Detecting Sleigh Fatigue

Monitoring Ageing in Embedded Systems for Predictive Maintenance at the North Pole

Leandro Lanzieri
leandro.lanzieri@haw-hamburg.de

Extra

Extra: Onboarding Material

Incident Reports

Incident Reports

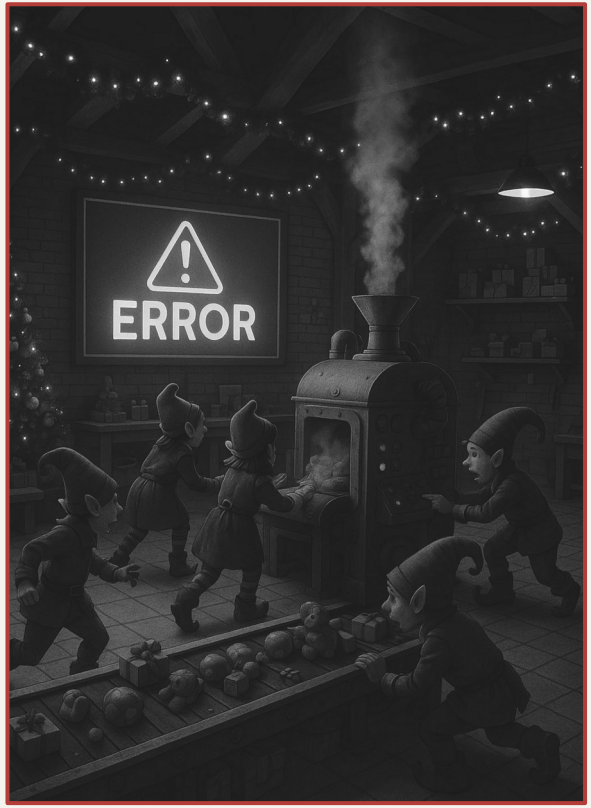

Incident Reports

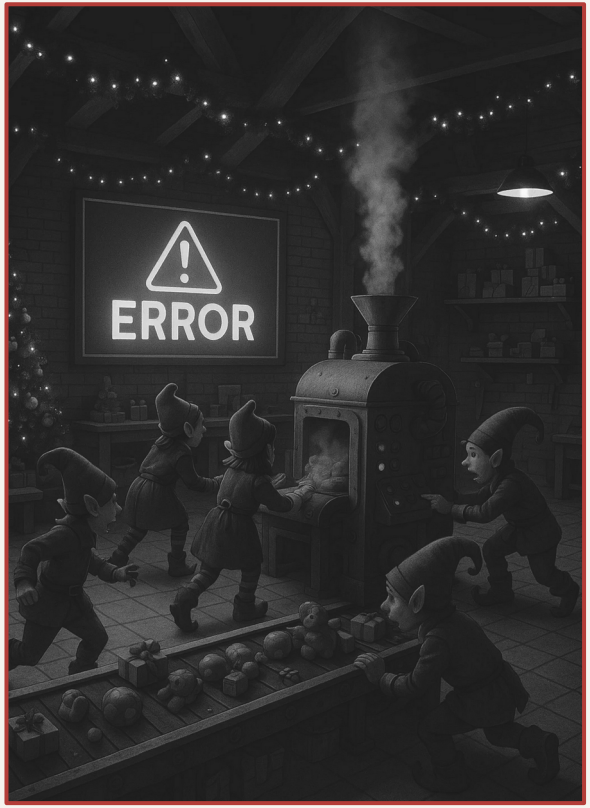

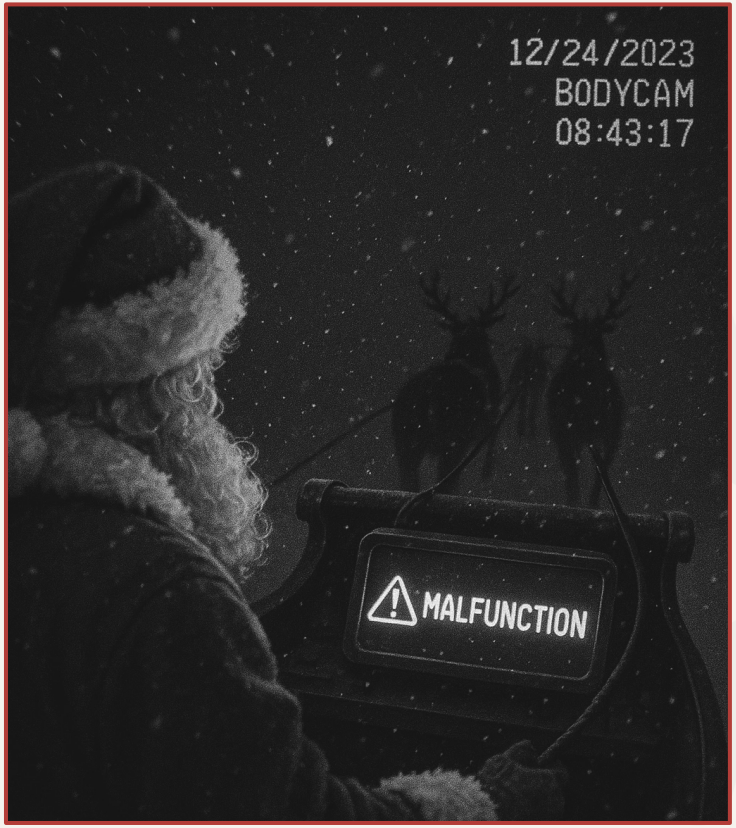

Why Monitor Hardware?

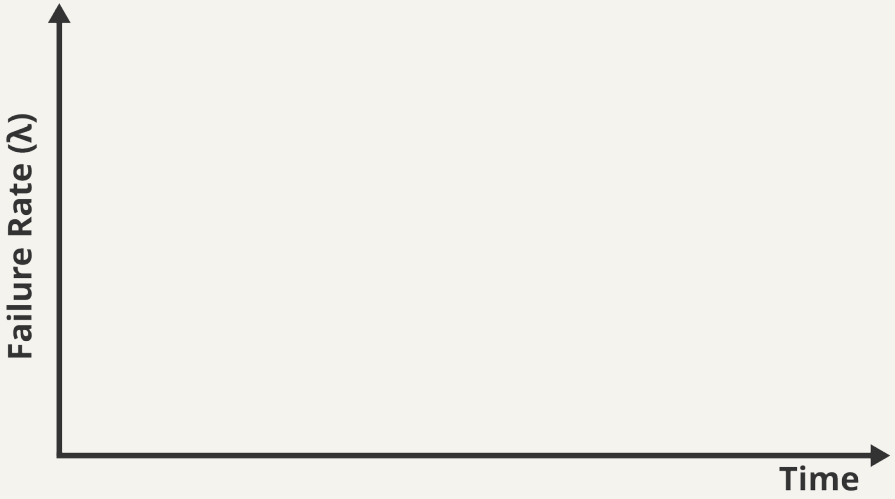

Why Monitor Hardware?

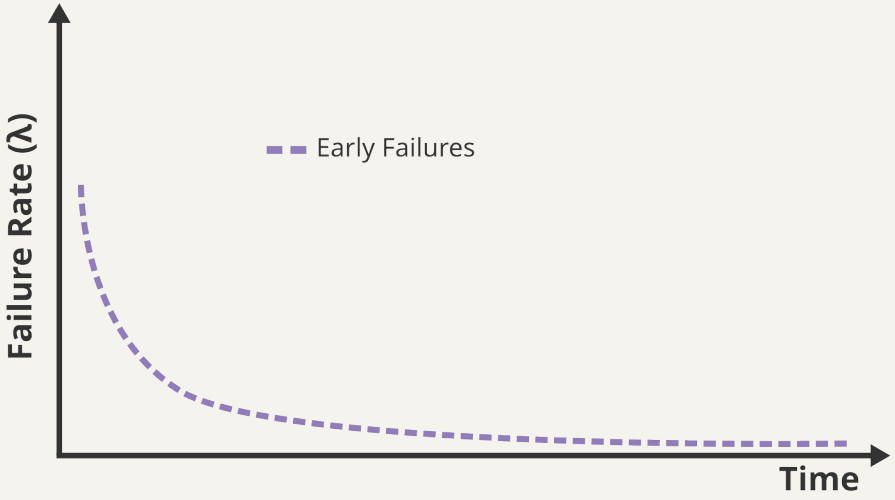

Why Monitor Hardware?

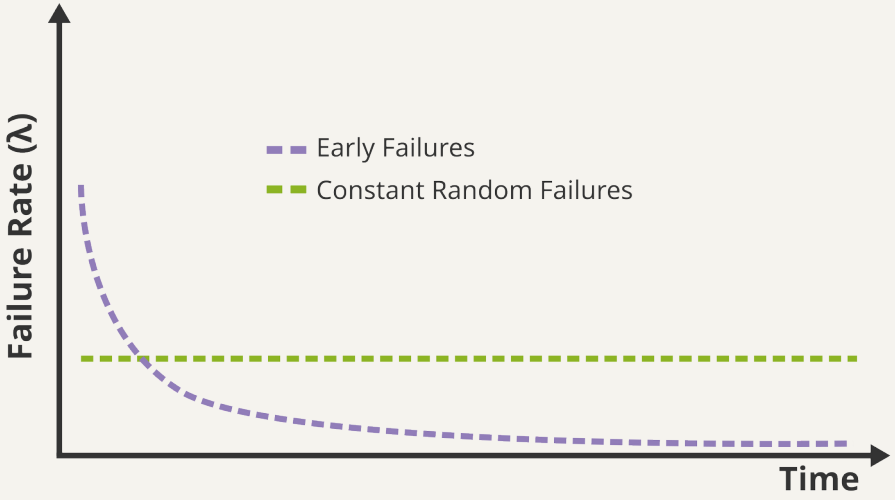

Why Monitor Hardware?

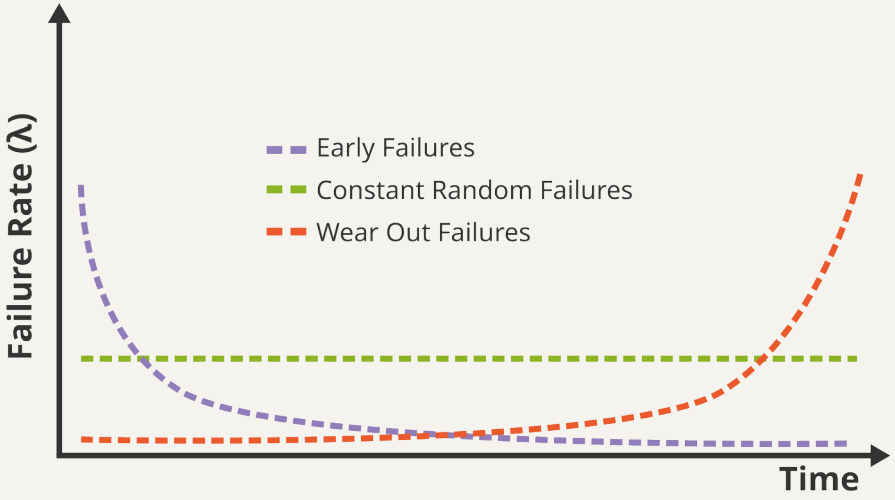

Why Monitor Hardware?

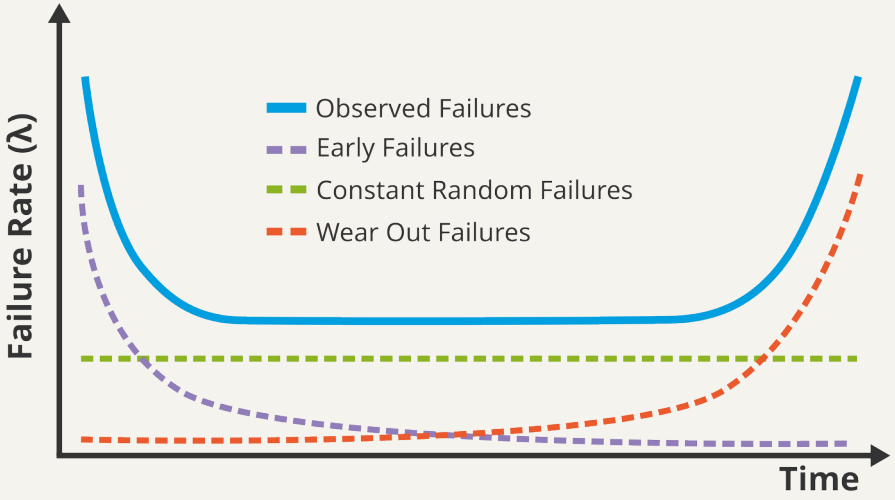

Why Monitor Hardware?

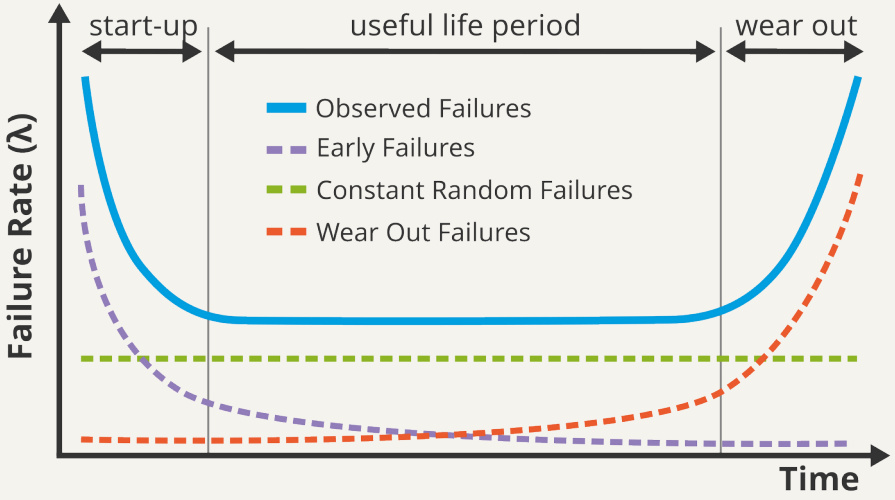

Why Monitor Hardware?

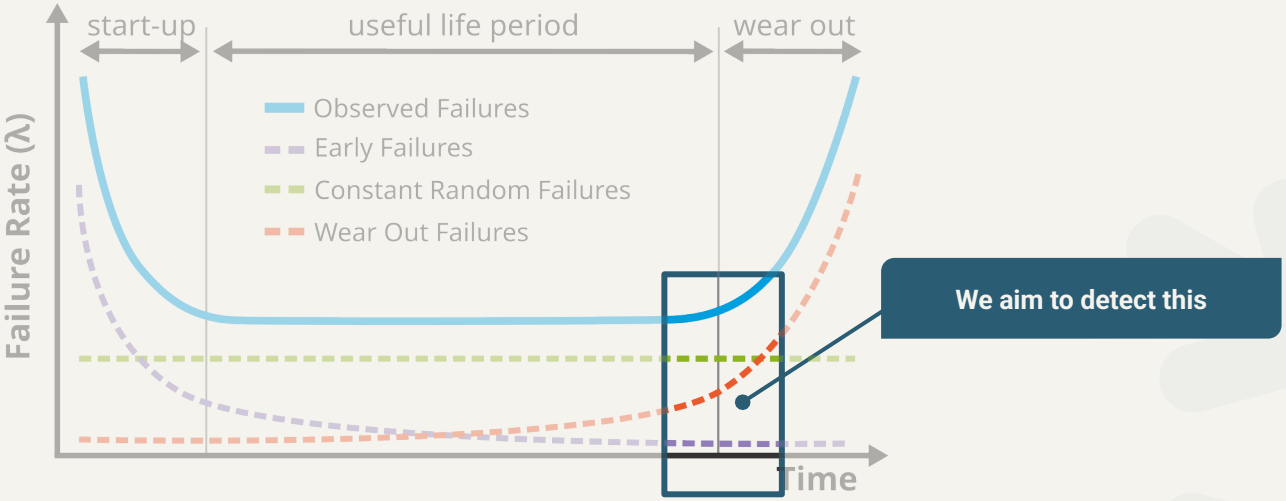

Why Monitor Hardware?

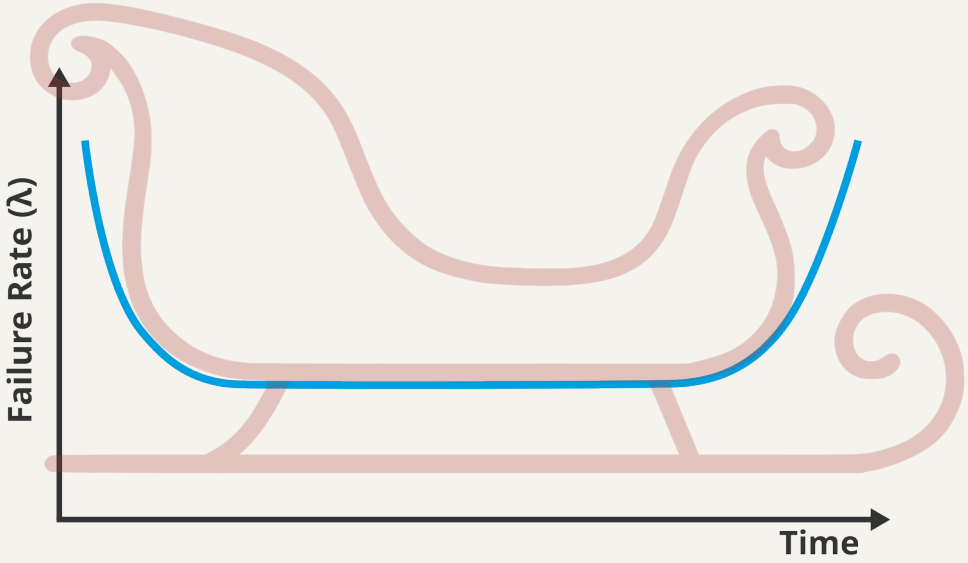

Degradation of Silicon Devices

[1]: L. Lanzieri, G. Martino, G. Fey, H. Schlarb, T. C. Schmidt, and M. Wählisch. 2024. “A Review of Techniques for Ageing Detection and Monitoring on Embedded Systems”. ACM Computing Surveys 57, 1, Article 24 (January 2025), 34 pages. DOI:10.1145/3695247.

Degradation of Silicon Devices

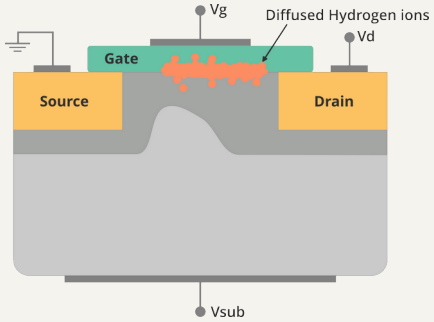

Bias Temperature Instability (BTI)

[1]: L. Lanzieri, G. Martino, G. Fey, H. Schlarb, T. C. Schmidt, and M. Wählisch. 2024. “A Review of Techniques for Ageing Detection and Monitoring on Embedded Systems”. ACM Computing Surveys 57, 1, Article 24 (January 2025), 34 pages. DOI:10.1145/3695247.

Degradation of Silicon Devices

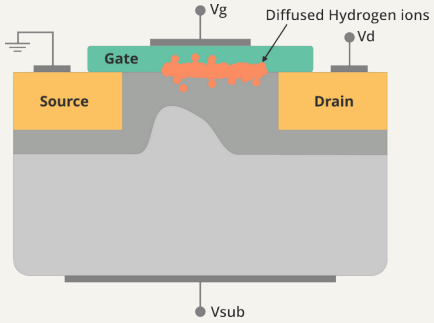

Bias Temperature Instability (BTI)

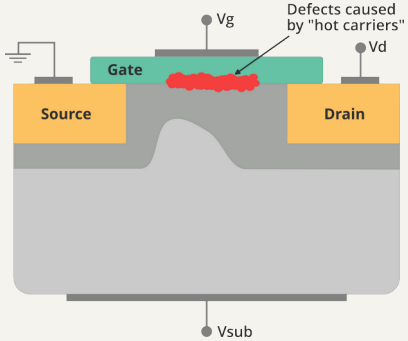

Hot Carrier Injection (HCI)

[1]: L. Lanzieri, G. Martino, G. Fey, H. Schlarb, T. C. Schmidt, and M. Wählisch. 2024. “A Review of Techniques for Ageing Detection and Monitoring on Embedded Systems”. ACM Computing Surveys 57, 1, Article 24 (January 2025), 34 pages. DOI:10.1145/3695247.

Degradation of Silicon Devices

BTI and HCI increase the threshold voltage (V_{th}) of transistors.

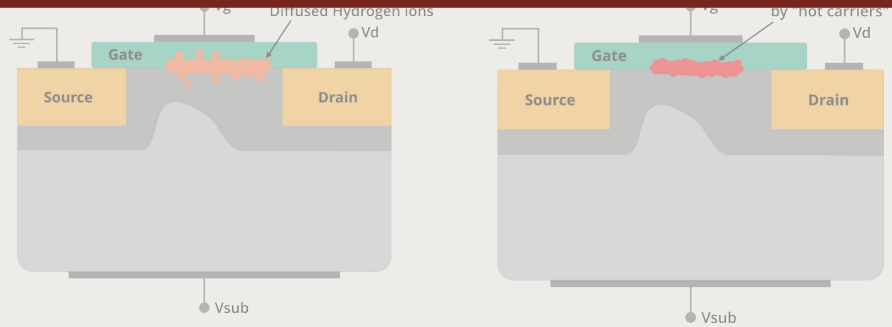

Bias Temperature Instability (BTI) | Hot Carrier Injection (HCI)

[1]: L. Lanzieri, G. Martino, G. Fey, H. Schlarb, T. C. Schmidt, and M. Wählisch. 2024. “A Review of Techniques for Ageing Detection and Monitoring on Embedded Systems”. ACM Computing Surveys 57, 1, Article 24 (January 2025), 34 pages. DOI:10.1145/3695247.

Degradation of Silicon Devices

BTI and HCI increase the threshold voltage (V_{th}) of transistors.

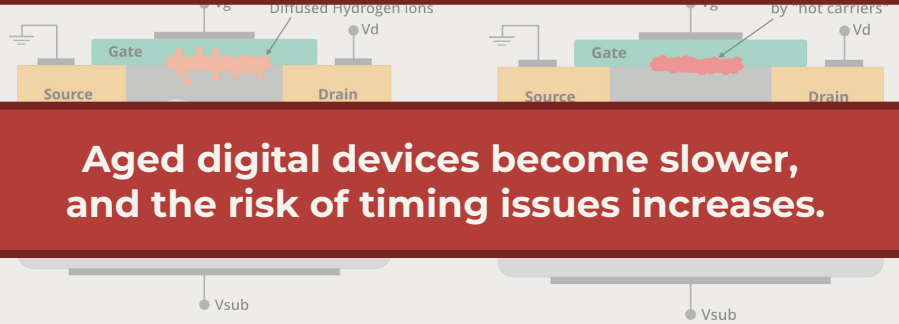

Aged digital devices become slower, and the risk of timing issues increases.

Bias Temperature Instability (BTI) | Hot Carrier Injection (HCI)

[1]: L. Lanzieri, G. Martino, G. Fey, H. Schlarb, T. C. Schmidt, and M. Wählisch. 2024. “A Review of Techniques for Ageing Detection and Monitoring on Embedded Systems”. ACM Computing Surveys 57, 1, Article 24 (January 2025), 34 pages. DOI:10.1145/3695247.

Degradation of Silicon Devices

BTI and HCI increase the threshold voltage (V_{th}) of transistors.

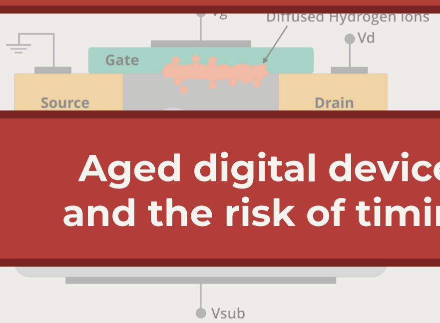

Bias Temperature Instability (BTI)

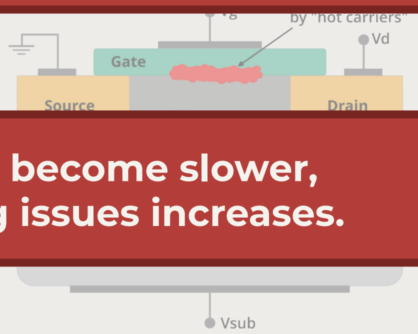

Hot Carrier Injection (HCI)

**Aged digital devices become slower,
and the risk of timing issues increases.**

For more details on ageing mechanisms, check ref. [1].

[1]: L. Lanzieri, ... Systems". ACM Computing Surveys 57, 1, Article 24 (January 2025), 34 pages. DOI:10.1145/3695247.

Considered Components

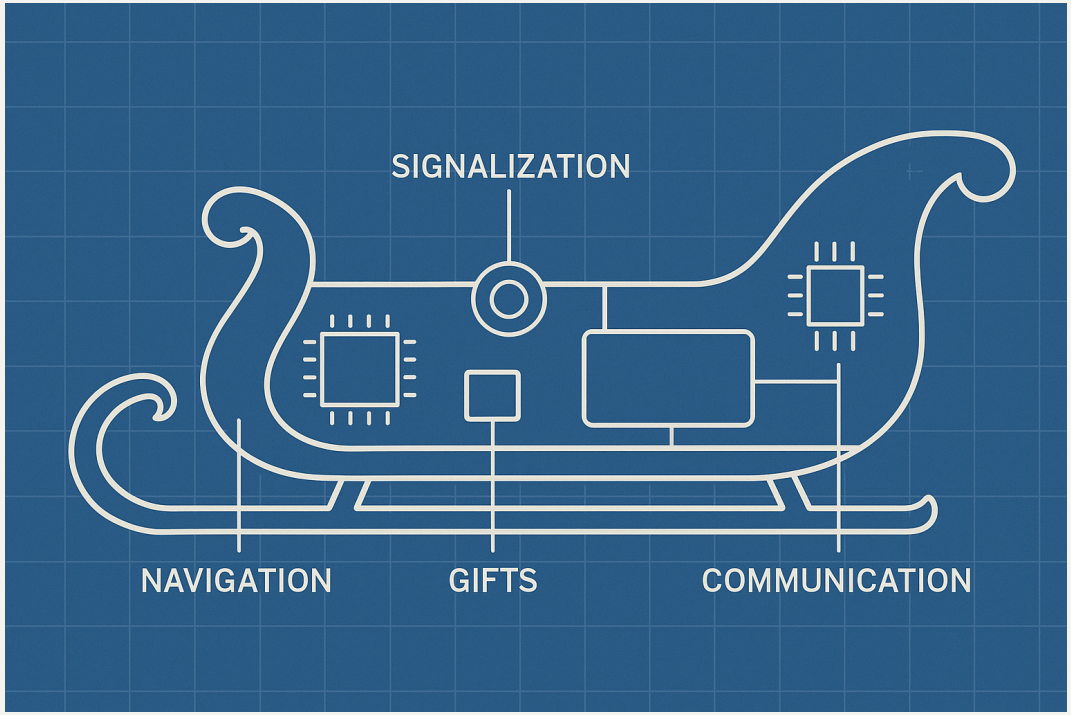

Considered Components

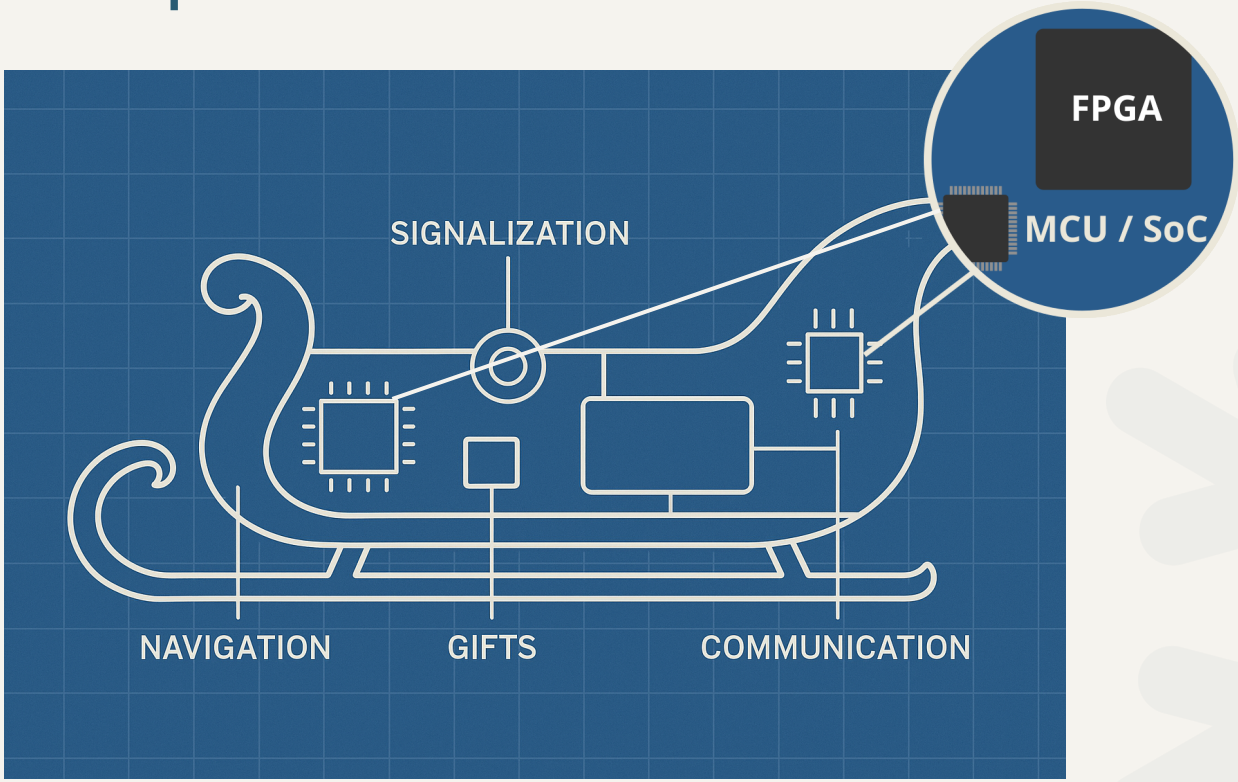

Considered Components

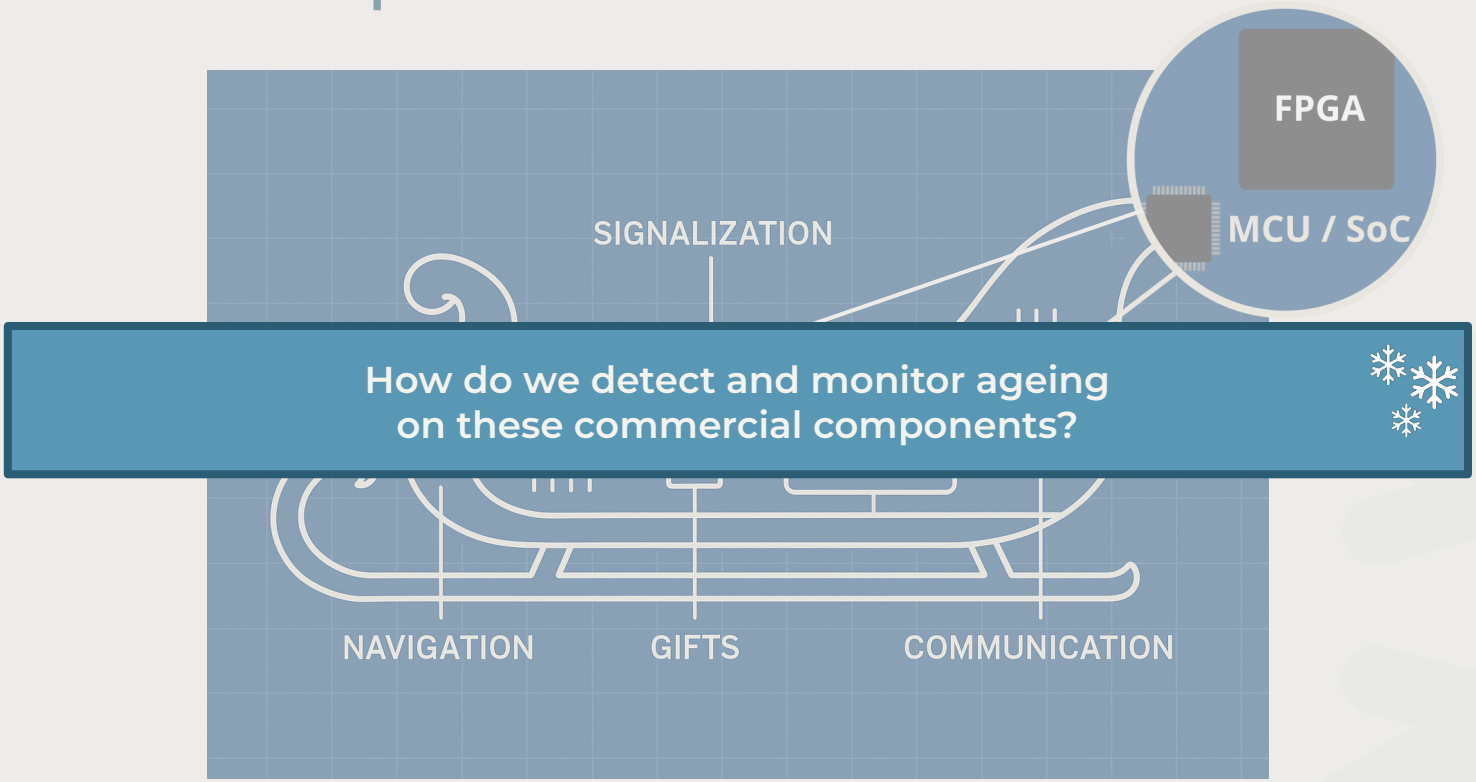

How do we detect and monitor ageing on these commercial components?

Agenda

- Ageing Detection on Embedded SRAMs
- Monitoring Degradation of Propagation Delay on FPGAs
- Ageing Monitoring for Commercial MCUs

Agenda

- **Ageing Detection on Embedded SRAMs**
- Monitoring Degradation of Propagation Delay on FPGAs
- Ageing Monitoring for Commercial MCUs

Static RAM (SRAM) 6T Cell

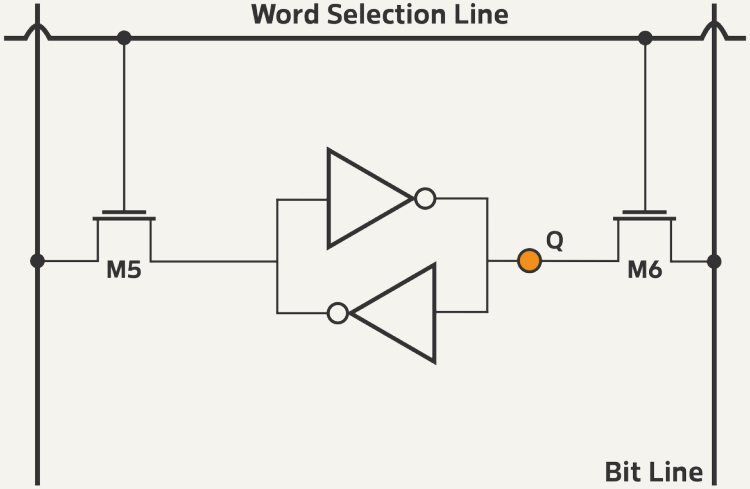

Static RAM (SRAM) 6T Cell

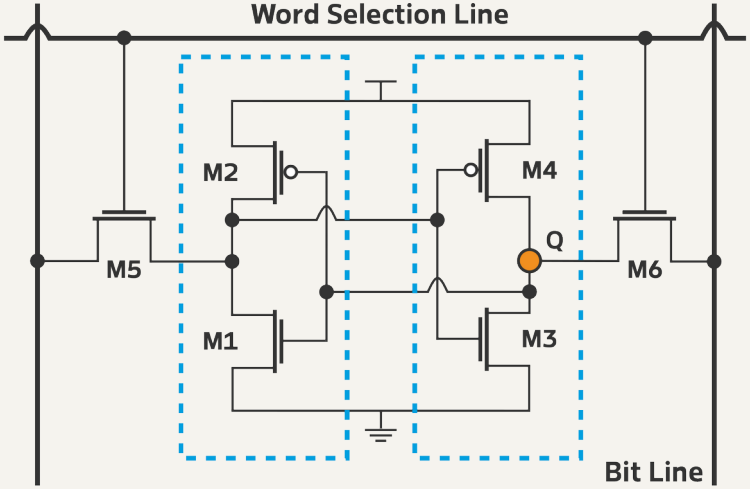

SRAM 6T Cell Startup

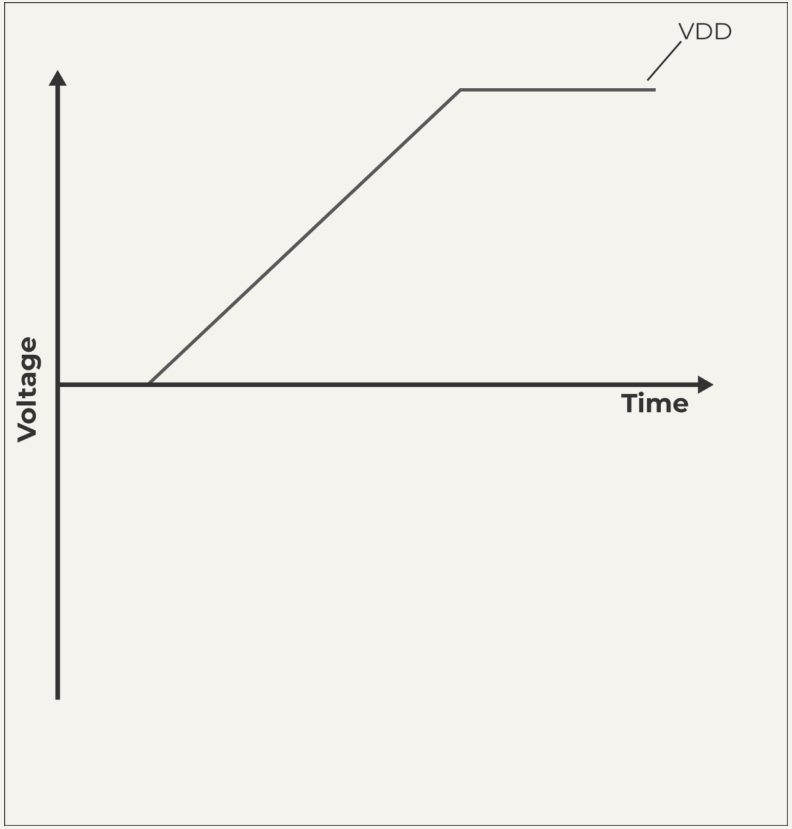

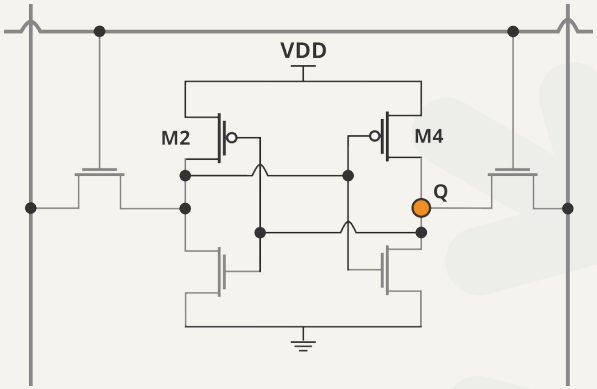

SRAM 6T Cell Startup

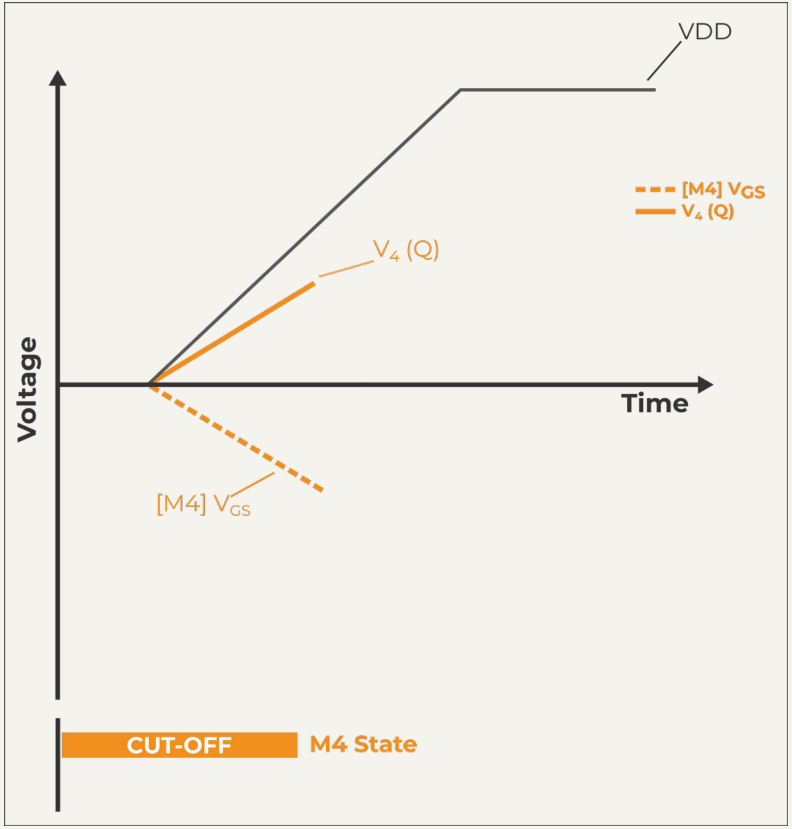

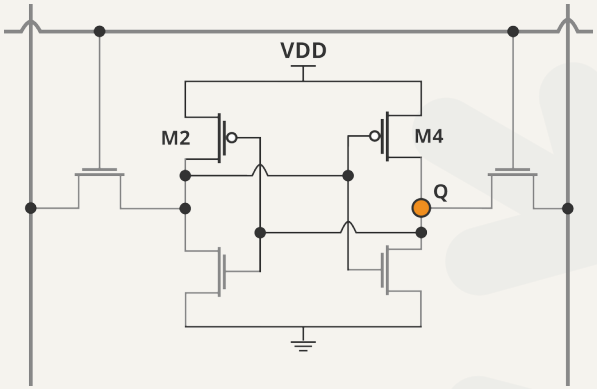

SRAM 6T Cell Startup

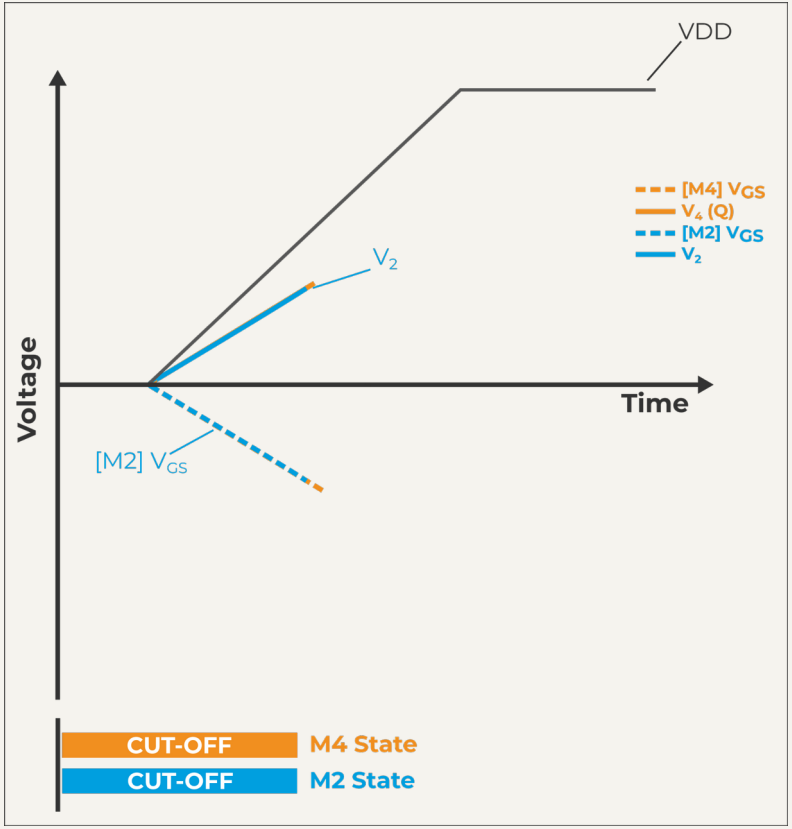

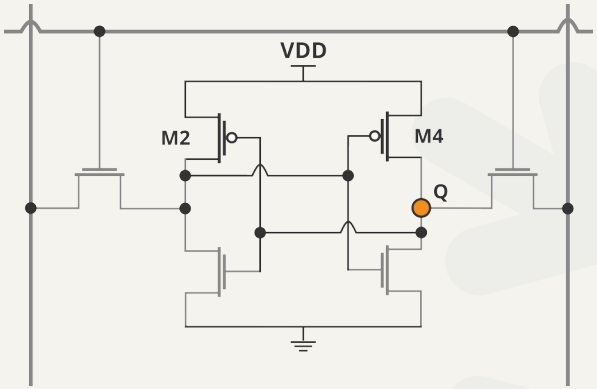

SRAM 6T Cell Startup

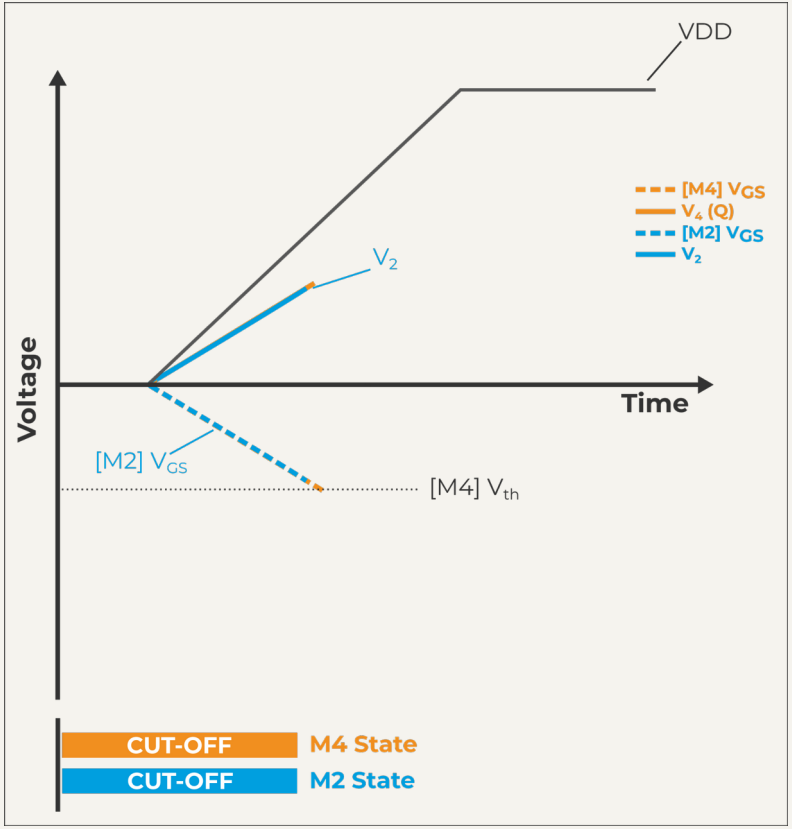

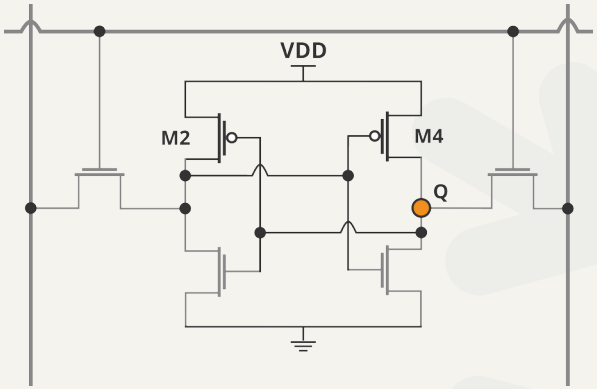

SRAM 6T Cell Startup

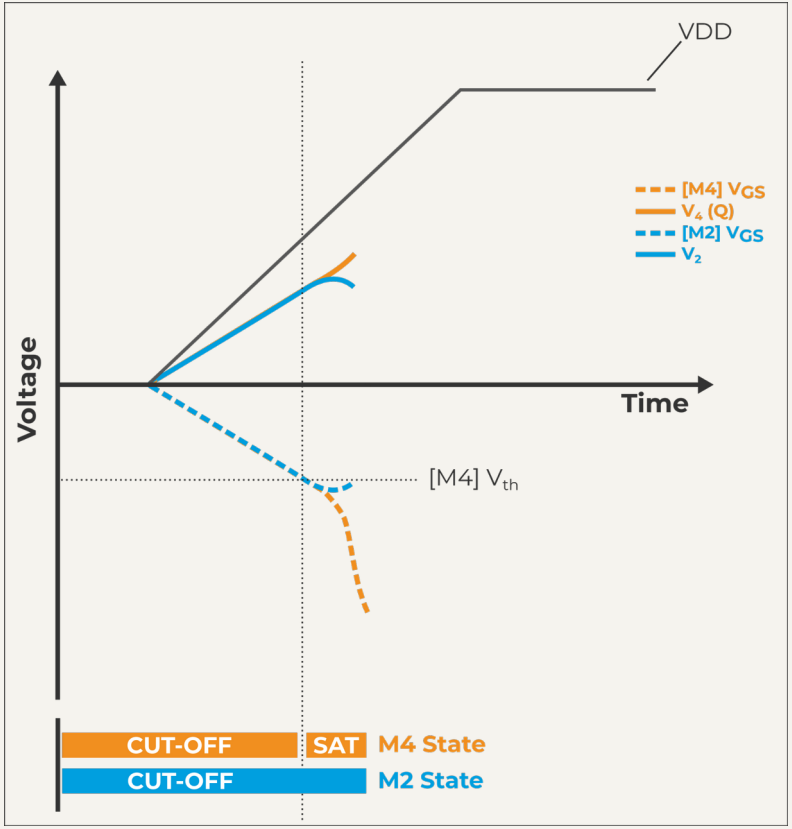

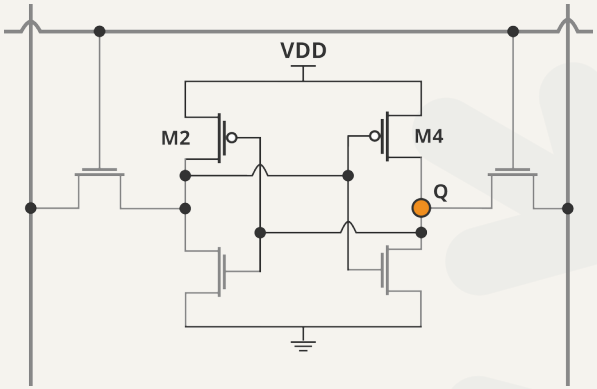

SRAM 6T Cell Startup

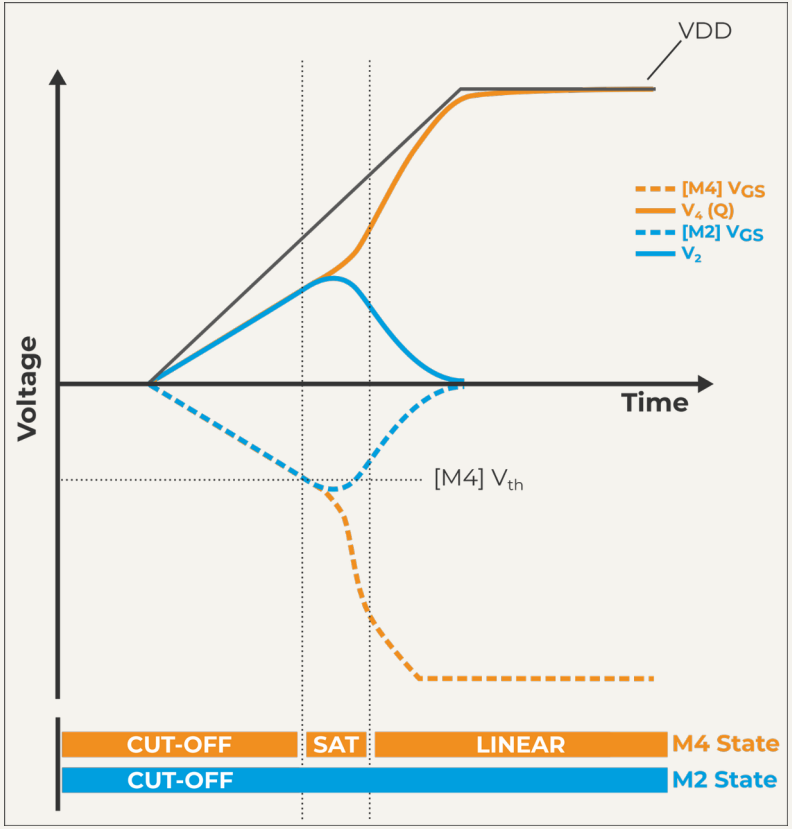

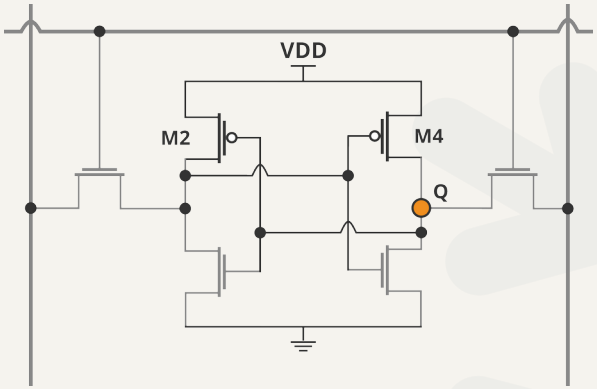

SRAM 6T Cell Startup

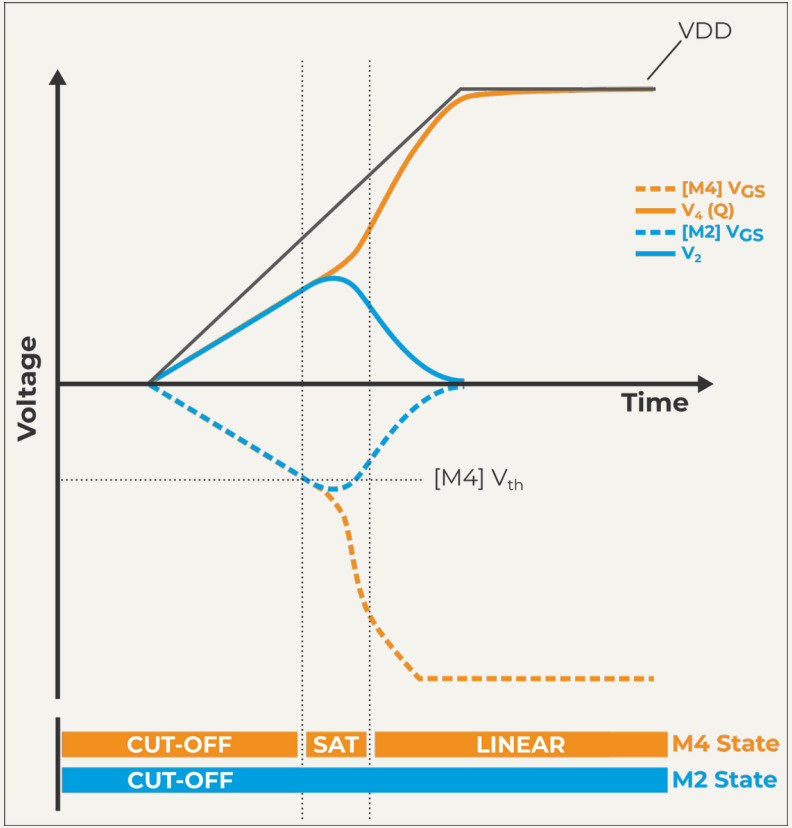

SRAM 6T Cell Startup

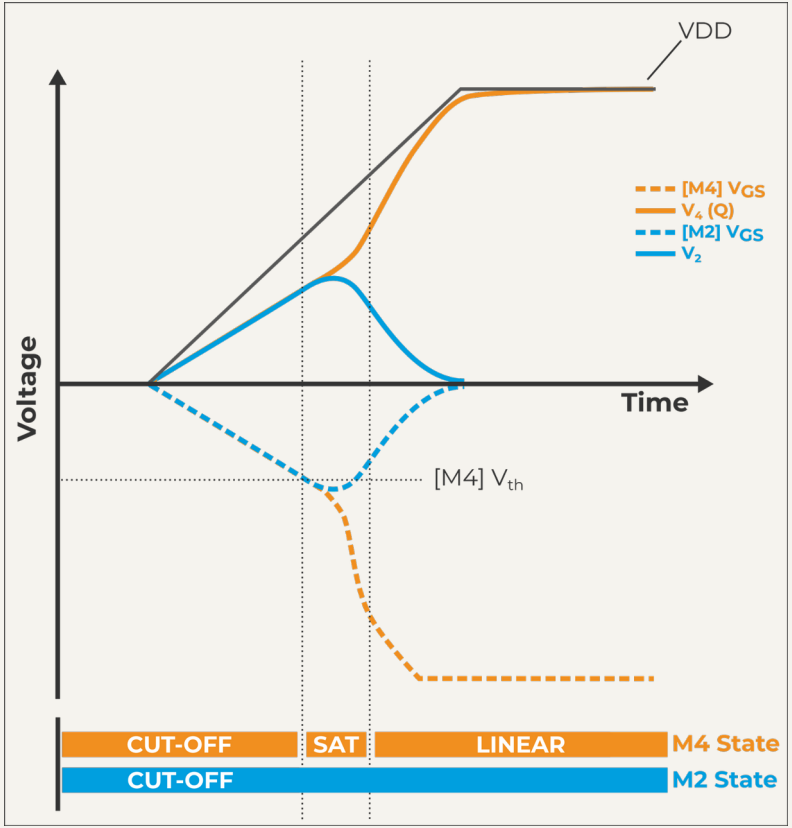

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:

SRAM 6T Cell Startup

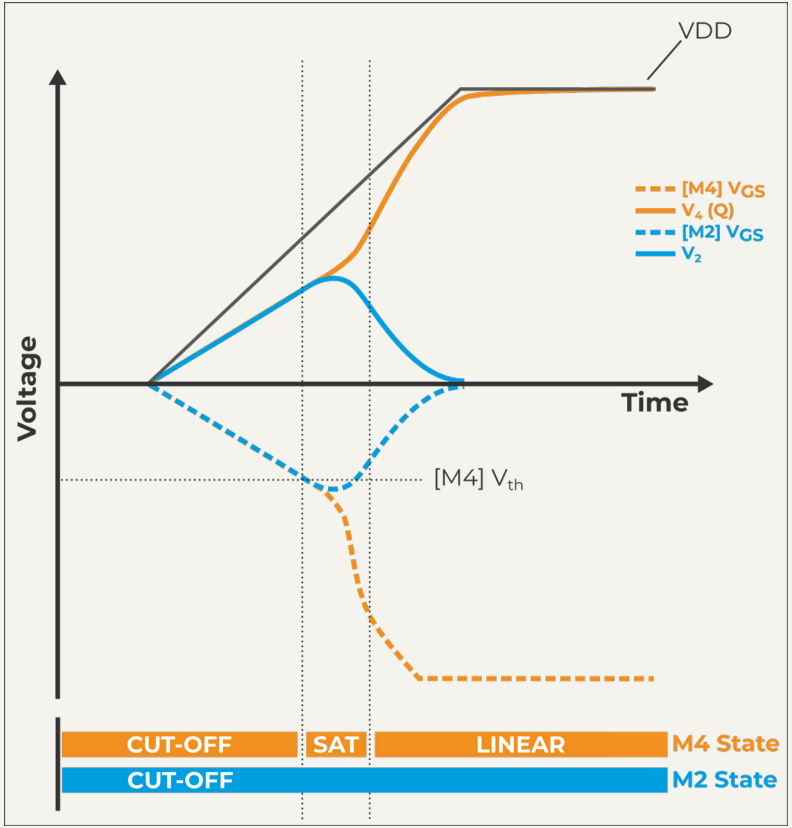

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:
  - **Small V_{th} difference ⇒ values ~ random.**

SRAM 6T Cell Startup

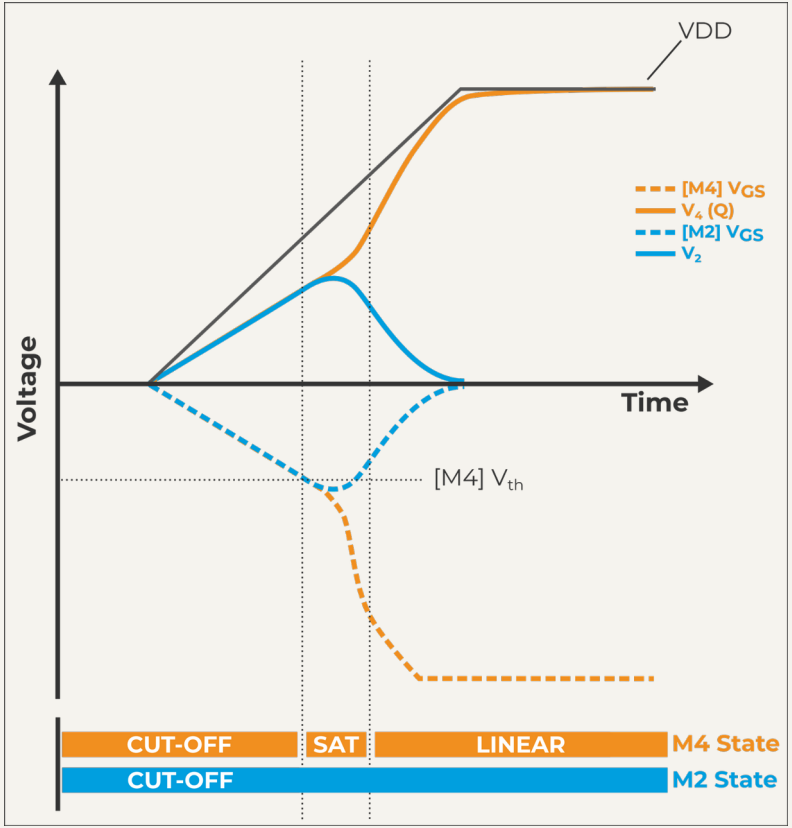

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:
  - **Small V_{th} difference ⇒ values ~ random.**
  - **Big V_{th} difference ⇒ values biased and ~ stable.**

SRAM 6T Cell Startup

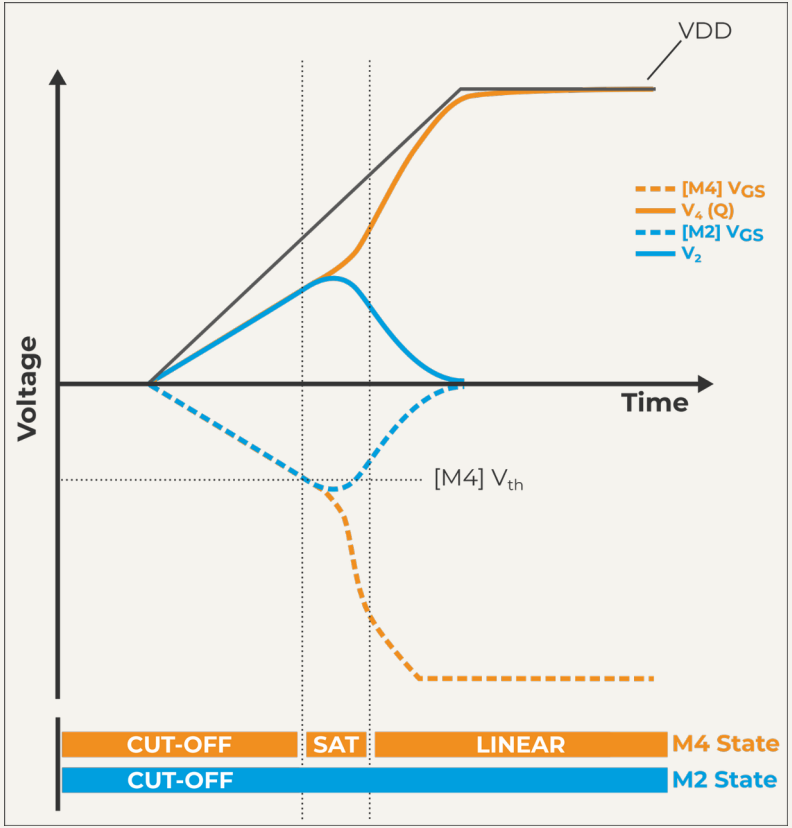

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:
  - **Small V_{th} difference ⇒ values ~ random.**
  - **Big V_{th} difference ⇒ values biased and ~ stable.**
- Stress (usage) progressively stabilizes unbiased cells:

SRAM 6T Cell Startup

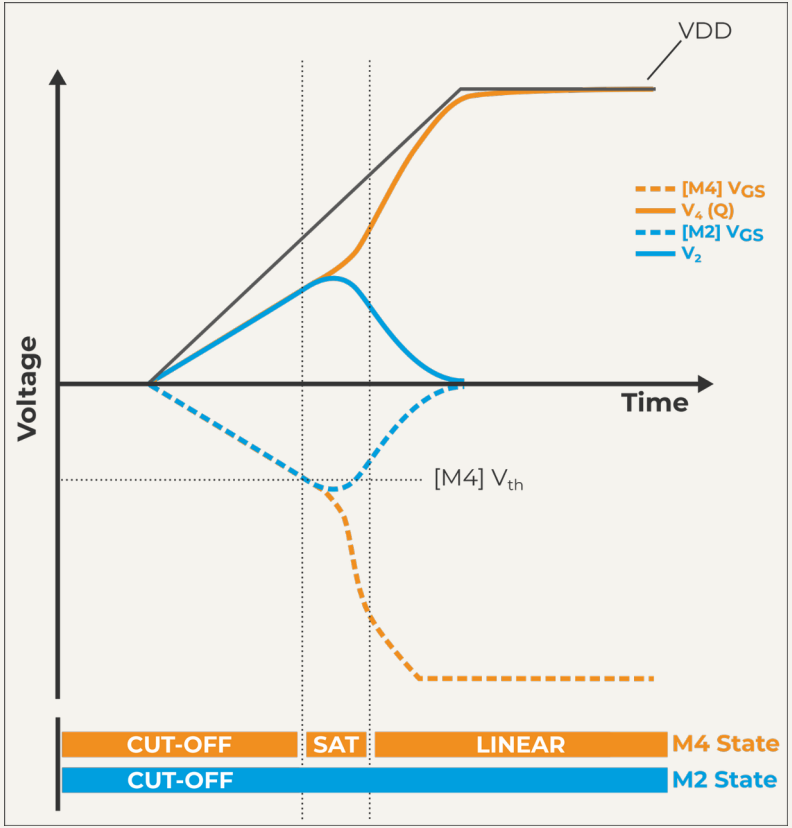

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:
  - **Small V_{th} difference ⇒ values ~ random.**
  - **Big V_{th} difference ⇒ values biased and ~ stable.**
- Stress (usage) progressively stabilizes unbiased cells:
  - **BTI** increases V_{th} when transistors are active.

SRAM 6T Cell Startup

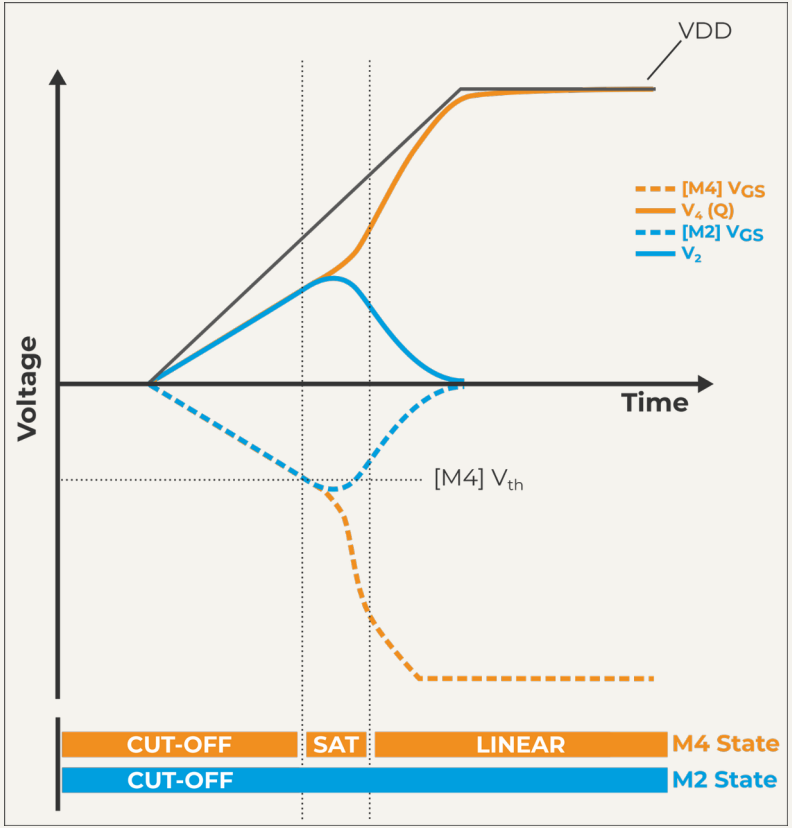

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:
  - **Small V_{th} difference ⇒ values ~ random.**
  - **Big V_{th} difference ⇒ values biased and ~ stable.**
- Stress (usage) progressively stabilizes unbiased cells:
  - **BTI** increases V_{th} when transistors are active.
  - Cells are biased towards the value opposite to stress.

SRAM 6T Cell Startup

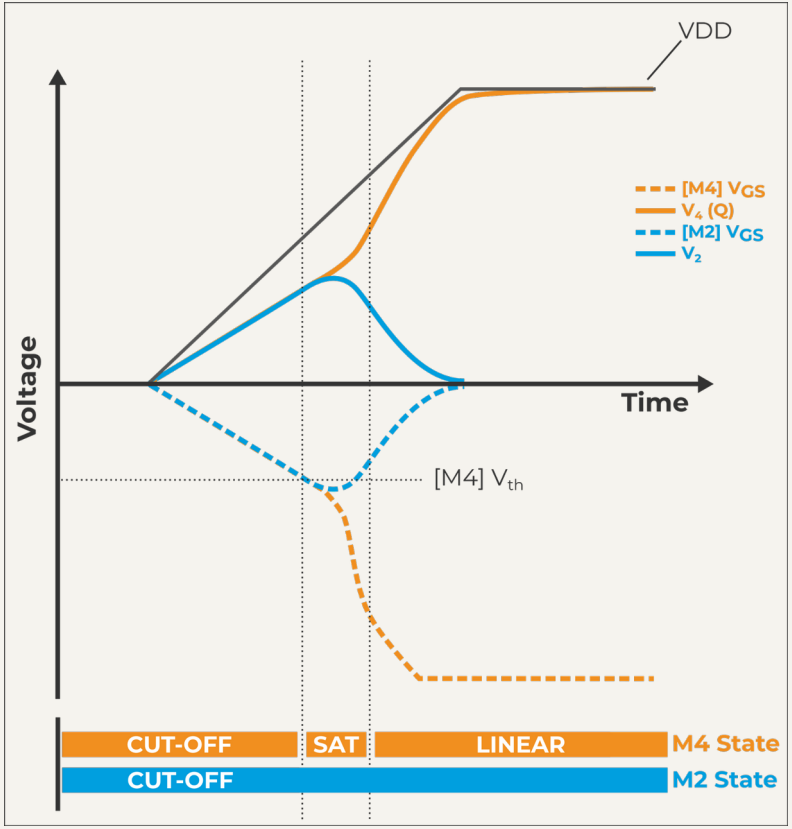

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:
  - **Small V_{th} difference ⇒ values ~ random.**
  - **Big V_{th} difference ⇒ values biased and ~ stable.**
- Stress (usage) progressively stabilizes unbiased cells:
  - **BTI** increases V_{th} when transistors are active.
  - Cells are biased towards the value opposite to stress.

SRAM 6T Cell Startup

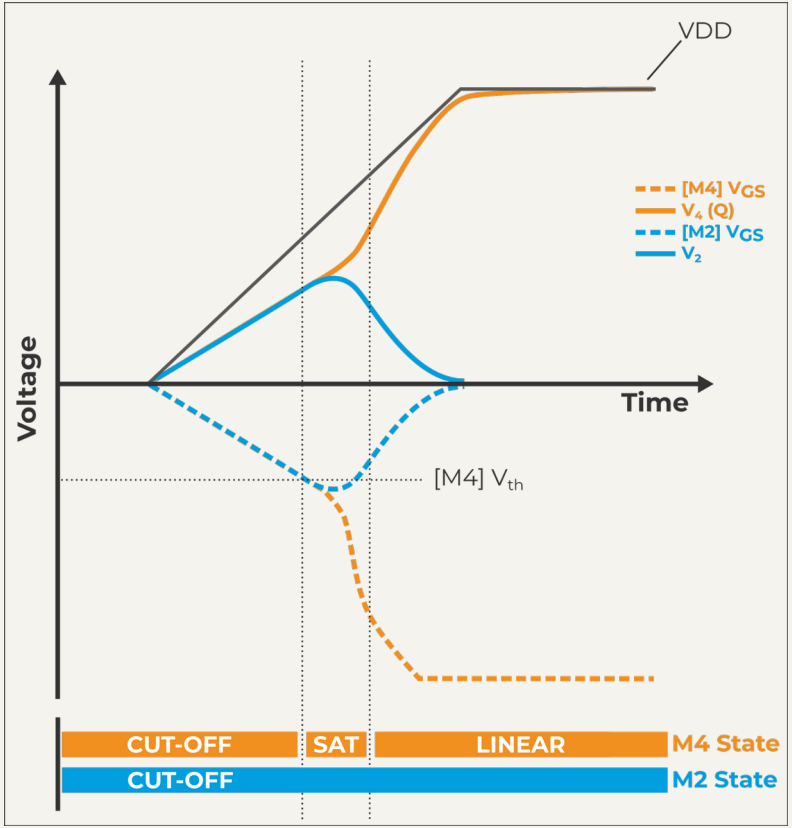

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:
  - **Small V_{th} difference ⇒ values ~ random.**
  - **Big V_{th} difference ⇒ values biased and ~ stable.**
- Stress (usage) progressively stabilizes unbiased cells:
  - **BTI** increases V_{th} when transistors are active.
  - Cells are biased towards the value opposite to stress.
- **'Percentage of 1s' on new SRAMs show ~50%.**

SRAM 6T Cell Startup

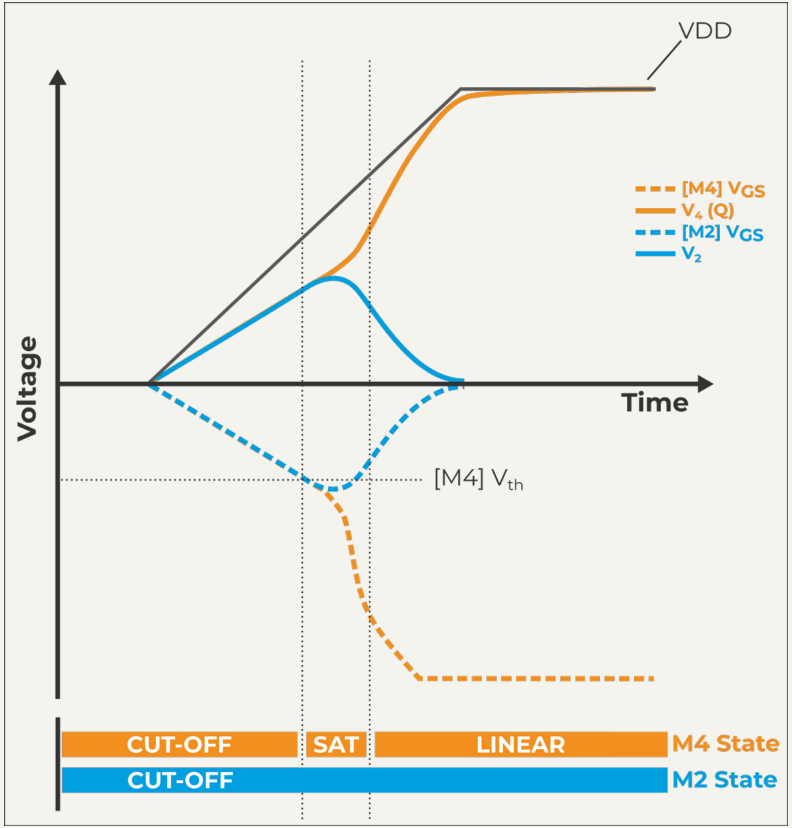

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:
  - **Small V_{th} difference ⇒ values ~ random.**
  - **Big V_{th} difference ⇒ values biased and ~ stable.**
- Stress (usage) progressively stabilizes unbiased cells:
  - **BTI** increases V_{th} when transistors are active.
  - Cells are biased towards the value opposite to stress.

- **'Percentage of 1s' on new SRAMs show ~50%.**
- **Aged SRAMs should drift over time.**

SRAM 6T Cell Startup

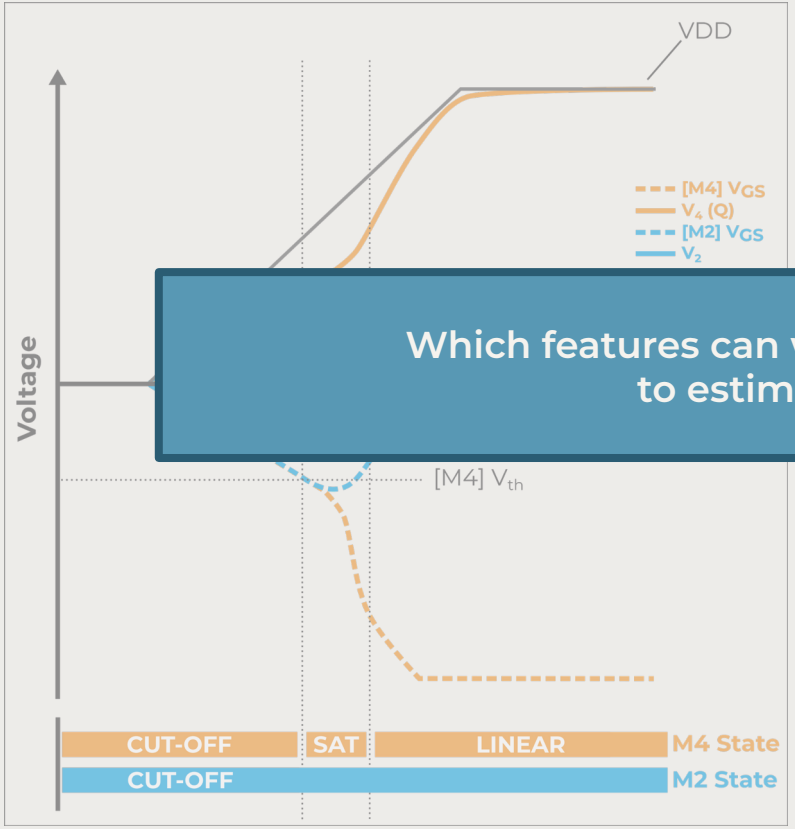

Which features can we extract from SRAM startup to estimate device usage?

- When initially powered, the value of the cell depends on the V_{th} of the inverter transistors:
  - **Small V_{th} difference ⇒ values ~ random.**
  - **table.**
  - cells:
  - **BTI** increases V_{th} when transistors are active.
  - Cells are biased towards the value opposite to stress.
- **'Percentage of 1s' on new SRAMs show ~50%.**
- **Aged SRAMs should drift over time.**

Evaluated Devices

- **154 ST Microelectronics MCUs** from public IoT testbeds across multiple locations.
- 7 Atmel MCUs from Eu-XFEL at DESY.
- Startup devices 1000 times, and observe un-initialized SRAM bits.

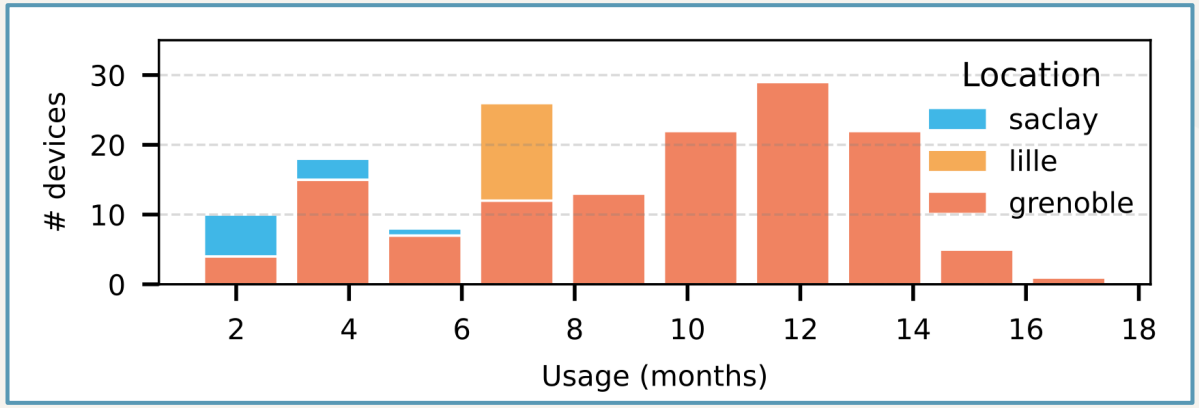

Percentage of Ones

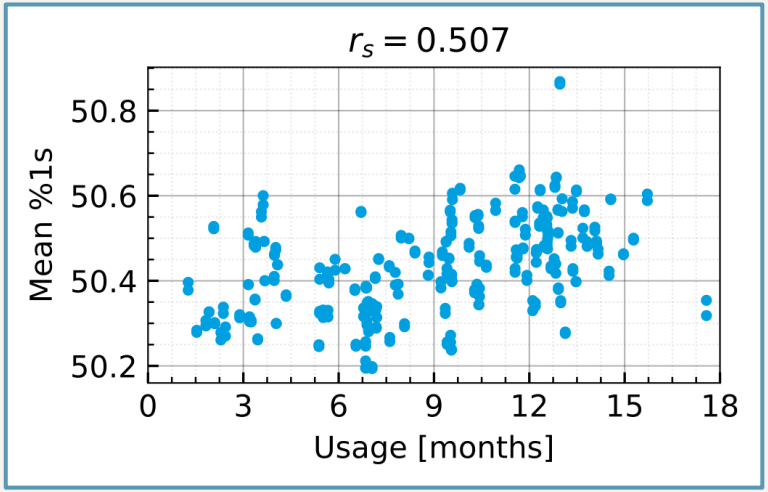

Percentage of Ones

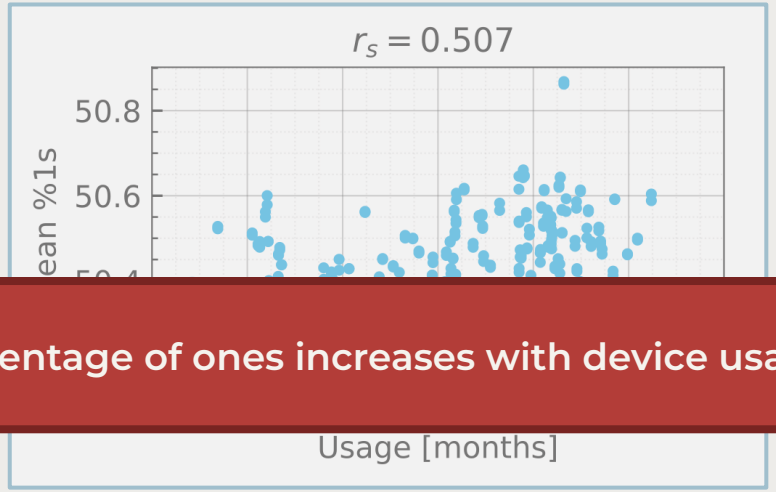

The percentage of ones increases with device usage time.

Probability of 1 (P_1) and Bit Activity

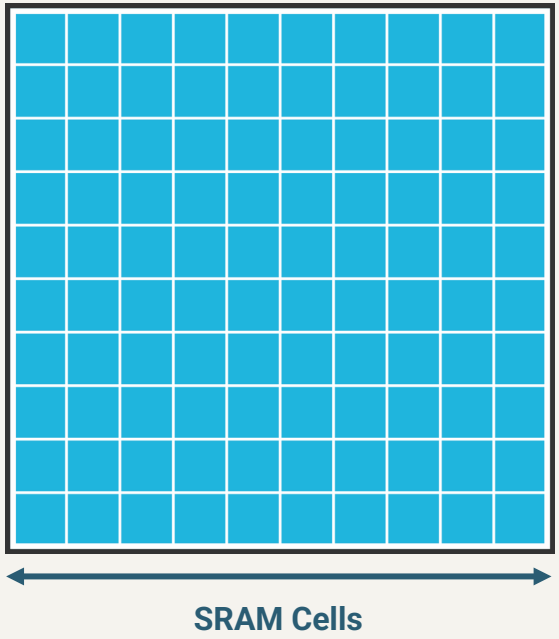

Probability of 1 (P_1) and Bit Activity

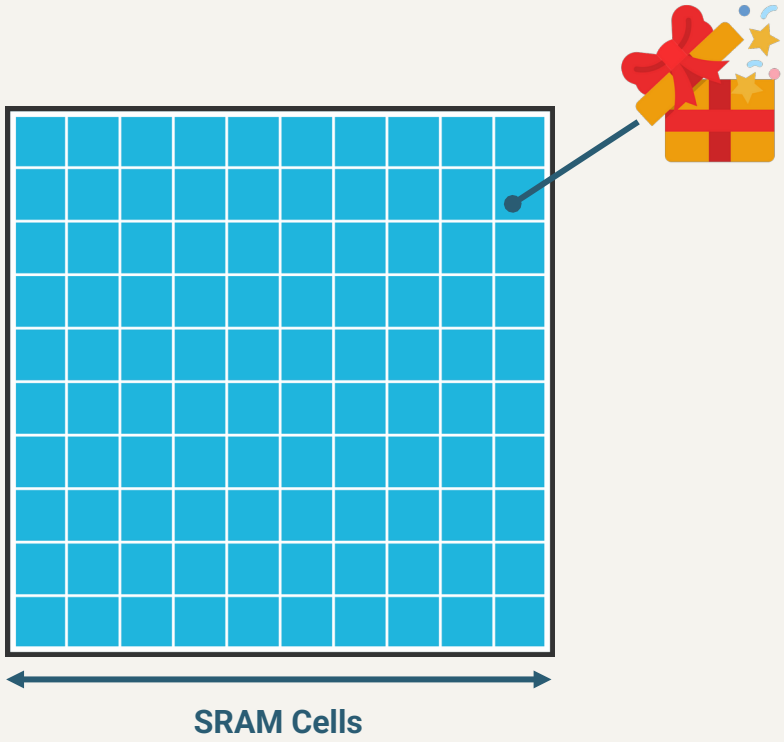

Probability of 1 (P_1) and Bit Activity

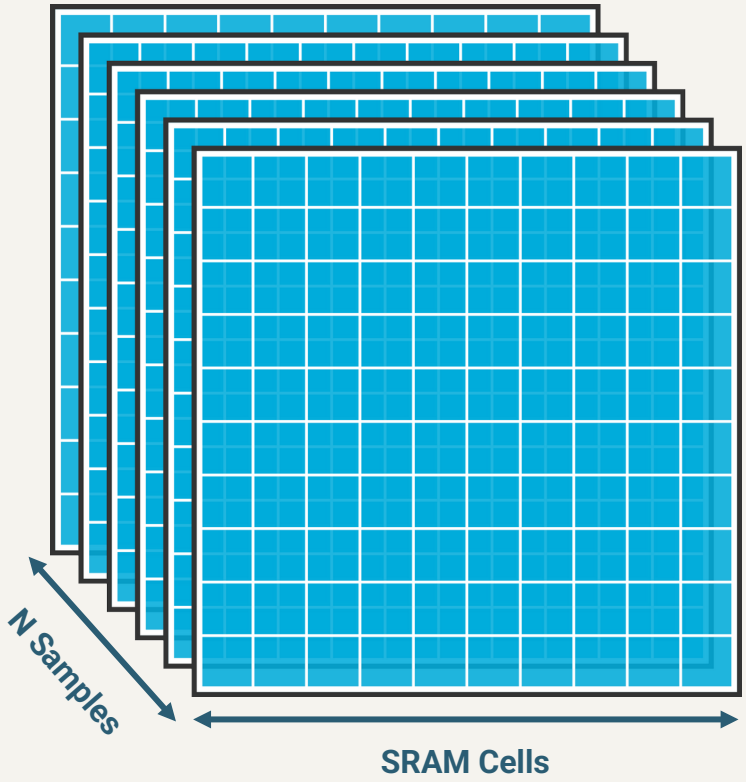

Probability of 1 (P_1) and Bit Activity

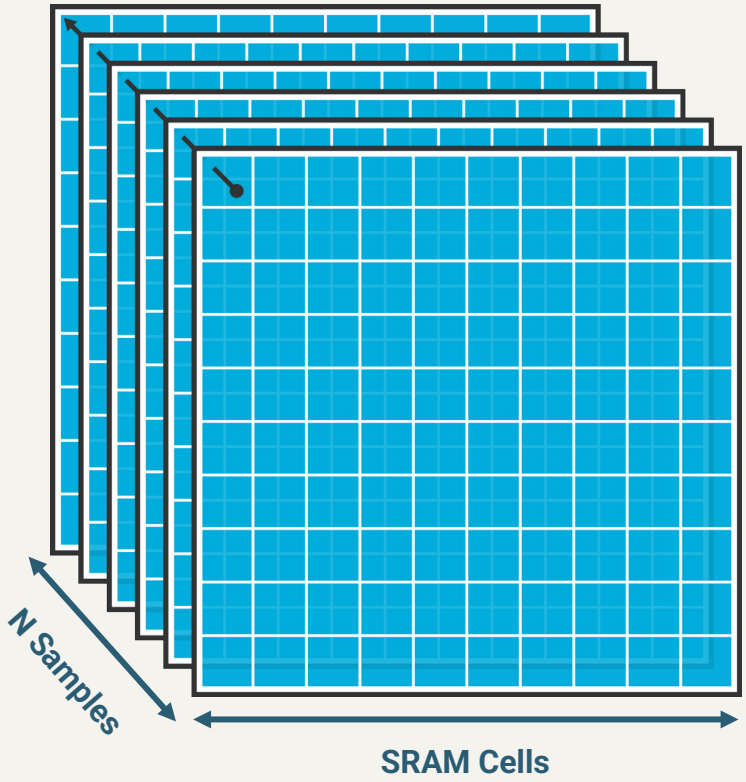

Probability of 1 (P_1) and Bit Activity

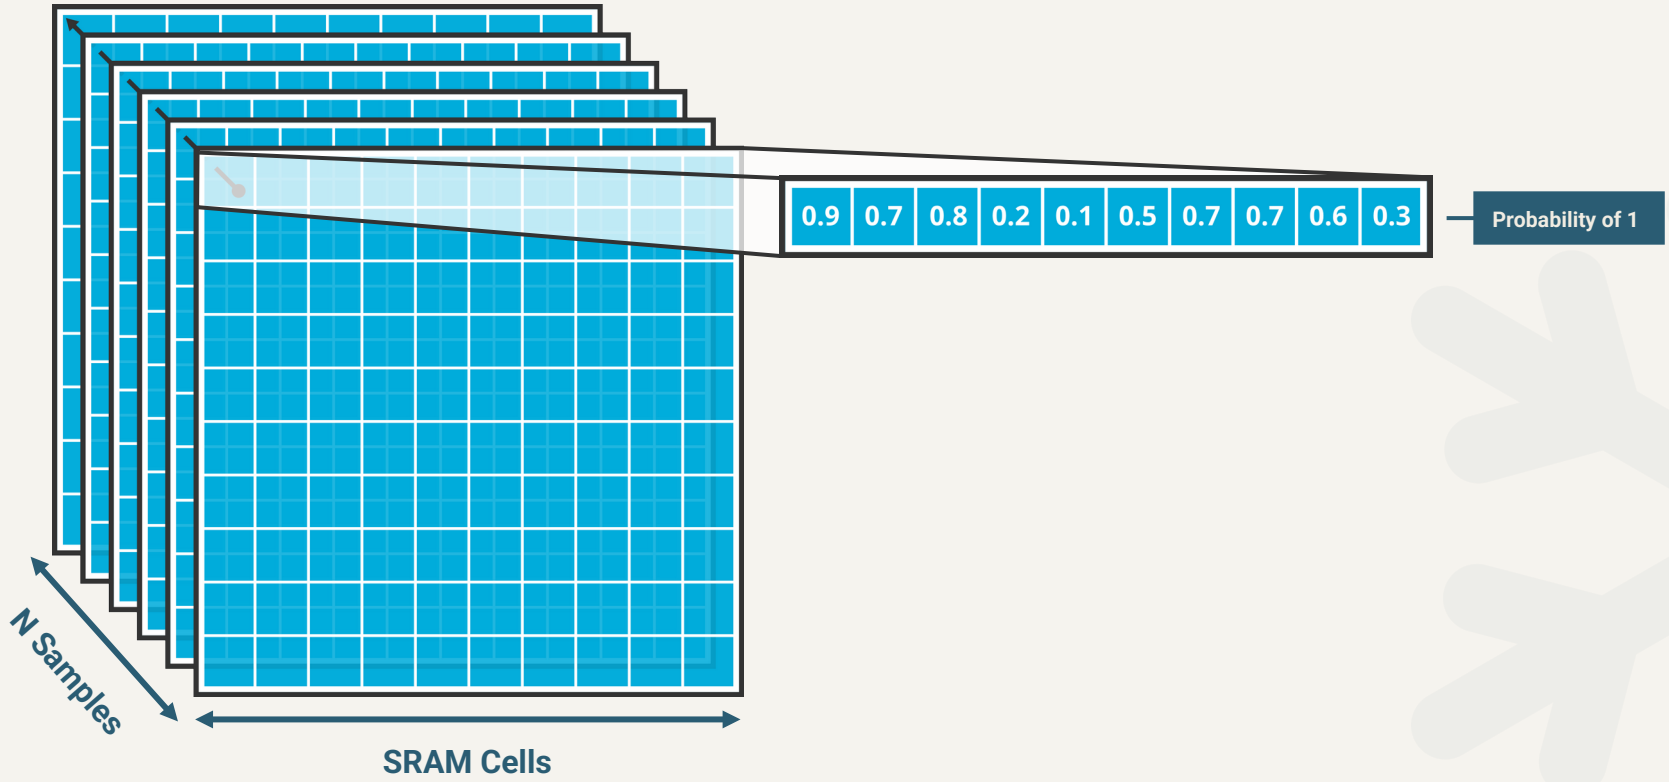

Probability of 1 (P_1) and Bit Activity

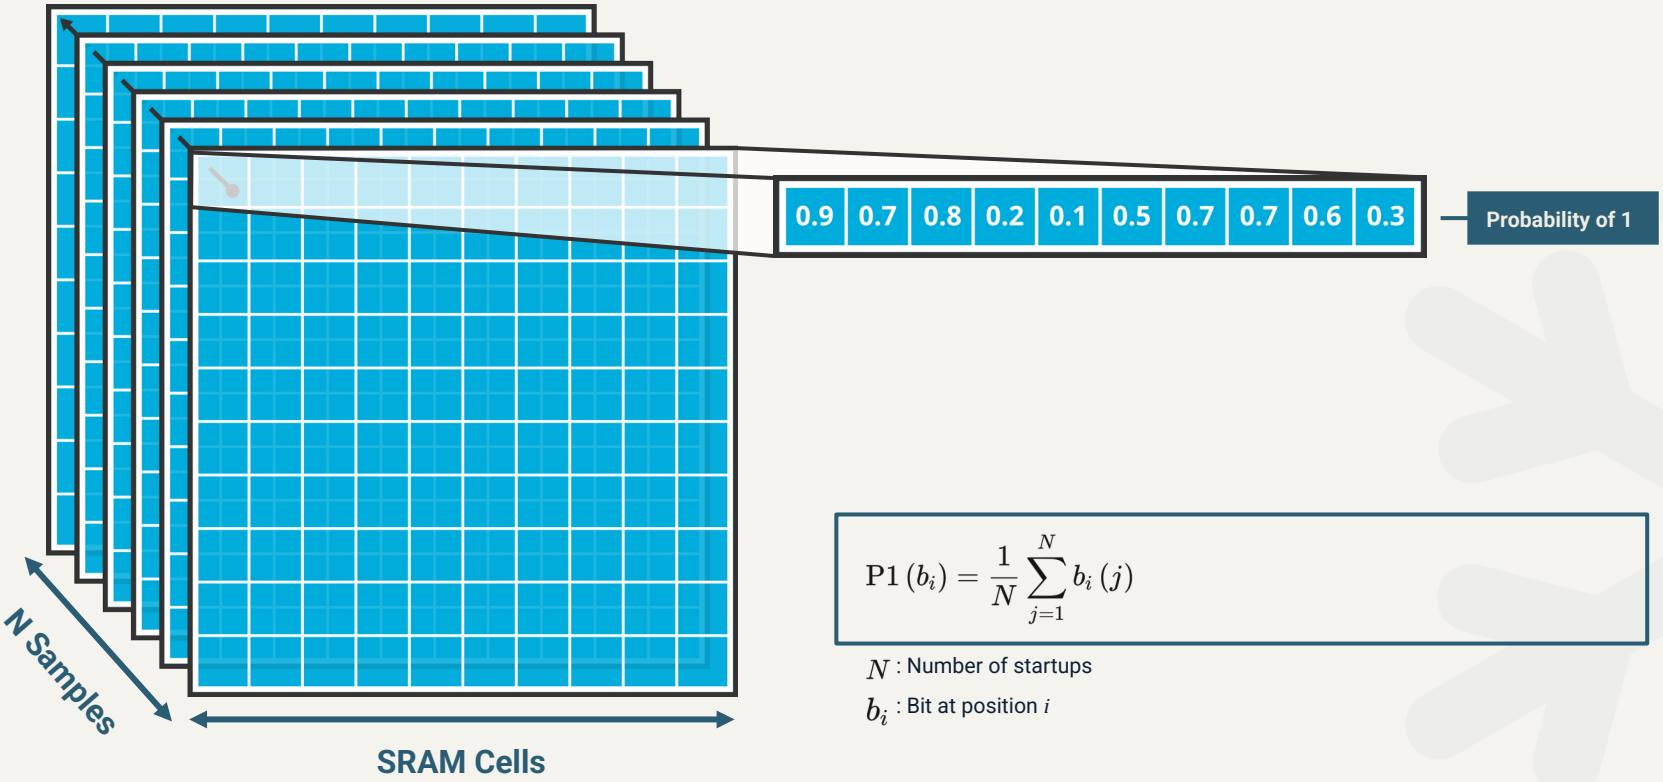

$$\mathrm{P1}\,(b_i) = \frac{1}{N}\sum_{j=1}^{N} b_i\,(j)$$

N : Number of startups

b_i : Bit at position i

Probability of 1 (P_1) and Bit Activity

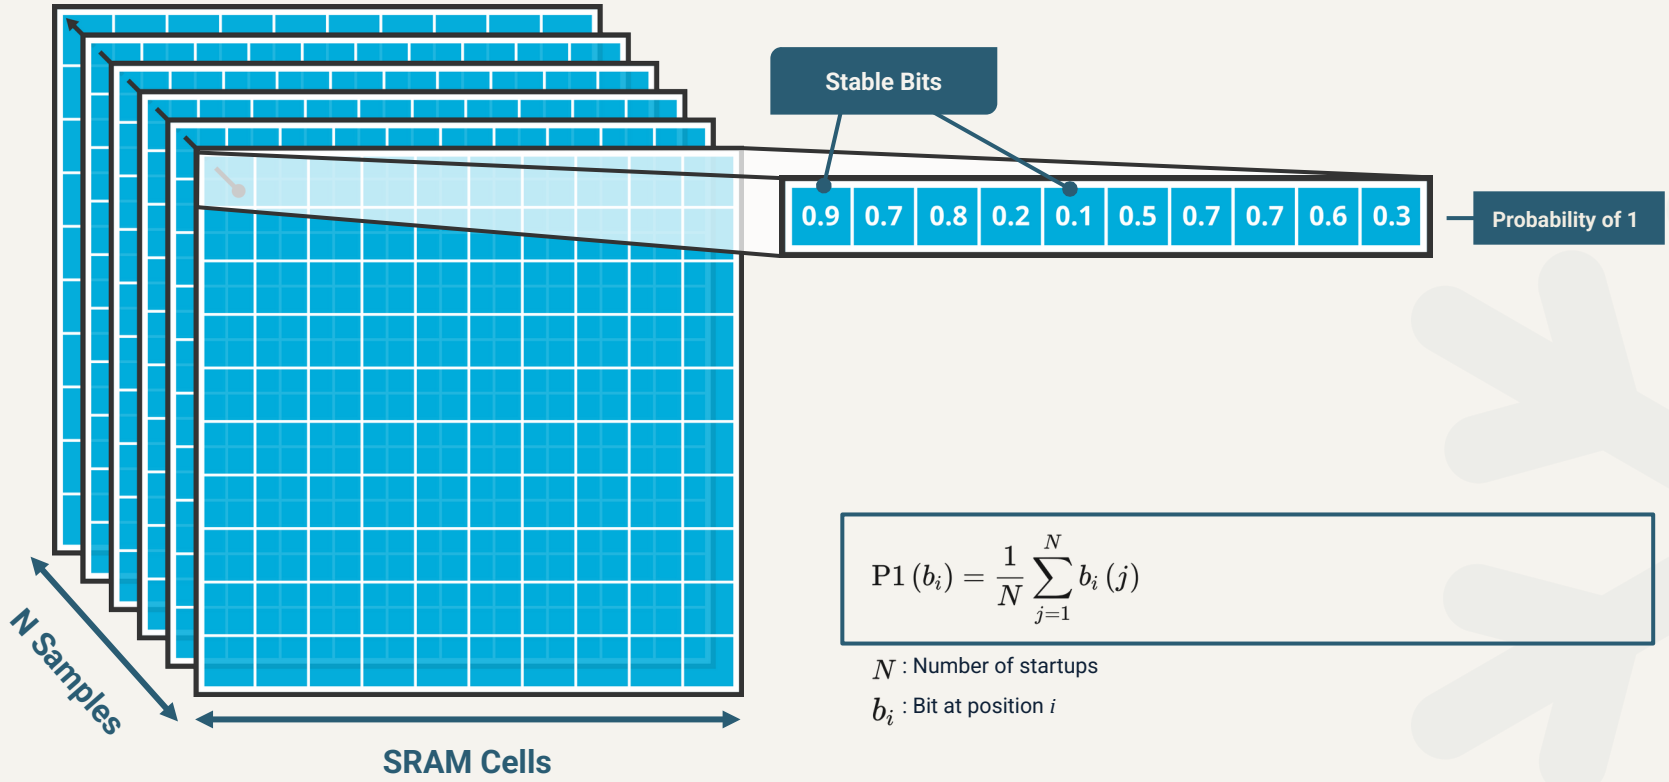

$$\mathrm{P1}\,(b_i) = \frac{1}{N}\sum_{j=1}^{N} b_i\,(j)$$

N : Number of startups

b_i : Bit at position i

Probability of 1 (P_1) and Bit Activity

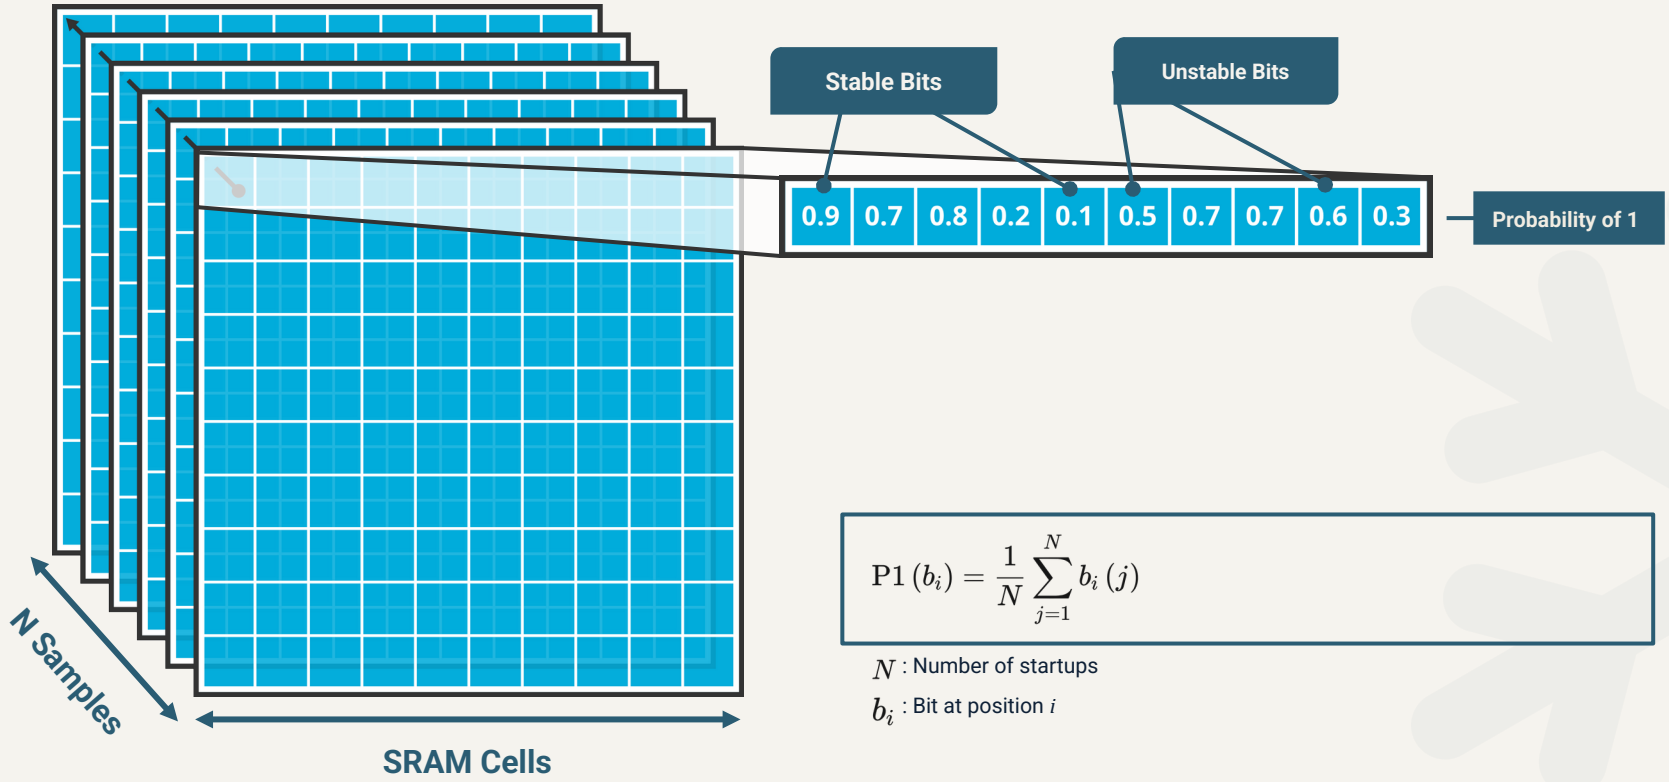

$$\mathrm{P1}\,(b_i) = \frac{1}{N}\sum_{j=1}^{N} b_i\,(j)$$

N : Number of startups

b_i : Bit at position i

Probability of 1 (P_1) and Bit Activity

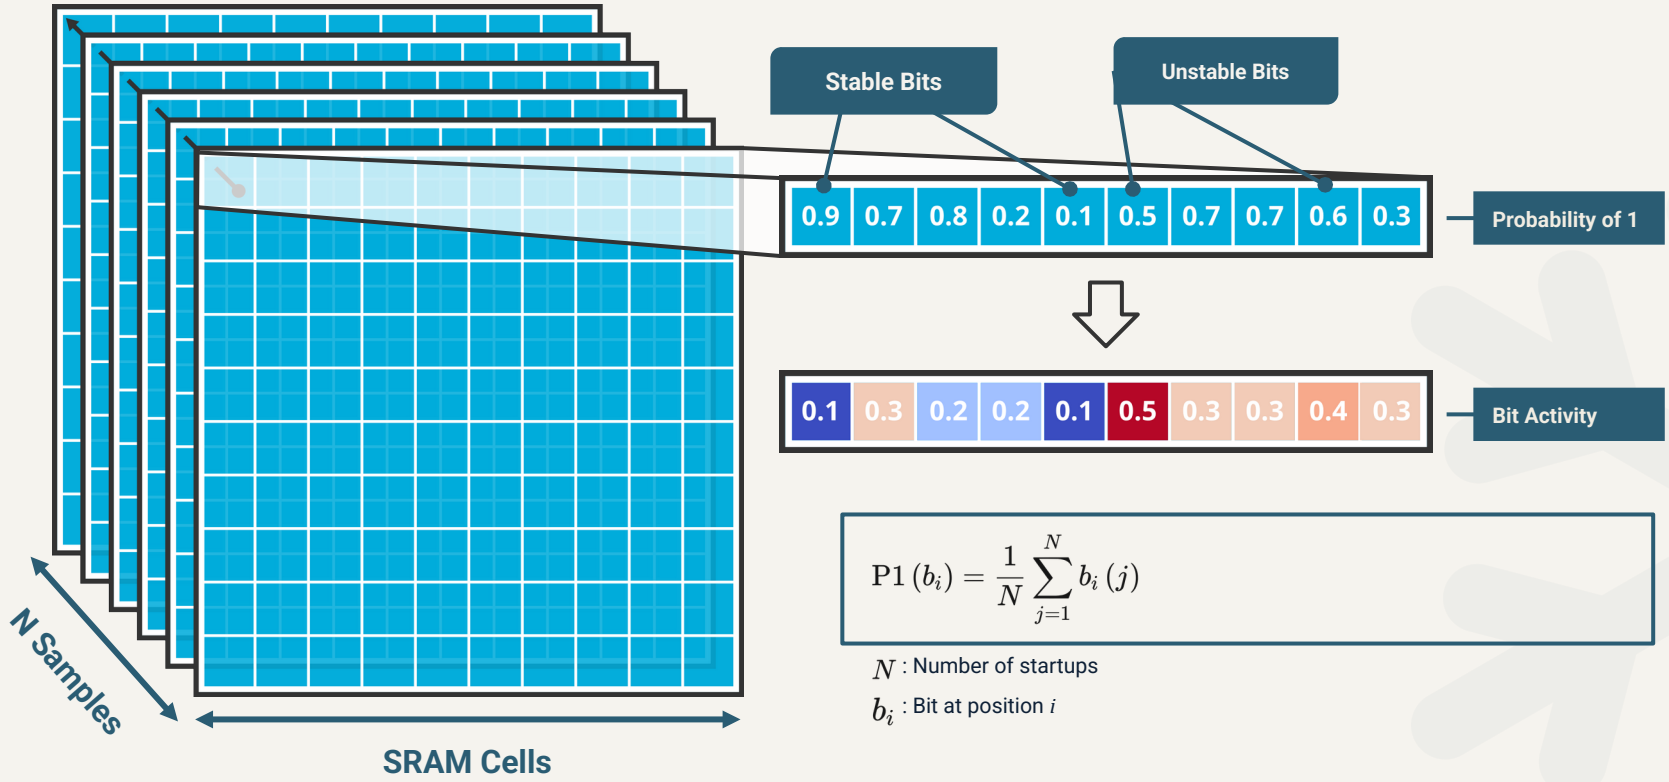

$$\mathrm{P1}\,(b_i) = \frac{1}{N}\sum_{j=1}^{N} b_i\,(j)$$

N : Number of startups

b_i : Bit at position i

Probability of 1 (P_1) and Bit Activity

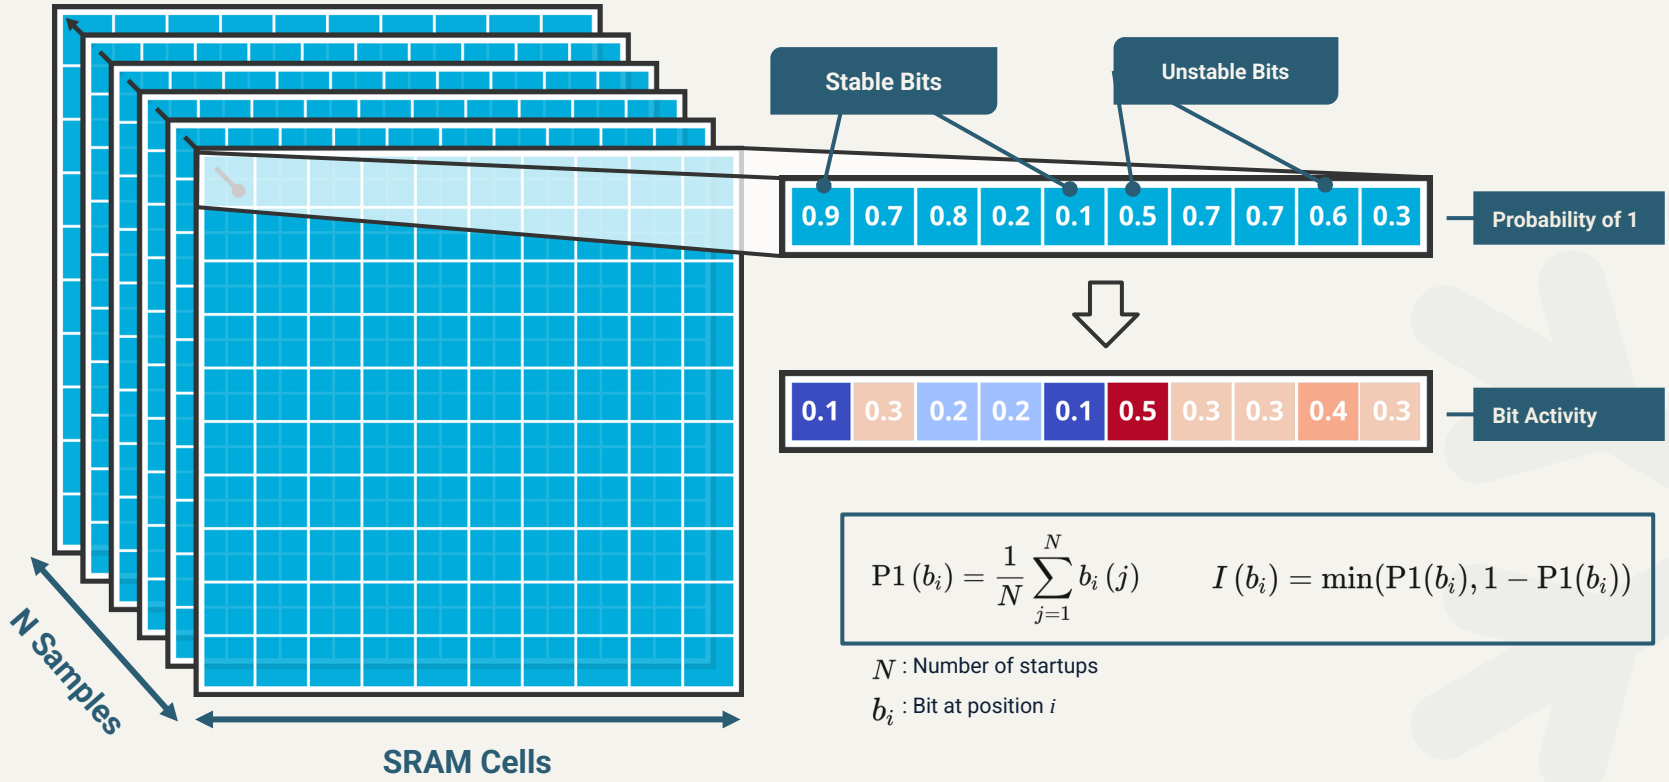

$$\mathrm{P1}\,(b_i) = \frac{1}{N}\sum_{j=1}^{N} b_i\,(j) \qquad I\,(b_i) = \min(\mathrm{P1}(b_i), 1 - \mathrm{P1}(b_i))$$

N : Number of startups

b_i : Bit at position i

Bit Activity Measurements

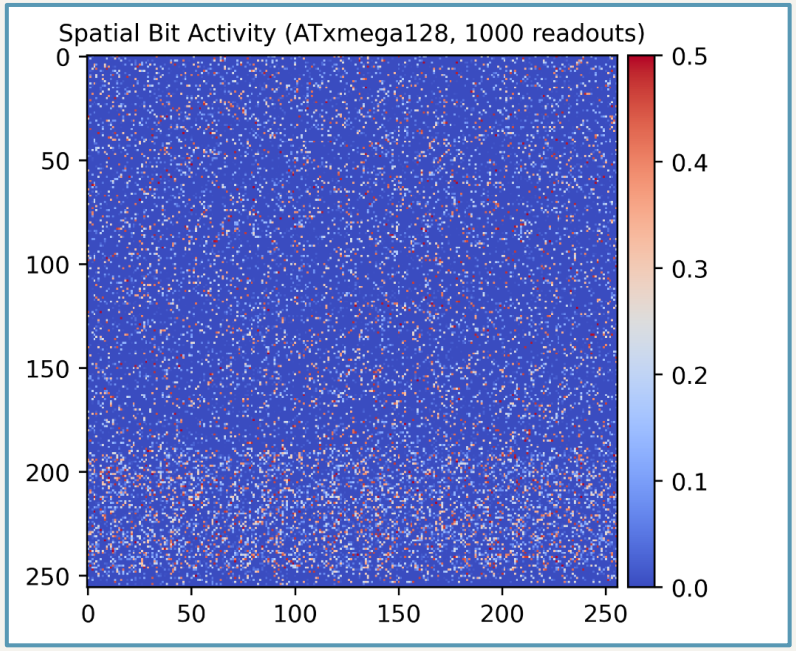

Bit Activity Measurements

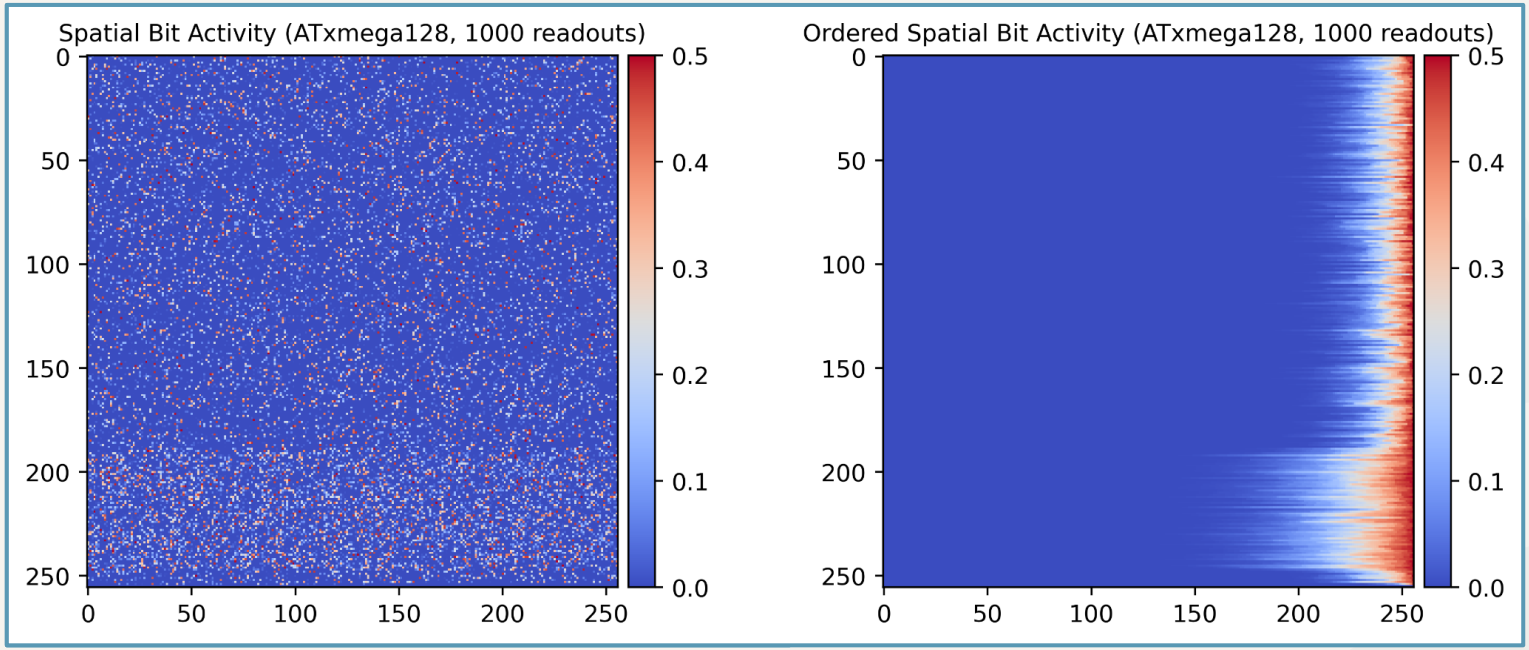

Bit Activity Measurements

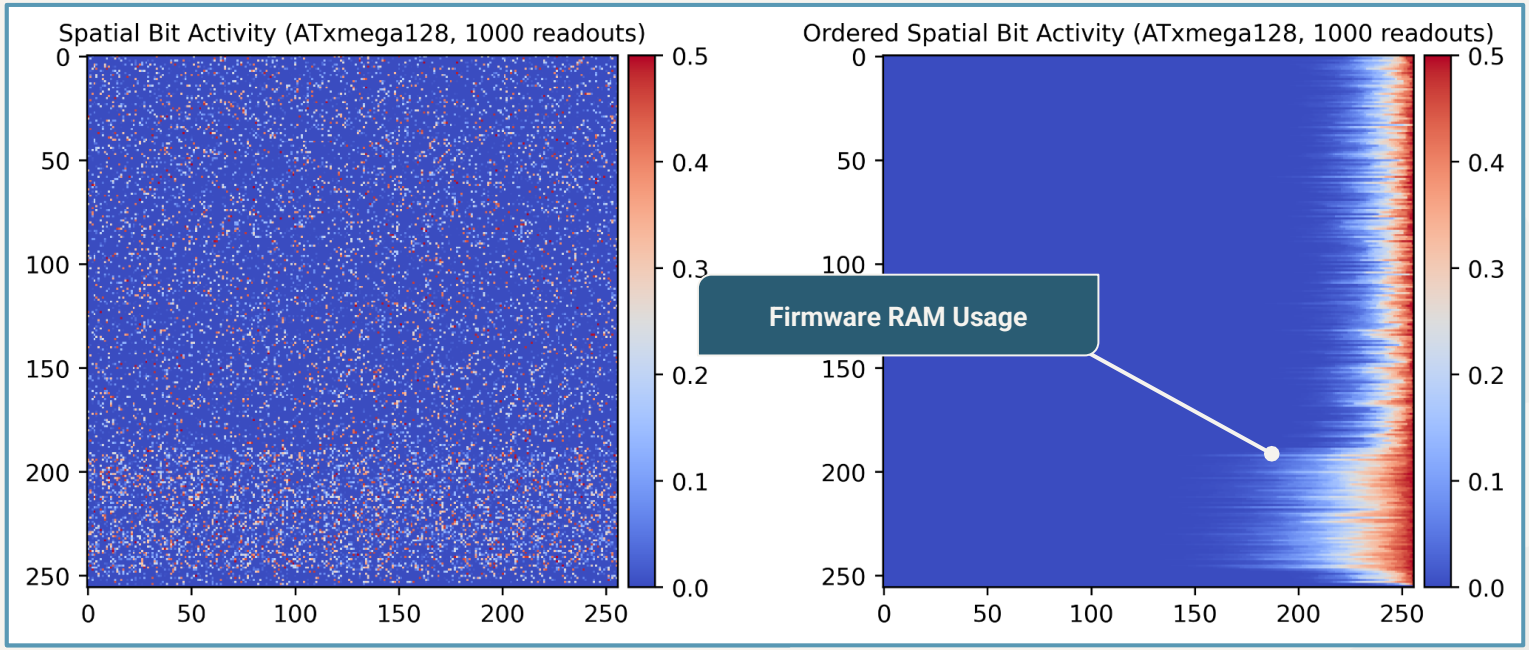

Bit Activity Measurements

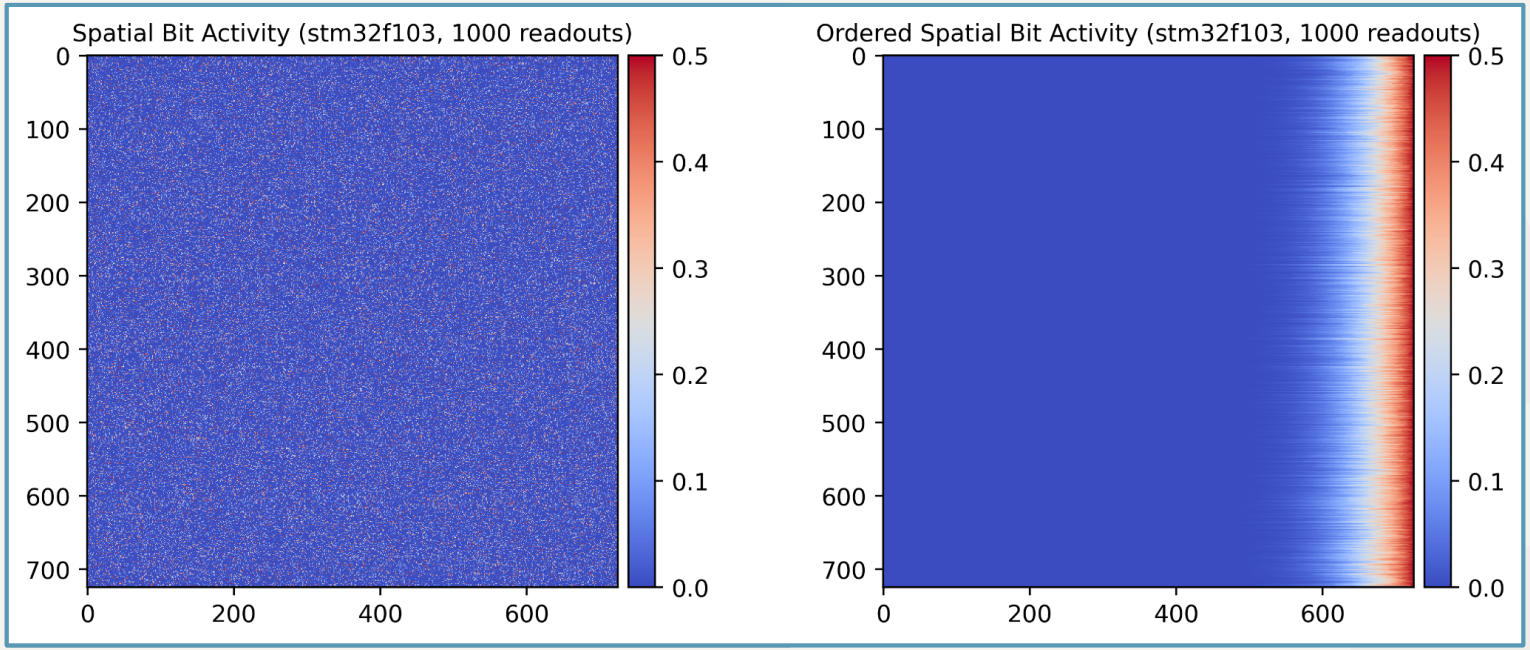

Correlation between P_1 and Address

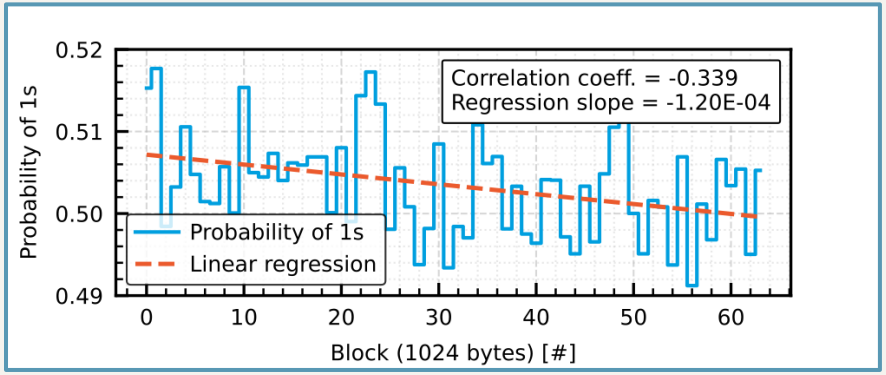

Correlation between P_1 and Address

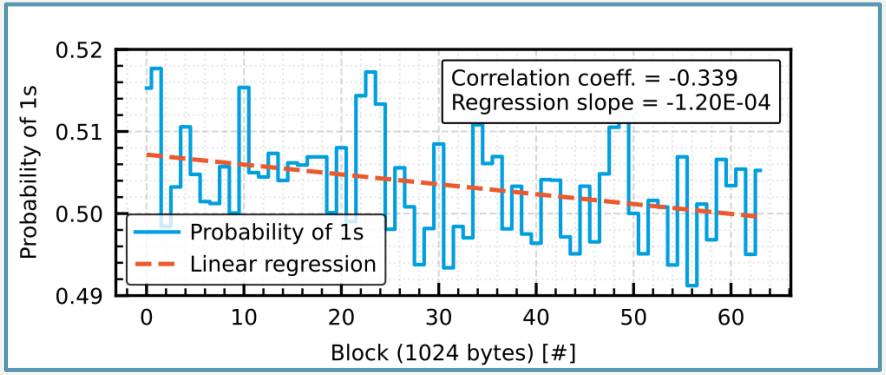

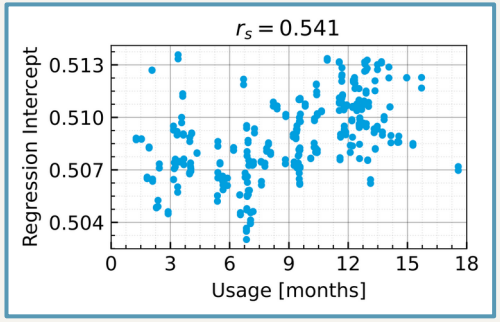

Correlation between P_1 and Address

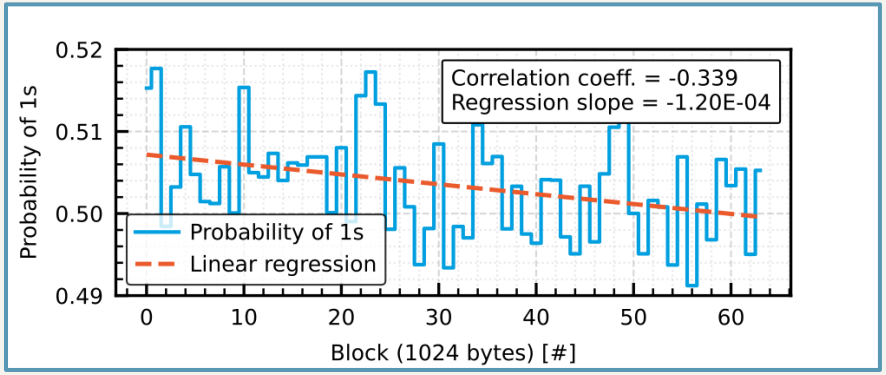

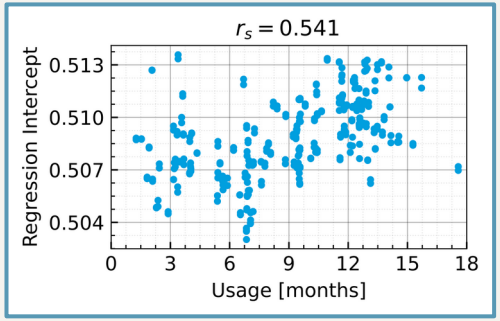

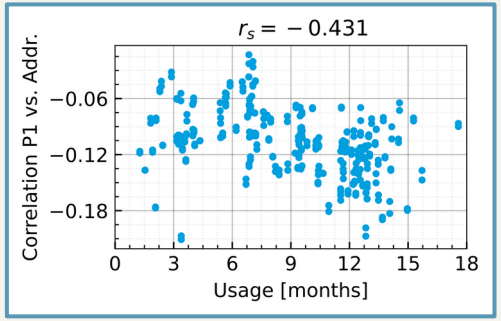

Correlation between P_1 and Address

Bits in lower memory addresses tend to start with 1s.

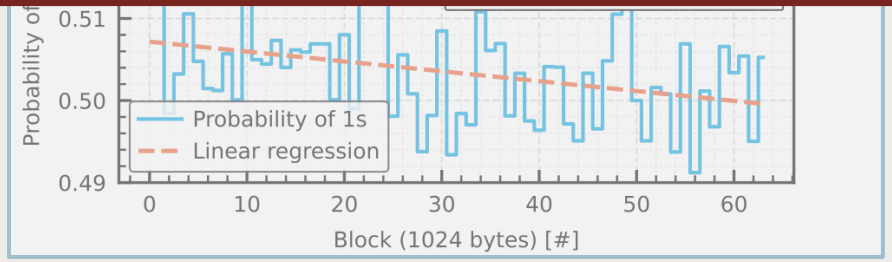

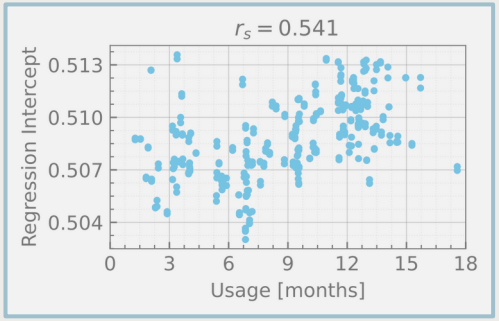

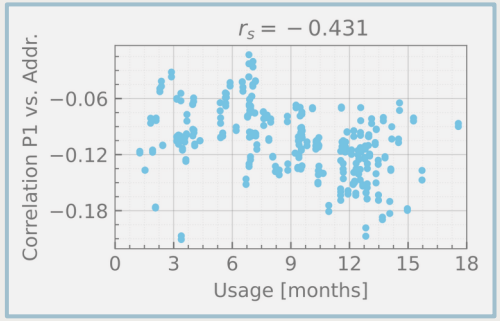

Correlation between P_1 and Address

Bits in lower memory addresses tend to start with 1s.

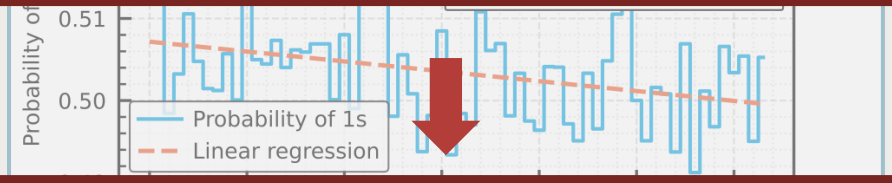

Lower memory addresses tend to hold 0s during operation.

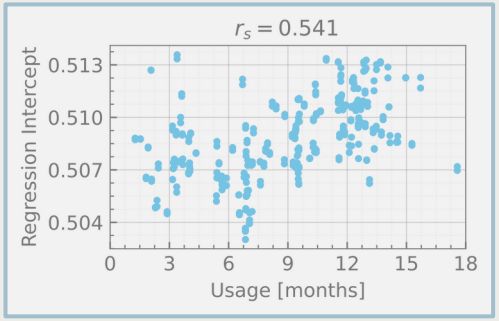

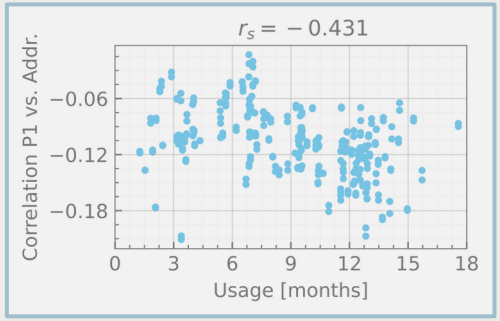

Correlation between P_1 and Address

Bits in lower memory addresses tend to start with 1s.

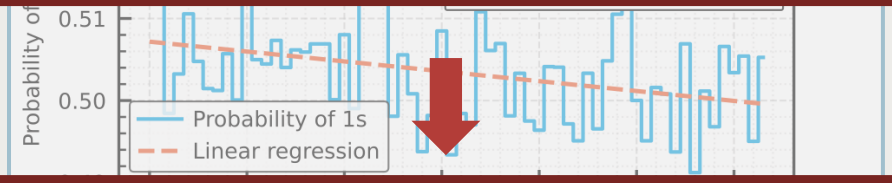

Lower memory addresses tend to hold 0s during operation.

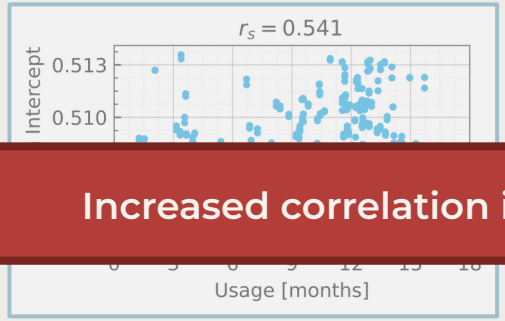

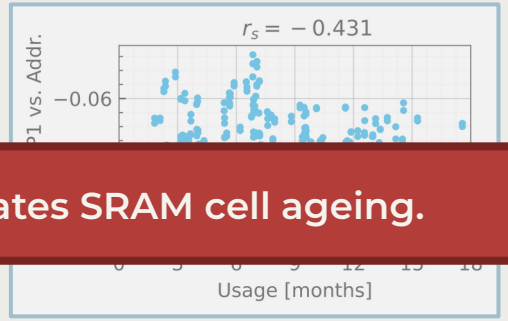

Increased correlation indicates SRAM cell ageing.

Spatial Frequency Spectrum of P_1

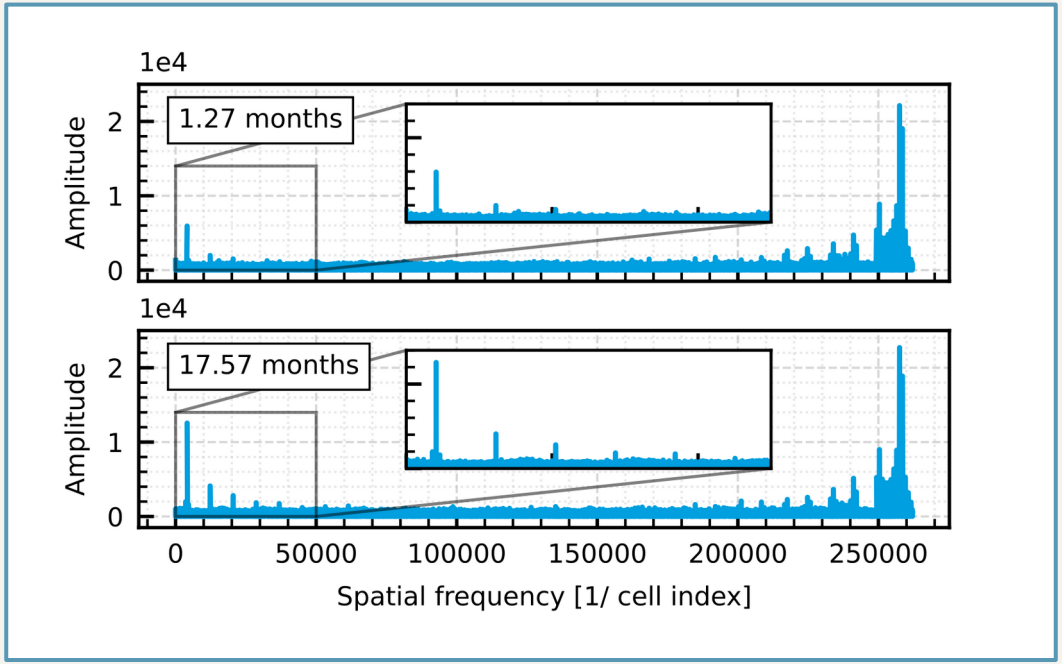

Spatial Frequency Spectrum of P_1

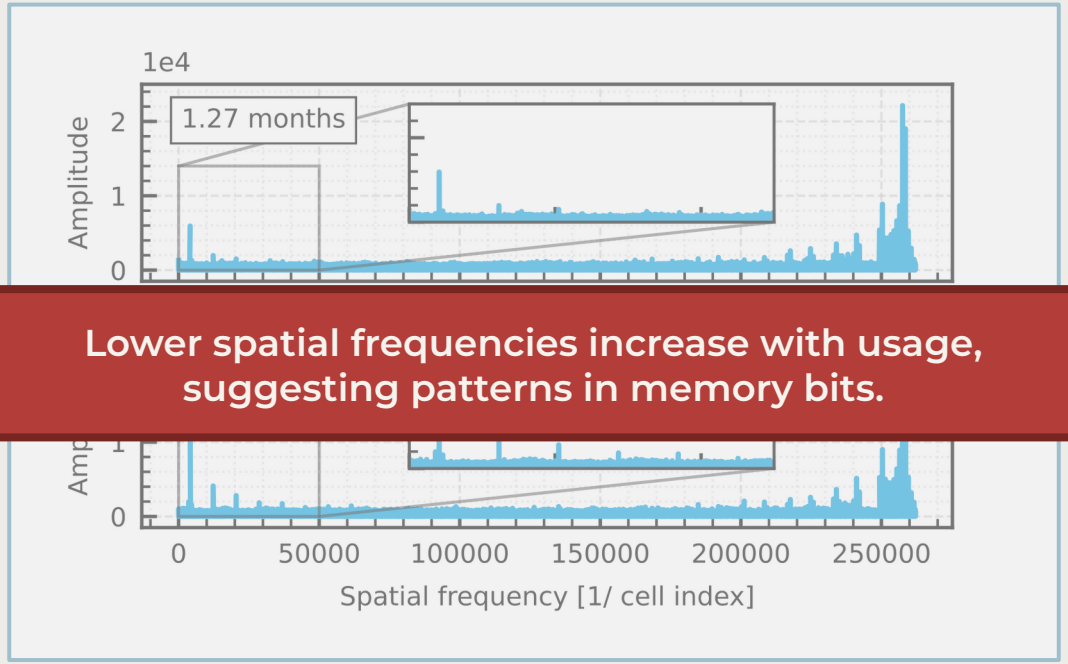

Lower spatial frequencies increase with usage, suggesting patterns in memory bits.

Evaluation of Usage Estimation

Evaluation of Usage Estimation

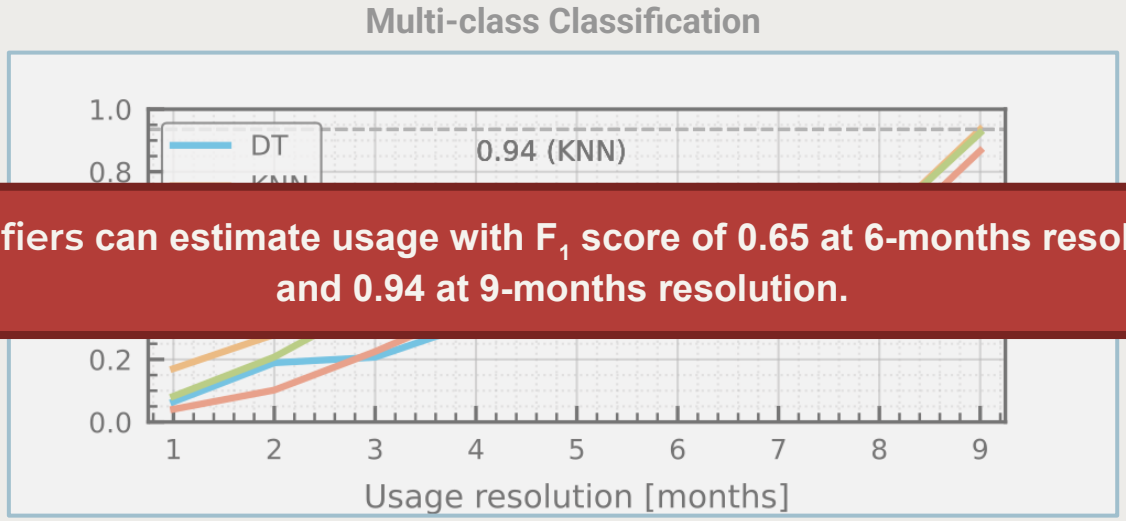

Classifiers can estimate usage with F_1 score of 0.65 at 6-months resolution, and 0.94 at 9-months resolution.

Evaluation of Usage Estimation

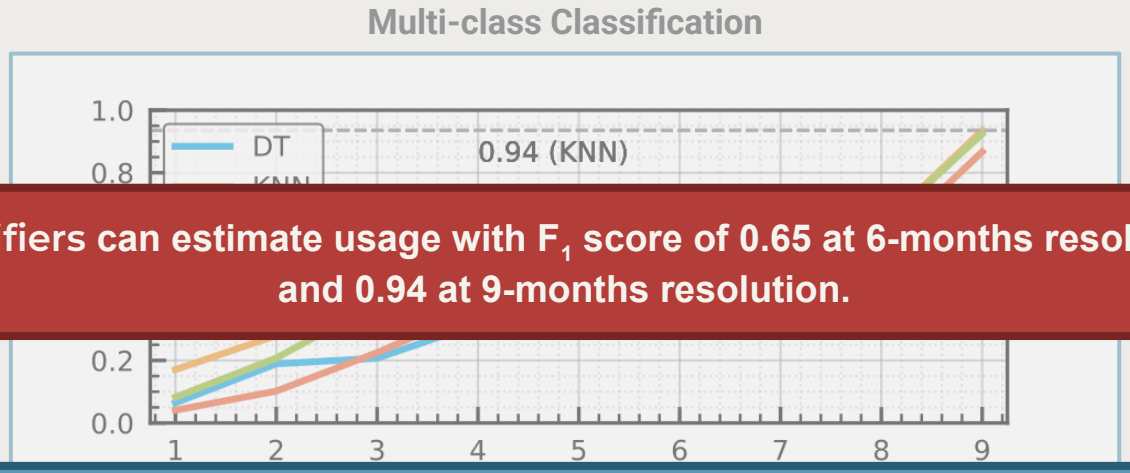

Classifiers can estimate usage with F_1 score of 0.65 at 6-months resolution, and 0.94 at 9-months resolution.

For more details on usage estimation, check ref. [2].

In Summary

[2]: L. Lanzieri, P. Kietzmann, G. Fey, H. Schlarb and T. C. Schmidt, "Ageing Analysis of Embedded SRAM on a Large-Scale Testbed Using Machine Learning", 2023. In proceedings of the 26th Euromicro Conference on Digital System Design (DSD), Golem, Albania, 2023, pp. 335-342. DOI: 10.1109/DSD60849.2023.00054.

In Summary

- Identified SRAM effects as a product of ageing:
  - The **share of ones** in patterns **increases** with device usage time.
  - Bits in **lower memory** addresses tend to **hold 0s** while running.
  - Lower spatial frequencies increase with usage, indicating patterns in memory cells.

[2]: L. Lanzieri, P. Kietzmann, G. Fey, H. Schlarb and T. C. Schmidt, "Ageing Analysis of Embedded SRAM on a Large-Scale Testbed Using Machine Learning", 2023. In proceedings of the 26th Euromicro Conference on Digital System Design (DSD), Golem, Albania, 2023, pp. 335-342. DOI: 10.1109/DSD60849.2023.00054.

In Summary

- Identified SRAM effects as a product of ageing:
  - The **share of ones** in patterns **increases** with device usage time.
  - Bits in **lower memory** addresses tend to **hold 0s** while running.
  - Lower spatial frequencies increase with usage, indicating patterns in memory cells.
- Usage time estimation based on SRAM startup features is feasible.

[2]: L. Lanzieri, P. Kietzmann, G. Fey, H. Schlarb and T. C. Schmidt, "Ageing Analysis of Embedded SRAM on a Large-Scale Testbed Using Machine Learning", 2023. In proceedings of the 26th Euromicro Conference on Digital System Design (DSD), Golem, Albania, 2023, pp. 335-342. DOI: 10.1109/DSD60849.2023.00054.

Agenda

- Ageing Detection on Embedded SRAMs
- **Monitoring Degradation of Propagation Delay on FPGAs**
- Ageing Monitoring for Commercial MCUs

Ring Oscillator Sensor

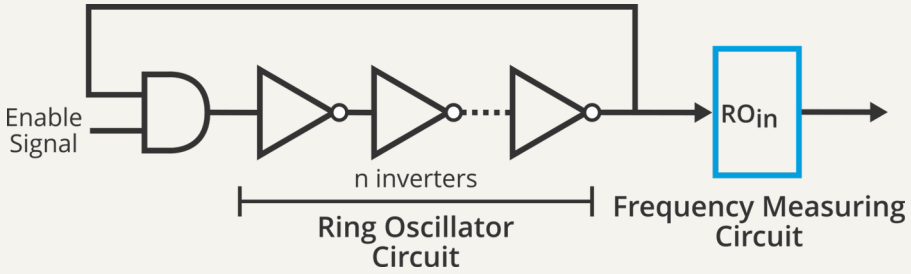

- Odd number of inverters in a closed loop.
- Signal to enable oscillation.
- Frequency **proportional** to propagation delay and size of the oscillator.

Ring Oscillator Sensor

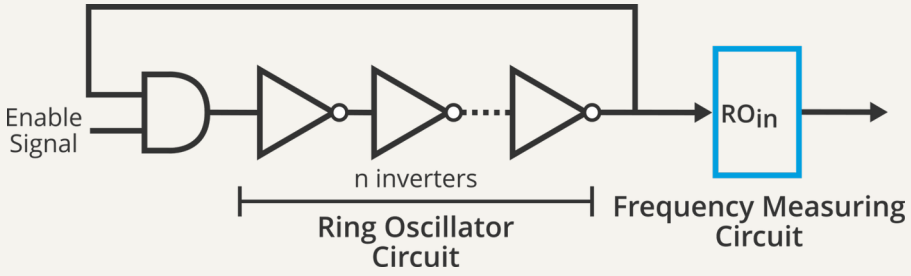

- Odd number of inverters in a closed loop.
- Signal to enable oscillation.
- Frequency **proportional** to propagation delay and size of the oscillator.

$$f = \frac{1}{2 \cdot n \cdot t_p}$$

Concept

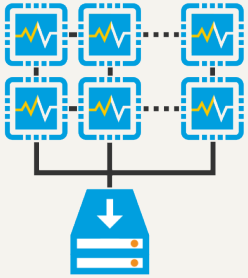

Online Self-Tests
Propagation Delay

Concept

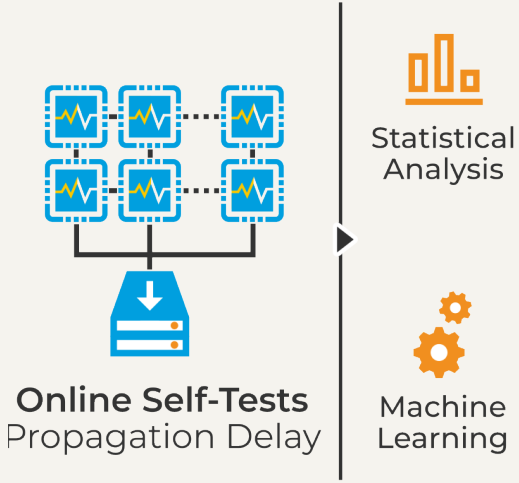
Online Self-Tests
Propagation Delay
Statistical Analysis
Machine Learning

Concept

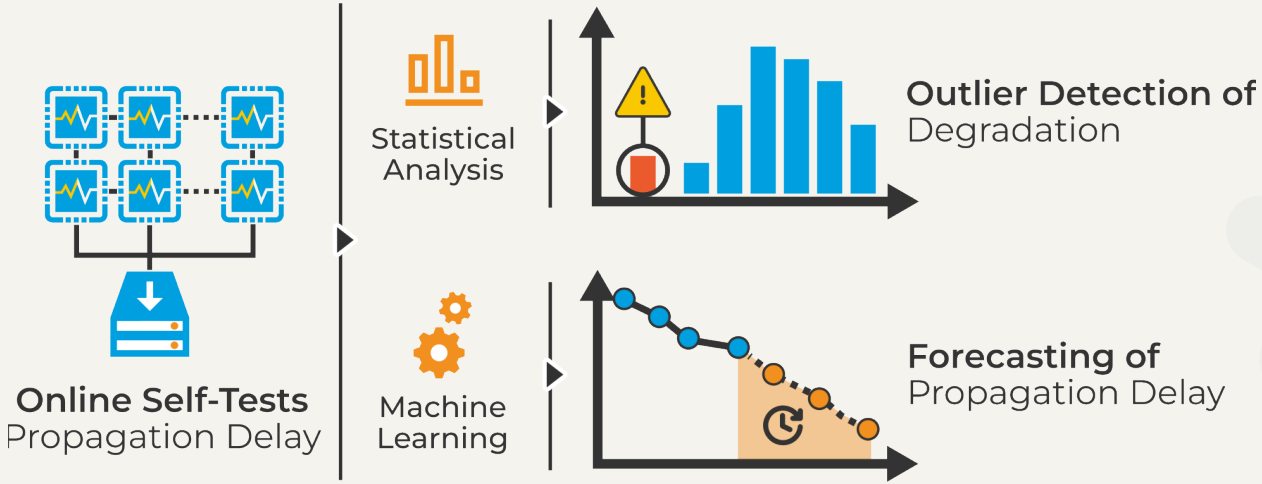
Online Self-Tests
Propagation Delay
Statistical
Analysis
Machine
Learning
Outlier Detection of
Degradation
Forecasting of
Propagation Delay

Concept

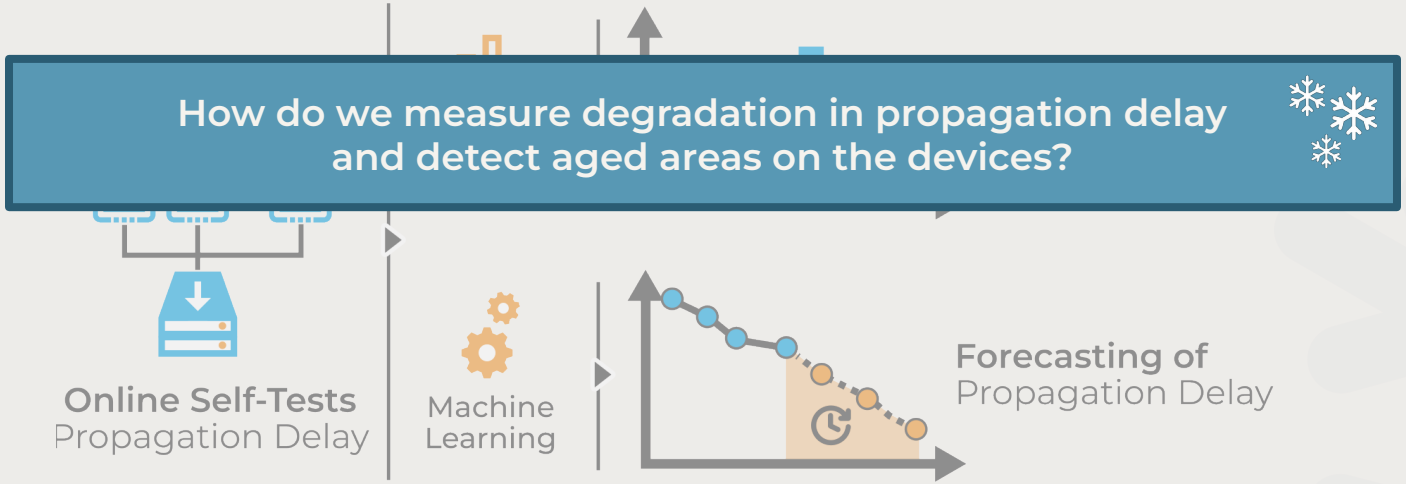
How do we measure degradation in propagation delay and detect aged areas on the devices?
Online Self-Tests
Propagation Delay
Machine Learning
Forecasting of
Propagation Delay

Concept


How do we measure degradation in propagation delay and detect aged areas on the devices?
How can we estimate the future propagation delay based on historical and environmental data?

Measuring the Propagation Delay

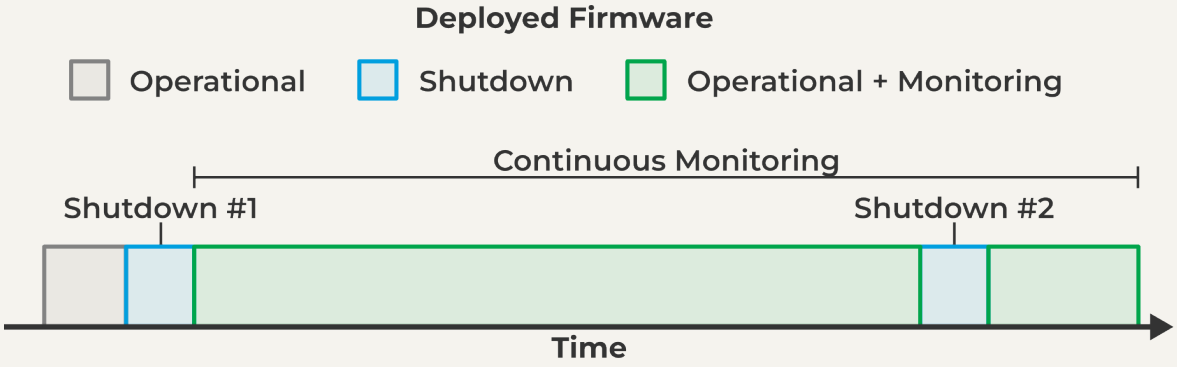

[3]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, and T. C. Schmidt. “Studying the Degradation of Propagation Delay on FPGAs at the European XFEL”, 2024. In Proceedings of the 27th Euromicro Conference on Digital System Design (DSD), Paris, France, 2024. DOI: 10.1109/DSD64264.2024.00018

Measuring the Propagation Delay

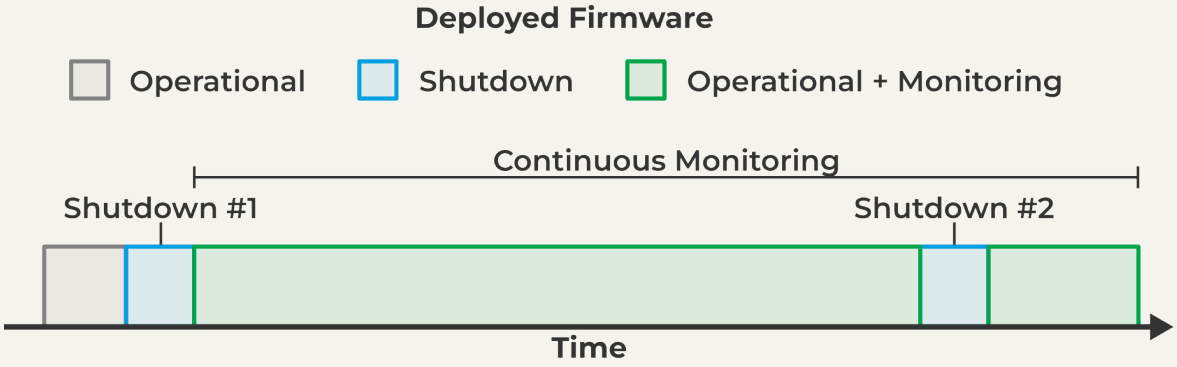

- **“High-resolution”** measurements during facility shutdowns (every 6 months).

[3]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, and T. C. Schmidt. “Studying the Degradation of Propagation Delay on FPGAs at the European XFEL”, 2024. In Proceedings of the 27th Euromicro Conference on Digital System Design (DSD), Paris, France, 2024. DOI: 10.1109/DSD64264.2024.00018

Measuring the Propagation Delay

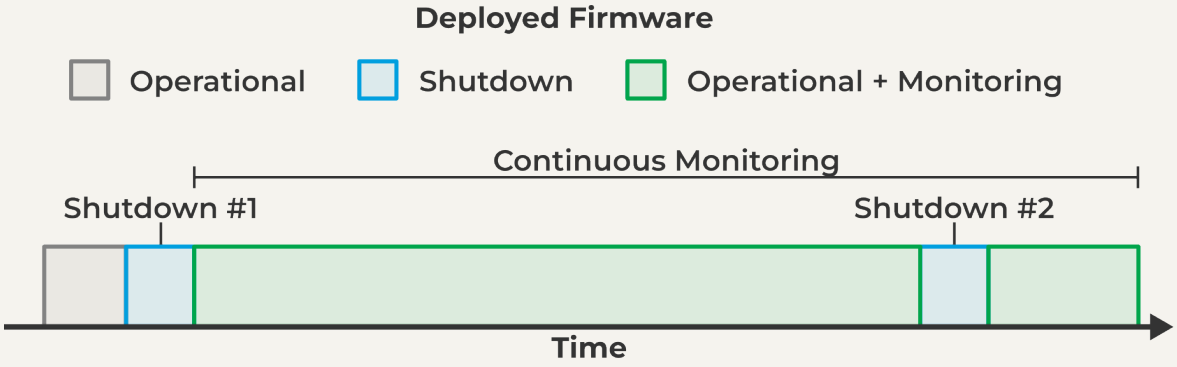

- **“High-resolution”** measurements during facility shutdowns (every 6 months).
- **“Low-resolution”** continuous measurements during facility operation (every 2 hours).

[3]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, and T. C. Schmidt. “Studying the Degradation of Propagation Delay on FPGAs at the European XFEL”, 2024. In Proceedings of the 27th Euromicro Conference on Digital System Design (DSD), Paris, France, 2024. DOI: 10.1109/DSD64264.2024.00018

European X-Ray Free-Electron Laser

European X-Ray Free-Electron Laser

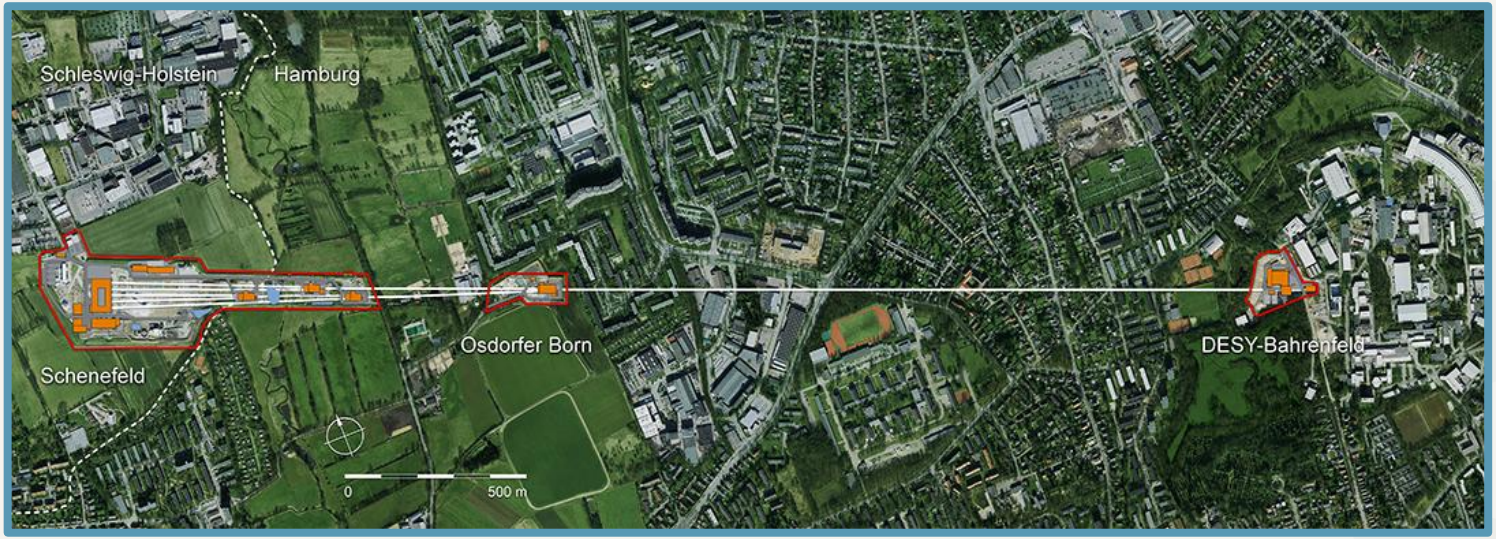

Monitoring Degradation of Propagation Delay on FPGAs

Shutdown Measurements

Detecting Degraded Regions

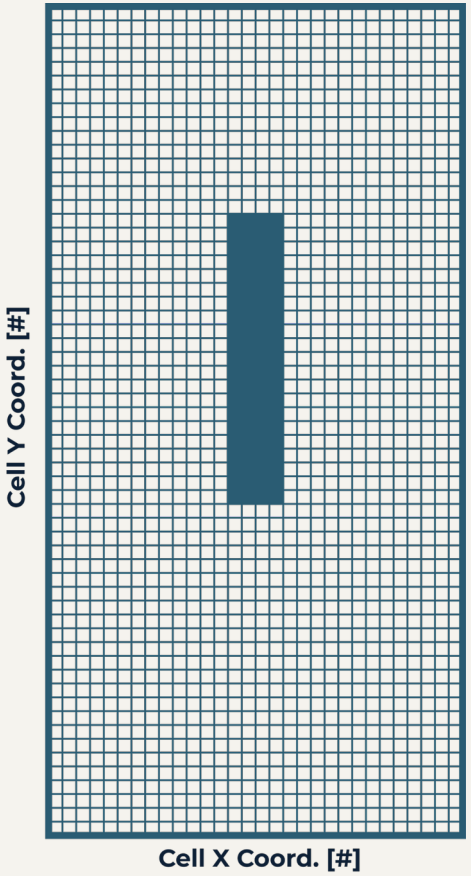

Detecting Degraded Regions

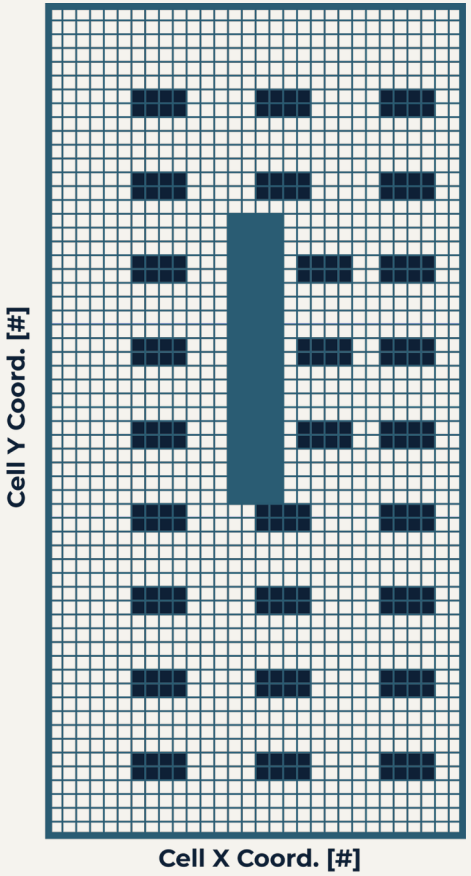

Detecting Degraded Regions

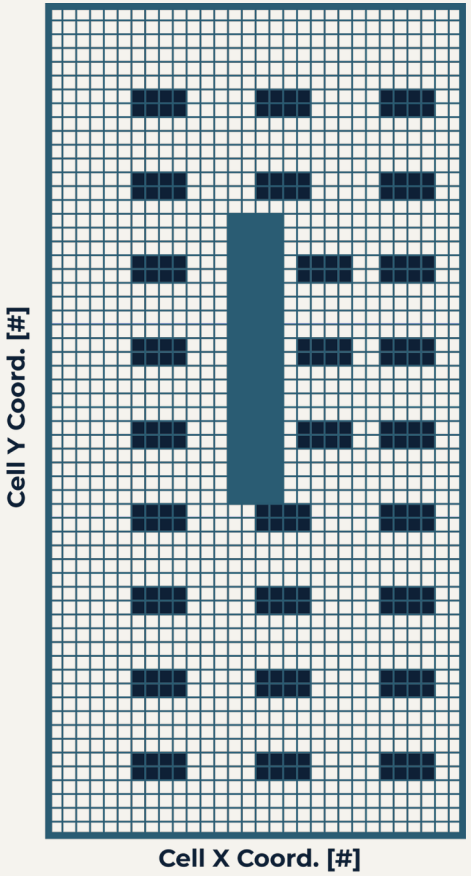

- Each RO covers a known region on the device.
- We compare the frequency of ROs between shutdowns.

Median Frequency Difference per Location

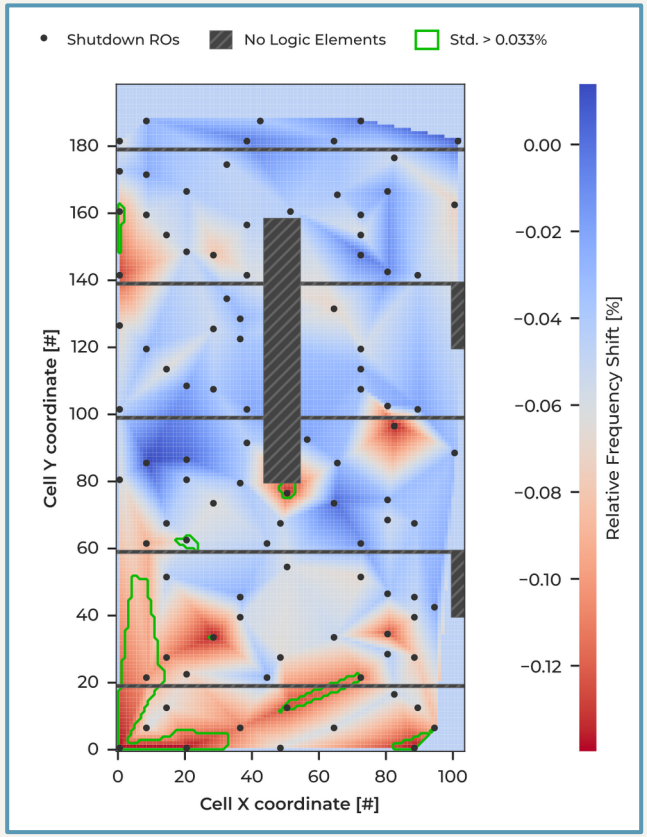

[4]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, T. C. Schmidt, “Switching Frequency as FPGA Monitor: Studying Degradation and Ageing Prognosis at Large Scale”, 2024. Preprint. DOI: 10.48550/arXiv.2412.15720

Median Frequency Difference per Location

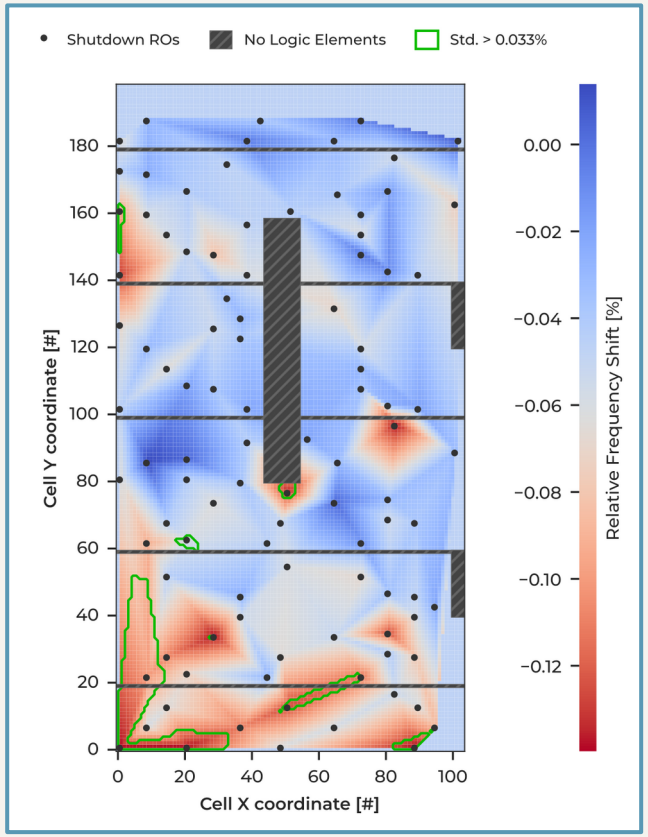

Certain areas present denser degradation.

[4]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, T. C. Schmidt, “Switching Frequency as FPGA Monitor: Studying Degradation and Ageing Prognosis at Large Scale”, 2024. Preprint. DOI: 10.48550/arXiv.2412.15720

Median Frequency Difference per Location

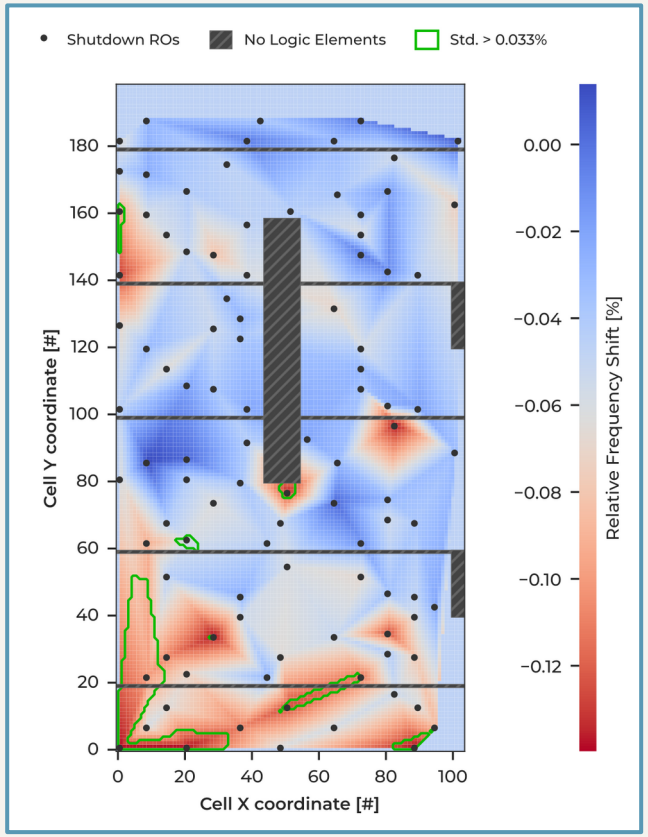

Certain areas present denser degradation.

Several hotspots with low deviation point to consistent degradation areas across devices.

[4]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, T. C. Schmidt, “Switching Frequency as FPGA Monitor: Studying Degradation and Ageing Prognosis at Large Scale”, 2024. Preprint. DOI: 10.48550/arXiv.2412.15720

Median Frequency Difference per Location

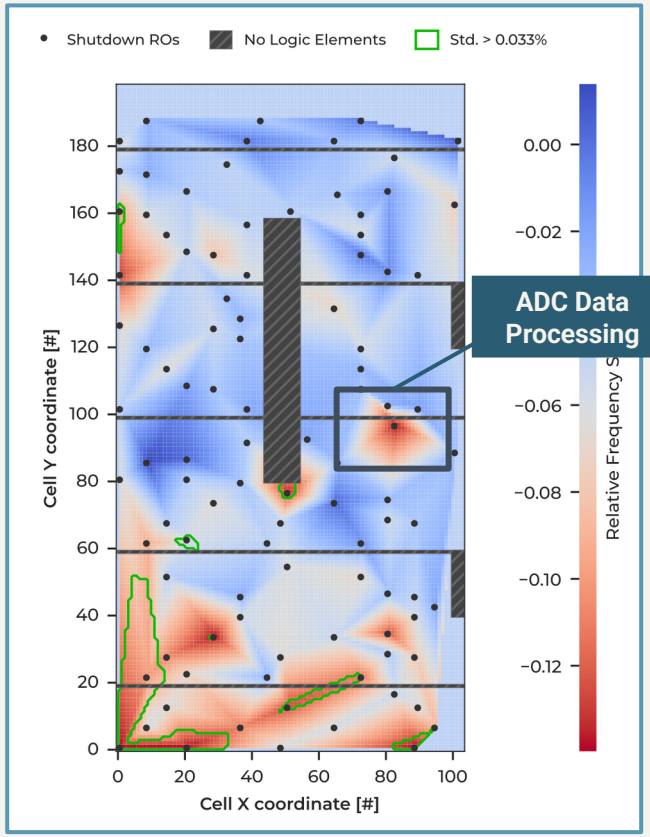

Certain areas present denser degradation.

Several hotspots with low deviation point to consistent degradation areas across devices.

[4]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, T. C. Schmidt, “Switching Frequency as FPGA Monitor: Studying Degradation and Ageing Prognosis at Large Scale”, 2024. Preprint. DOI: 10.48550/arXiv.2412.15720

Median Frequency Difference per Location

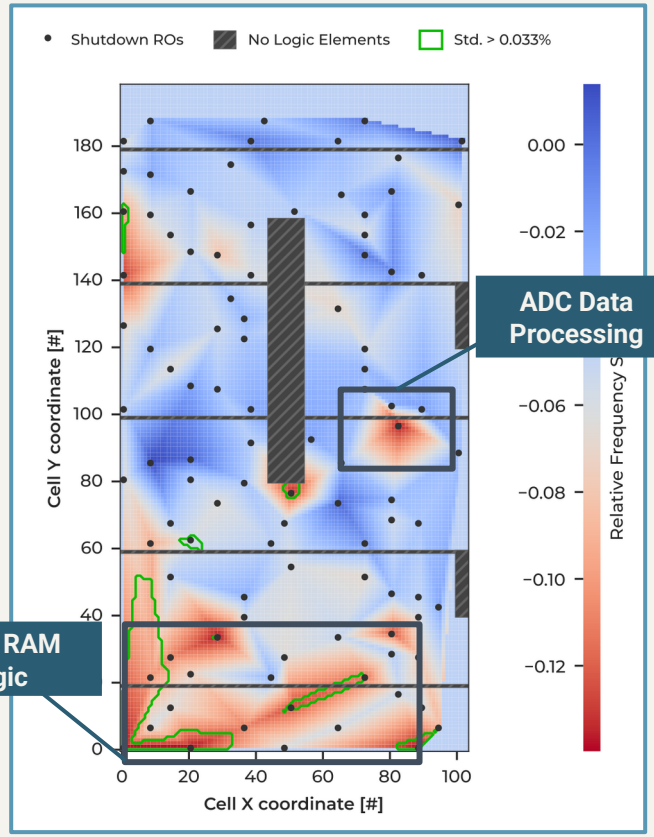

Certain areas present denser degradation.

Several hotspots with low deviation point to consistent degradation areas across devices.

[4]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, T. C. Schmidt, “Switching Frequency as FPGA Monitor: Studying Degradation and Ageing Prognosis at Large Scale”, 2024. Preprint. DOI: 10.48550/arXiv.2412.15720

Monitoring Degradation of Propagation Delay on FPGAs

Continuous Measurements

Frequency Shift of ‘Low-resolution’ measurements

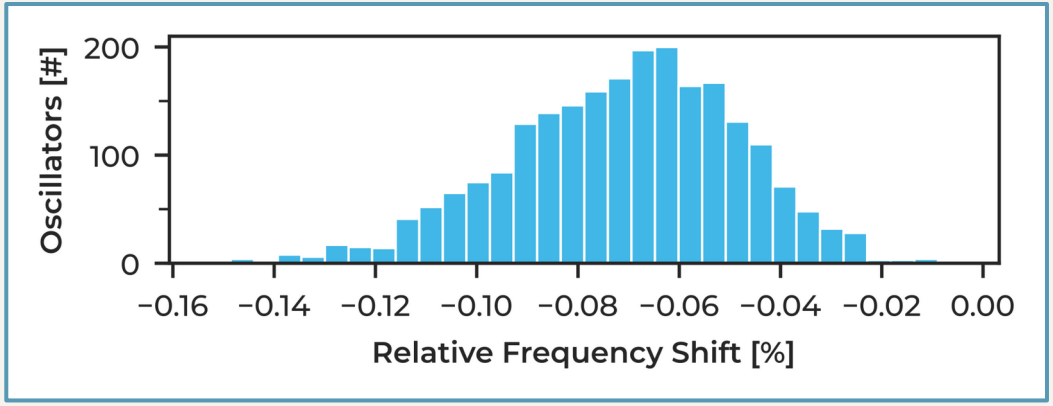

Frequency Shift of ‘Low-resolution’ measurements

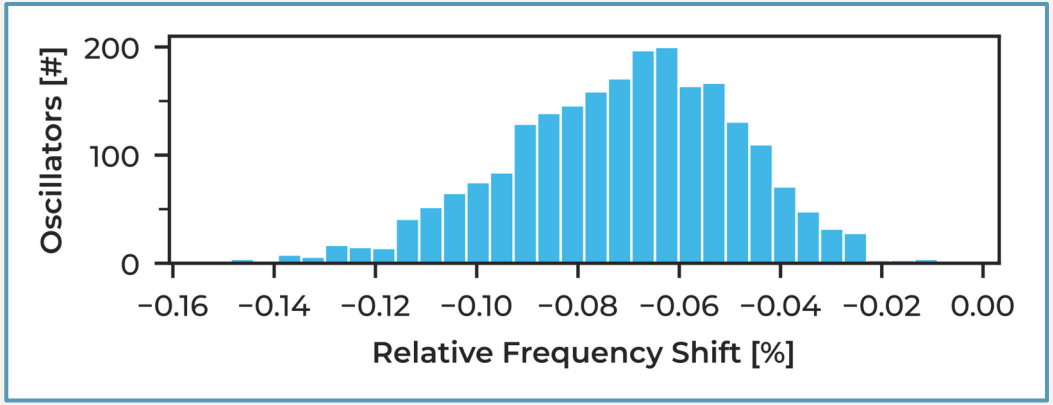

Median shift in RO frequency of -0.069%.

Frequency Shift of ‘Low-resolution’ measurements

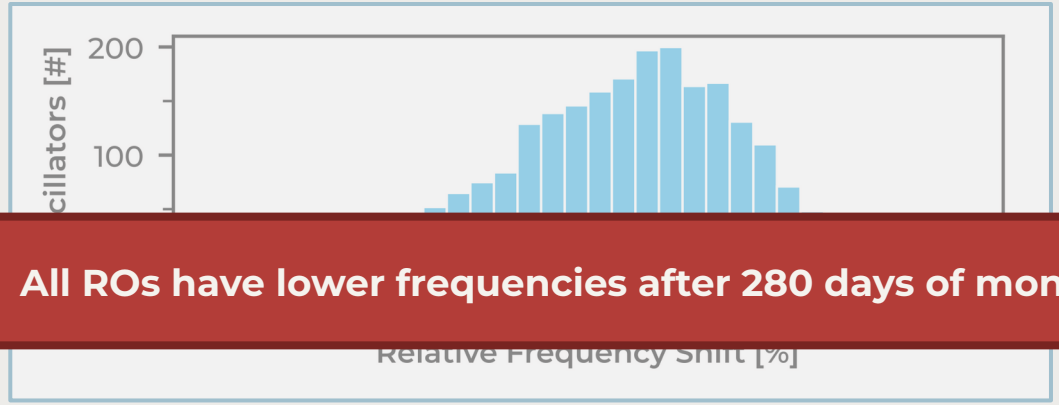

All ROs have lower frequencies after 280 days of monitoring.

Median shift in RO frequency of -0.069%.

Frequency Trends

Exponentially Weighted Moving Average

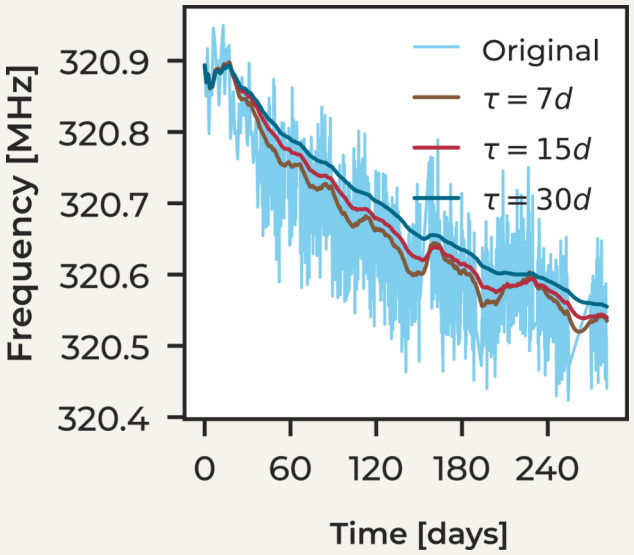

Seasonal-Trend Decomposition Based on LOESS

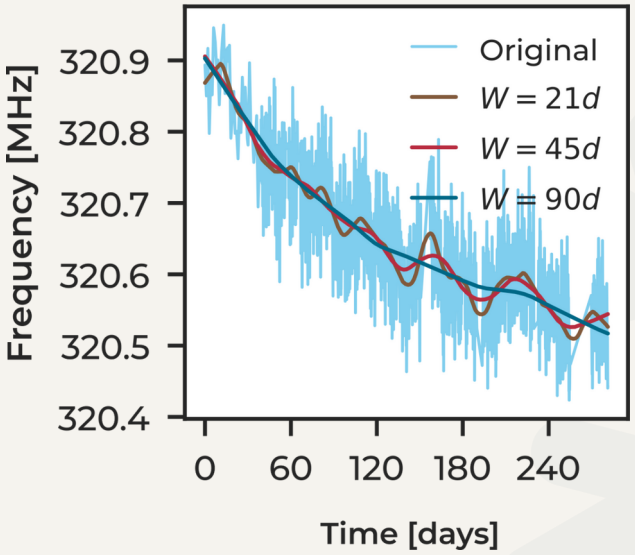

Frequency Trends

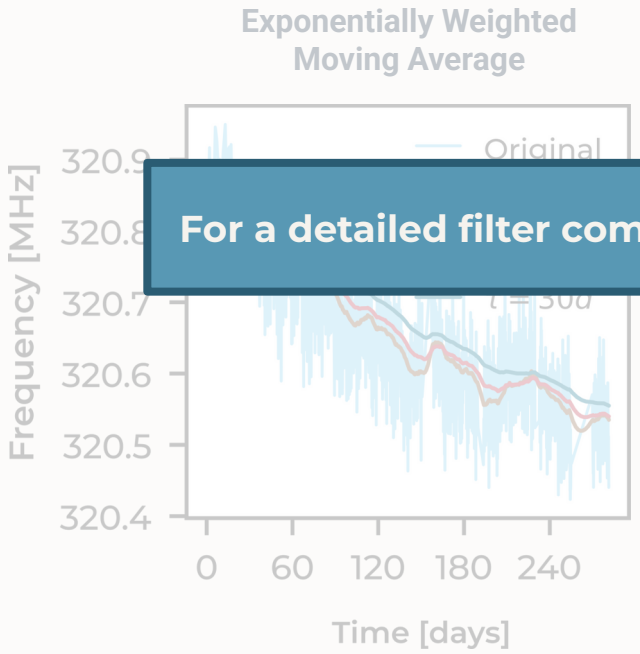

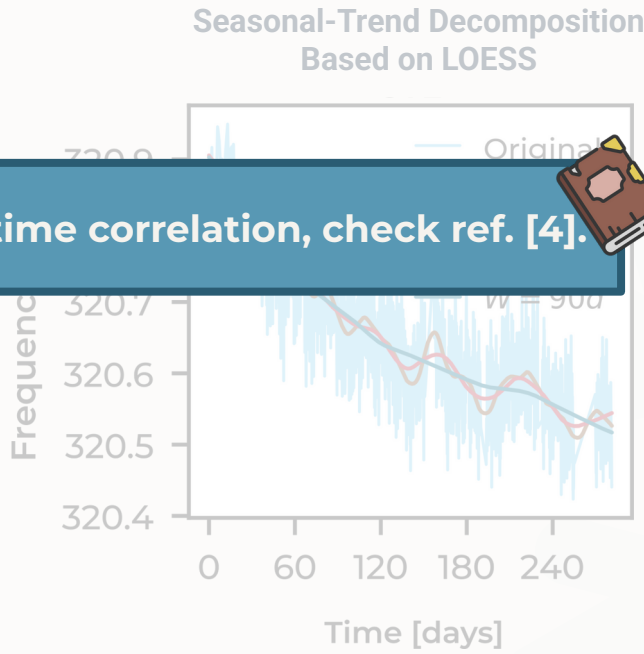

For a detailed filter comparison and time correlation, check ref. [4].

Frequency Trends

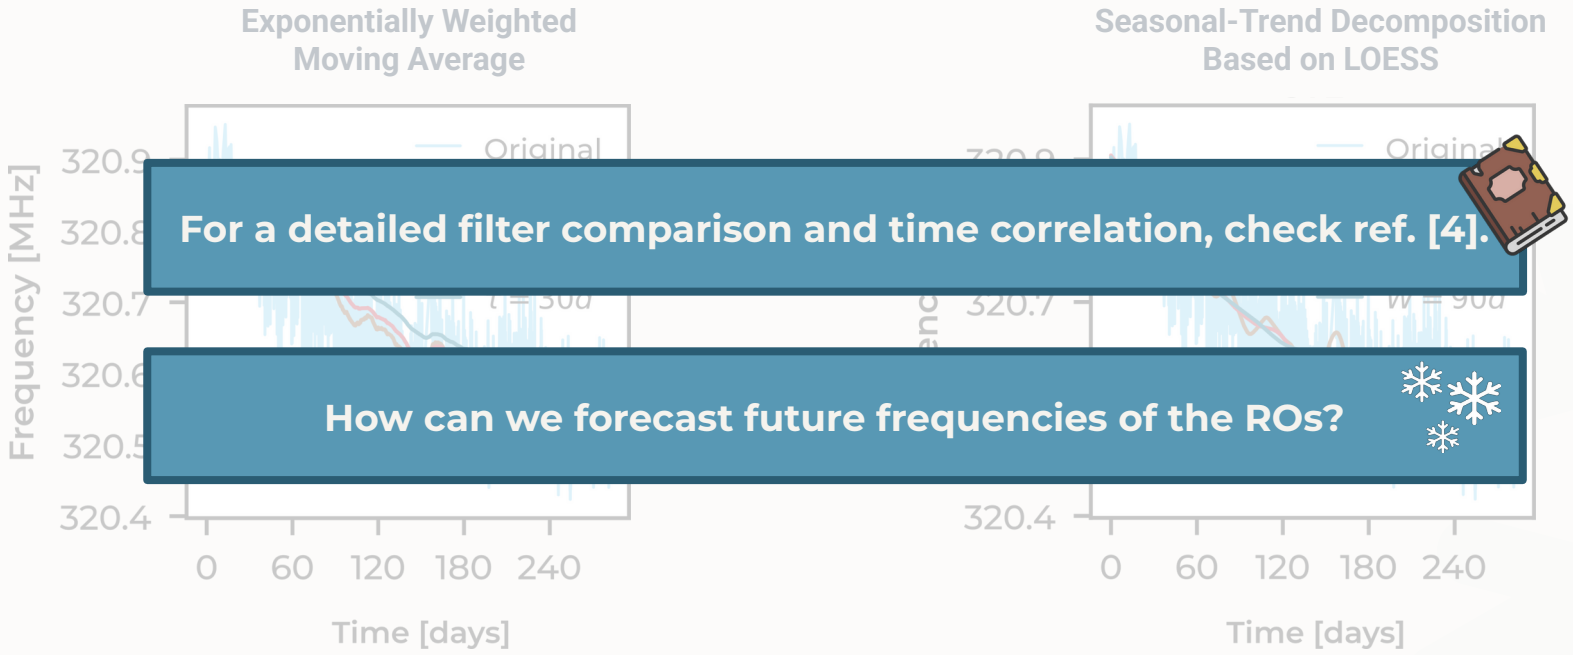

For a detailed filter comparison and time correlation, check ref. [4].

How can we forecast future frequencies of the ROs?

Forecasting Models for Propagation Delay

Forecasting Models for Propagation Delay

- Comparative evaluation of ML models on **forecasting RO frequencies.**

Forecasting Models for Propagation Delay

- Comparative evaluation of ML models on **forecasting RO frequencies.**
- Evaluate the effect of covariates in the training (temperature, voltage, current, etc.).

Forecasting Models for Propagation Delay

- Comparative evaluation of ML models on **forecasting RO frequencies.**
- Evaluate the effect of covariates in the training (temperature, voltage, current, etc.).

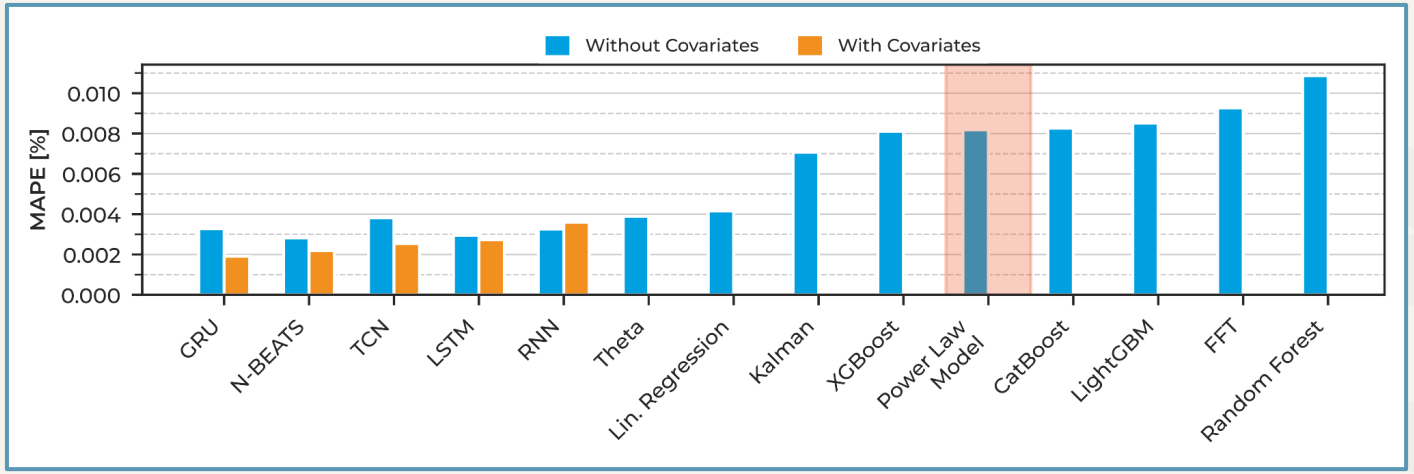

Forecasting Models for Propagation Delay

- Comparative evaluation of ML models on **forecasting RO frequencies.**
- Evaluate the effect of covariates in the training (temperature, voltage, current, etc.).

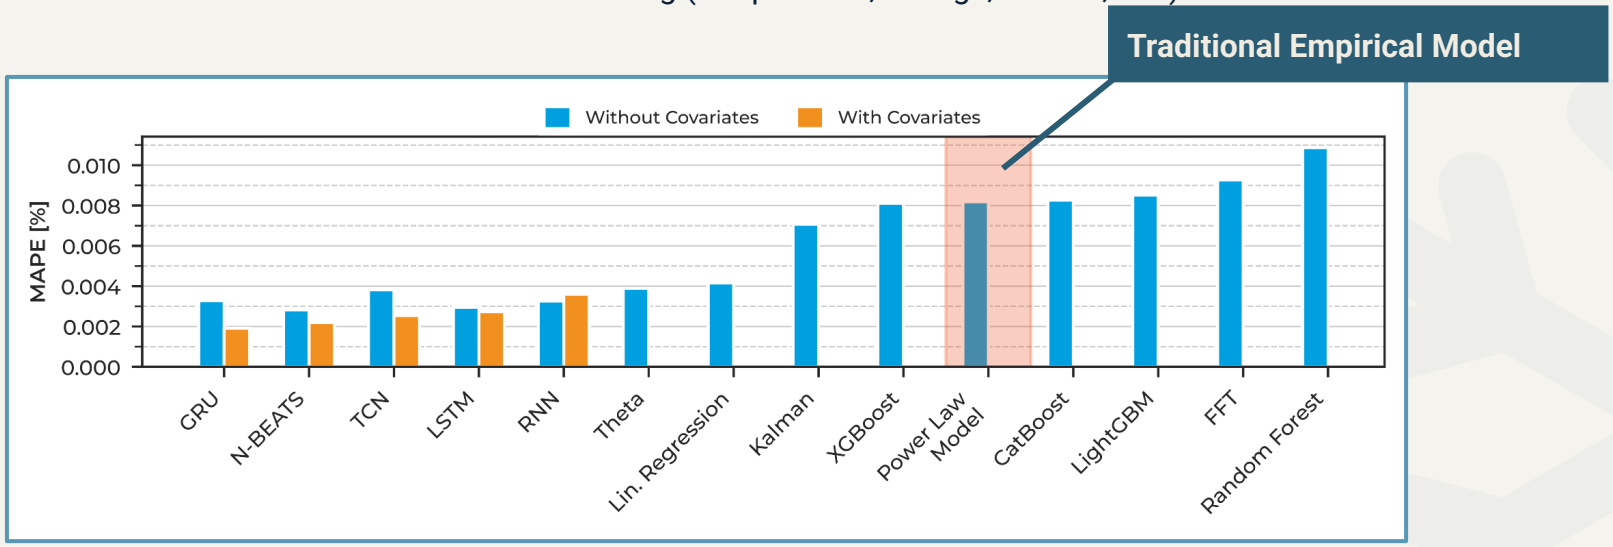

Backtesting

Backtesting

- Method to evaluate how models would have performed historically, had they been deployed.

Backtesting

- Method to evaluate how models would have performed historically, had they been deployed.
- The best non-DL model, **Theta**, is evaluated.

Backtesting

- Method to evaluate how models would have performed historically, had they been deployed.
- The best non-DL model, **Theta**, is evaluated.

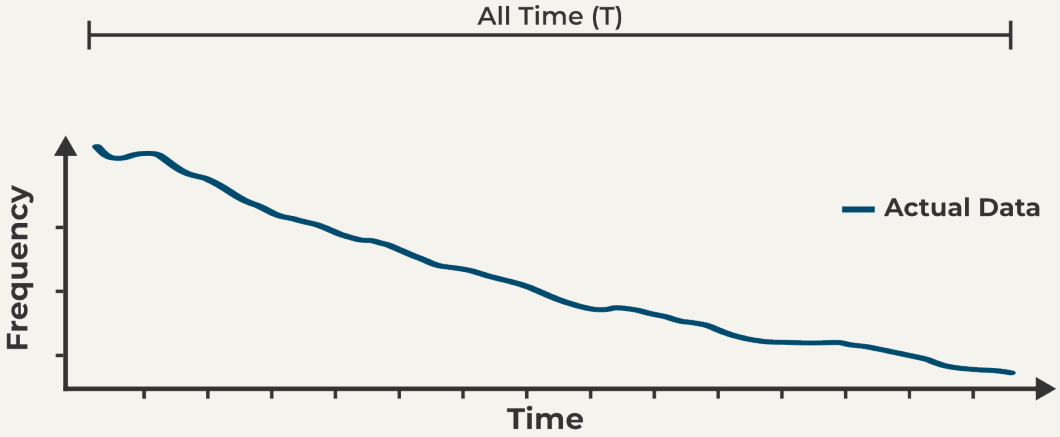

Backtesting

- Method to evaluate how models would have performed historically, had they been deployed.
- The best non-DL model, **Theta**, is evaluated.

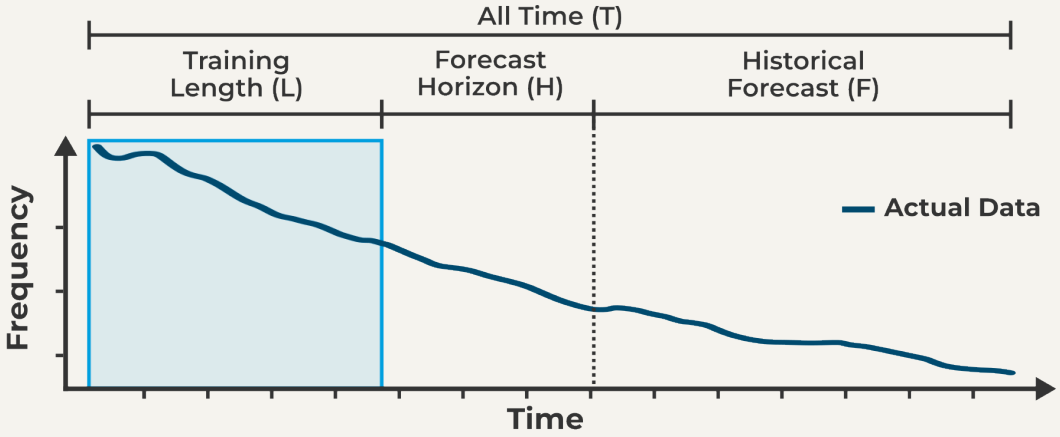

Backtesting

- Method to evaluate how models would have performed historically, had they been deployed.
- The best non-DL model, **Theta**, is evaluated.

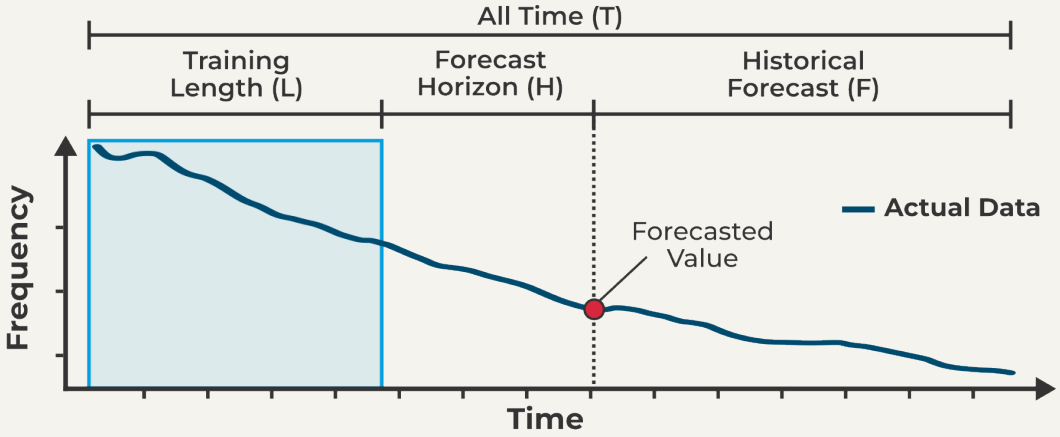

Backtesting

- Method to evaluate how models would have performed historically, had they been deployed.
- The best non-DL model, **Theta**, is evaluated.

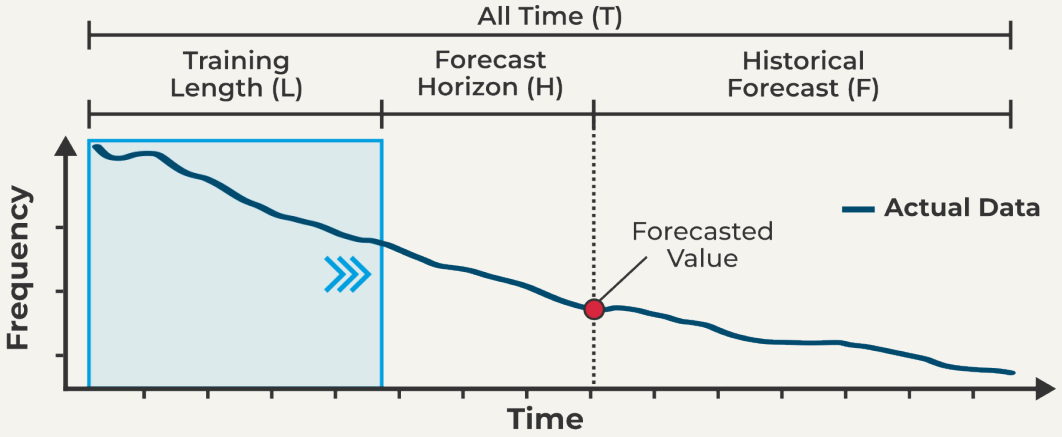

Backtesting

- Method to evaluate how models would have performed historically, had they been deployed.
- The best non-DL model, **Theta**, is evaluated.

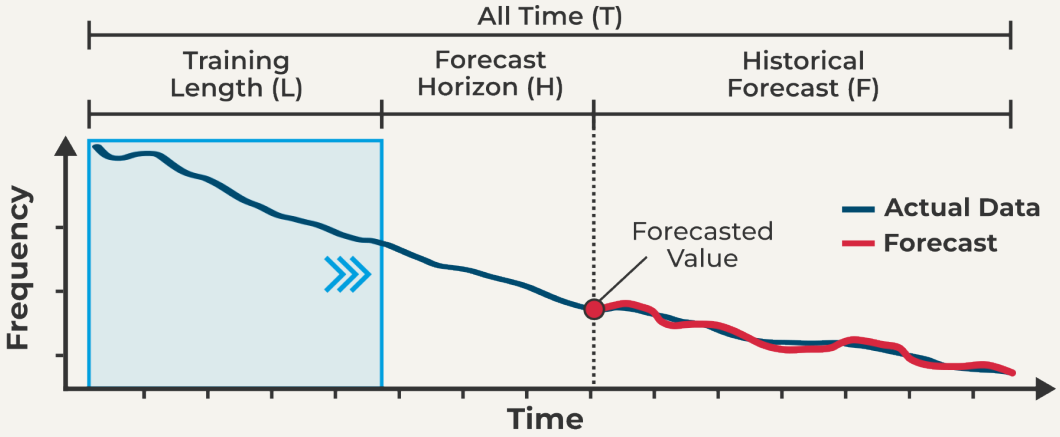

Backtesting for Predictive Maintenance

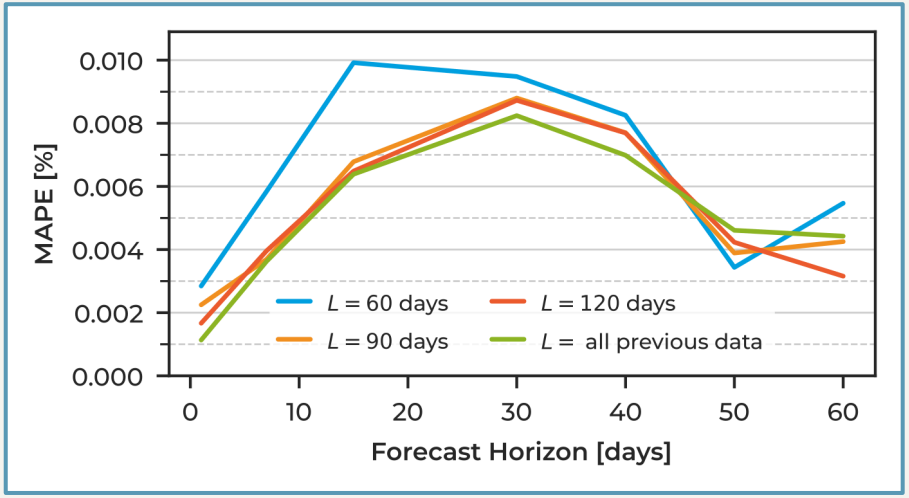

Backtesting for Predictive Maintenance

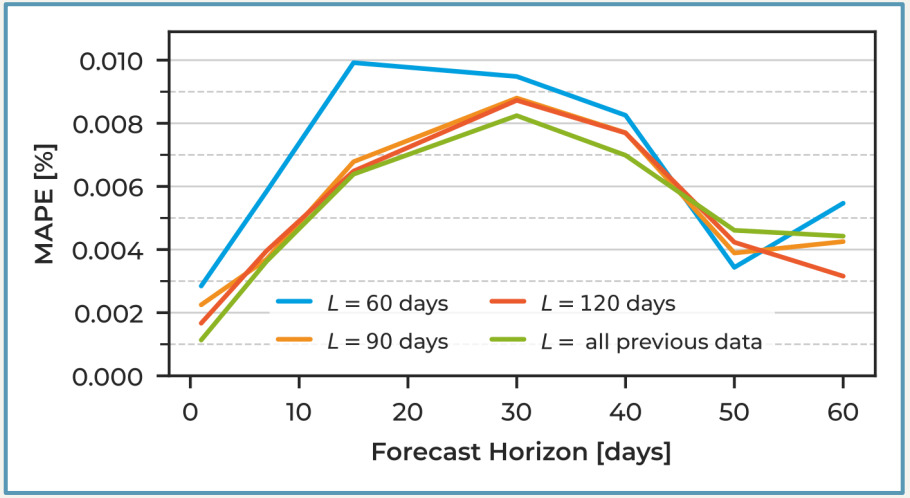

30-days breaking point due to leftover seasonality in the data?

Backtesting for Predictive Maintenance

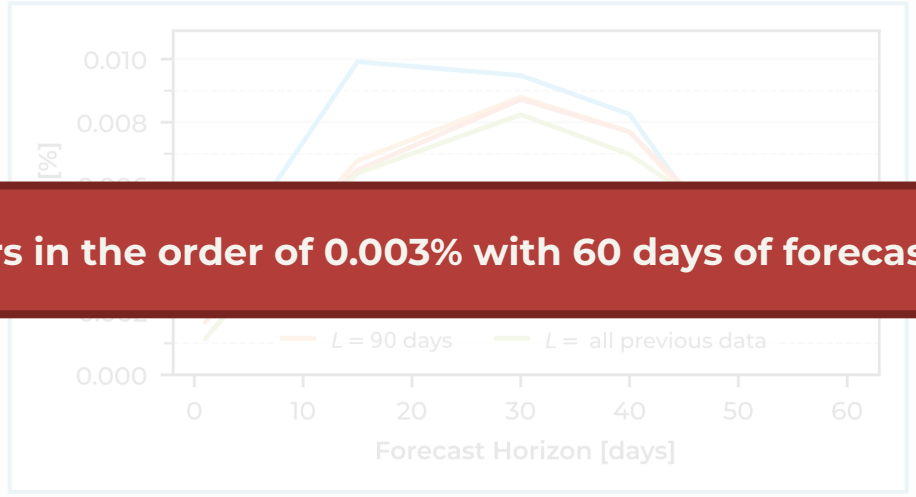

Errors in the order of 0.003% with 60 days of forecast horizon.

30-days breaking point due to leftover seasonality in the data?

Theta Model Backtest on all ROs

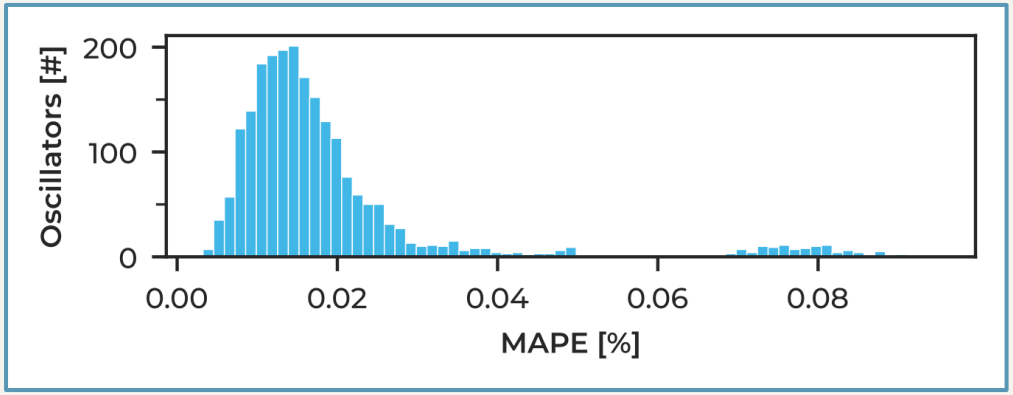

Theta Model Backtest on all ROs

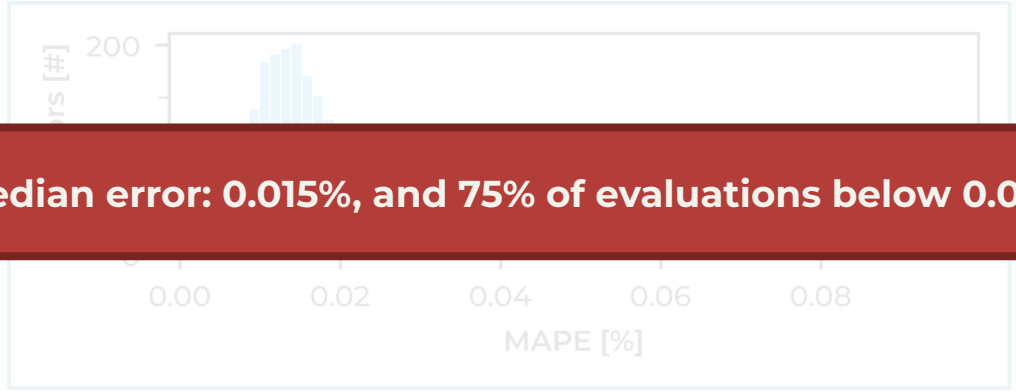

Median error: 0.015%, and 75% of evaluations below 0.020%.

In Summary

In Summary

- Devices underwent a noticeable degradation over the 280 days of study:

In Summary

- Devices underwent a noticeable degradation over the 280 days of study:
  - A median of **0.069% (and up to 0.25%) reduction** in RO oscillation frequency**.**

In Summary

- Devices underwent a noticeable degradation over the 280 days of study:
  - A median of **0.069% (and up to 0.25%) reduction** in RO oscillation frequency**.**
  - Areas with high-speed signals and heavy load showed increased degradation.

In Summary

- Devices underwent a noticeable degradation over the 280 days of study:
  - A median of **0.069% (and up to 0.25%) reduction** in RO oscillation frequency**.**
  - Areas with high-speed signals and heavy load showed increased degradation.

In Summary

- Devices underwent a noticeable degradation over the 280 days of study:
  - A median of **0.069% (and up to 0.25%) reduction** in RO oscillation frequency**.**
  - Areas with high-speed signals and heavy load showed increased degradation.

- The application of ML for predictive maintenance is possible:

In Summary

- Devices underwent a noticeable degradation over the 280 days of study:
  - A median of **0.069% (and up to 0.25%) reduction** in RO oscillation frequency**.**
  - Areas with high-speed signals and heavy load showed increased degradation.

- The application of ML for predictive maintenance is possible:
  - Forecasts with horizons of 60 days yield errors in the order of **0.003%**.

In Summary

- Devices underwent a noticeable degradation over the 280 days of study:
  - A median of **0.069% (and up to 0.25%) reduction** in RO oscillation frequency**.**
  - Areas with high-speed signals and heavy load showed increased degradation.

- The application of ML for predictive maintenance is possible:
  - Forecasts with horizons of 60 days yield errors in the order of **0.003%**.
  - ML models performed an order of magnitude better than the traditional power model.

Agenda

- Ageing Detection on Embedded SRAMs
- Monitoring Degradation of Propagation Delay on FPGAs
- **Ageing Monitoring for Commercial MCUs**

Timing Windows and Degradation

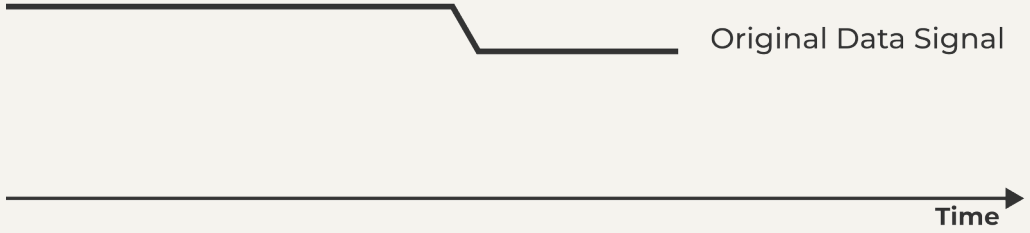

Timing Windows and Degradation

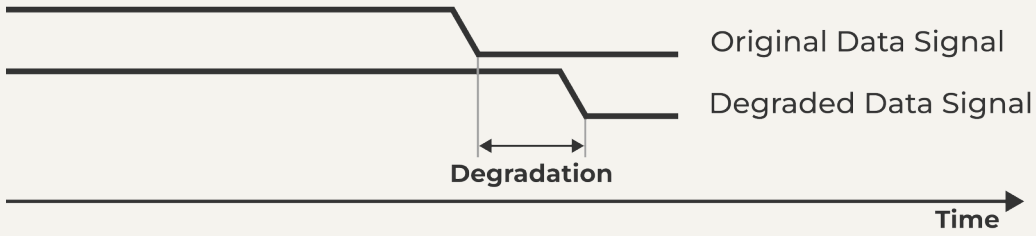

Timing Windows and Degradation

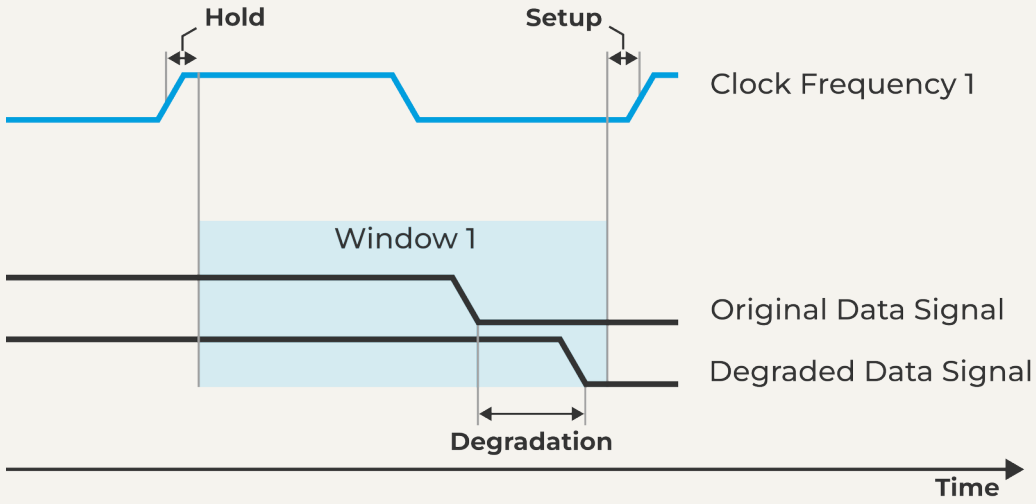

Timing Windows and Degradation

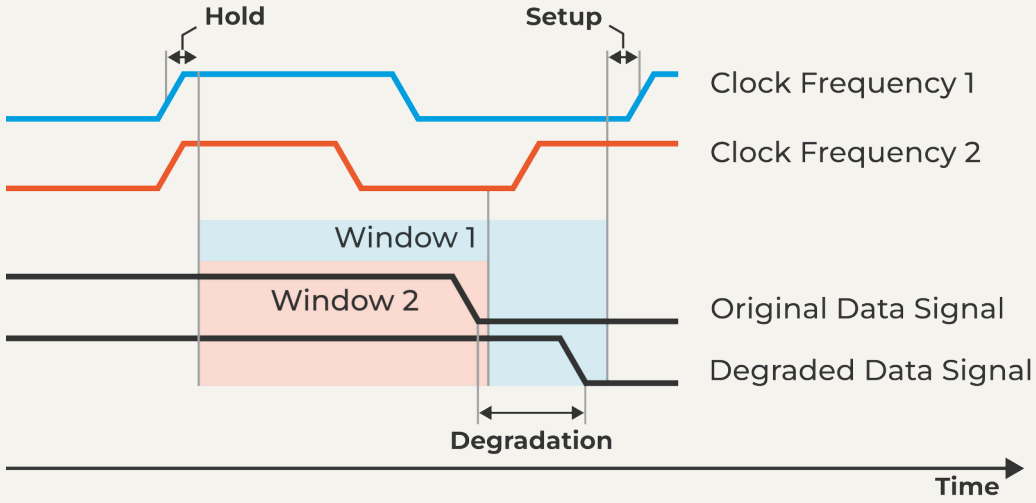

Timing Windows and Degradation

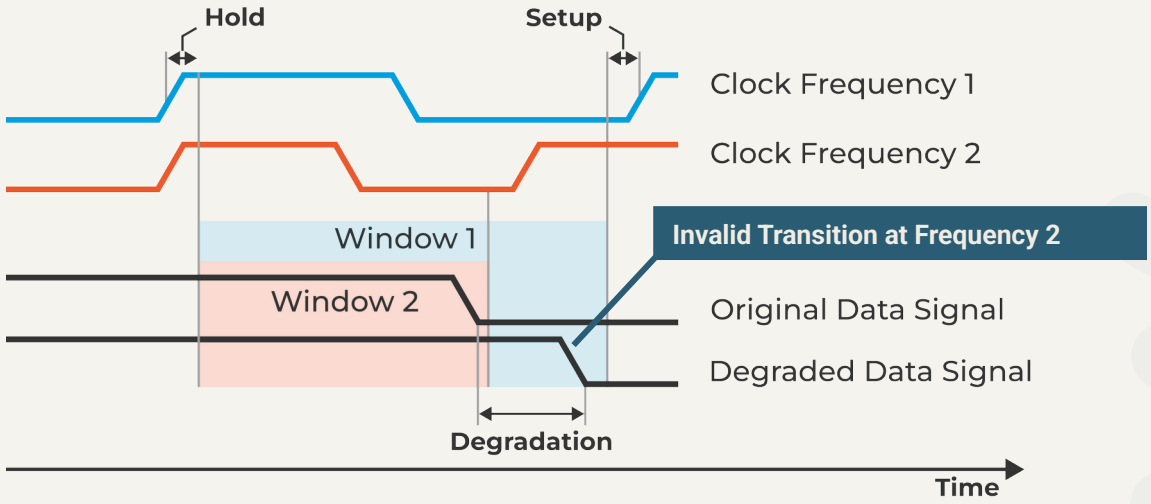

Our Goal

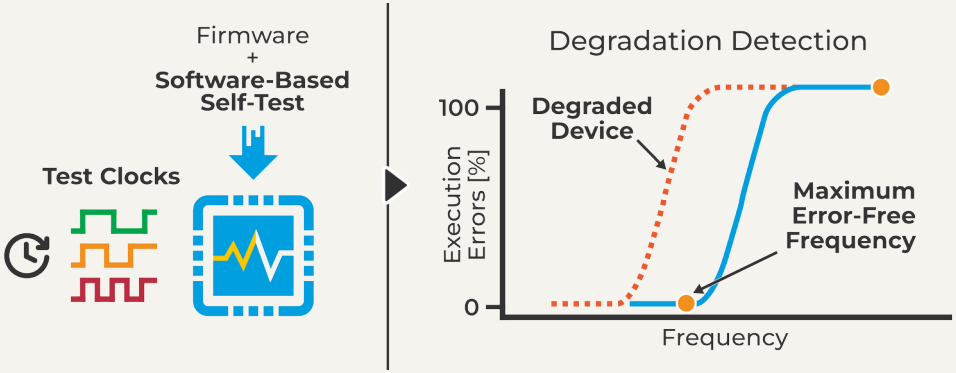
Firmware
+
Software-Based
Self-Test
Test Clocks
Degradation Detection
100
0
Execution
Errors [%]
Degraded
Device
Maximum
Error-Free
Frequency
Frequency

Our Goal

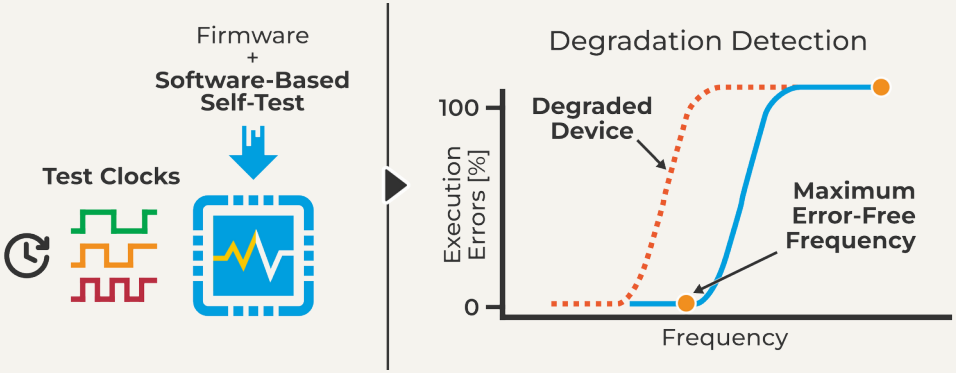

- Periodically execute a **software-based self-test** on MCUs to estimate ageing.
- Validate approach using high temperature as ageing proxy.
- Compare firmware computing payloads.

Experiment Setup

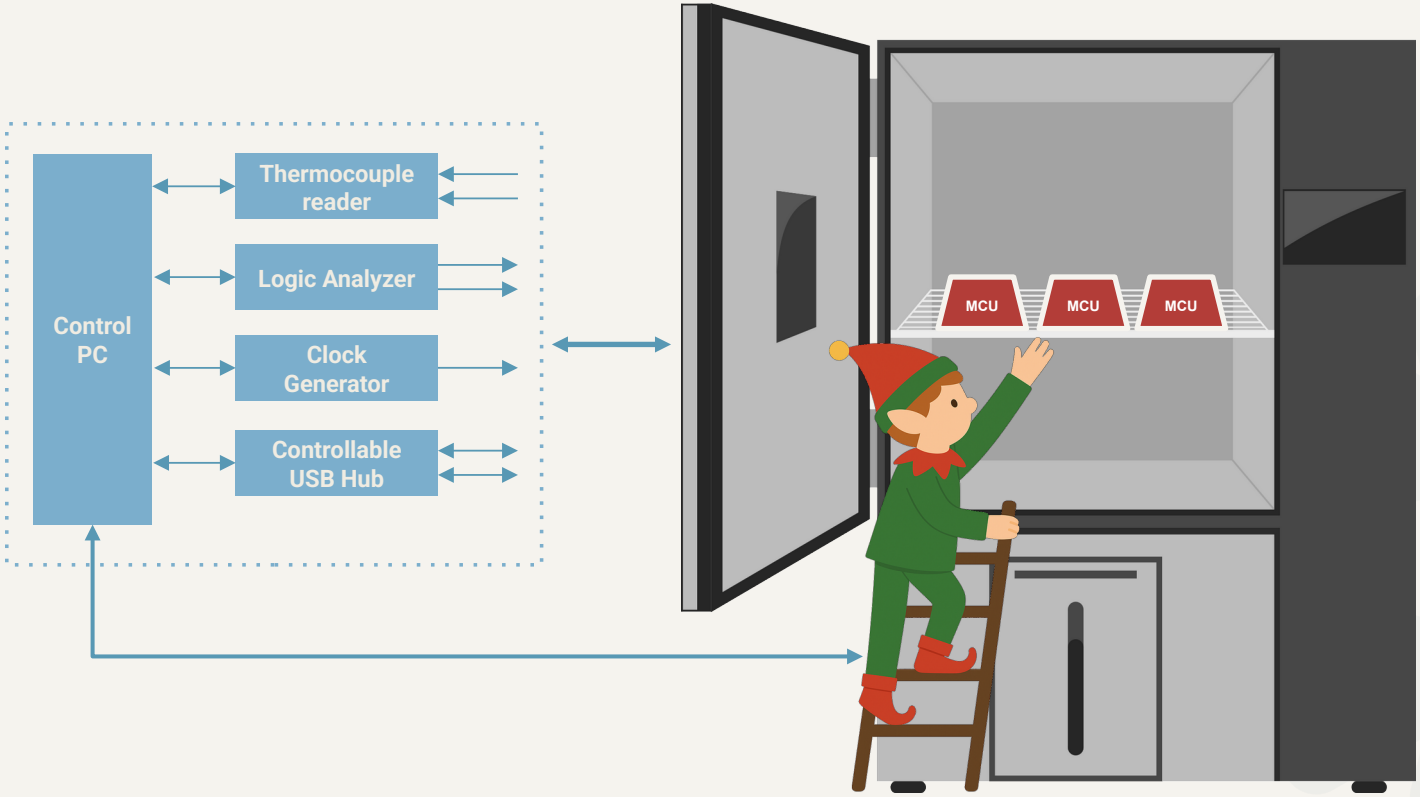

Experiment Setup

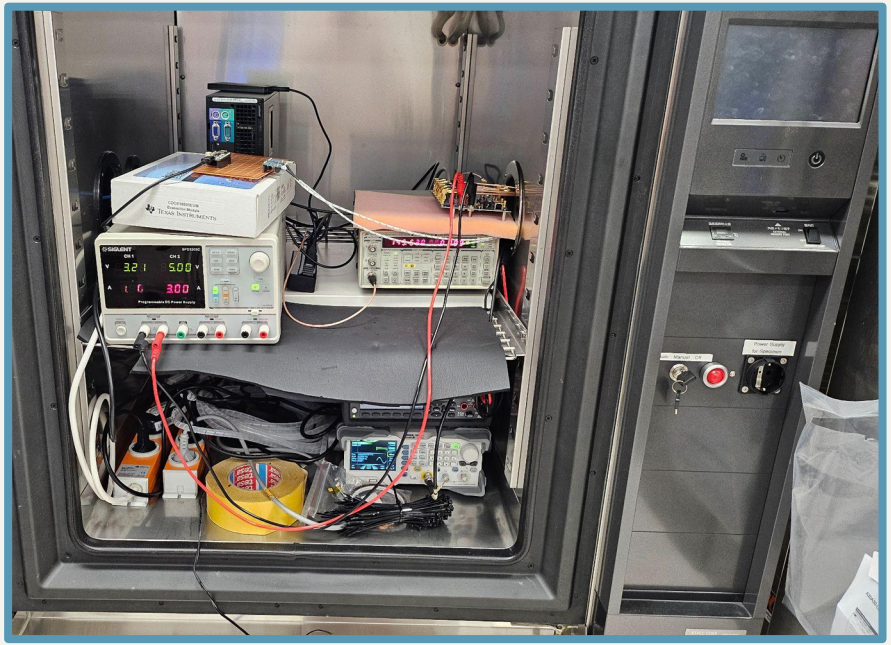

Experiment Setup

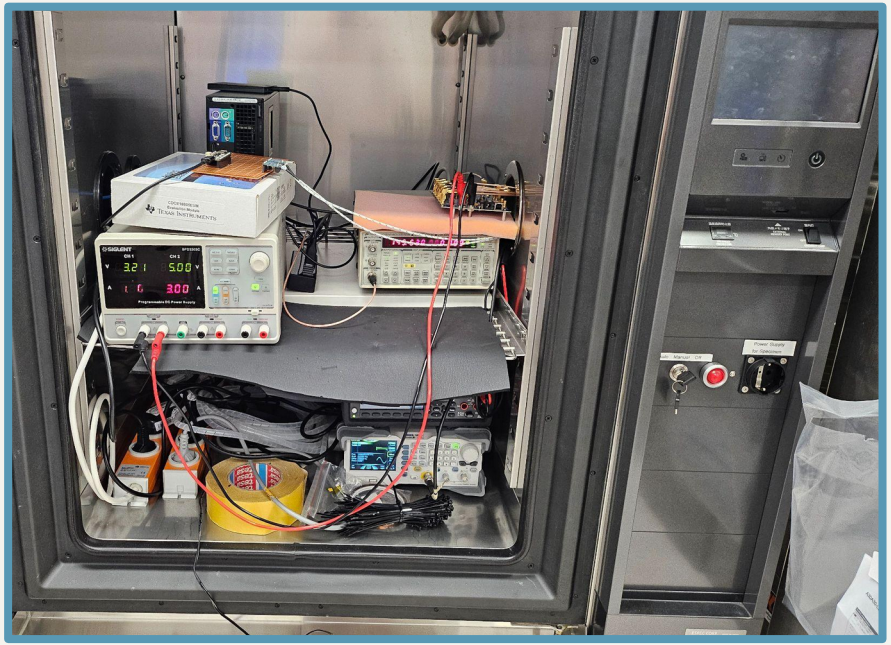

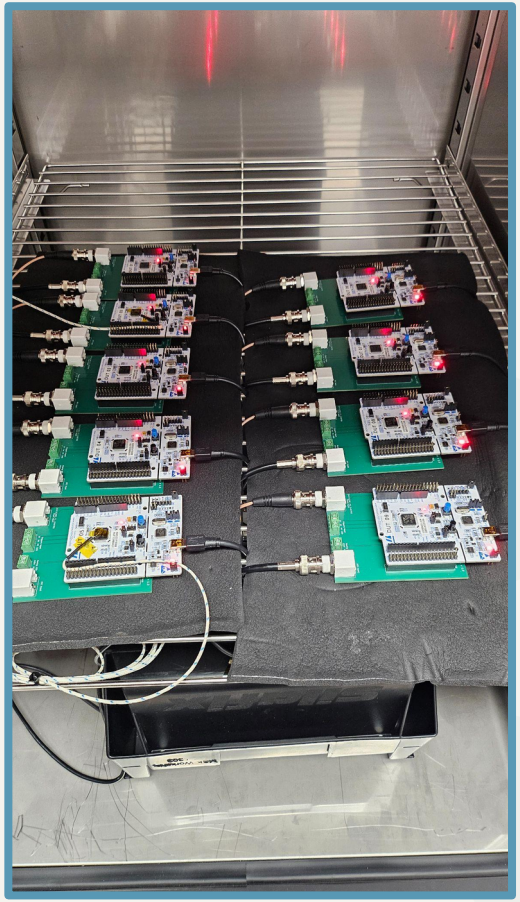

Evaluated Payloads and Configurations

Evaluated Payloads and Configurations

- **Payloads:**

Evaluated Payloads and Configurations

- **Payloads:**
  - Matrix Multiplication.

Evaluated Payloads and Configurations

- **Payloads:**
  - Matrix Multiplication.
  - Flash Read.

Evaluated Payloads and Configurations

- **Payloads:**
  - Matrix Multiplication.
  - Flash Read.
  - RAM Read and Write.

Evaluated Payloads and Configurations

- **Payloads:**
  - Matrix Multiplication.
  - Flash Read.
  - RAM Read and Write.
  - RAM March-C (code in RAM).

Evaluated Payloads and Configurations

- **Payloads:**
  - Matrix Multiplication.
  - Flash Read.
  - RAM Read and Write.
  - RAM March-C (code in RAM).
  - CPU Vendor Test (code in RAM).

Evaluated Payloads and Configurations

- **Payloads:**
  - Matrix Multiplication.
  - Flash Read.
  - RAM Read and Write.
  - RAM March-C (code in RAM).
  - CPU Vendor Test (code in RAM).

Evaluated Payloads and Configurations

- **Payloads:**
  - Matrix Multiplication.
  - Flash Read.
  - RAM Read and Write.
  - RAM March-C (code in RAM).
  - CPU Vendor Test (code in RAM).

- **Configuration Parameters:**

Evaluated Payloads and Configurations

- **Payloads:**
  - Matrix Multiplication.
  - Flash Read.
  - RAM Read and Write.
  - RAM March-C (code in RAM).
  - CPU Vendor Test (code in RAM).

- **Configuration Parameters:**
  - "*Buffered*": Pre-fetch buffer enabled + 2 Flash wait states.

Evaluated Payloads and Configurations

- **Payloads:**
  - Matrix Multiplication.
  - Flash Read.
  - RAM Read and Write.
  - RAM March-C (code in RAM).
  - CPU Vendor Test (code in RAM).

- **Configuration Parameters:**
  - "*Buffered*": Pre-fetch buffer enabled + 2 Flash wait states.
  - "*Unbuffered*": Pre-fetch buffered disabled + 0 Flash wait states.

Relative Degradation of MEF with Temperature

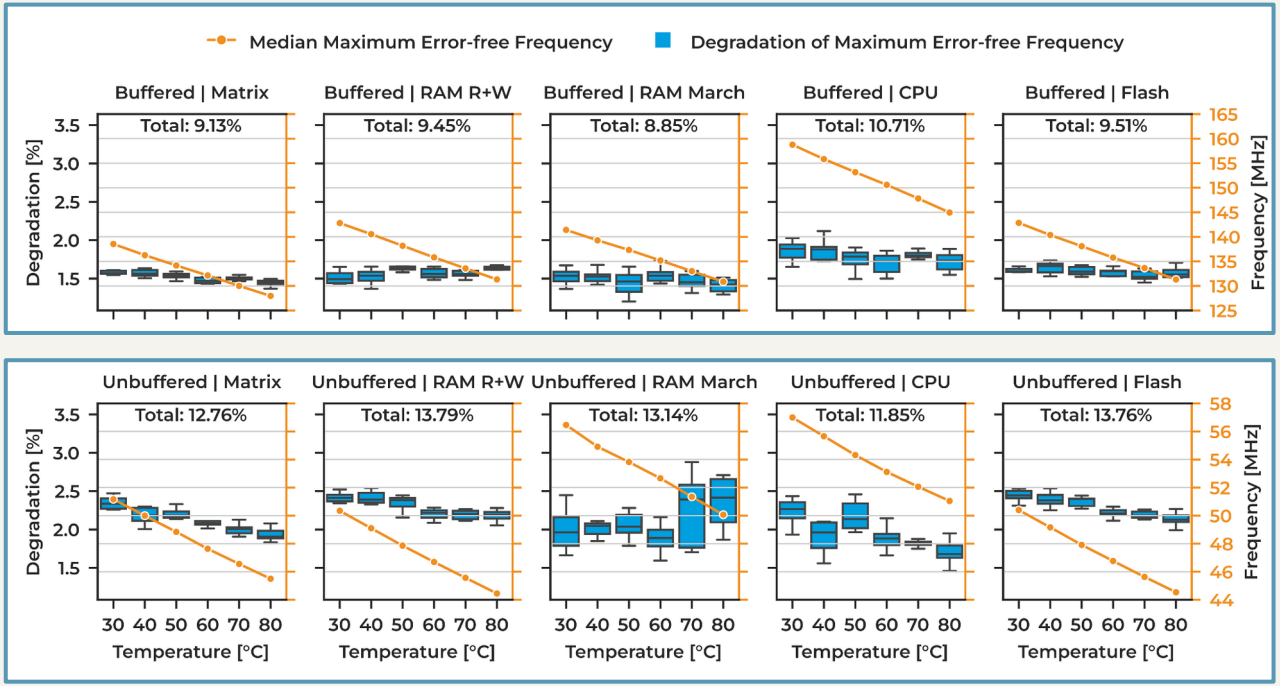

Relative Degradation of MEF with Temperature

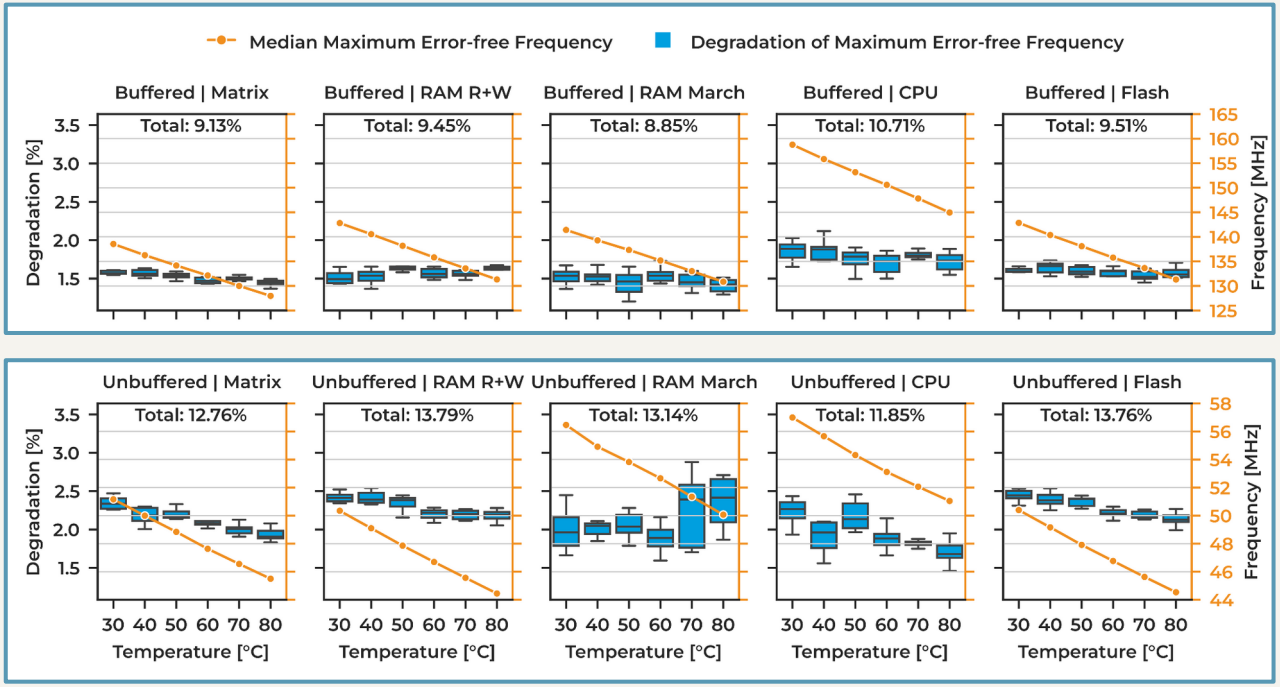

Unbuffered: payloads with heavy memory usage fail at lower frequencies.

Buffered: matrix multiplication fails at lower frequencies. Possibly due to non-sequential memory access and ALU usage.

In general: ~2% of MEF degradation per 10 ◦C.

Transition in the Number of Errors

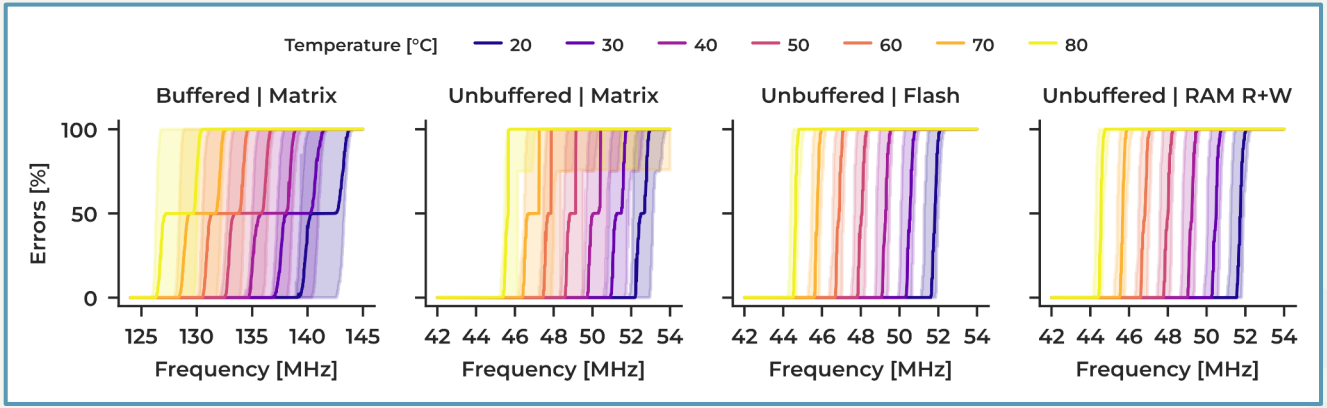

- Errors of 100% correspond to all 500 payload executions failing.
- Solid lines are median values, and shades are the interquartile range.

Transition in the Number of Errors

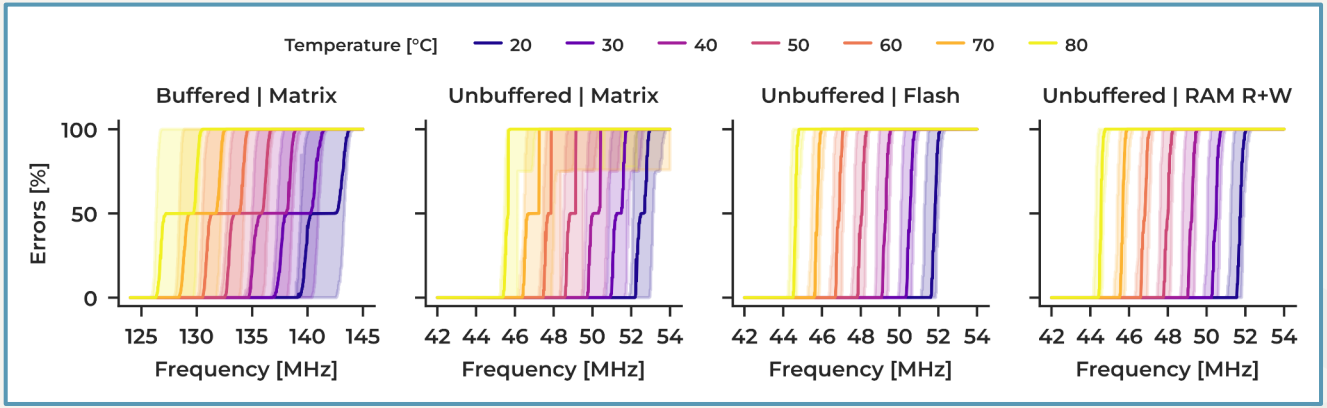

- Errors of 100% correspond to all 500 payload executions failing.
- Solid lines are median values, and shades are the interquartile range.

Flash and RAM R+W have smaller spreads than Matrix payload.

Matrix is the only payload with transitions on both configurations.

Qualitative Comparison of Payloads

Which one gets more cookies?

Payload	MEF	Execution Time	Error Transition
CPU			
RAM March			
RAM R+W			
Flash			
Matrix			

Qualitative Comparison of Payloads

Which one gets more cookies?

Payload	MEF	Execution Time	Error Transition
CPU			
RAM March			
RAM R+W			
Flash			
Matrix			

Qualitative Comparison of Payloads

Which one gets more cookies?

Payload	MEF	Execution Time	Error Transition
CPU			
RAM March			
RAM R+W			
Flash			
Matrix			

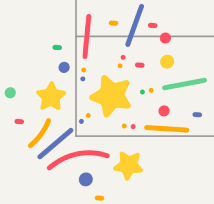

Qualitative Comparison of Payloads

Which one gets more cookies?

Payload	MEF	Execution Time	Error Transition
RAM March			
RAM R+W			
Flash			
Matrix			

On the wishlist: automate the generation of payloads

Agenda

- Ageing Detection on Embedded SRAMs
- Monitoring Degradation of Propagation Delay on FPGAs
- Ageing Monitoring for Commercial MCUs

Automatic Self-Test Generation for Ageing Detection on Commercial MCUs

Test Generation Process

Test Generation Process

Initial Payload Population

Test Generation Process

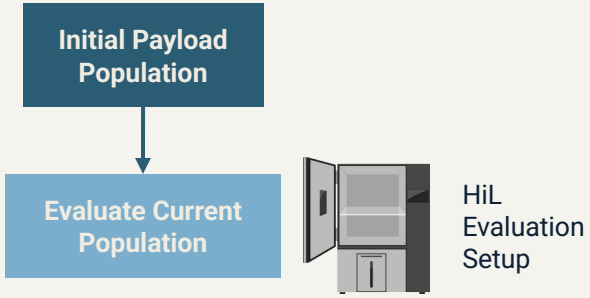

Test Generation Process

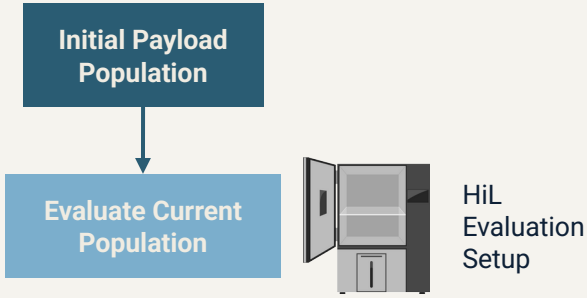

- For each firmware, find MEF at two temperatures.
- Measure MEF degradation and execution time.

Test Generation Process

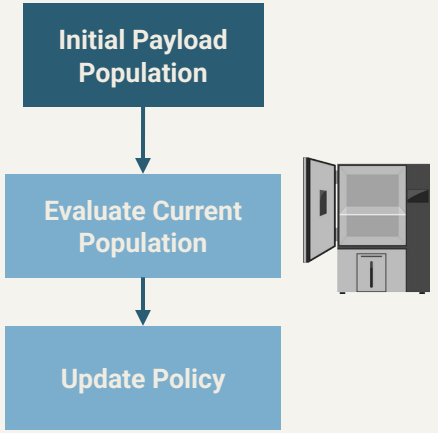

HiL Evaluation Setup

- For each firmware, find MEF at two temperatures.
- Measure MEF degradation and execution time.

Test Generation Process

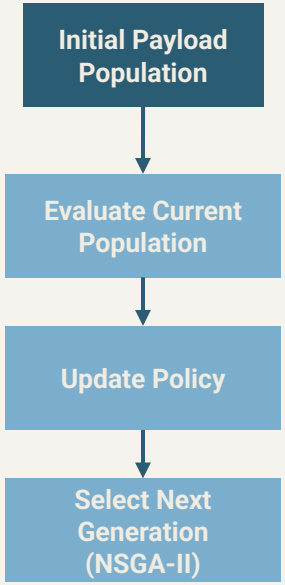

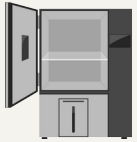

HiL Evaluation Setup

- For each firmware, find MEF at two temperatures.
- Measure MEF degradation and execution time.

Test Generation Process

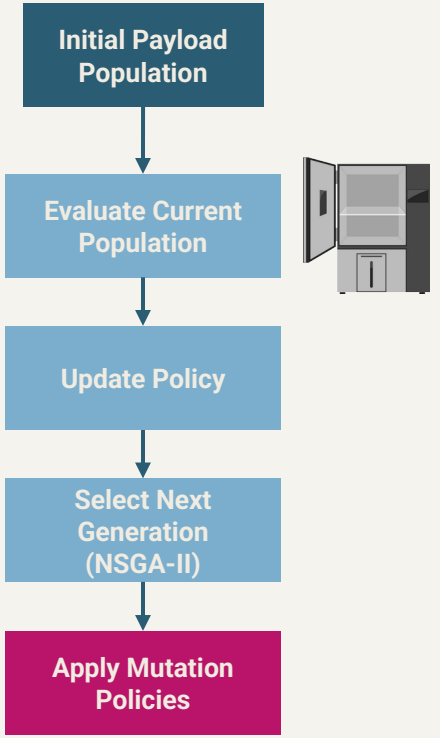

HiL Evaluation Setup

- For each firmware, find MEF at two temperatures.
- Measure MEF degradation and execution time.

Test Generation Process

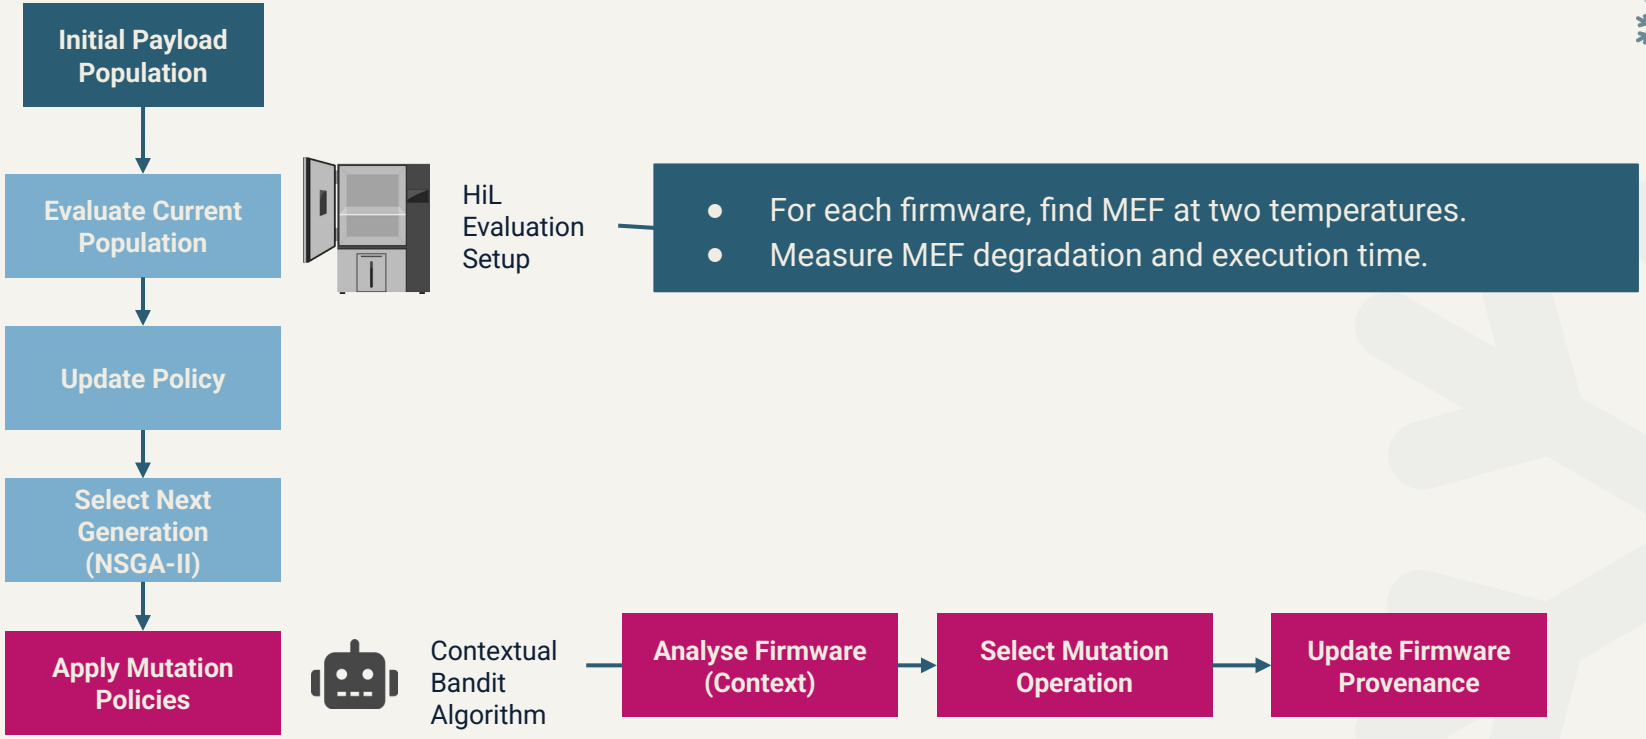

Test Generation Process

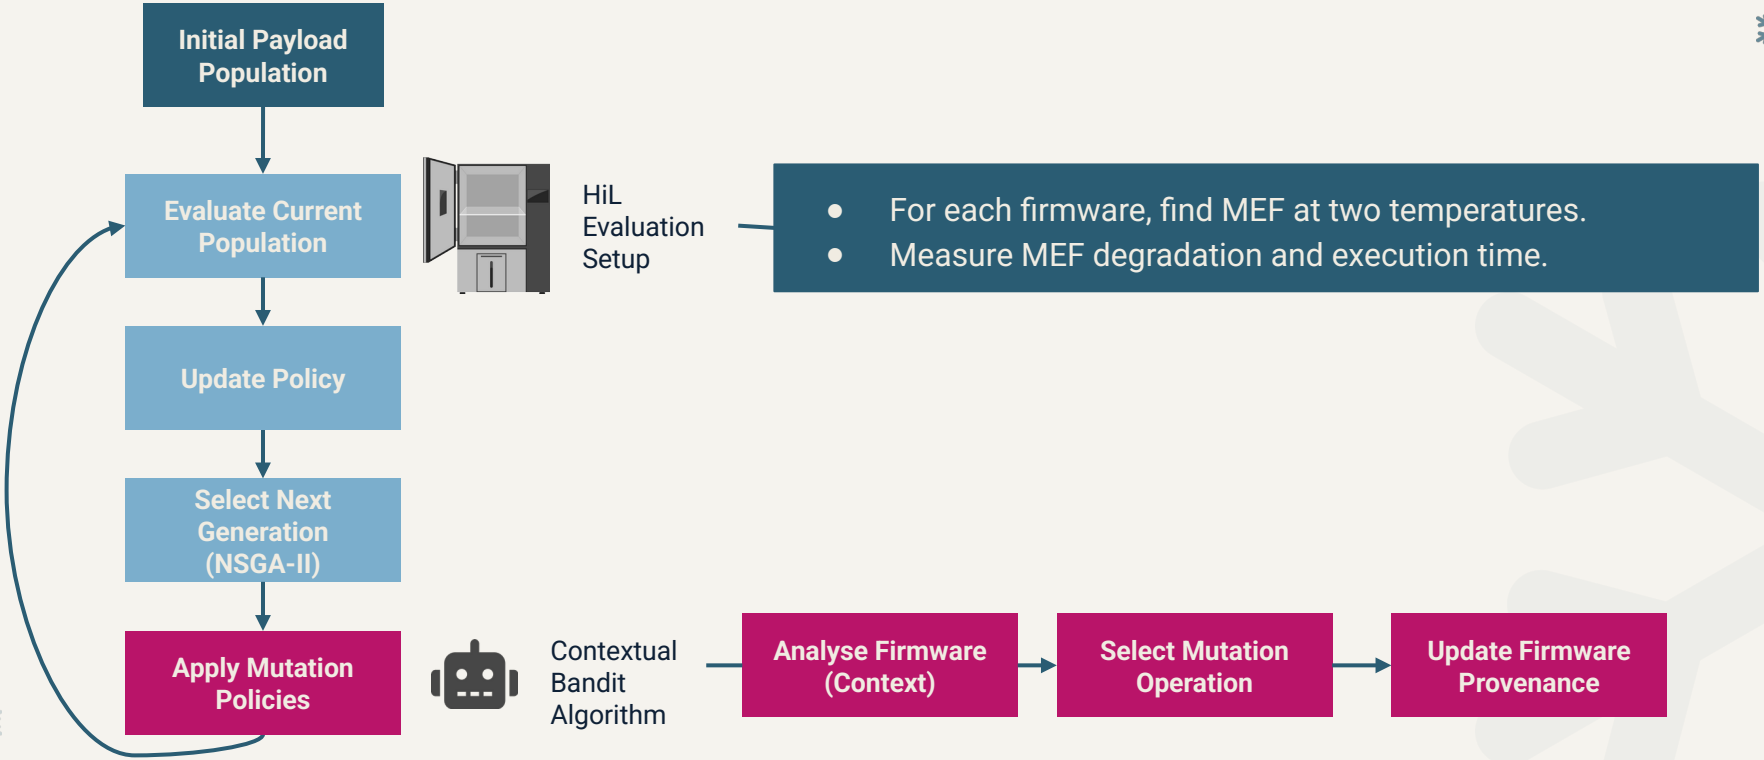

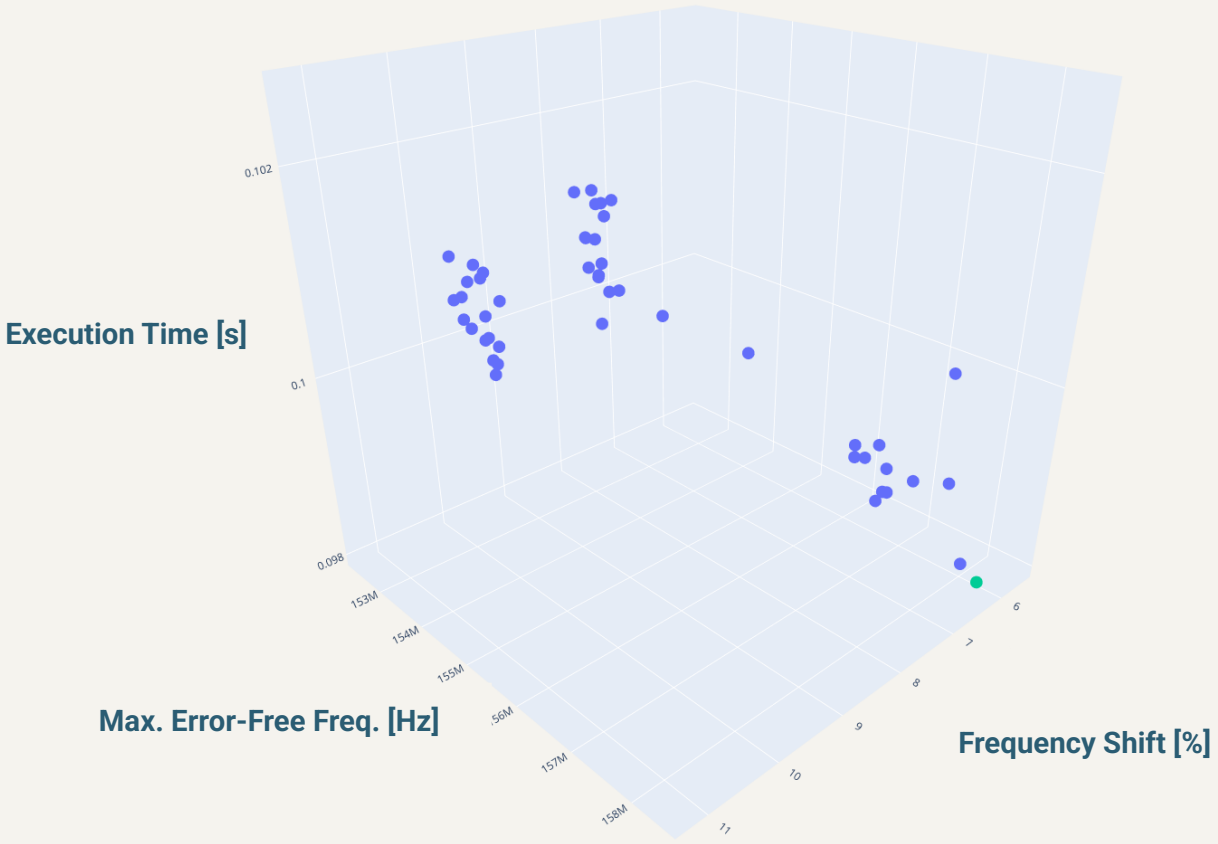
Execution Time [s]
0.102
0.1
0.098
Max. Error-Free Freq. [Hz]
153M
154M
155M
156M
157M
158M
Frequency Shift [%]
6
7
8
9
10
11

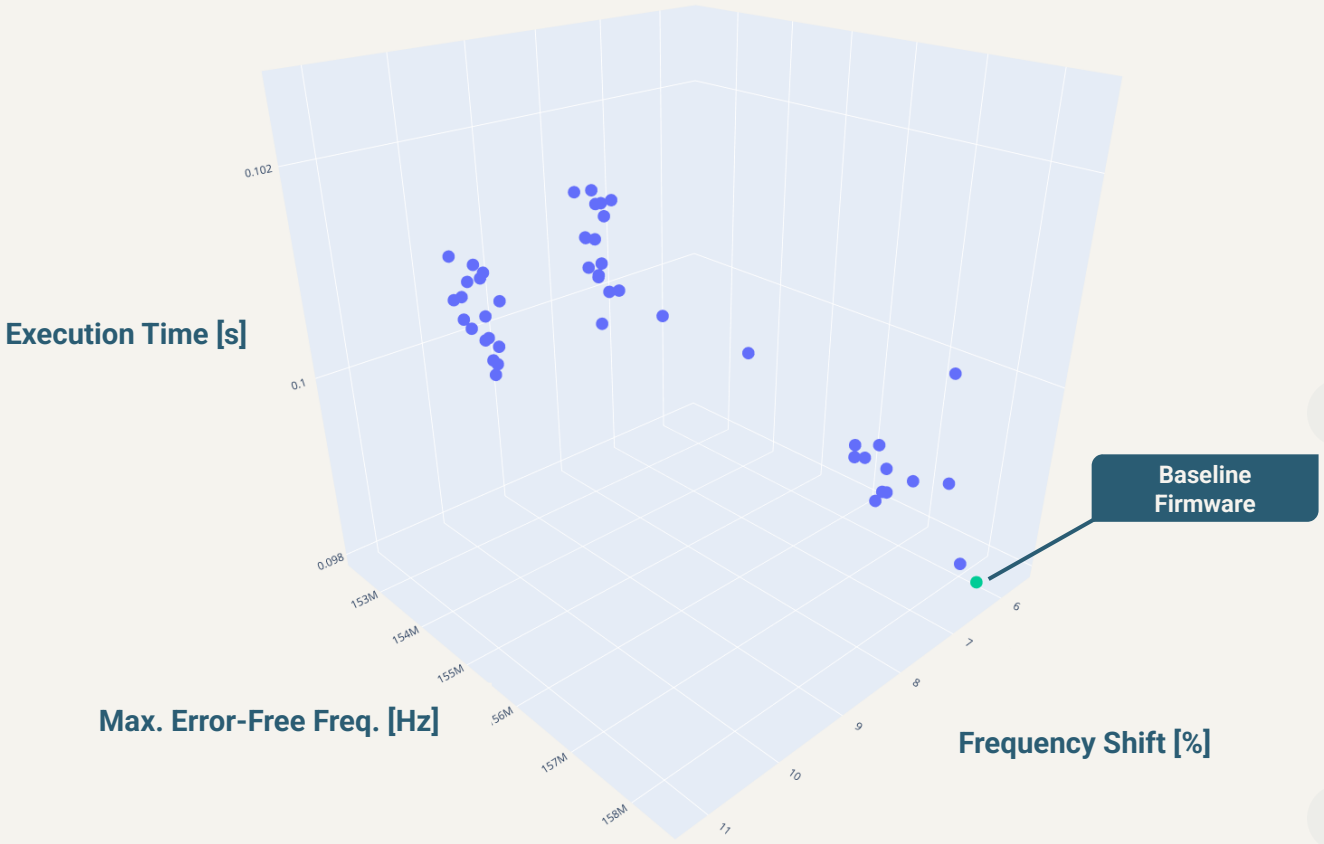
Execution Time [s]
0.102
0.1
0.098
Max. Error-Free Freq. [Hz]
153M
154M
155M
156M
157M
158M
Frequency Shift [%]
6
7
8
9
10
11
Baseline Firmware

So far, more than double sensitivity to degradation than the baseline firmware.

Conclusions

Hardware ages…even in Santa’s sleigh

Hardware ages…even in Santa's sleigh

- Observed shift of SRAM startup patterns, and predicted usage time.

Hardware ages…even in Santa's sleigh

- Observed shift of SRAM startup patterns, and predicted usage time.
- Measured and forecasted FPGAs slowdown due to long-term operation.

Hardware ages…even in Santa's sleigh

- Observed shift of SRAM startup patterns, and predicted usage time.
- Measured and forecasted FPGAs slowdown due to long-term operation.
- Detected MCUs degradation caused by temperature.

Hardware ages…even in Santa's sleigh

- Observed shift of SRAM startup patterns, and predicted usage time.
- Measured and forecasted FPGAs slowdown due to long-term operation.
- Detected MCUs degradation caused by temperature.
- **ML enables detection and forecasting, supporting predictive maintenance.**

Hardware ages…even in Santa's sleigh

- Observed shift of SRAM startup patterns, and predicted usage time.
- Measured and forecasted FPGAs slowdown due to long-term operation.
- Detected MCUs degradation caused by temperature.
- **ML enables detection and forecasting, supporting predictive maintenance.**

Ageing is observable, measurable, and modelable.

Hardware ages…even in Santa's sleigh

- Observed shift of SRAM startup patterns, and predicted usage time.
- Measured and forecasted FPGAs slowdown due to long-term operation.
- Detected MCUs degradation caused by temperature.
- **ML enab**

Test today. → Deliver presents tomorrow.

Ageing is observable, measurable, and modelable.

References

- [1]: L. Lanzieri, G. Martino, G. Fey, H. Schlarb, T. C. Schmidt, and M. Wählisch. 2024. “A Review of Techniques for Ageing Detection and Monitoring on Embedded Systems”. ACM Computing Surveys 57, 1, Article 24 (January 2025), 34 pages. DOI:10.1145/3695247.
- [2]: L. Lanzieri, P. Kietzmann, G. Fey, H. Schlarb and T. C. Schmidt, "Ageing Analysis of Embedded SRAM on a Large-Scale Testbed Using Machine Learning", 2023. In proceedings of the 26th Euromicro Conference on Digital System Design (DSD), Golem, Albania, 2023, pp. 335-342. DOI: 10.1109/DSD60849.2023.00054.
- [3]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, and T. C. Schmidt. “Studying the Degradation of Propagation Delay on FPGAs at the European XFEL”, 2024. In Proceedings of the 27th Euromicro Conference on Digital System Design (DSD), Paris, France, 2024. DOI: 10.1109/DSD64264.2024.00018
- [4]: L. Lanzieri, L. Butkowski, J. Kral, G. Fey, H. Schlarb, T. C. Schmidt, “Switching Frequency as FPGA Monitor: Studying Degradation and Ageing Prognosis at Large Scale”, 2024. Preprint. DOI: 10.48550/arXiv.2412.15720

Icon Attributions

- Old icons created by juicy_fish - Flaticon
- Christmas icons created by surang - Flaticon
- Wish icons created by photo3idea_studio - Flaticon
- Cookie icons created by Smashicons - Flaticon
- Coal icons created by Freepik - Flaticon
- Gift box icons created by photo3idea_studio - Flaticon
- Christmas icons created by Freepik - Flaticon
- Reindeer icons created by Leremy - Flaticon

HAW HAMBURG

DESY.

TUHH
Hamburg
University of
Technology

Thank you

Leandro Lanzieri
leandro.lanzieri@haw-hamburg.de