



# Extending Flash Memory Through Adaptive Parameter Tuning

Conor Ryan

CTO - Software

[Conor.Ryan@NVMdurance.com](mailto:Conor.Ryan@NVMdurance.com)

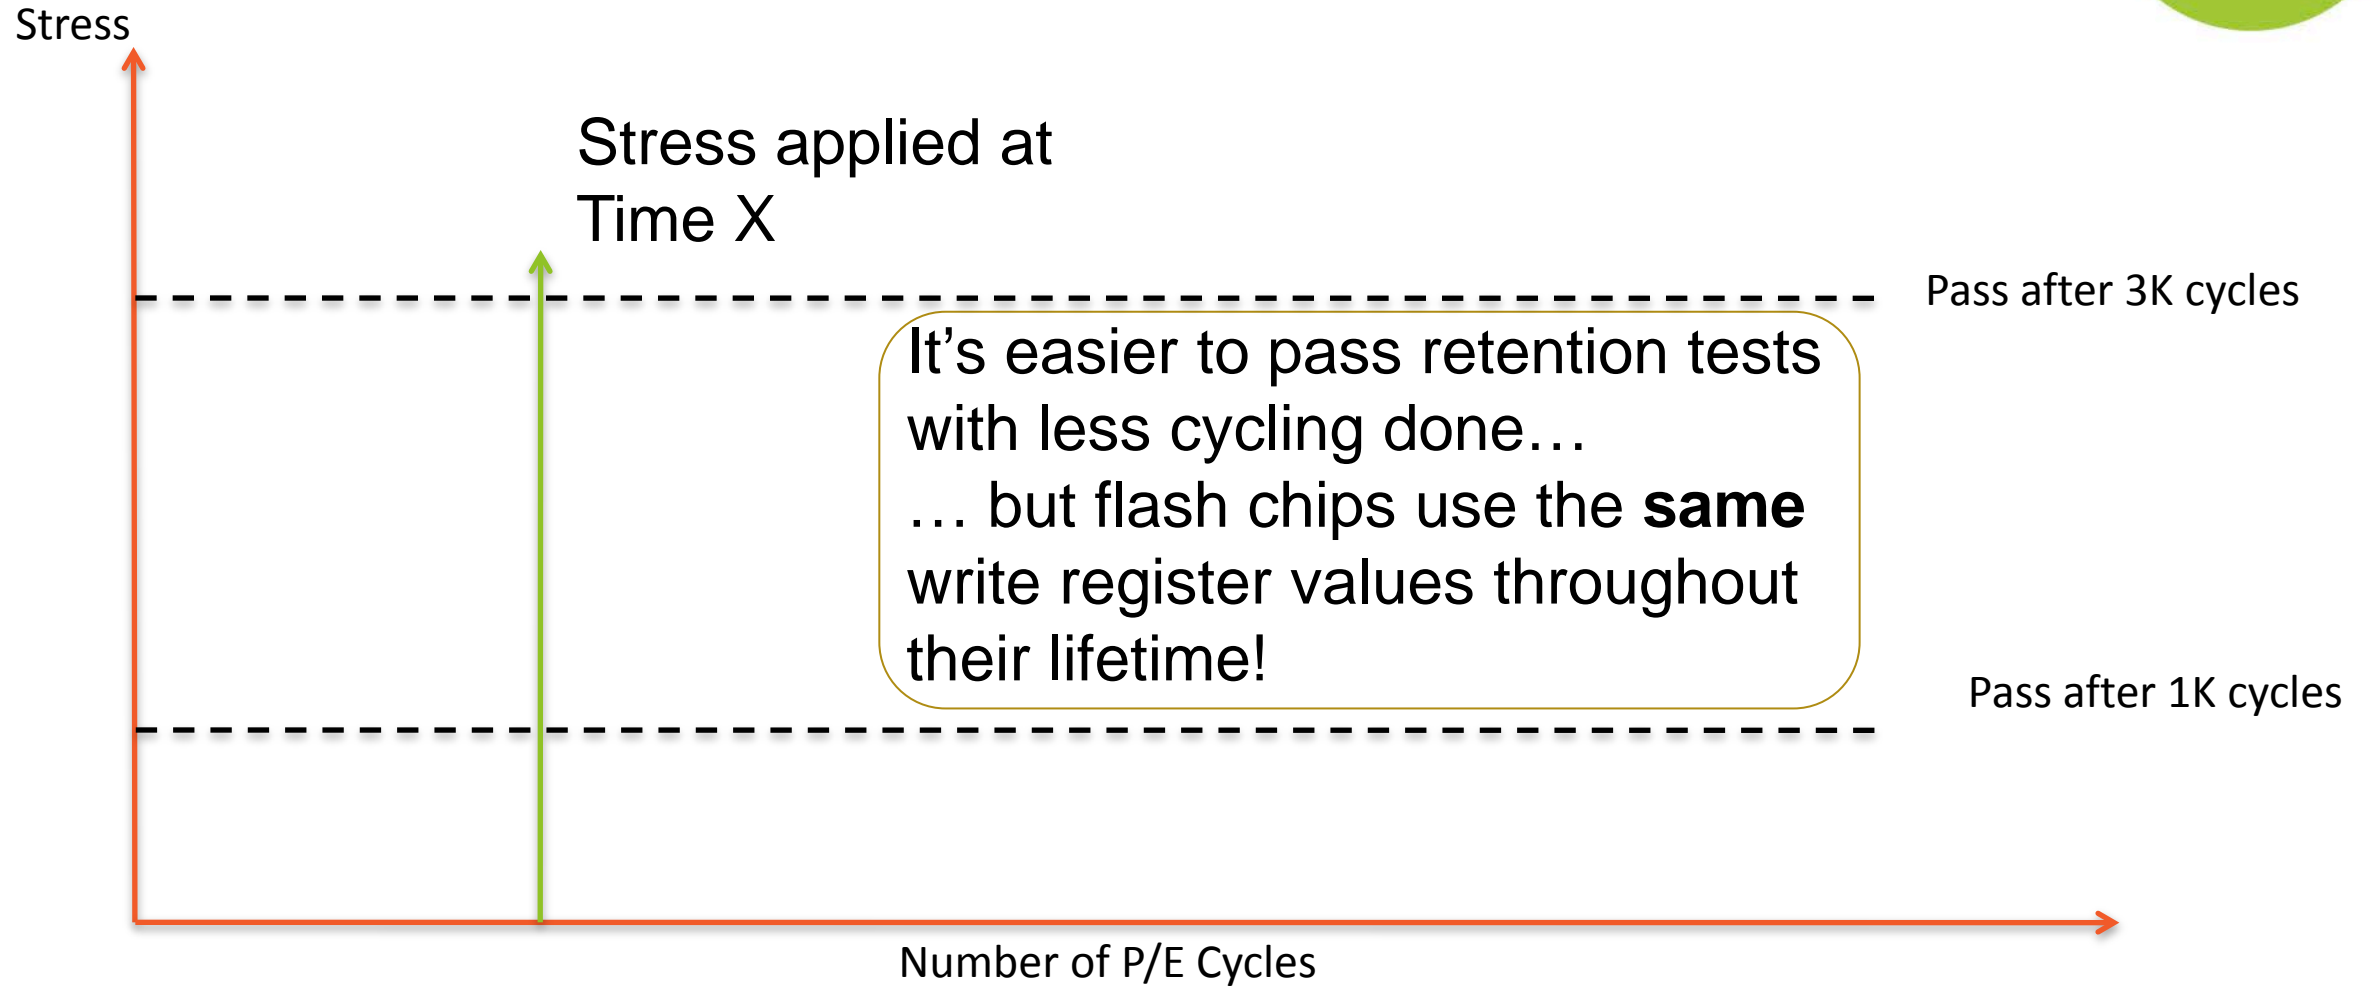
# Take Home Messages




- ▶ 25x increase in endurance on 1X nm devices
  - ▶ Machine learning software fully characterizes NAND chips before they go into a product
  - ▶ Lightweight software running on the product autonomically manages its health, using machine learning knowledge
  - ▶ 40-bit ECC gives 10x increase; 120-bit ECC gives 25x increase
- ▶ Uses the least *stress* possible
- ▶ Raw flash runs *faster*

- ▶ NVMdurance Pathfinder
  - ▶ Delivers most of the endurance gain
  - ▶ Suite of Machine Learning Algorithms
  - ▶ Determines sets of optimal register values for NAND chips before they go into a product
  - ▶ “Heavy lifting” done before shipping
- ▶ NVMdurance Navigator
  - ▶ Exploits endurance gains enabled by Pathfinder
  - ▶ Autonomic system, runs on the controller
  - ▶ Chooses register values at run-time from those discovered by Pathfinder

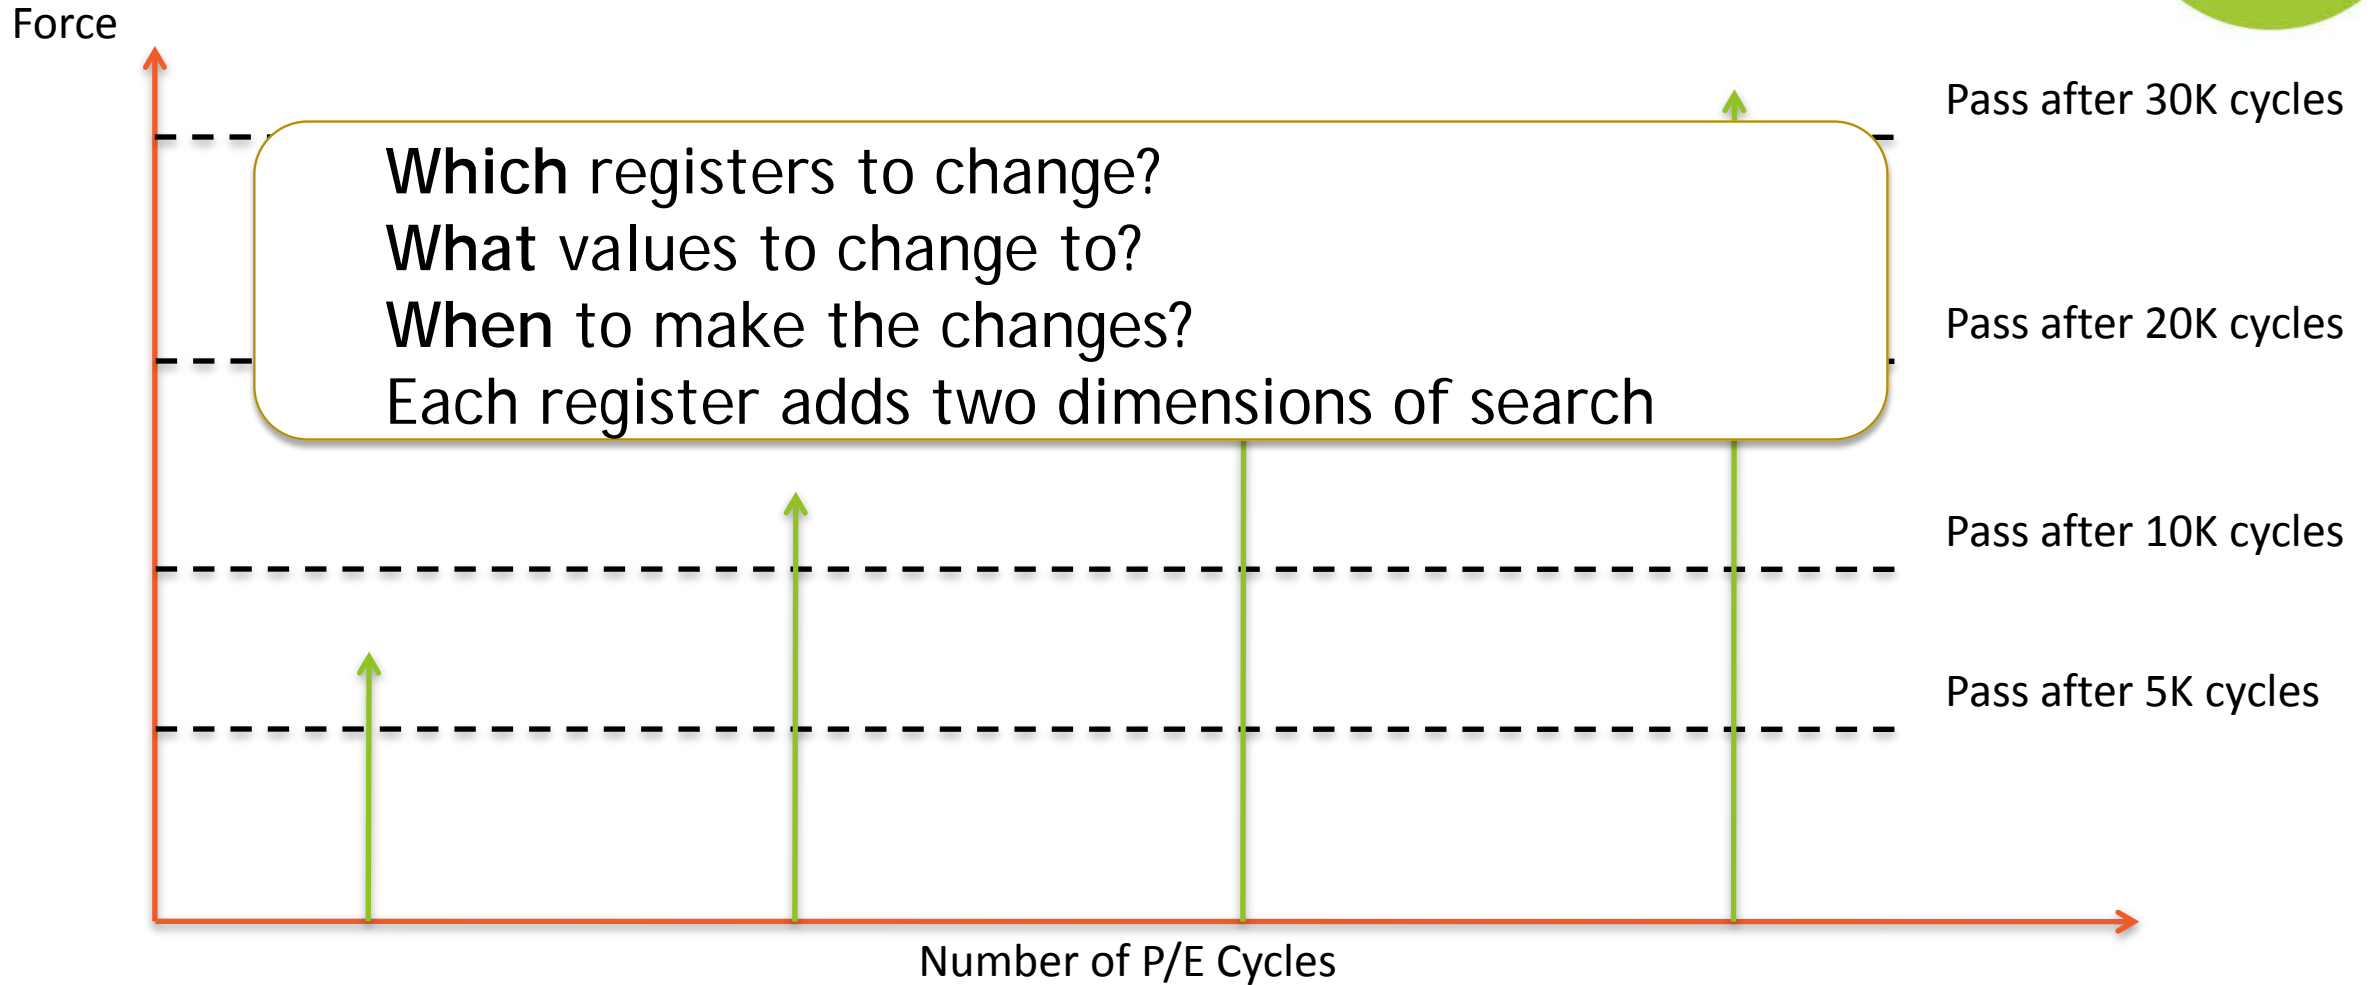
# The Problem



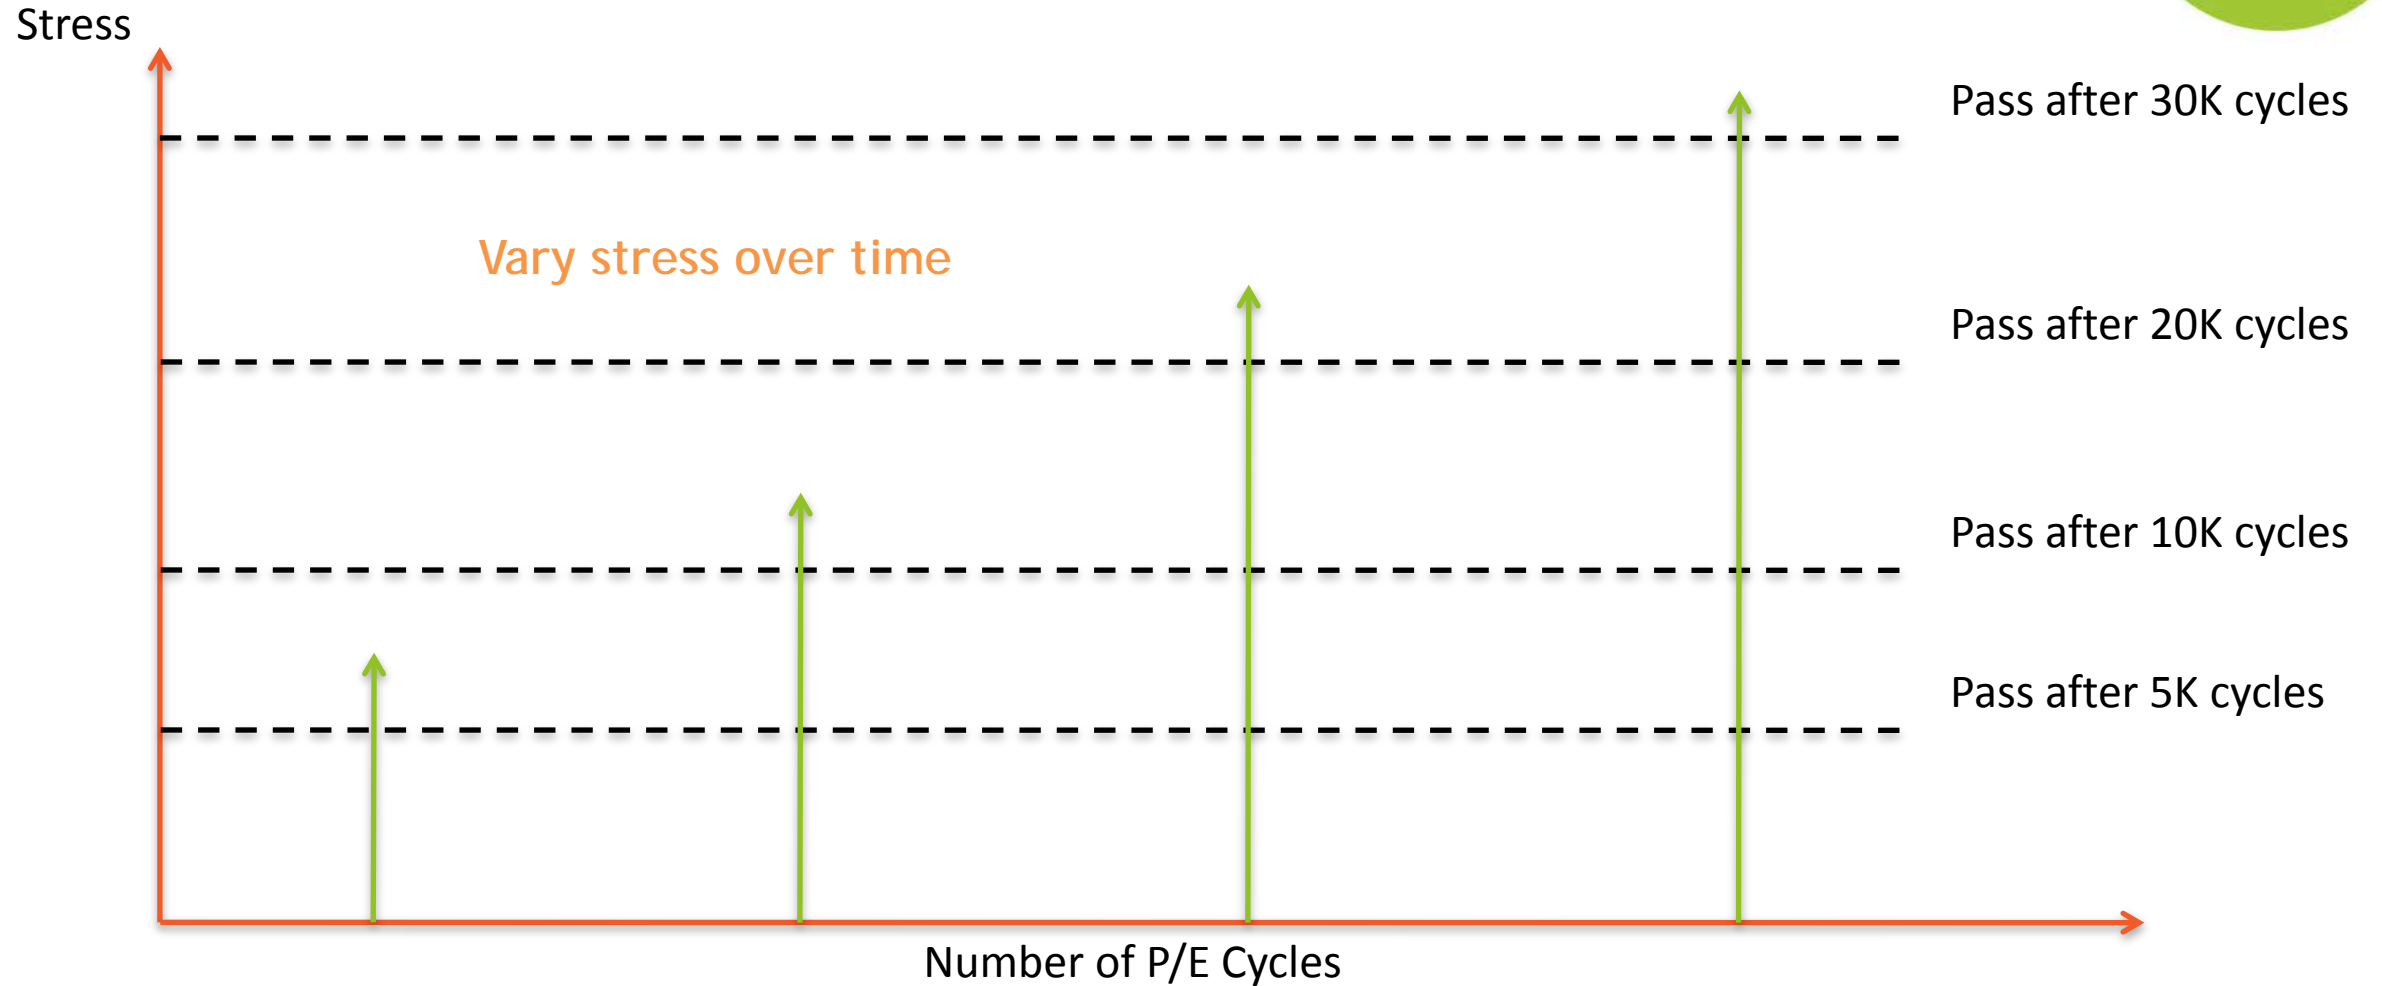
# Solution




# Solution



# Solution

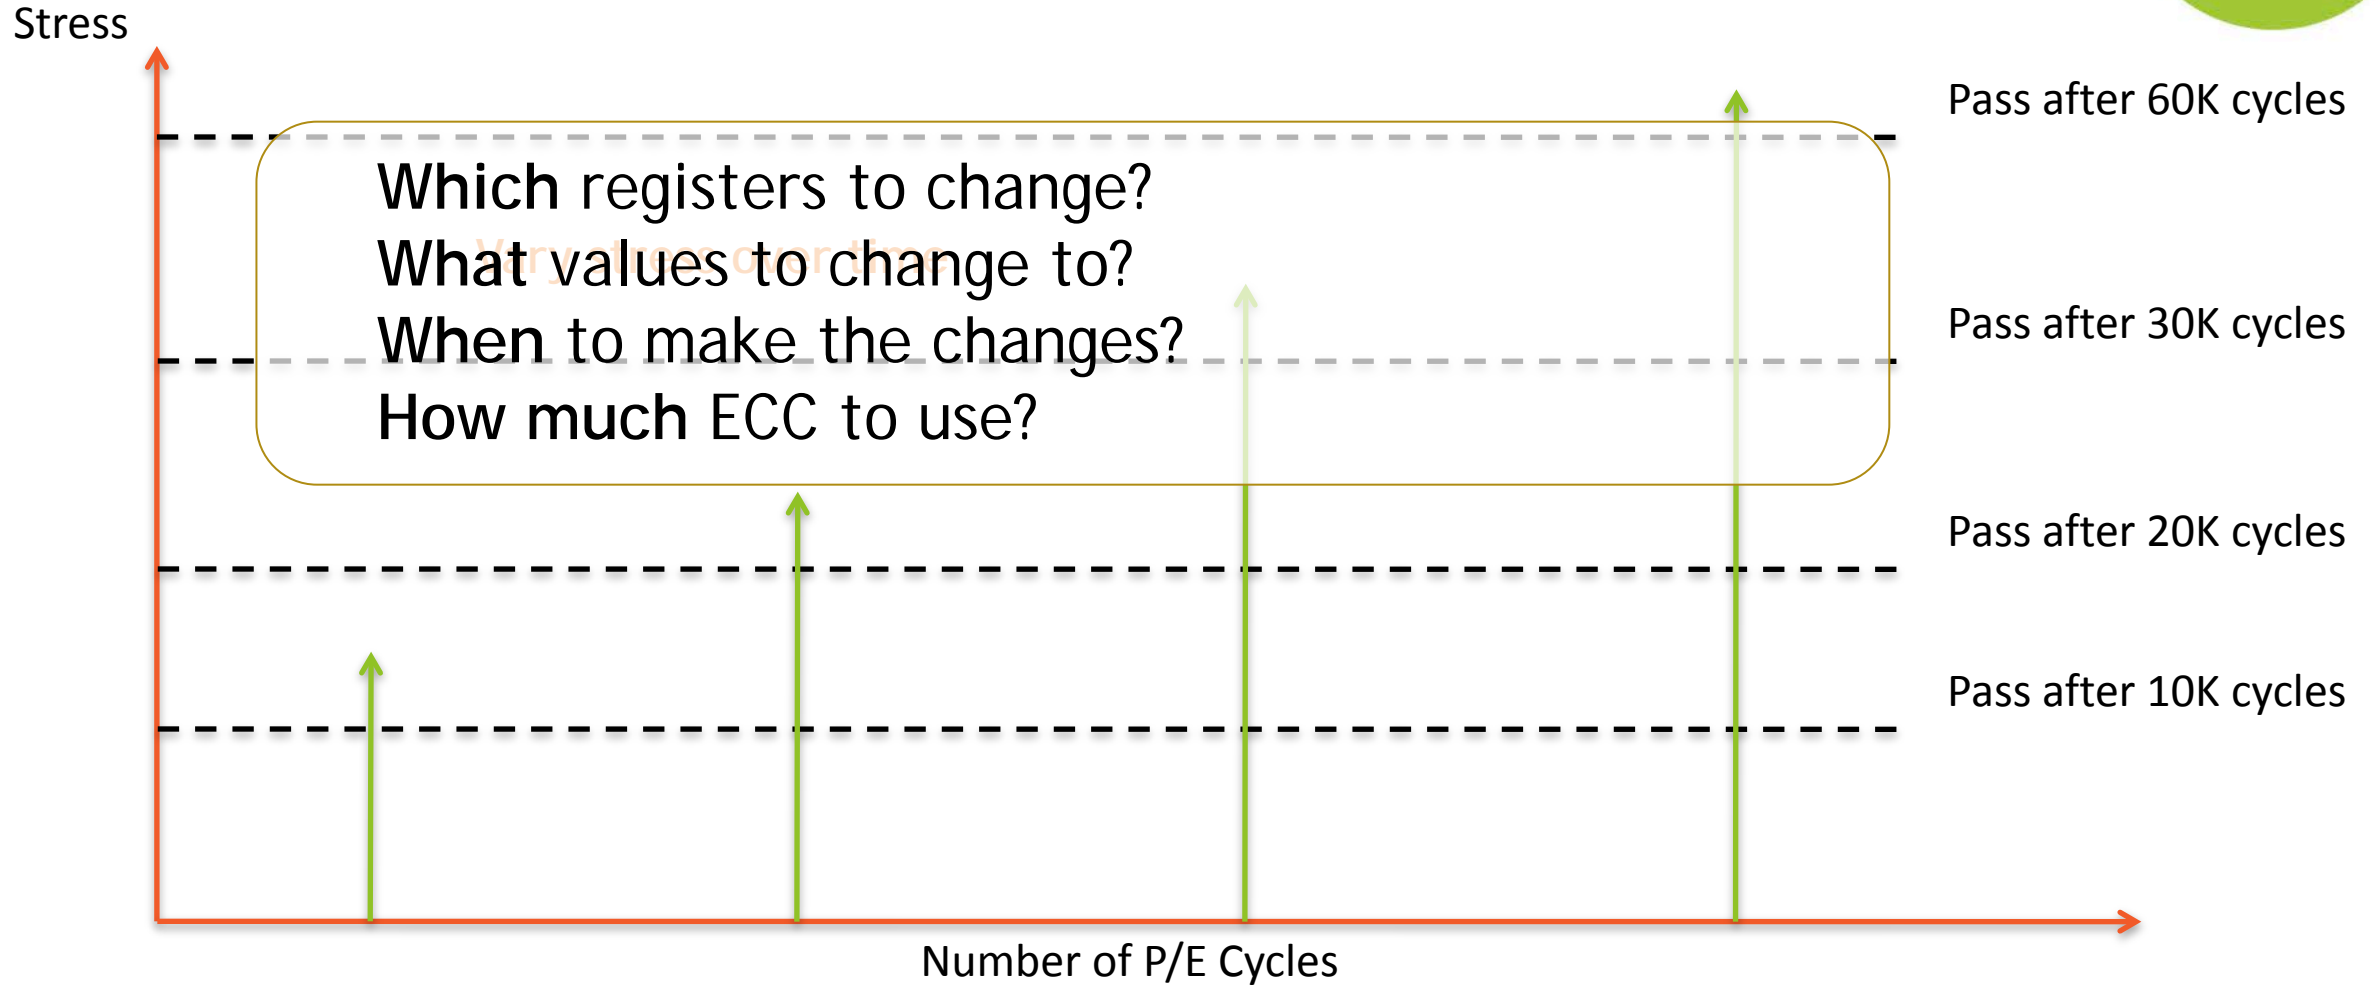


# Solution with variable ECC





# Solution with variable ECC



# The Not-So-Secret Sauce



- ▶ Reduce dimensionality of the problem
  - ▶ Reduce the number of independent variables
  - ▶ Understand the silicon; vary as few registers as possible
- ▶ Guide the search
  - ▶ Be sensible, if not insightful
- ▶ Test only what has to be tested
  - ▶ Or at least know what NOT to test
- ▶ ...this is still an astronomically difficult problem!

# The Secret Sauce



- ▶ Plot “safe” paths through massively highly dimensional space
  - ▶ Each register adds two dimensions (more than 50 write registers!)
- ▶ Tune paths on the fly
  - ▶ Not all blocks degrade at the same rate
  - ▶ Retention may impact various blocks differently
  - ▶ Anticipate health issues before they happen
- ▶ Treat the flash as though it is a dynamic, living thing

# NVMdurance Pathfinder



Start


Imagine the lifetime of a device to be a journey through space...  
... touch an “asteroid” and we have unrecoverable data

Destination

# NVMdurance Pathfinder



Start



Destination

# NVMdurance Pathfinder



Start



Destination


# NVMdurance Pathfinder



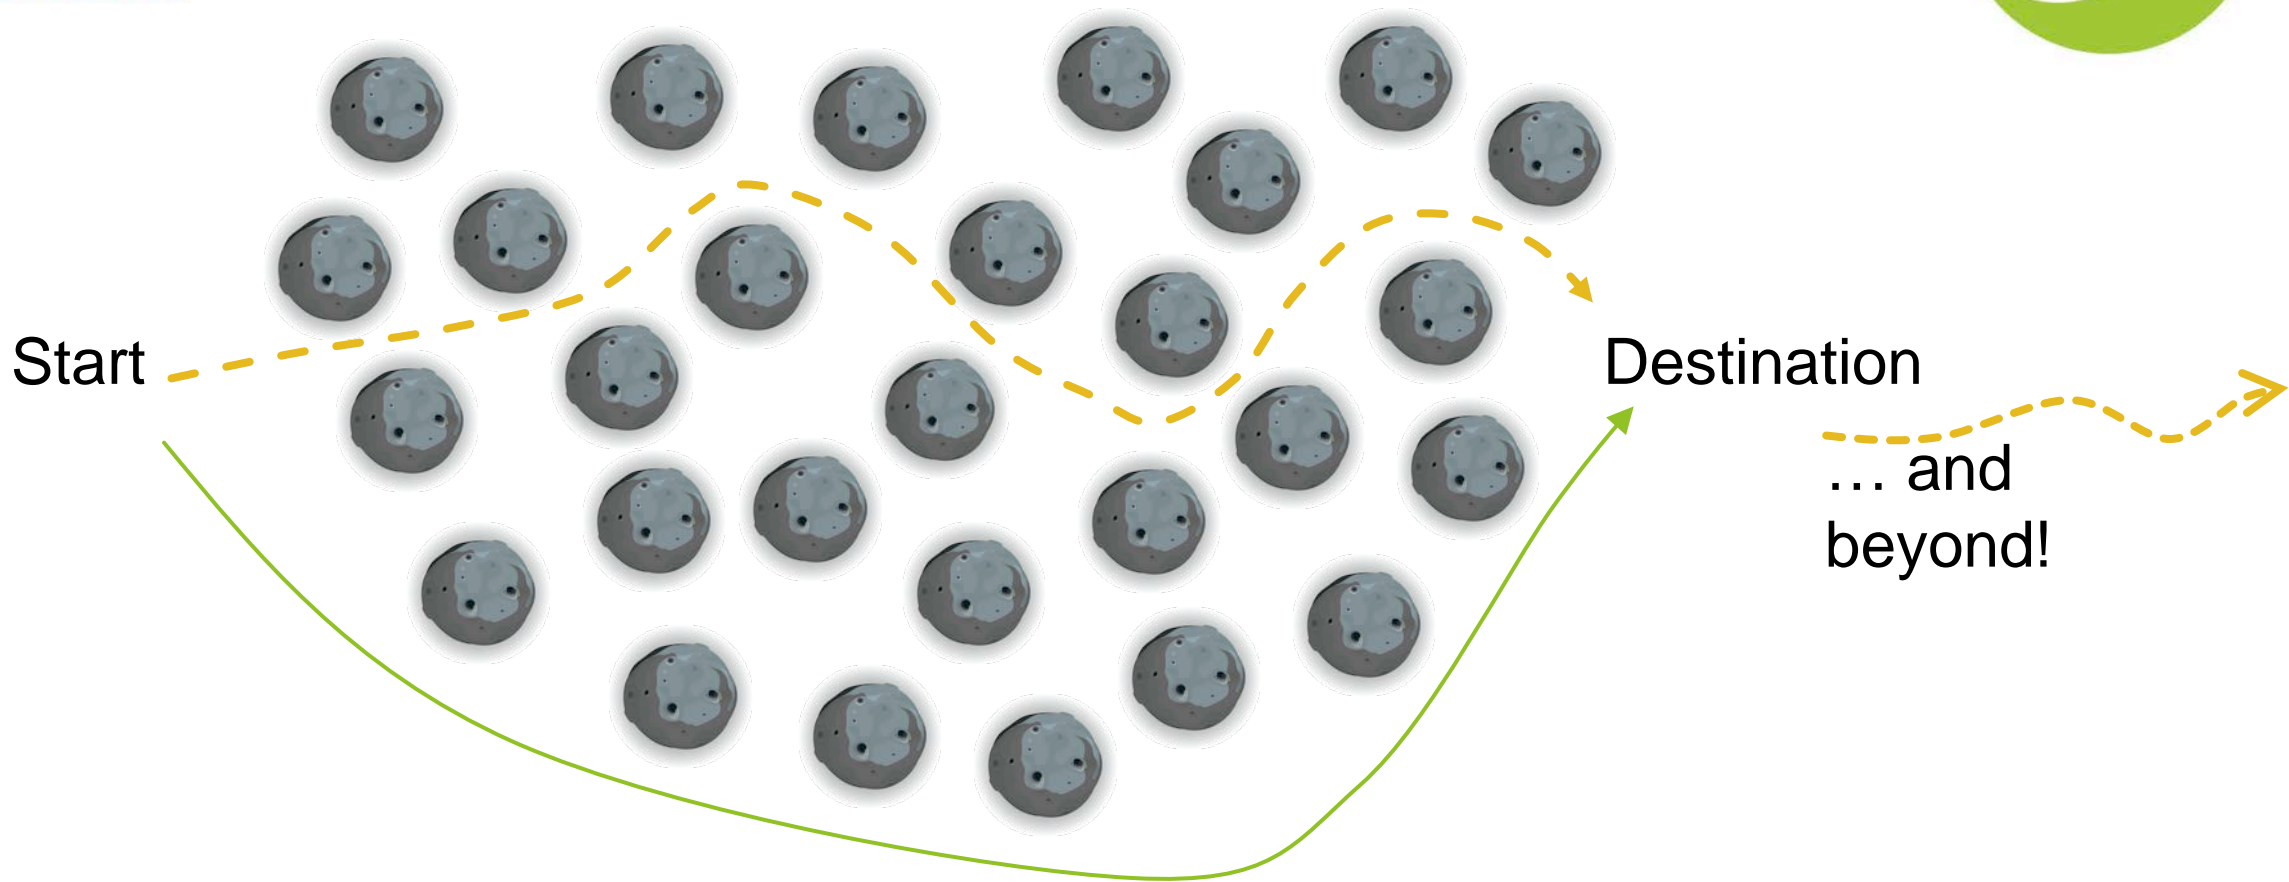
Start



Destination



# NVMdurance Pathfinder






# NVMdurance Pathfinder



Start




Destination

# NVMdurance Pathfinder



Start




Destination

# NVMdurance Pathfinder



Start




Destination

# NVMdurance Pathfinder




Start



Destination

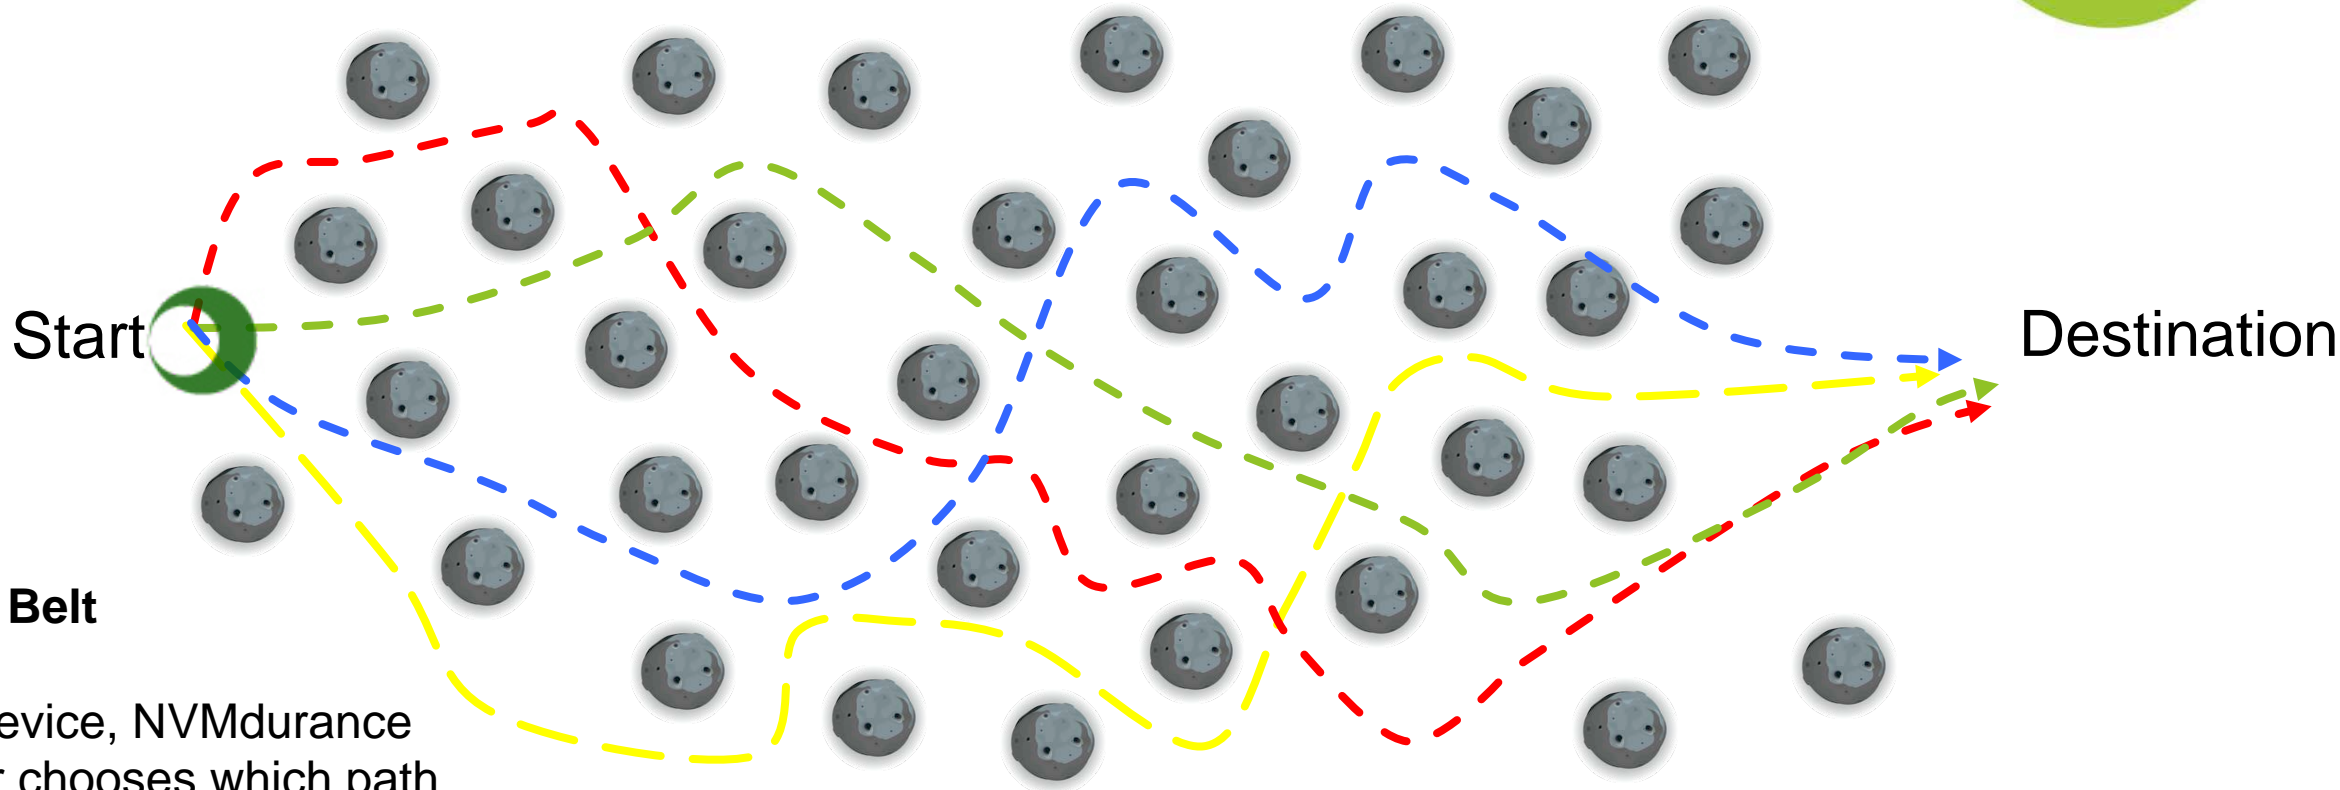
# NVMdurance Pathfinder



Asteroid Belt


NVMdurance Pathfinder determines sets of optimal register values (analogous to safe paths through the asteroid belt) for NAND chips before they go into a product

# NVMdurance Navigator



## Asteroid Belt

Live on device, NVMdurance Navigator chooses which path through the space to use, based on the “health” of the device.




Start

Destination


## Asteroid Belt

Live on device, NVMdurance Navigator chooses which path through the space to use, based on the "health" of the device.

# NVMdurance Navigator








Start

Destination

## Asteroid Belt

NVMdurance Navigator constantly monitors the health, and jumps and jumps from path to path as the use-case requires.




Start

Destination

## Asteroid Belt

NVMdurance Navigator constantly monitors the health, and jumps and jumps from path to path as the use-case requires.




Start

Destination

## Asteroid Belt

NVMdurance Navigator constantly monitors the health, and jumps and jumps from path to path as the use-case requires.



Start

Destination

## Asteroid Belt

NVMdurance Navigator constantly monitors the health, and jumps and jumps from path to path as the use-case requires.

# NVMdurance system



## Machine Learning Parameter Discovery


*NVMdurance Pathfinder*: Discover routes through multidimensional space such that every parameter set passes retention for that point of life (by fully characterizing NAND chips before they go into a product)

## Autonomic (runs live on the SSD controller)

*NVMdurance Navigator*: Observes deterioration of the SSD; chooses when to change parameters (using the knowledge delivered by Pathfinder)




- ▶ Sets of Write register values for that time in life
  - ▶ 30-60 write registers, e.g. start voltage, step size, MSB, LSB, odd, even, etc.
  - ▶ More registers means more control, but larger search space




- ▶ Sets of Write register values for that time in life
  - ▶ 30-60 write registers, e.g. start voltage, step size, MSB, LSB, odd, even, etc.



Early Life

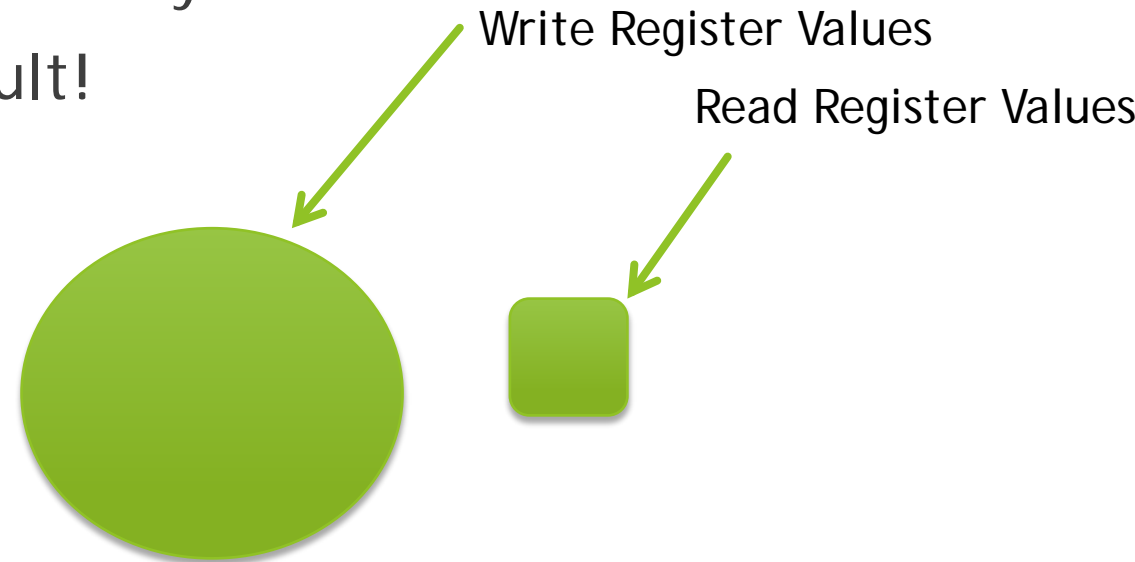


Mid Life



Late Life

- ▶ Each stage has multiple *waypoints* (each a set of read register values) to guide it to the next stage
  - ▶ Often only **one** set in early life - no read retry!
  - ▶ First stage often lasts longer than default!
  - ▶ Never more than **eight** waypoints






# Waypoints


- ▶ More waypoints required later in life
  - ▶ Higher wear uncovers higher variability



Early Life




Mid Life




Late Life

- ▶ NVMdurance Pathfinder (machine learning offline)
  - ▶ Automatically discovers and proves viability of stages
  - ▶ Each stage passes retention test



P/E Cycle Count as Multiple of Default Rating

- ▶ NVMdurance Navigator (run time)
  - ▶ Runs on controller (autonomic)
  - ▶ Monitors “Health” of devices
  - ▶ Determines when to progress to next stage
  - ▶ Chooses which waypoint to use during each stage




P/E Cycle Count as Multiple of Default Rating


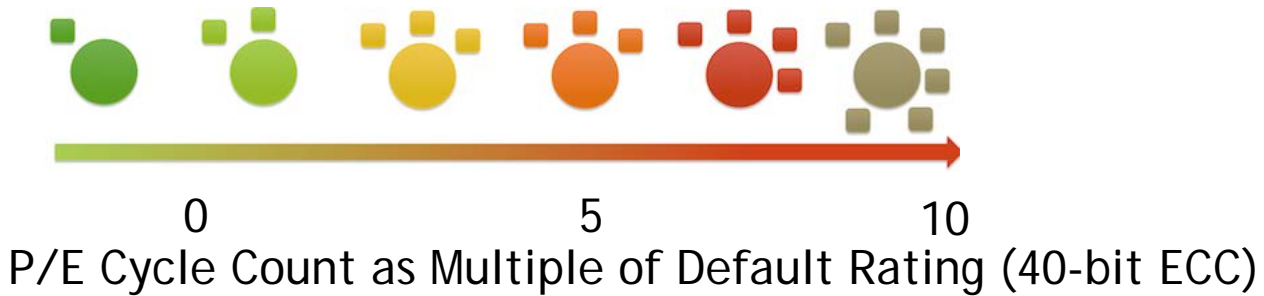
# Increasing ECC



- ▶ System can tolerate higher BER
  - ▶ Stages last longer
  - ▶ Weaker stages possible earlier in life
  - ▶ NVMdurance enables a truly variable/adaptive ECC



# One approach - many use cases




## Target Device

Device	1X nm
ECC	Up to 40 bits per sector
Retention	12 months
Baseline Endurance	1
Intrinsic Endurance (3 months retention)	1.8x longer than rated endurance

## Summary


Number of stages	8
Stage length	1x - 2x longer than rated endurance
Retention	3 months
ECC	Up to 40 bits per sector
Minimum Window Stress (stage 1)	45%
Approximately equal stress	Stage 5
Maximum Window Stress (stage 8)	120%
Maximum P/E cycles	10x longer than rated endurance

### P/E Cycle Time & Write Stress Vs P/E Cycles



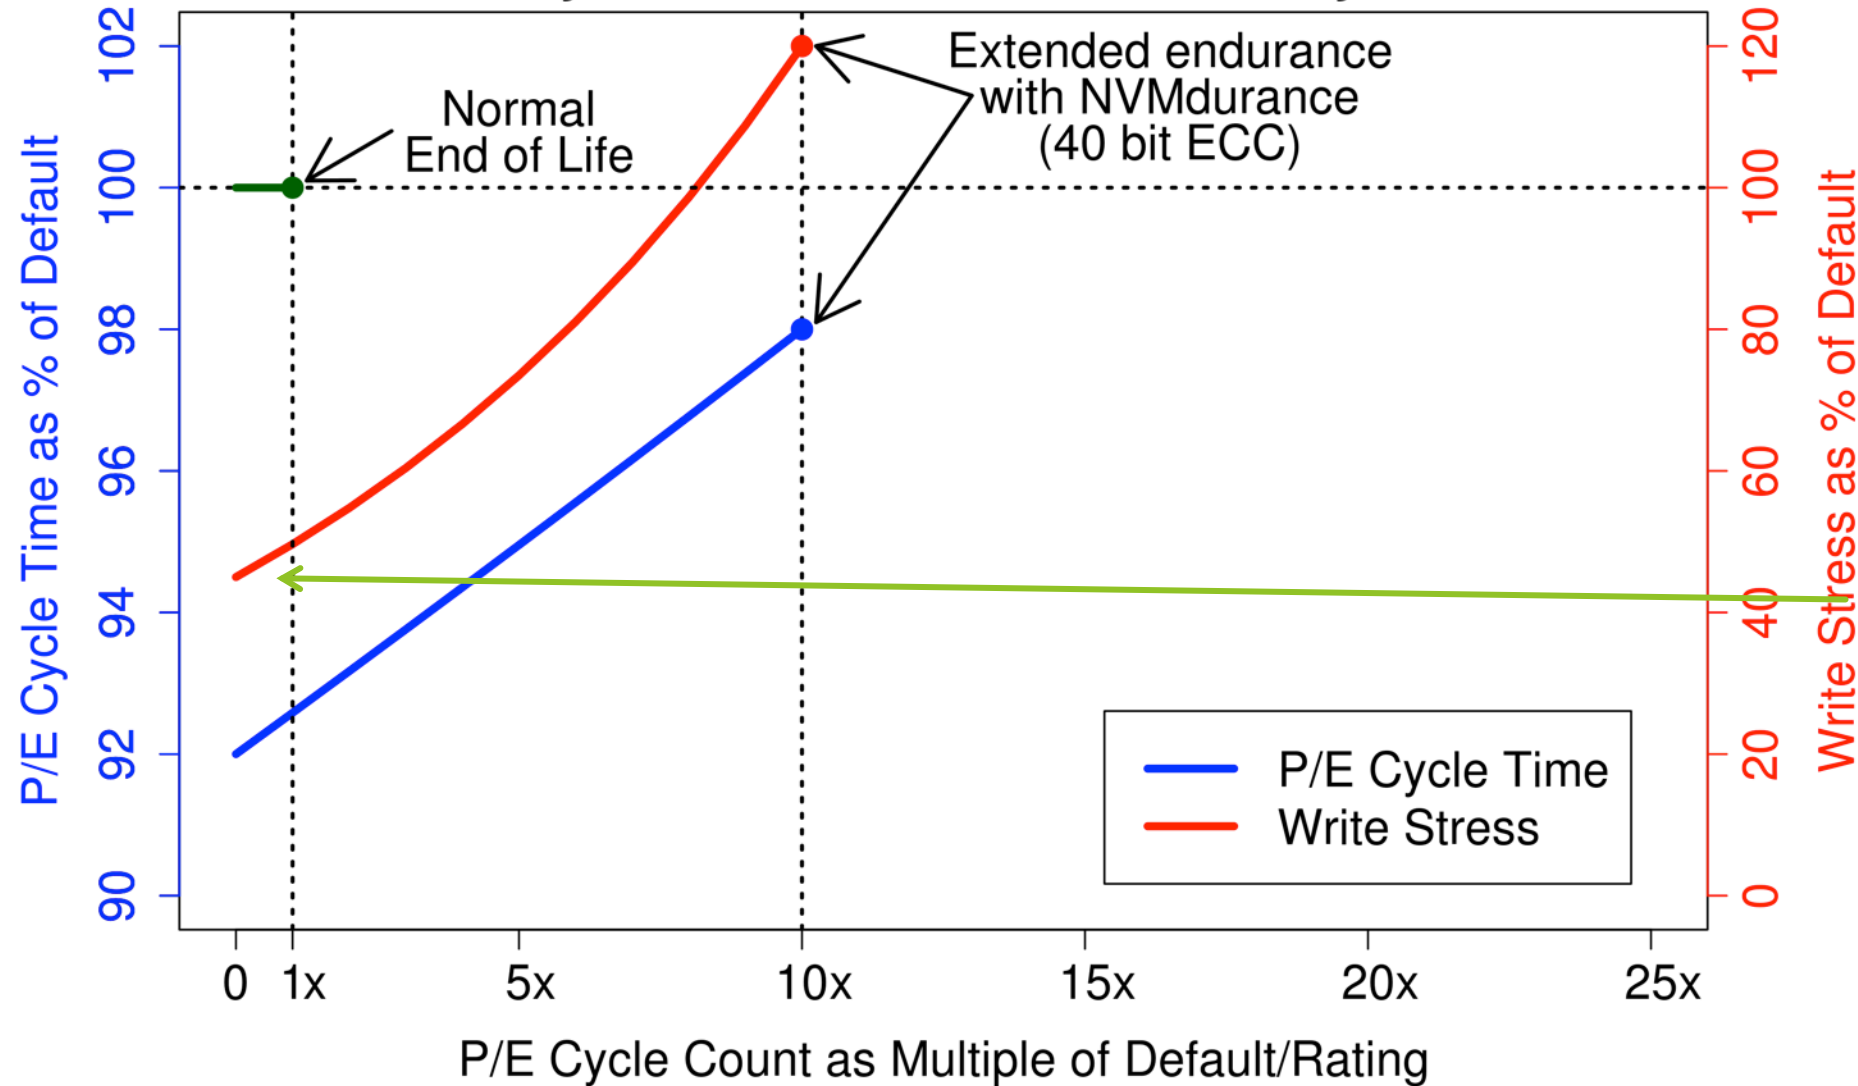


### P/E Cycle Time & Write Stress Vs P/E Cycles



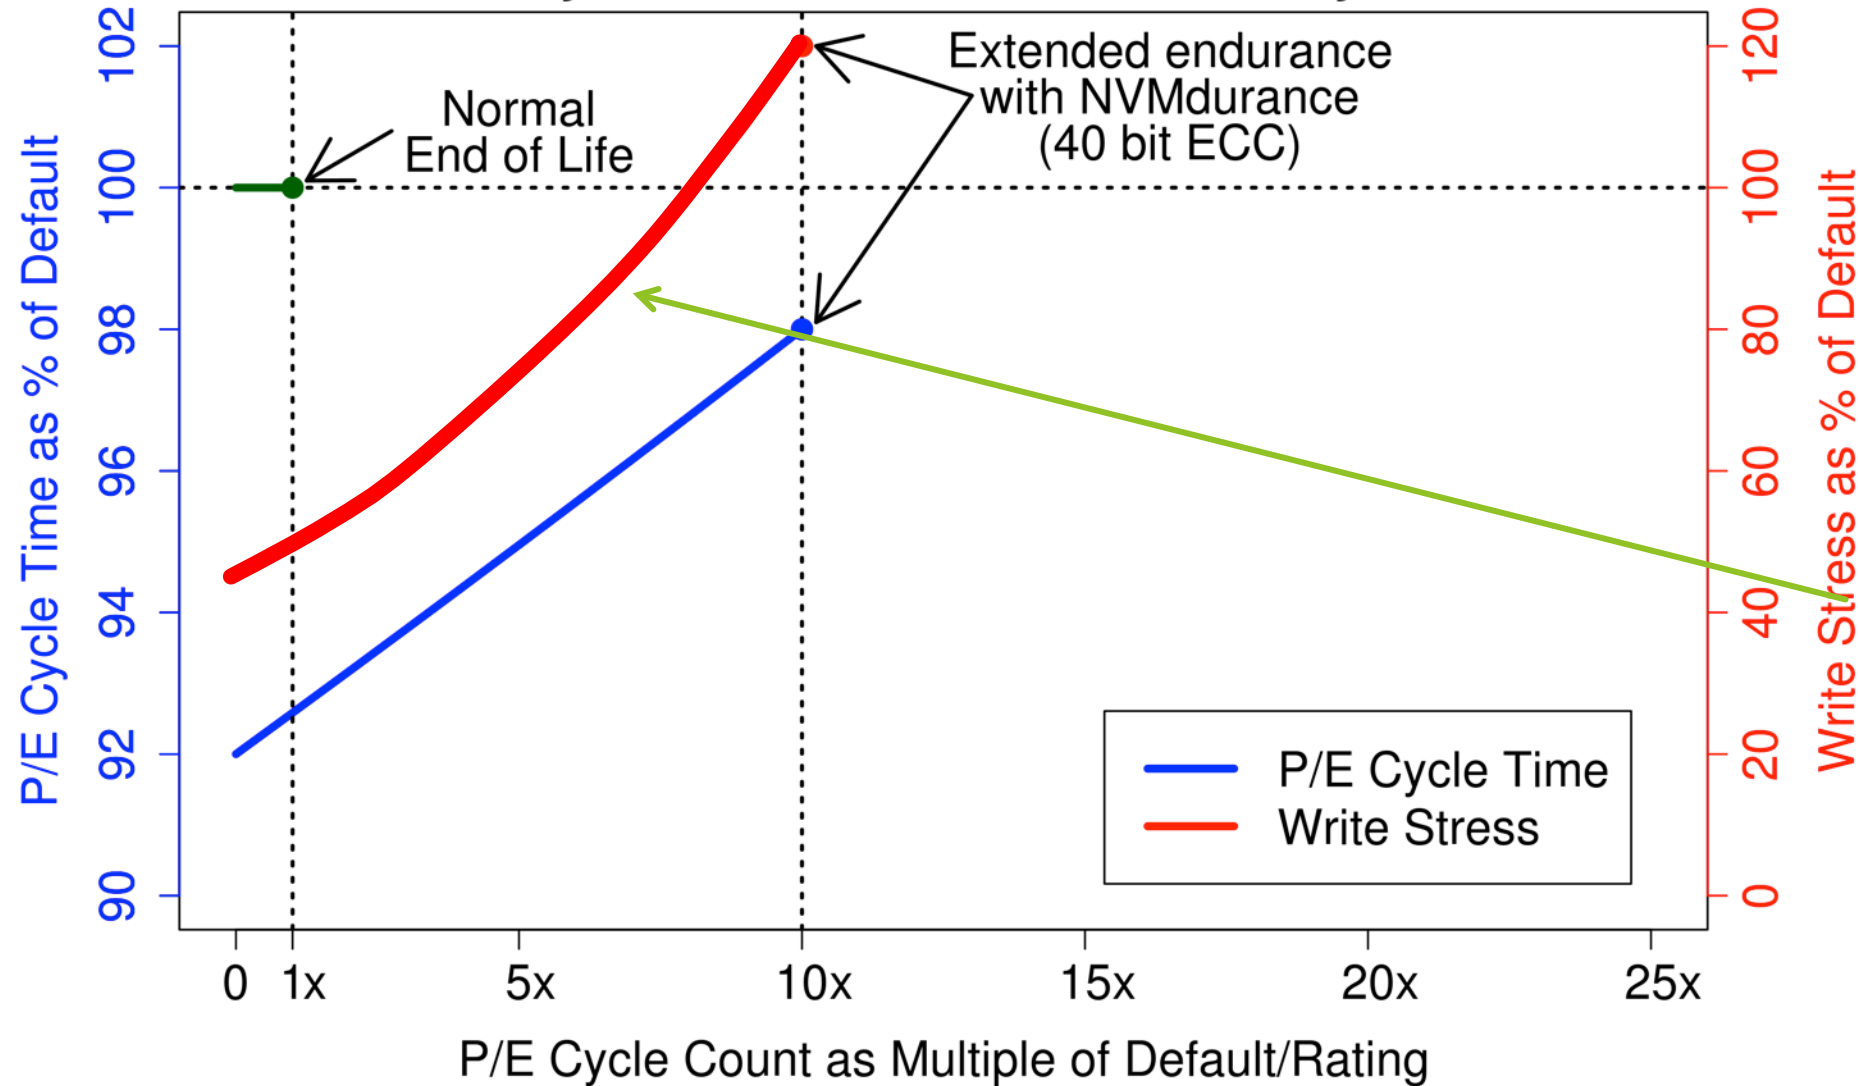
Default lifetime, normalized to one

**P/E Cycle Time & Write Stress Vs P/E Cycles**




Initial write stress starts at 45% of default

### P/E Cycle Time & Write Stress Vs P/E Cycles




Write stress slowly increases

### P/E Cycle Time & Write Stress Vs P/E Cycles




Write stress exceeds default level

### P/E Cycle Time & Write Stress Vs P/E Cycles



10x increase in endurance

### P/E Cycle Time & Write Stress Vs P/E Cycles




P/E cycle time as % of default; never exceeds default.

## Results - Increasing ECC



### P/E Cycle Time & Write Stress Vs P/E Cycles




Stages get longer as ECC increases

# Results - Increasing ECC



P/E Cycle Time & Write Stress Vs P/E Cycles



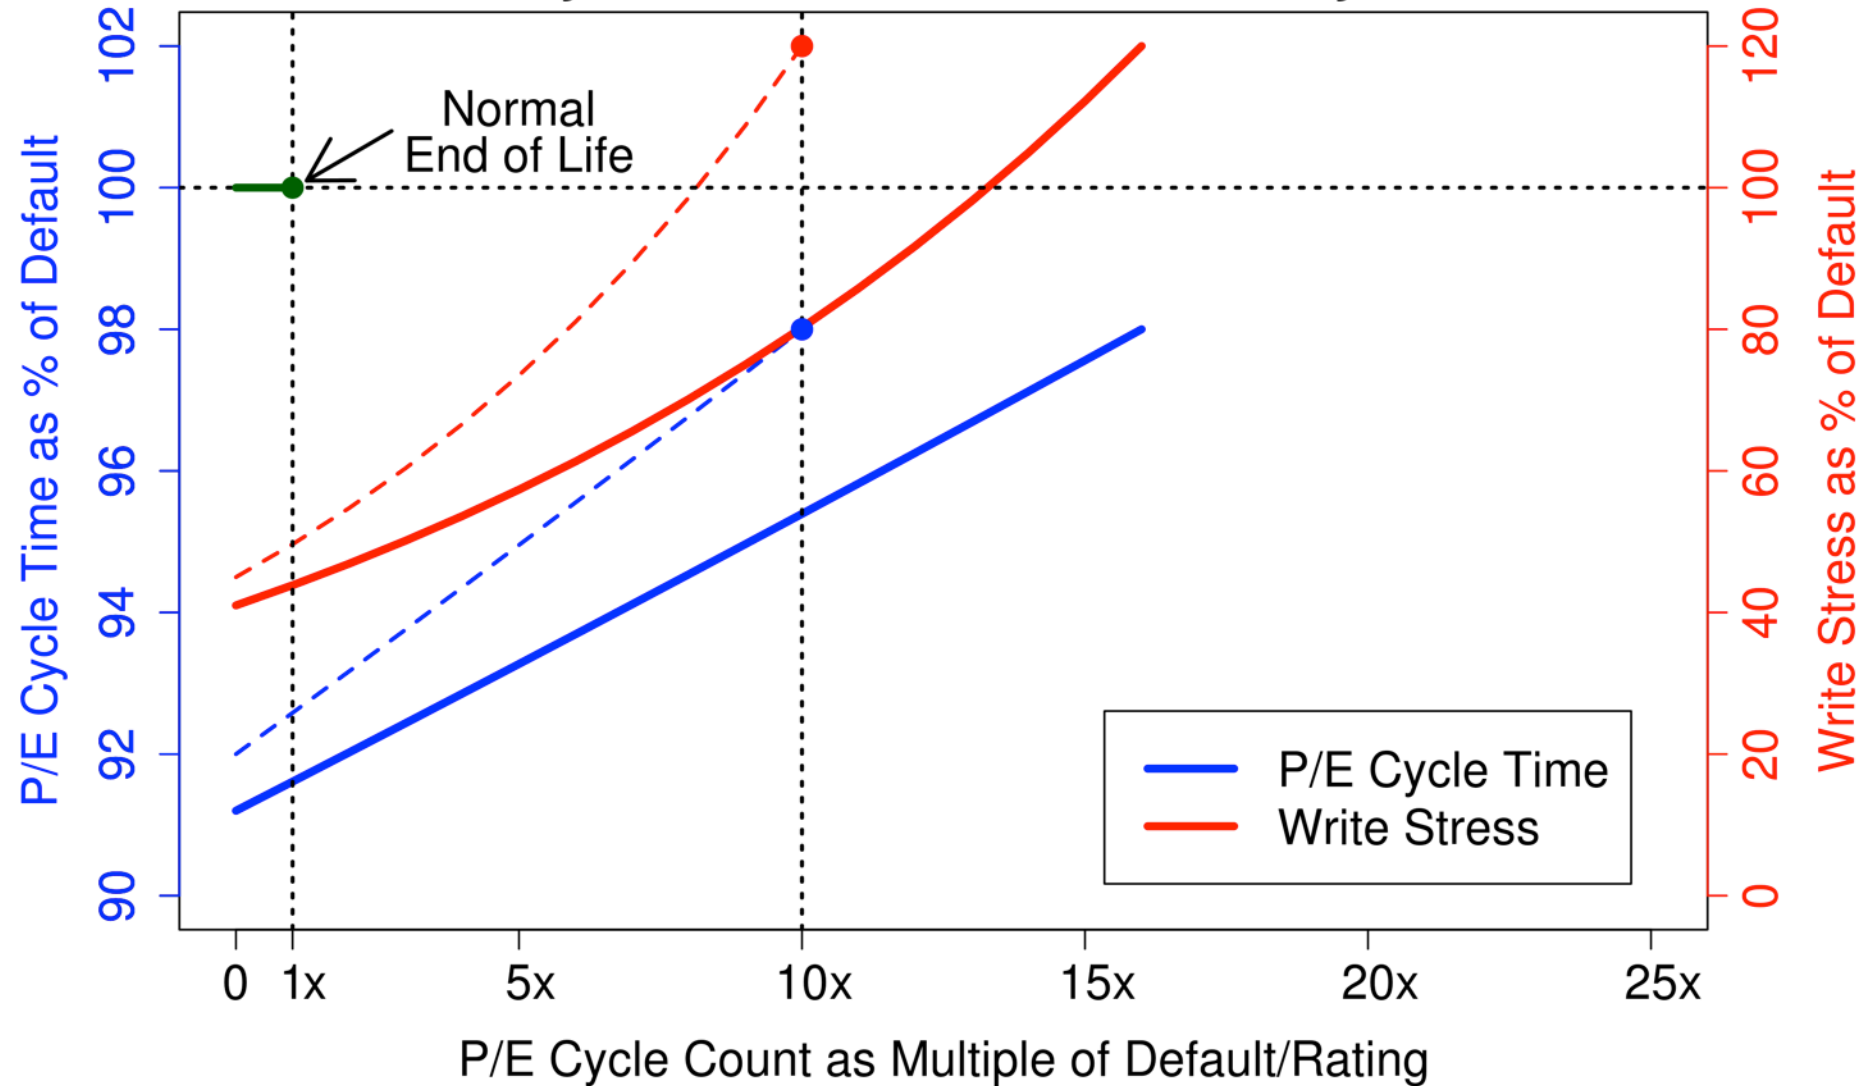
Stages get longer as ECC increases



# Results - Increasing ECC



P/E Cycle Time & Write Stress Vs P/E Cycles




Stages get longer as ECC increases

# Results - Increasing ECC



P/E Cycle Time & Write Stress Vs P/E Cycles




Stages get longer as ECC increases

# Results - Increasing ECC



P/E Cycle Time & Write Stress Vs P/E Cycles




Stages get longer as ECC increases

# Results - Increasing ECC




P/E Cycle Time & Write Stress Vs P/E Cycles




25x improvement in endurance with the same retention!

### P/E Cycle Time & Write Stress Vs P/E Cycles




Initial write stress starts at 25% of default

P/E Cycle Time & Write Stress Vs P/E Cycles




Write stress increases more slowly

### P/E Cycle Time & Write Stress Vs P/E Cycles



Write stress exceeds default level

### P/E Cycle Time & Write Stress Vs P/E Cycles



P/E cycle time as % of default; never exceeds default.



## Final Results

Number of stages	8
Stage length	1.2x - 7x longer than rated endurance
Retention	3 months
ECC	Up to 120 bits per sector
Minimum Window Stress (stage 1)	25%
Approximately equal stress	Stage 5
Maximum Window Stress (stage 8)	120%
Maximum P/E Cycles	25x longer than rated endurance

# Results Summary



- ▶ 25-fold increase in endurance on 1X nm
  - ▶ 10-fold increase on 1X with 40-bit ECC
- ▶ We avoid the problem of live optimization of parameters
  - ▶ Most work done before flash is put in product
- ▶ NVMdurance Navigator can predict imminent failure
  - ▶ Our “health” measure is very precise
- ▶ P/E operations run faster than defaults
- ▶ Results proven with current generation devices from multiple manufacturers

## Conclusion



- ▶ Industry leading endurance gains
- ▶ NVMdurance's technology is synergistic to existing flash controller technology
  - ▶ Machine-learning software fully characterizes NAND chips before they go into a product
  - ▶ Lightweight software running on the product autonomically manages its health, using that knowledge

Stop by and see us at booth #921  
Conor.Ryan@NVMdurance.com