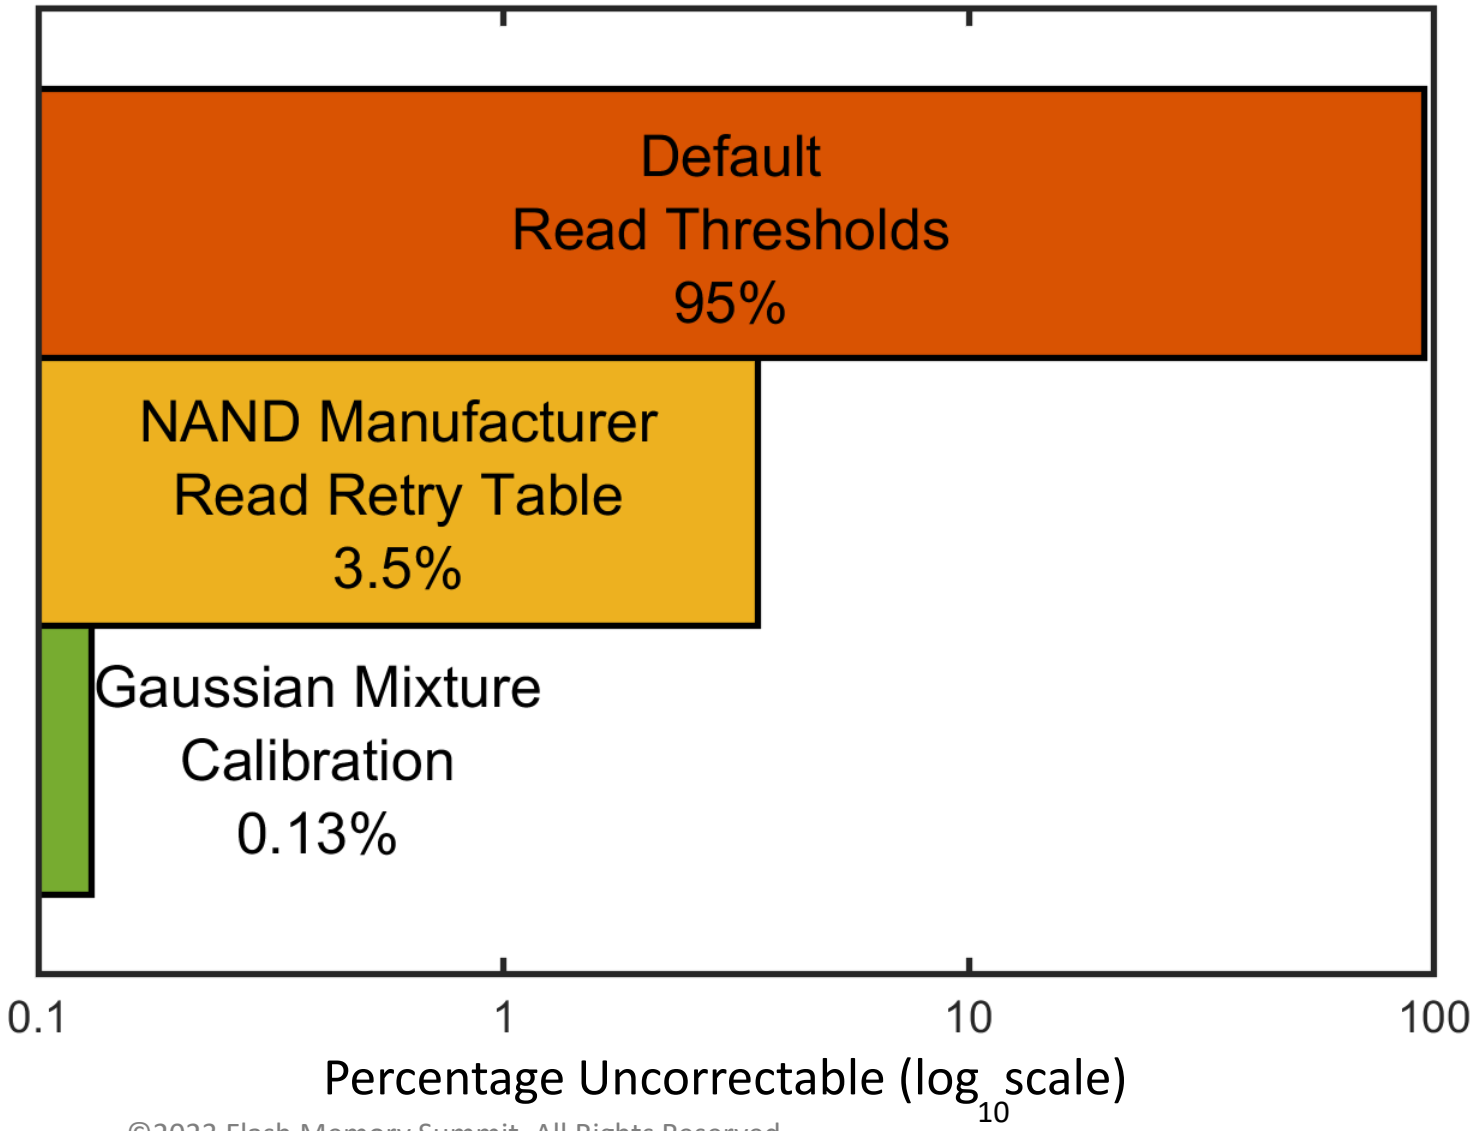

- **SSD device containing 3D TLC NAND**
  - Several logical sectors returning I/O error upon read request
  - Controller's ECC and data processing scheme known
  - NAND manufacturer provides recommended read retry table
- **Attempt "chip-off" and offline recovery**
- **Offline correction of images using the read retry table**
  - Reduced the amount of uncorrectable data from 95% to 3.5% compared with the offline correction of the image using the default read thresholds
- **Offline correction of the optimal image using Gaussian mixture calibration**
  - Further reduced the amount of uncorrectable data to 0.13%

Findings so far – QLC device

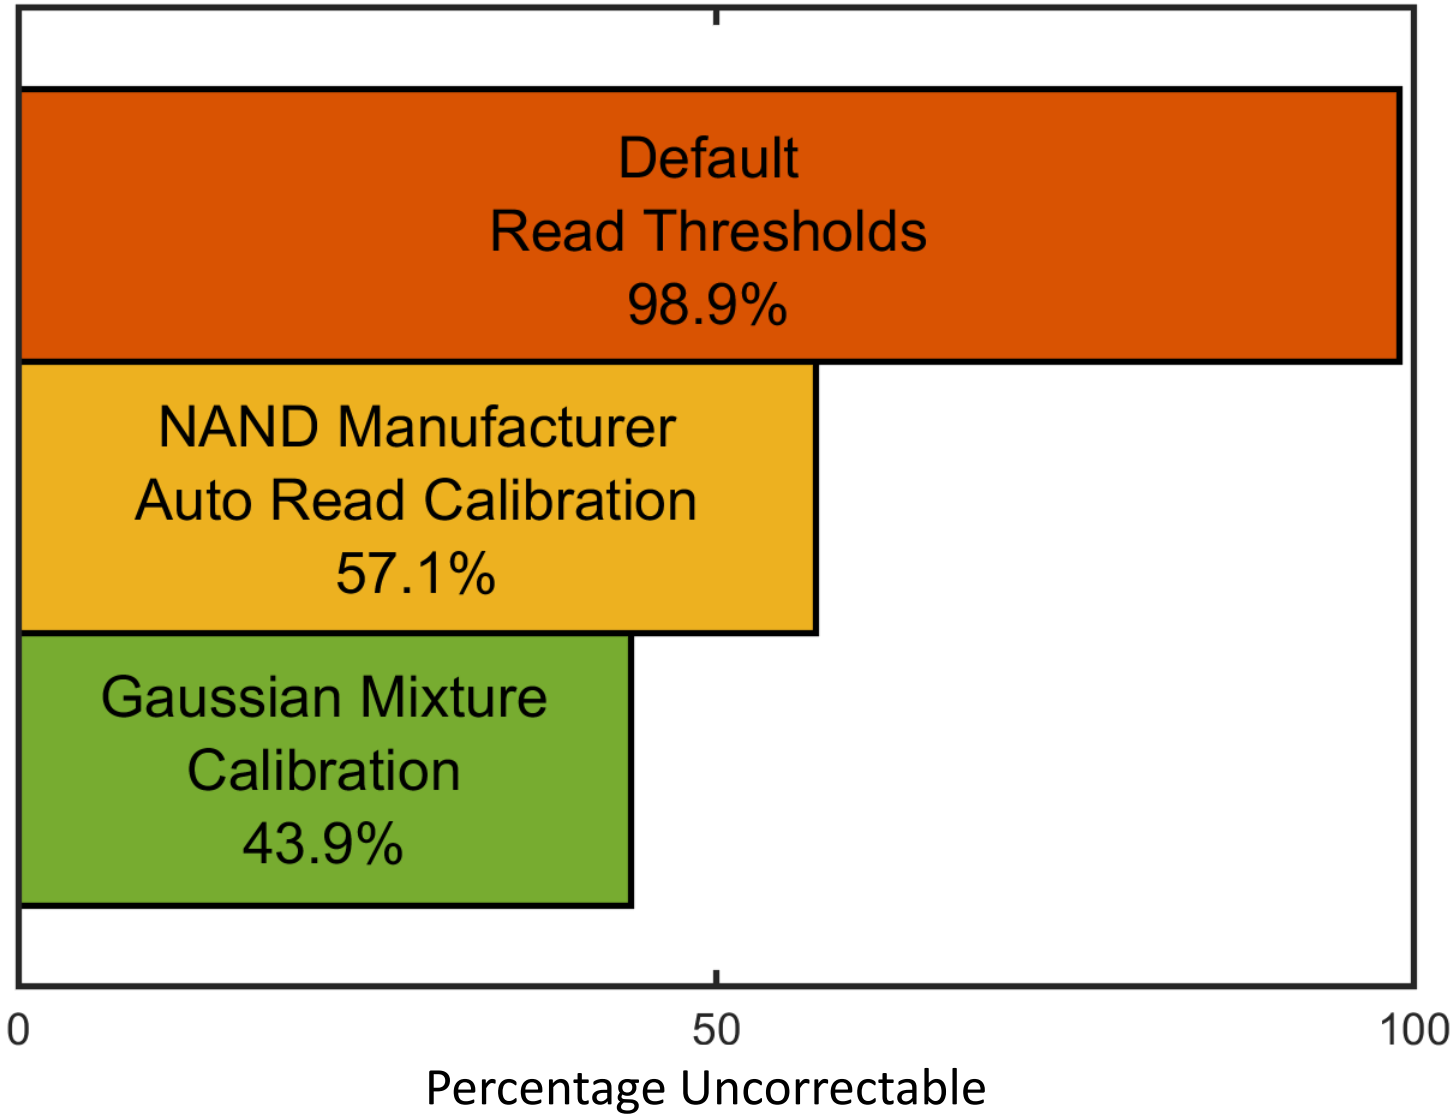

- **SSD device containing 3D QLC NAND**
  - Several logical sectors returning I/O error upon read request
  - Controller's ECC and data processing scheme known
  - NAND manufacturer provides auto read calibration feature
- **Attempt "chip-off" and offline recovery**
- **Offline correction of image using auto read calibration**
  - Reduced the amount of uncorrectable data from 99% to 57% compared with the offline correction of the image using the default read thresholds
- **Offline correction of the optimal image using Gaussian mixture calibration**
  - Further reduced the amount of uncorrectable data to 44%
- **Note that in this case, a significant amount of the uncorrectable data remained uncorrected**
  - The Gaussian mixture calibration solved the problem of incorrect trough determination, but could only correct data on the lower pages

Conclusions

Flash Memory Summit

Advantages of Gaussian mixture calibration for chip-off recovery

- **Gaussian mixture calibration gives a greater reduction in uncorrectable data compared to read retry and auto read calibration**
  - Determines more precise optimal read thresholds than read retry tables
  - Less vulnerable to incorrect trough determination compared to gradient methods such as auto read calibration
- **Gaussian mixture calibration is verifiable and adaptable**
  - Produces a visual check of the voltage distribution, the data fit and optimal thresholds
  - Optimal image can be adapted if a different data fitting model is available
  - Data fitting model can be customized to suit or different makes/models of NAND

Disadvantages of Gaussian Mixture calibration for chip-off recovery

- **Multiple reads of NAND required at incremental read thresholds**
  - Slow
  - Requires read threshold control
  - Introduces read disturb
- **Data fitting requires significant post-processing computation**

Ontrack® by KLDiscovery SIGLEAD

Next Steps

- Further experiments: introduce uncorrectable errors to sample devices with known data and where controller signal-processing and logic algorithms are known
- Optimise algorithm and processing speed for data fitting
- Examine behaviour of technique with different NAND vendor devices and measure effectiveness
- Further develop hardware (socketed NAND, controller, buffered I/O) for experiments with retuned NAND data on different controllers and SSD models
- Investigate possibility of using entire replacement NAND for tuned data versus rewriting just those blocks that have uncorrectable errors in original NAND
- Apply technique to SSDs where signal-processing and logic algorithms are unknown using original device controller to perform recovery

Ontrack® by KLDiscovery SIGLEAD

Thank You!

Please visit us at the Ontrack booth

#651

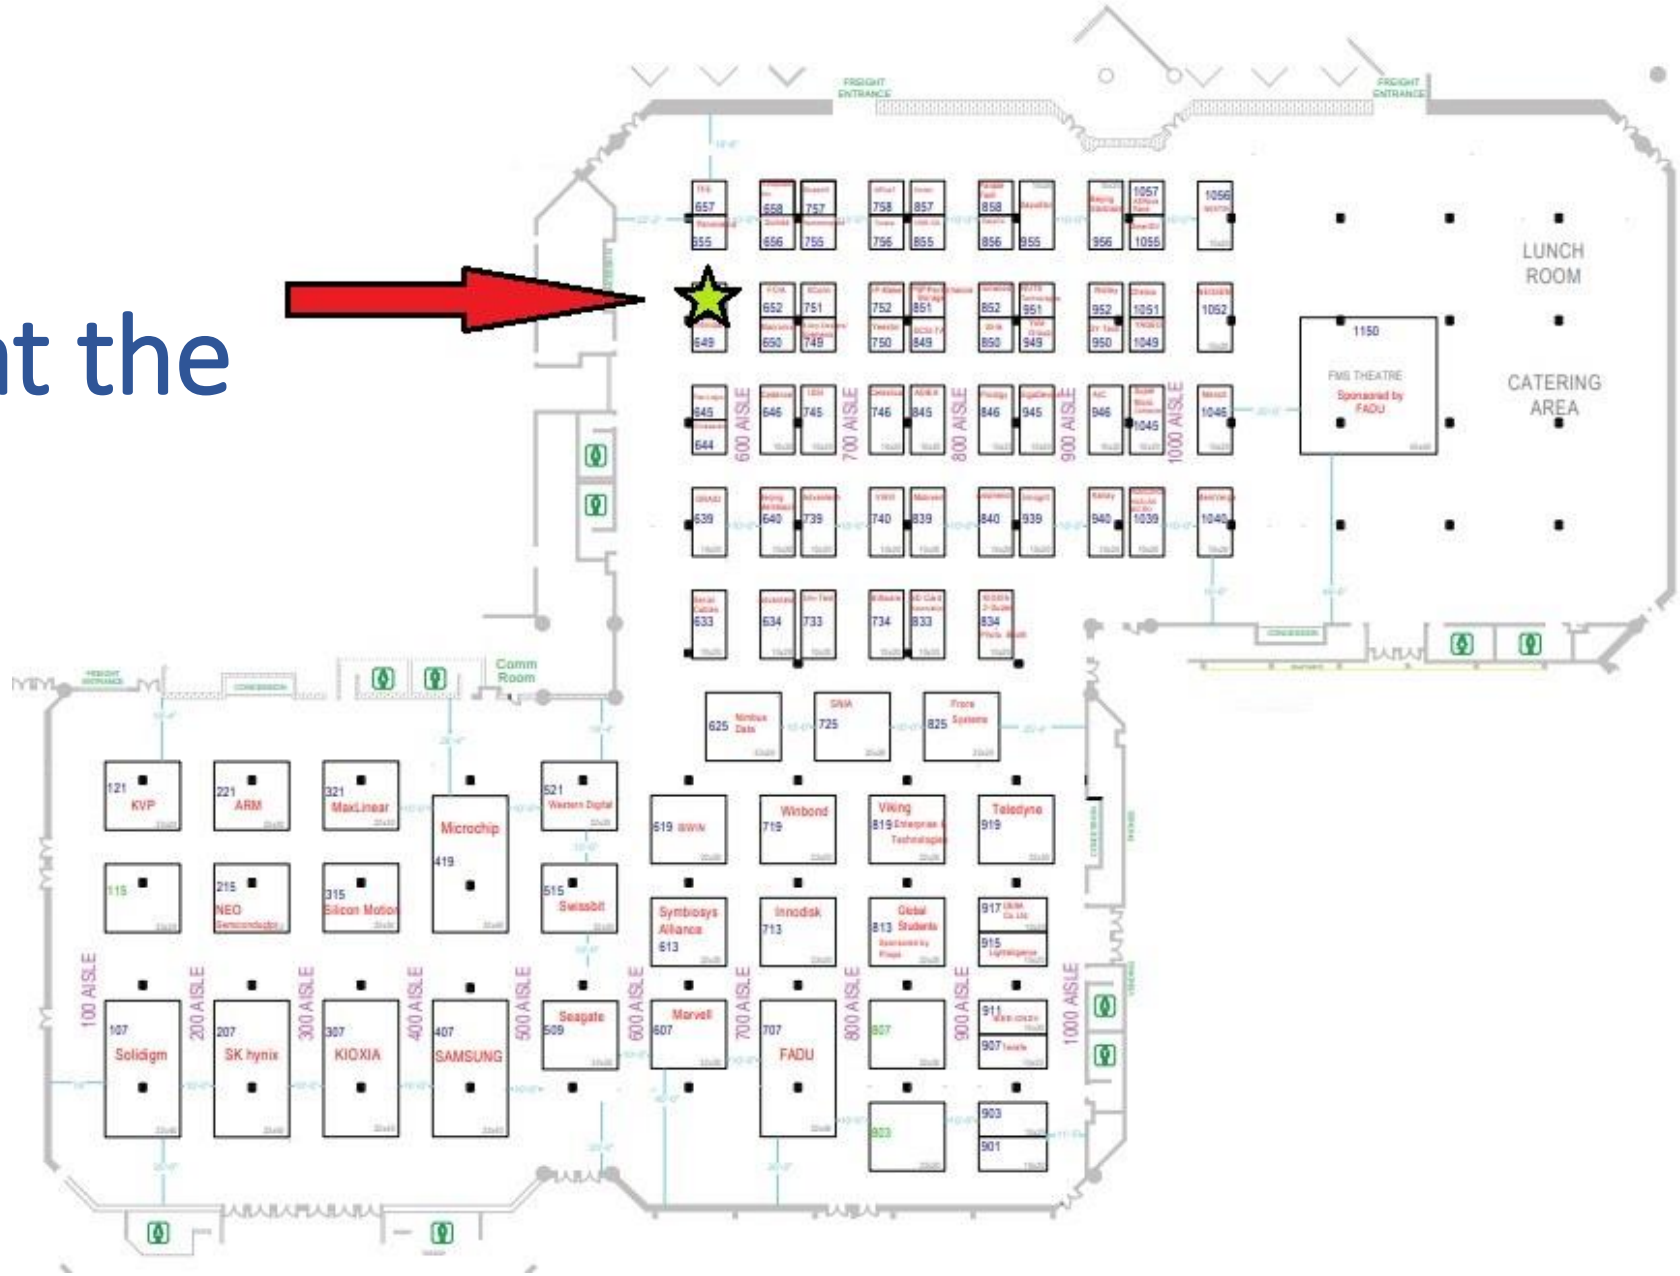

Ontrack® by KLDiscovery SIGLEAD