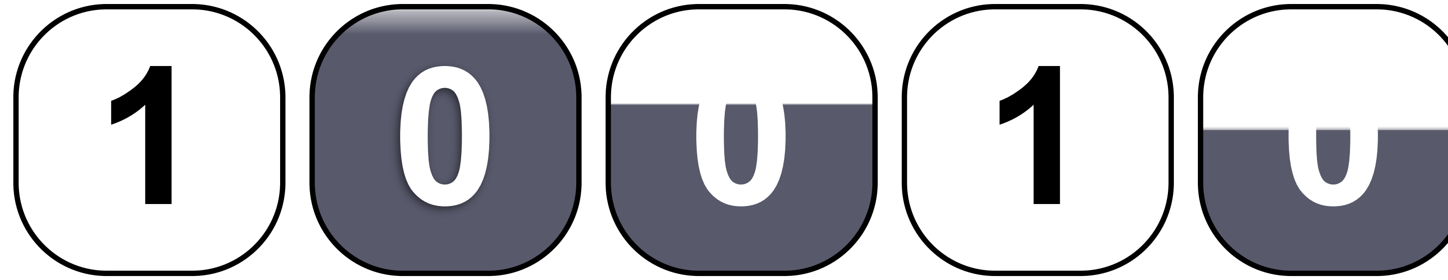

Exploiting Half-Wits: Smarter Storage for Low-Power Devices

Mastooreh (Negin) Salajegheh, Yue Wang
Kevin Fu, Andrew Jiang, Erik Learned-Miller

Supported in part by a Sloan Research Fellowship and NSF CCF-0747415, CNS-0627476, CNS-0627529, CNS-0845874, CNS-0923313, ECCS-0802107.
Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

figure: treehugger.com

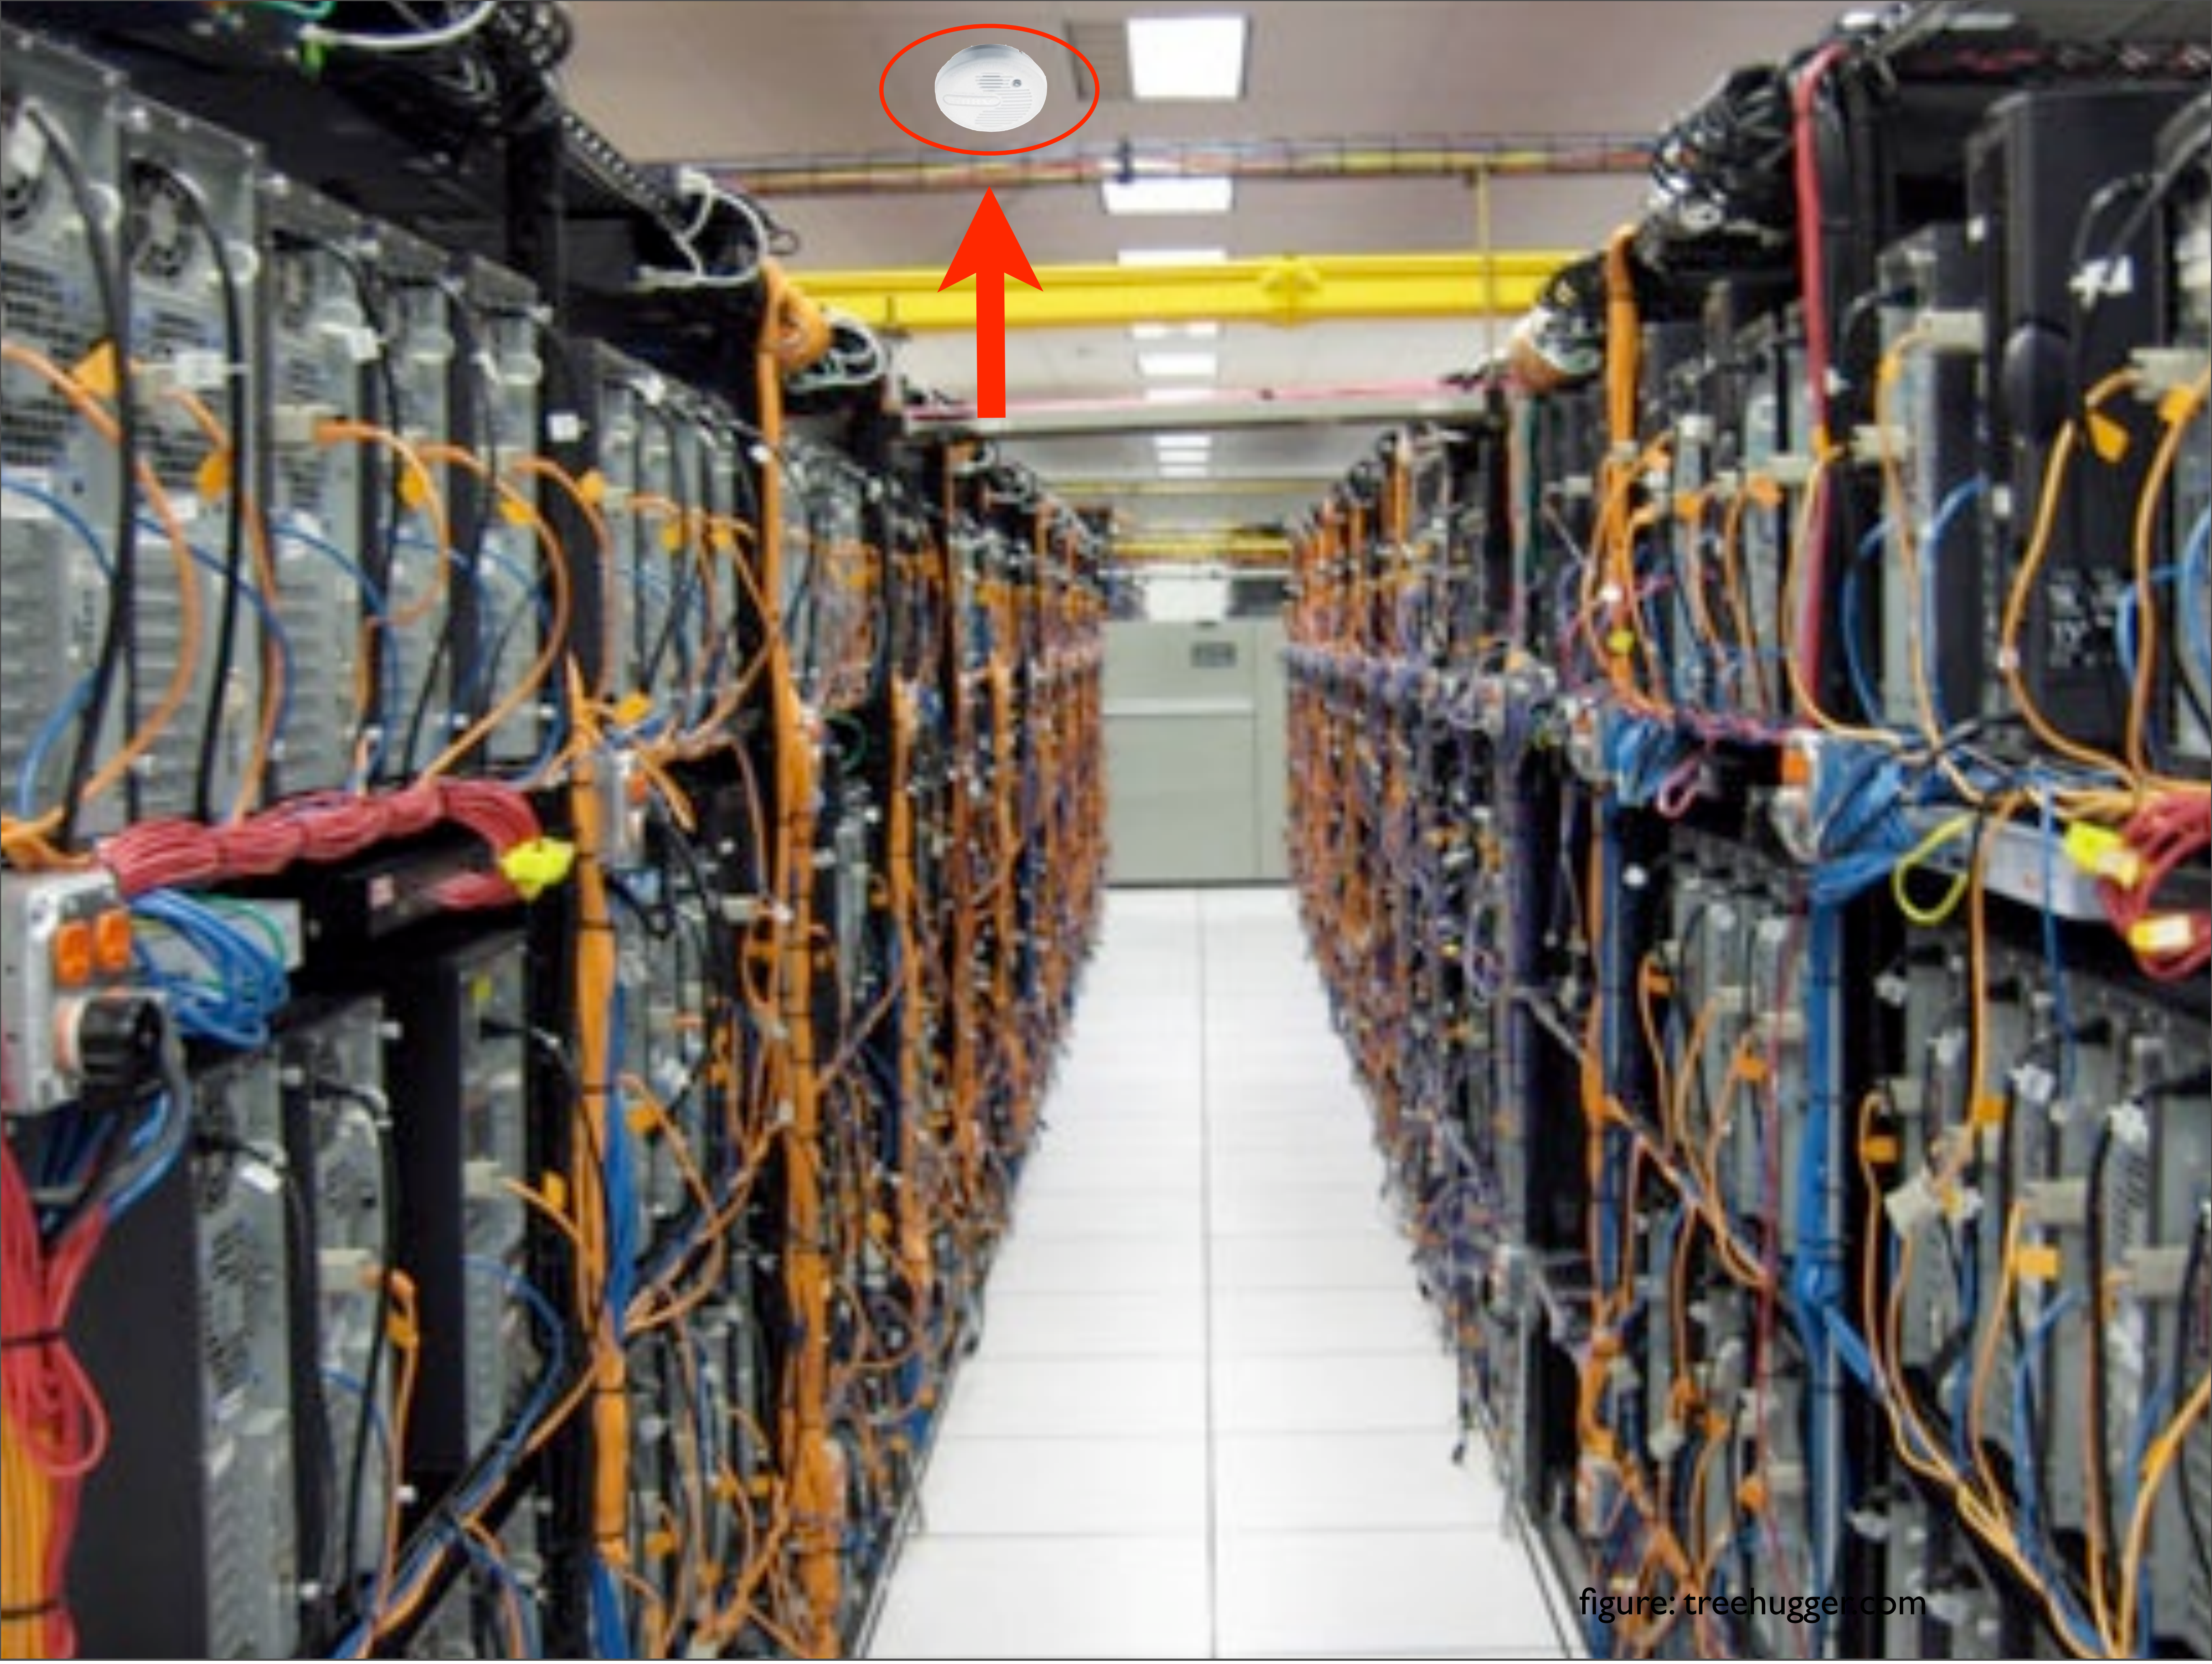
figure: treehugger.com

Storage on Embedded Devices

Storage on Embedded Devices

>$10 billion

Storage on Embedded Devices

>$10 billion

Storage on Embedded Devices

On-chip Flash

On-chip Flash

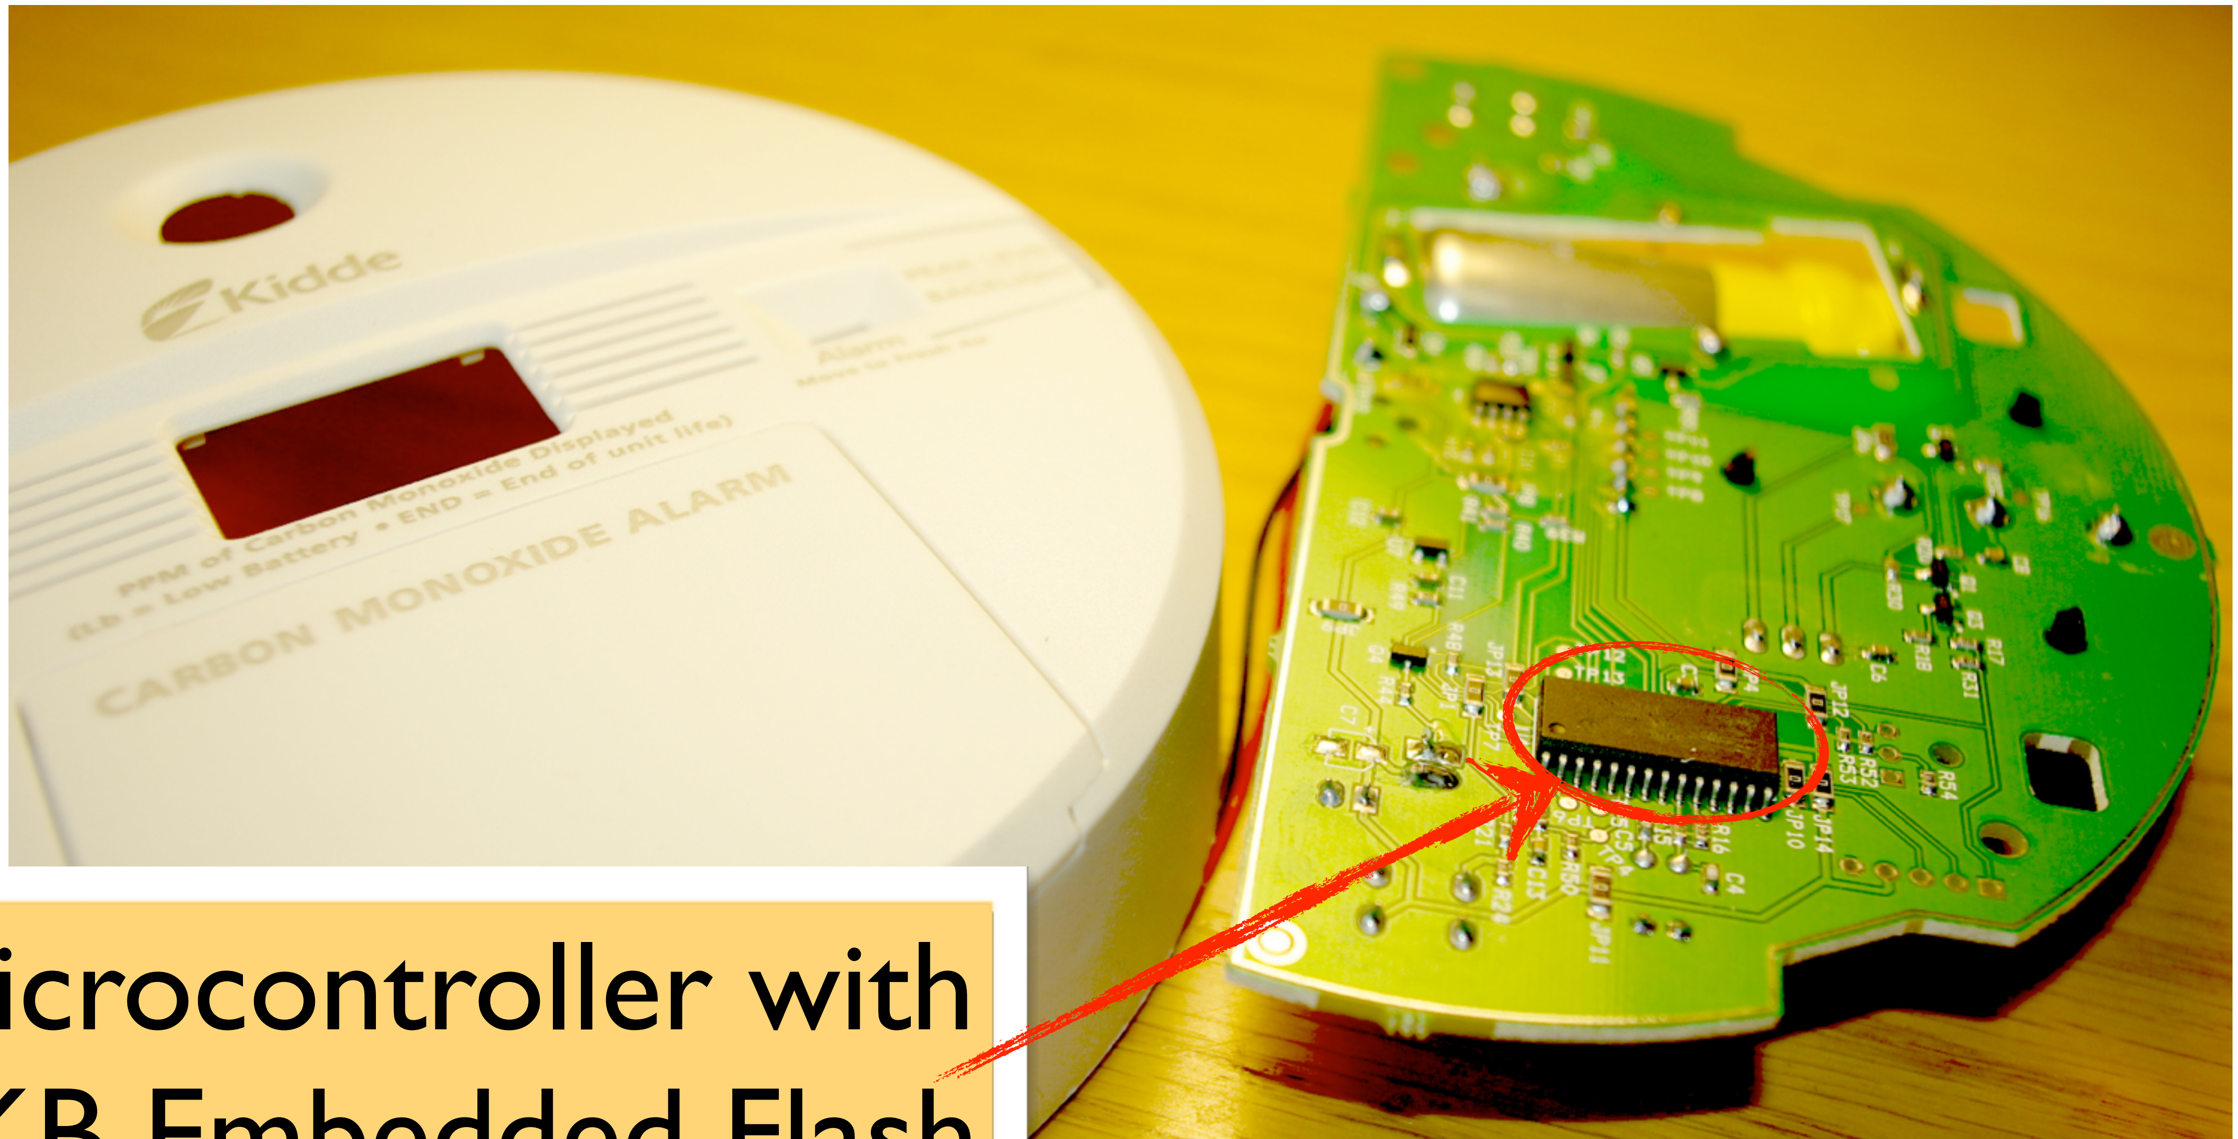

Microcontroller with 8KB Embedded Flash Memory

MICROCHIP

On-chip Flash

2.2 V vs. 4.5 V

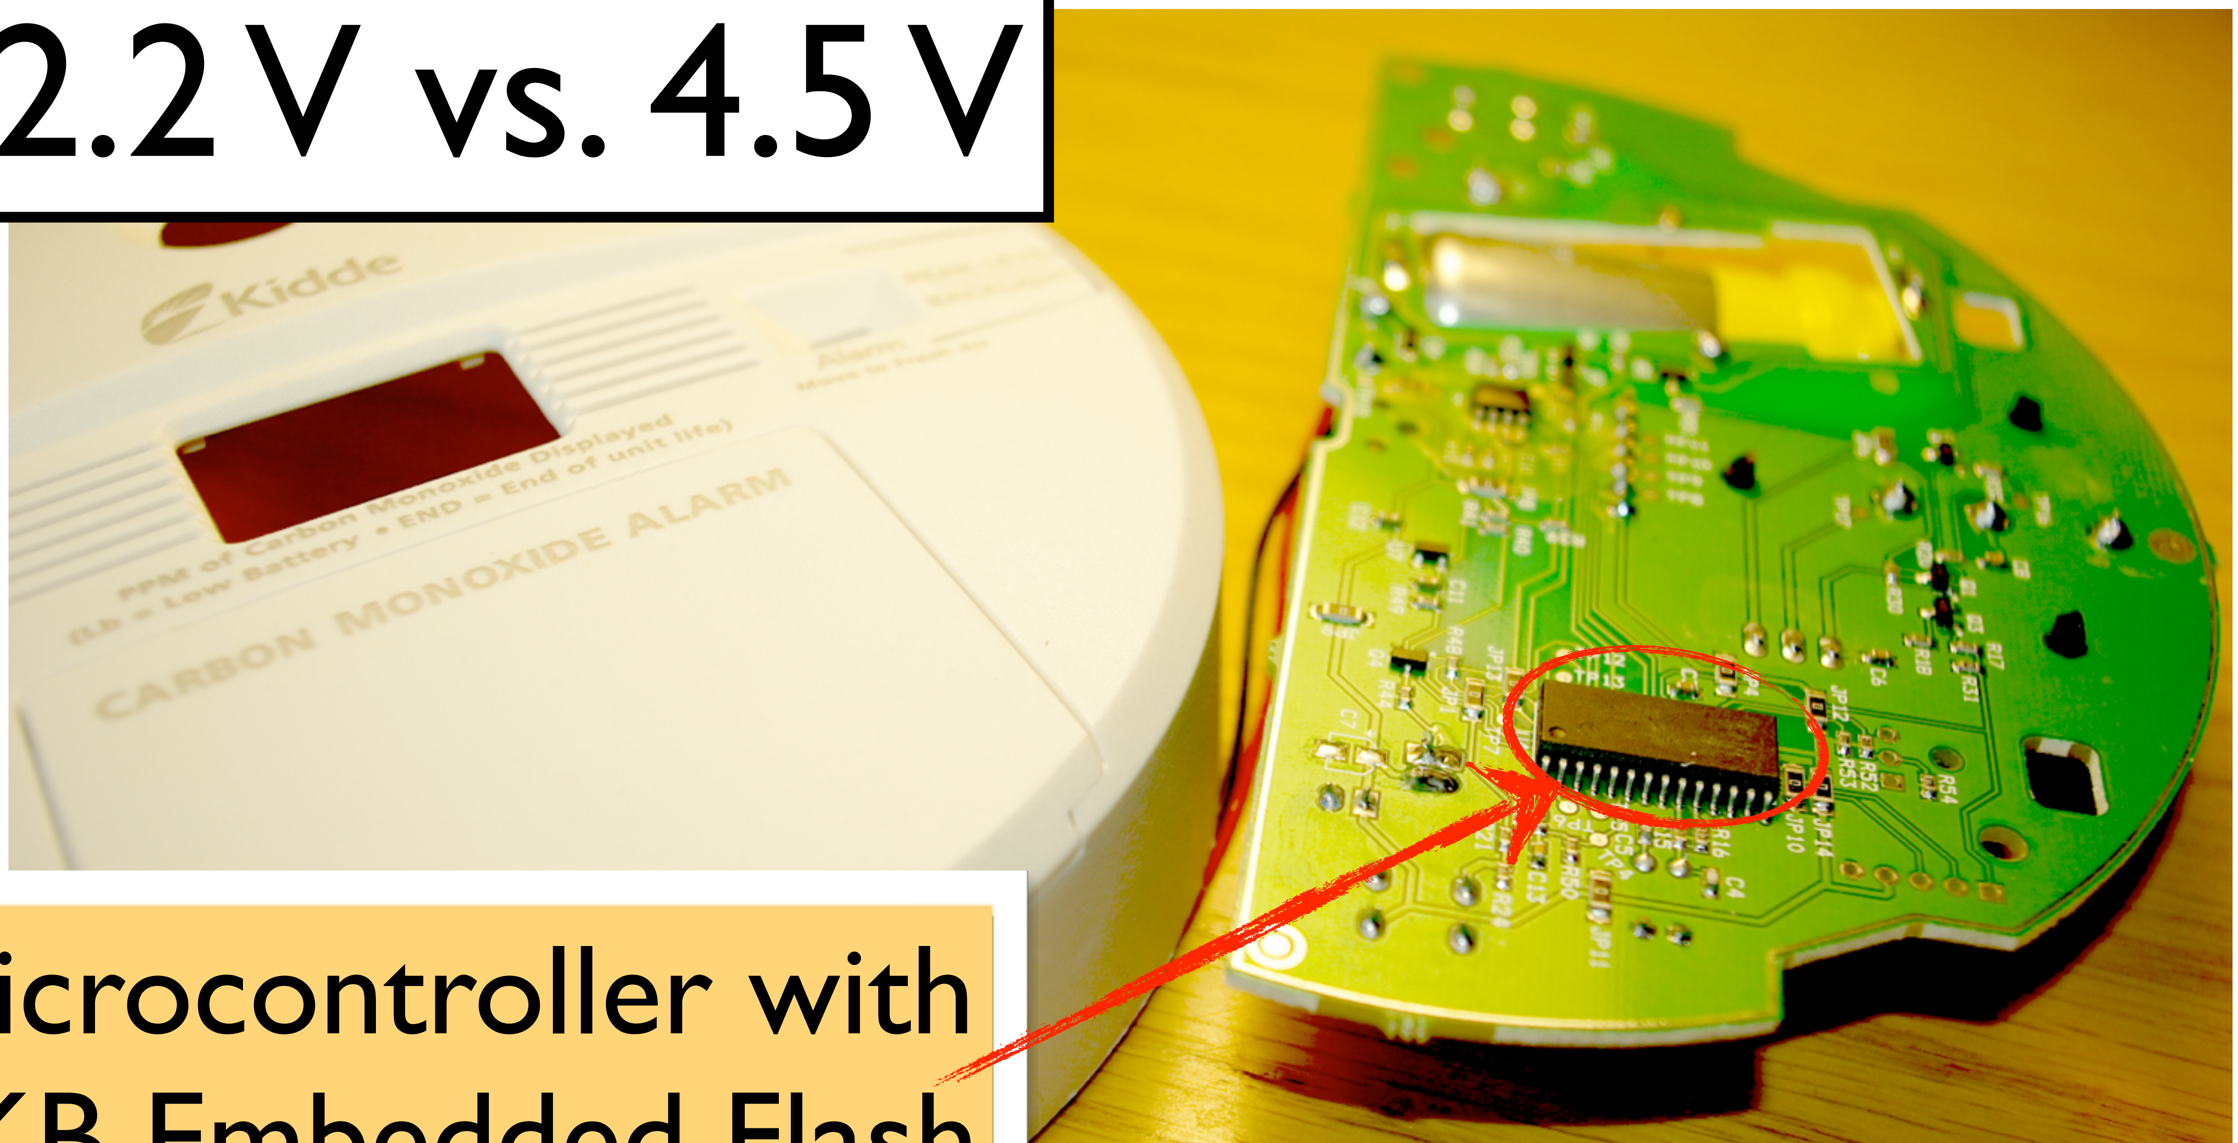

Microcontroller with 8KB Embedded Flash Memory

MICROCHIP

Ideal

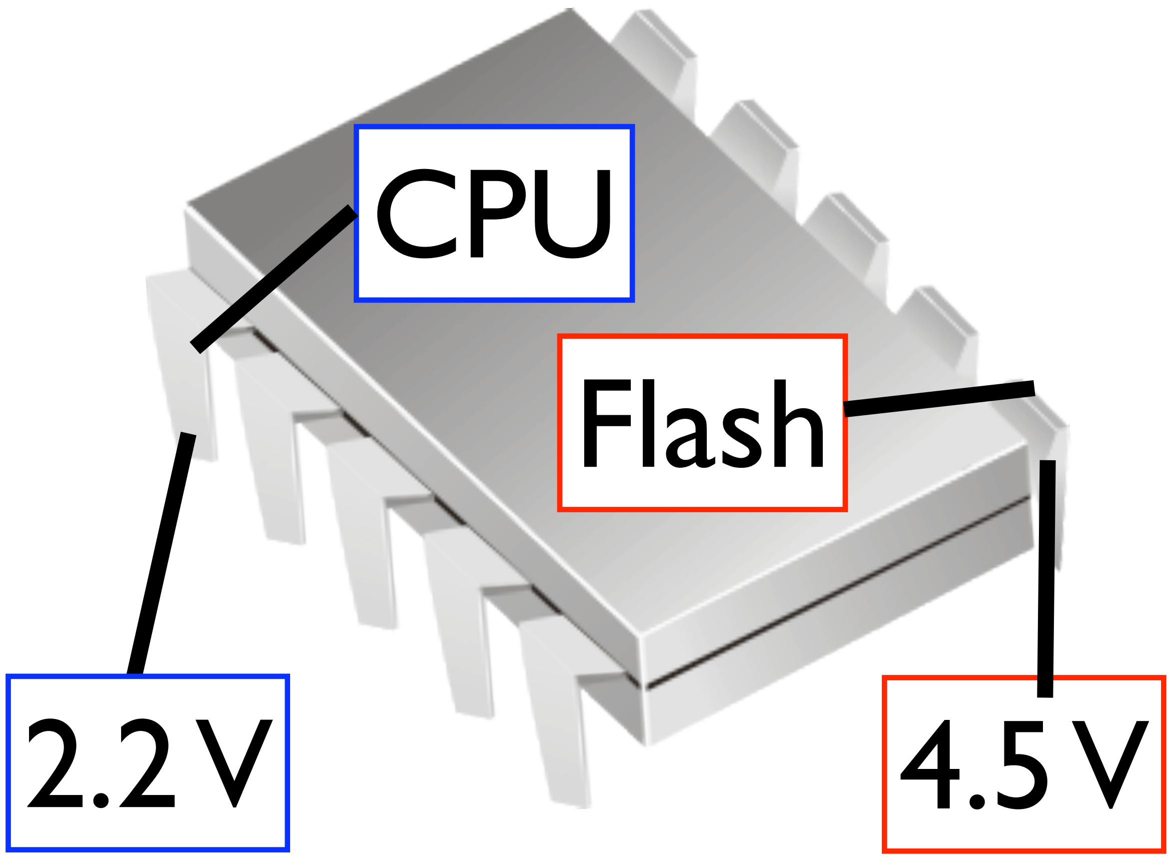

Energy $\propto$ Workload

Ideal	**Actual**

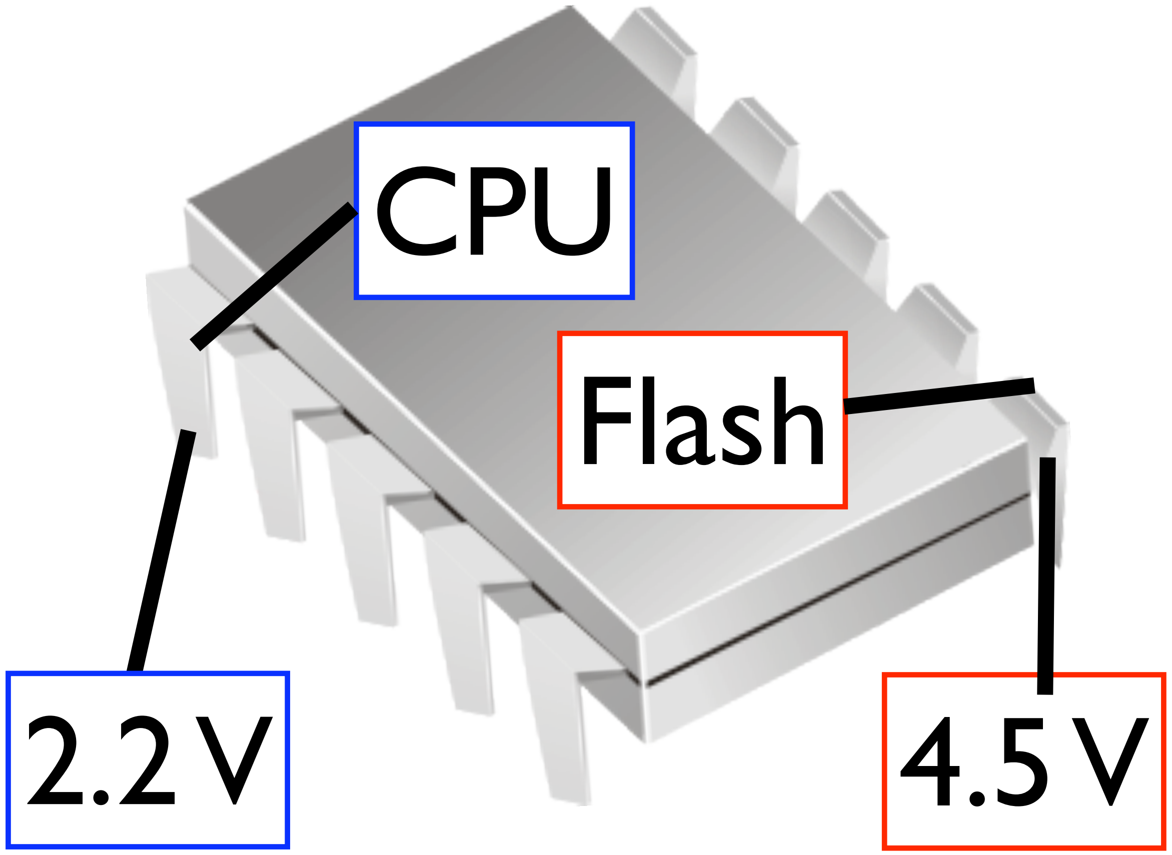

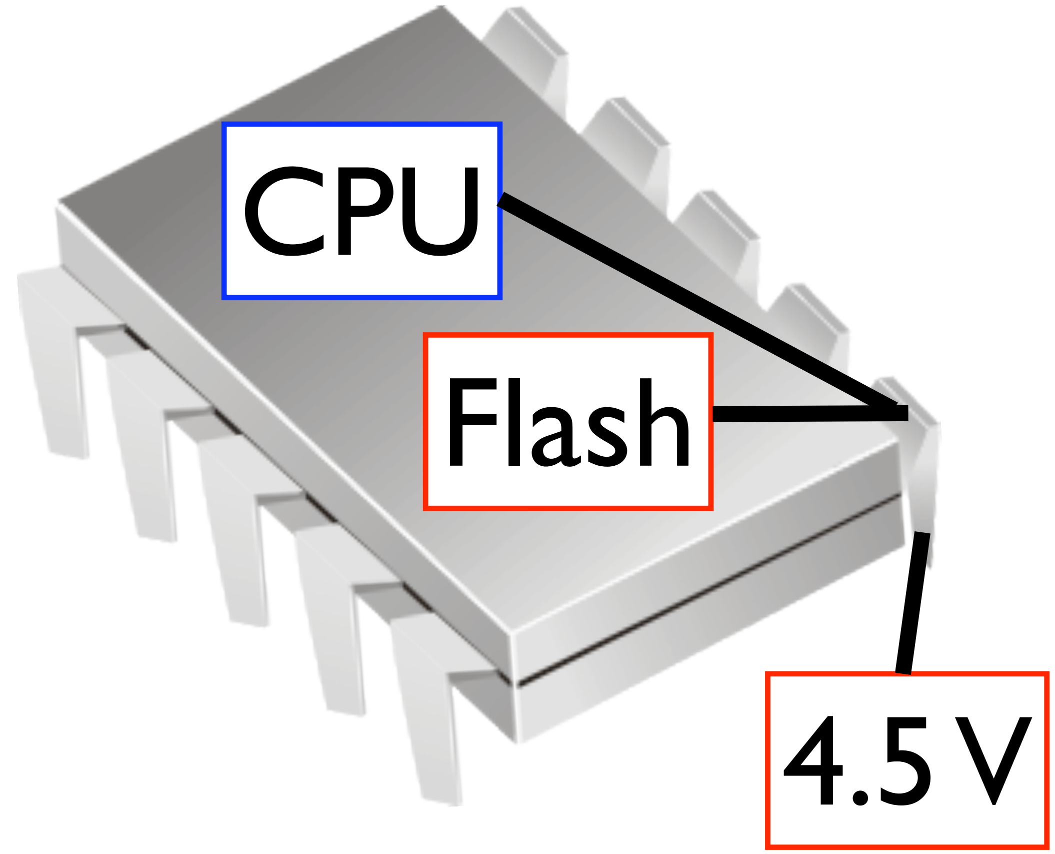

Energy $\propto$ Workload

Energy $\propto$ Worst case

Goal of this work

- To reduce the wasted energy consumption for embedded storage.

Contribution

- Software for using flash memory at low voltage
  - Quantifying the impact on reliability
  - Measuring the energy savings

State of the Art

- Pick the highest voltage...Excessive power

State of the Art

- Pick the highest voltage...Excessive power
- Add hardware...$$

State of the Art

- Pick the highest voltage...Excessive power
- Add hardware...$$
- Don't use flash memory...Ugh!

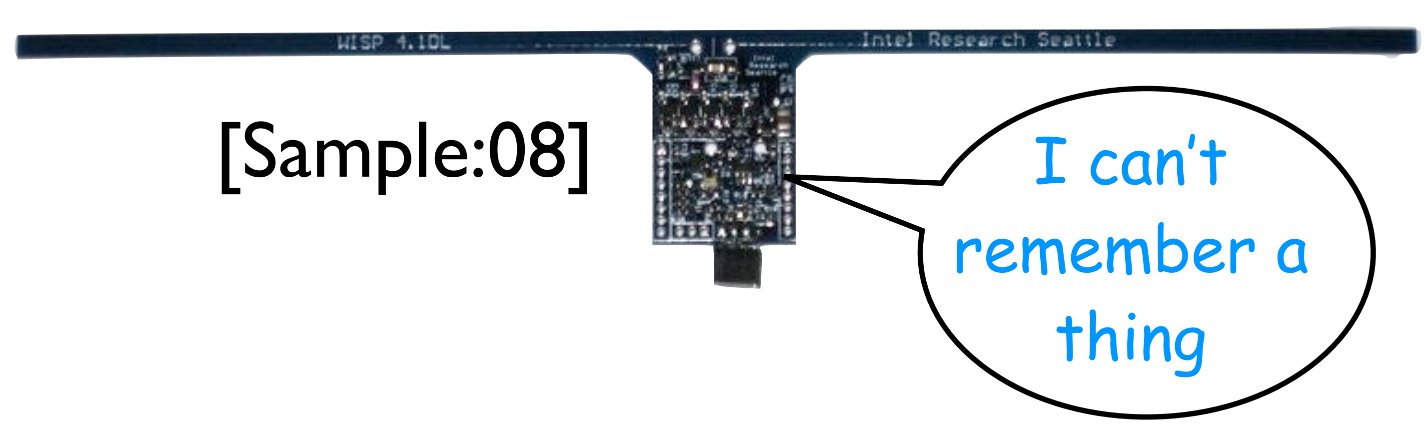

[Sample:08]

State of the Art

- Pick the highest voltage...Excessive power
- Add hardware...$$
- Don't use flash memory...Ugh!

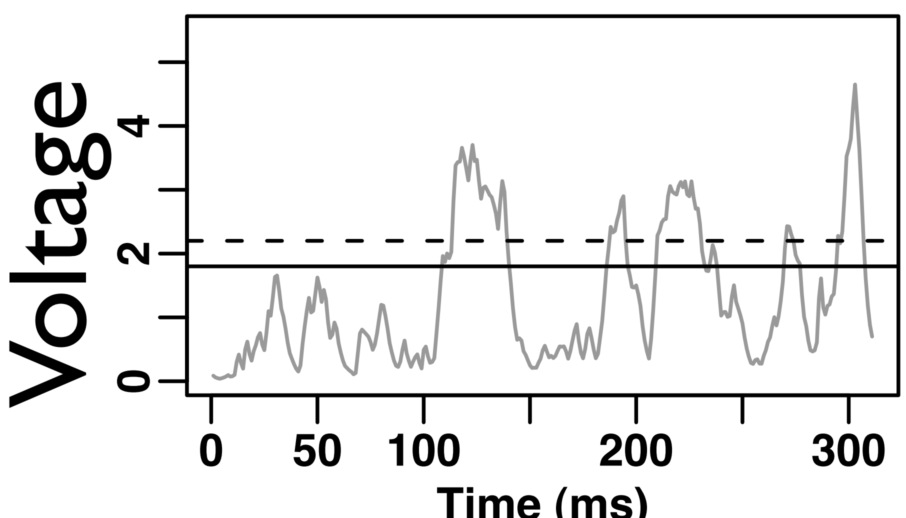

[Sample:08]

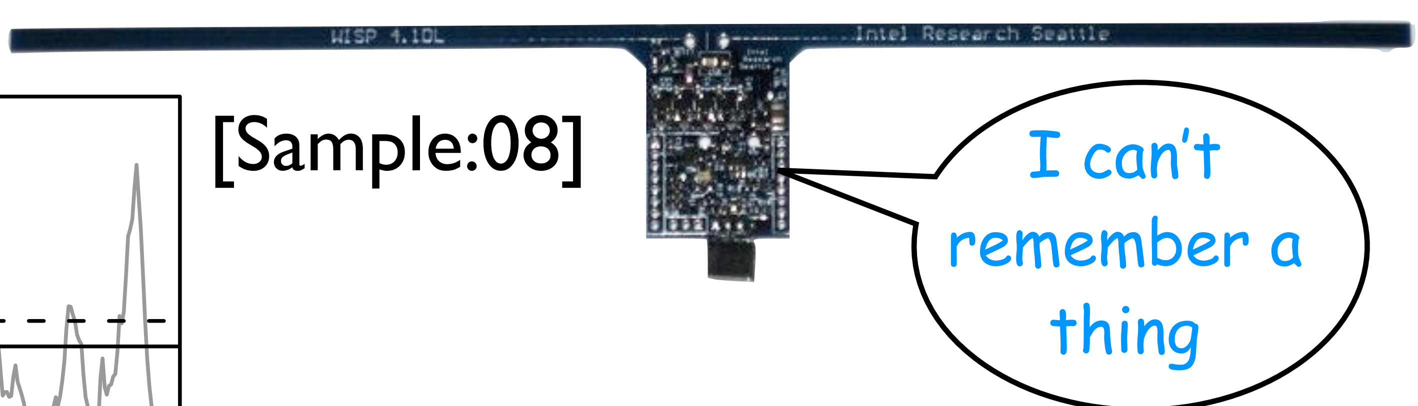

[Ransford: ASPLOS11]

Our Approach

Savings:
Low-voltage

Write to flash memory at low voltage.

Our Approach

Savings:
Low-voltage

Write to flash memory at low voltage.

Cost:
Errors

How hard is it to correct the errors?

Write once bits (Wits) [Rivest:82]

figure: http://arcweb.archives.gov

Partial Failure at Low Voltage

- Example:

Initialized: 1 1 1 1 1 1 1 1

Partial Failure at Low Voltage

- Example:

Initialized: **1111 1111**

Input: **1111 1100**

Partial Failure at Low Voltage

- Example:

Initialized: **1111 1111**

Input: **1111 1100**

Partial Failure at Low Voltage

- Example:

Initialized: 1111 1111

Input: 1111 1100

Result: 1111 1101

Partial Failure at Low Voltage

- Example:

Initialized: 1111 1111

Input: 1111 1100

Result: 1111 1101

Error

Transitions at low voltage

- 1→0 might fail with P≥0.
- 1→1 never fails (P=0).

[Klove:95]

What might influence the error rate

What might influence the error rate

✓ Operating voltage level

What might influence the error rate

✓ Operating voltage level

✓ Hamming weight of data

What might influence the error rate

✓ Operating voltage level

✓ Hamming weight of data

✓ Wear-out history

What might influence the error rate

- ✓ Operating voltage level
- ✓ Hamming weight of data
- ✓ Wear-out history
- – Neighbor cells

What might influence the error rate

✓ Operating voltage level

✓ Hamming weight of data

✓ Wear-out history

- Neighbor cells

- Permutation of 0s

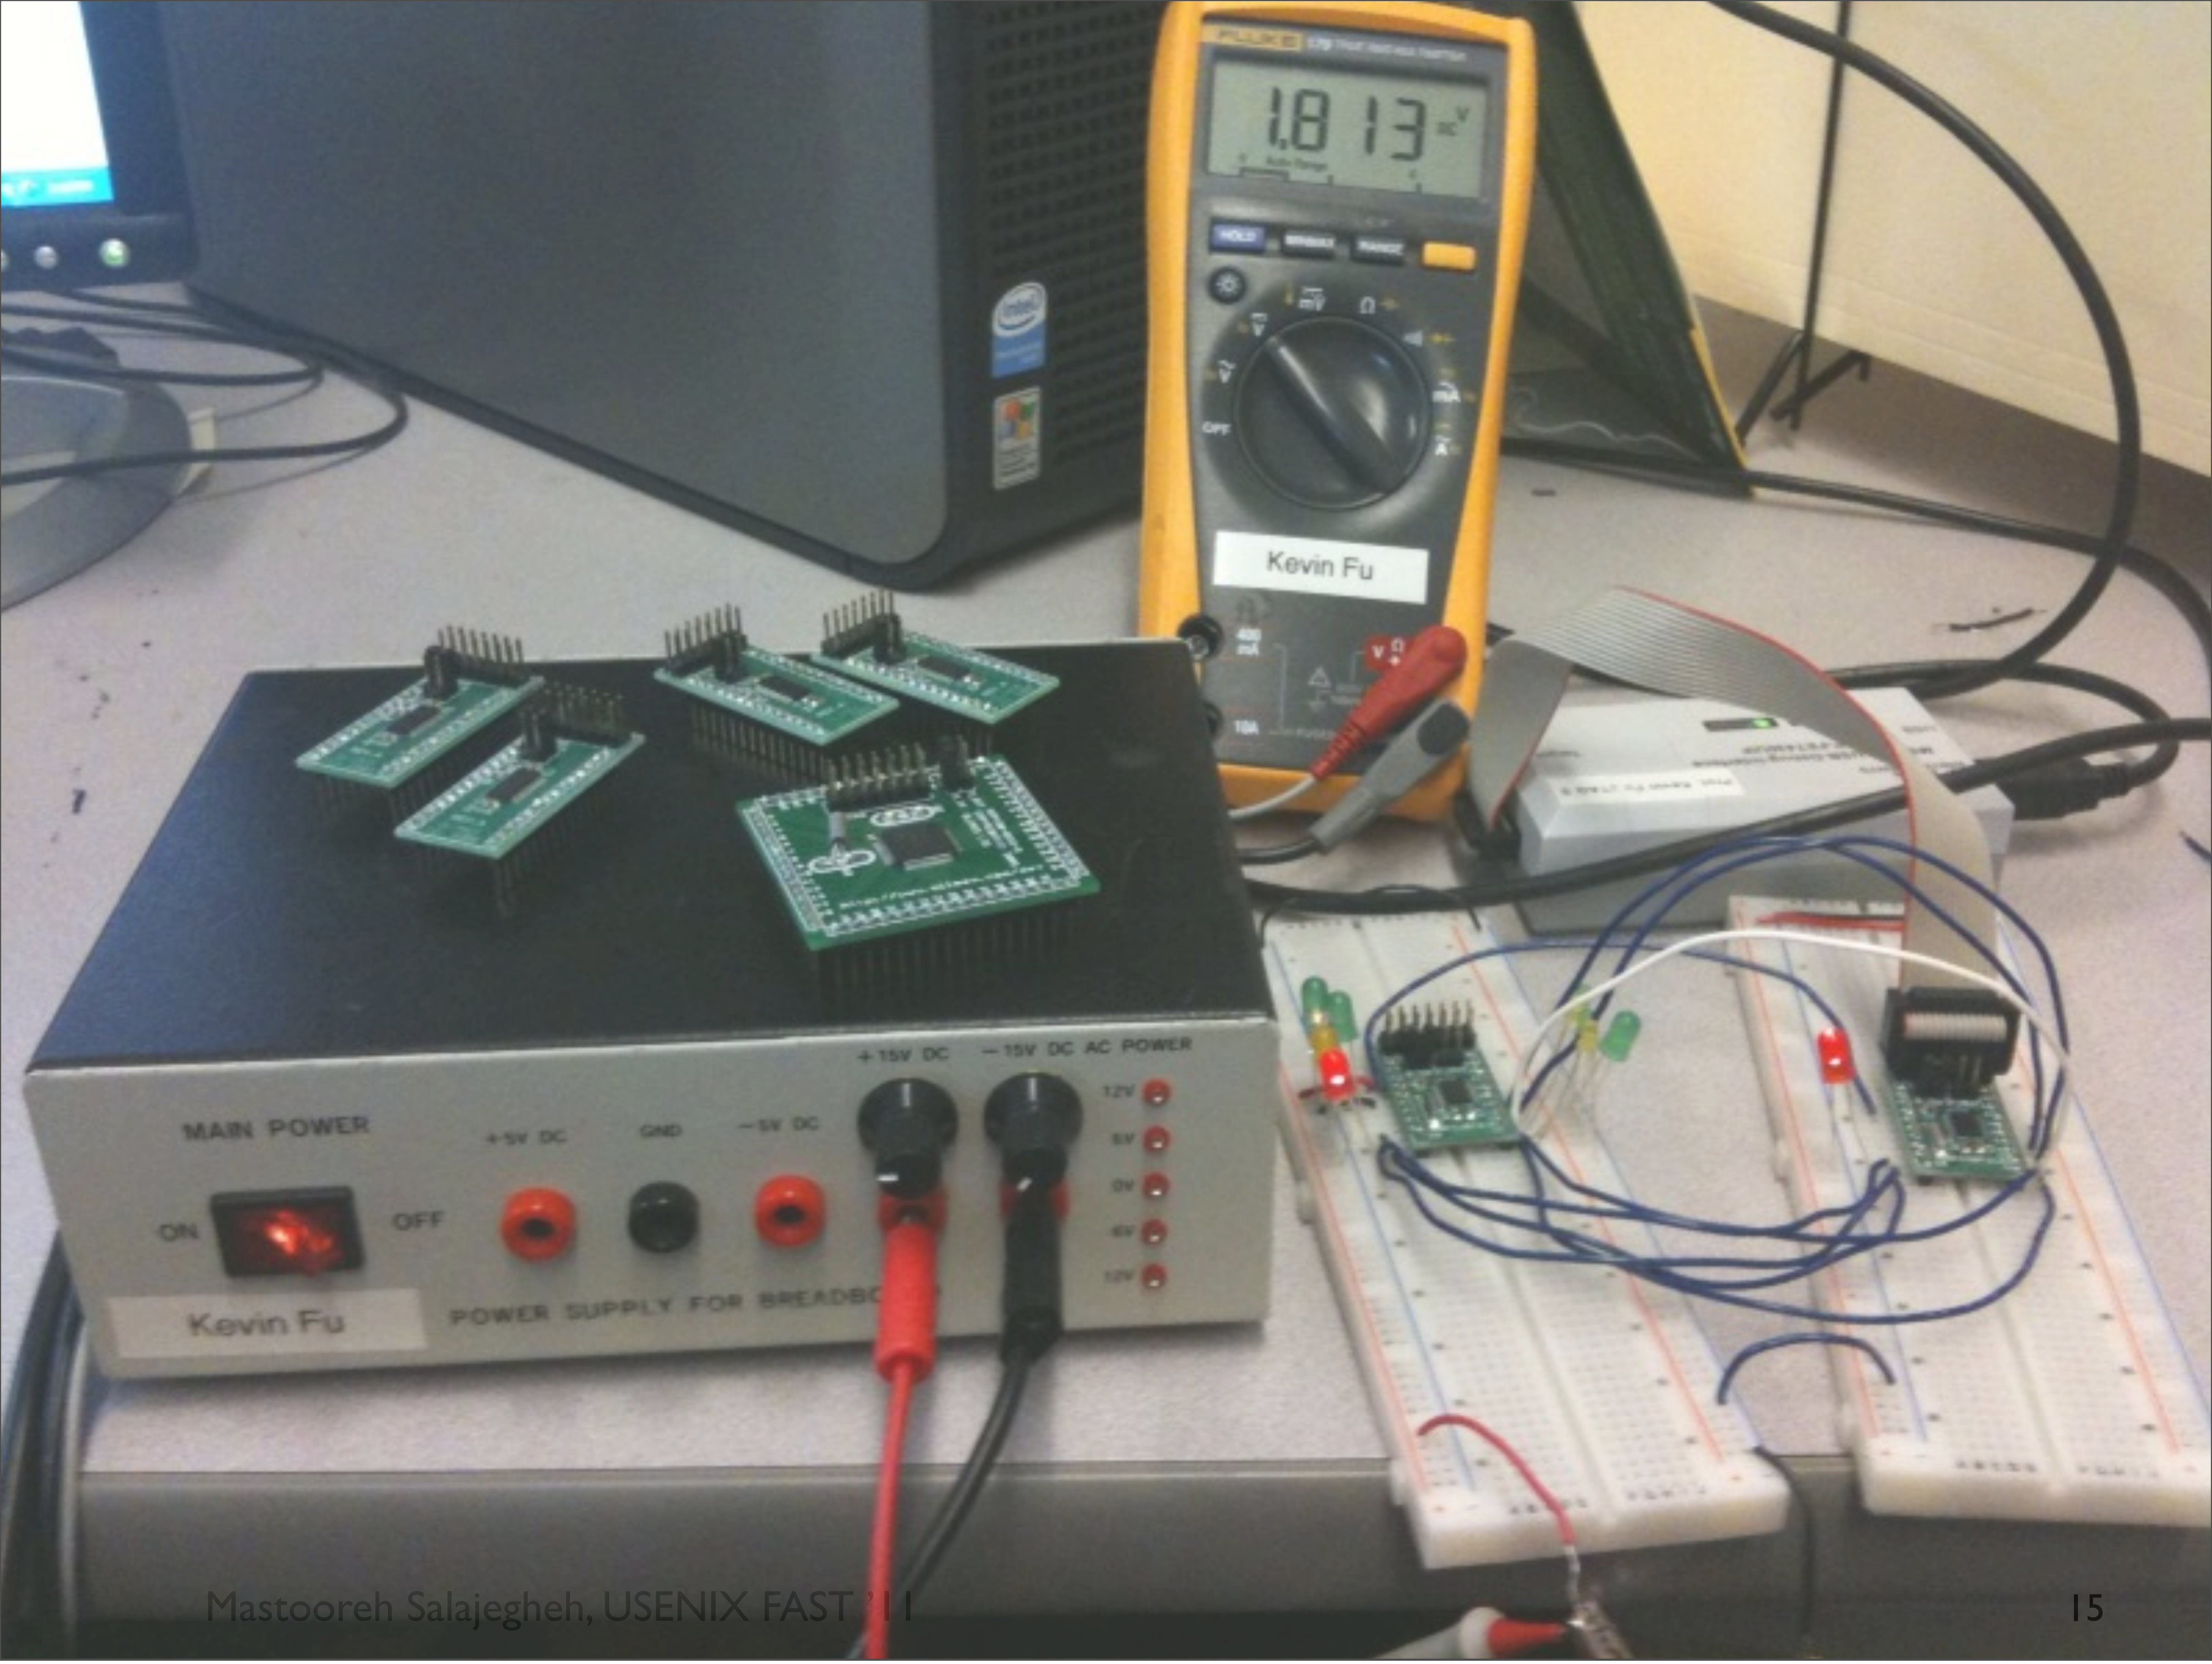
1.813
Kevin Fu
MAIN POWER
ON
OFF
+5V DC
GND
−5V DC
+15V DC
−15V DC AC POWER
Kevin Fu

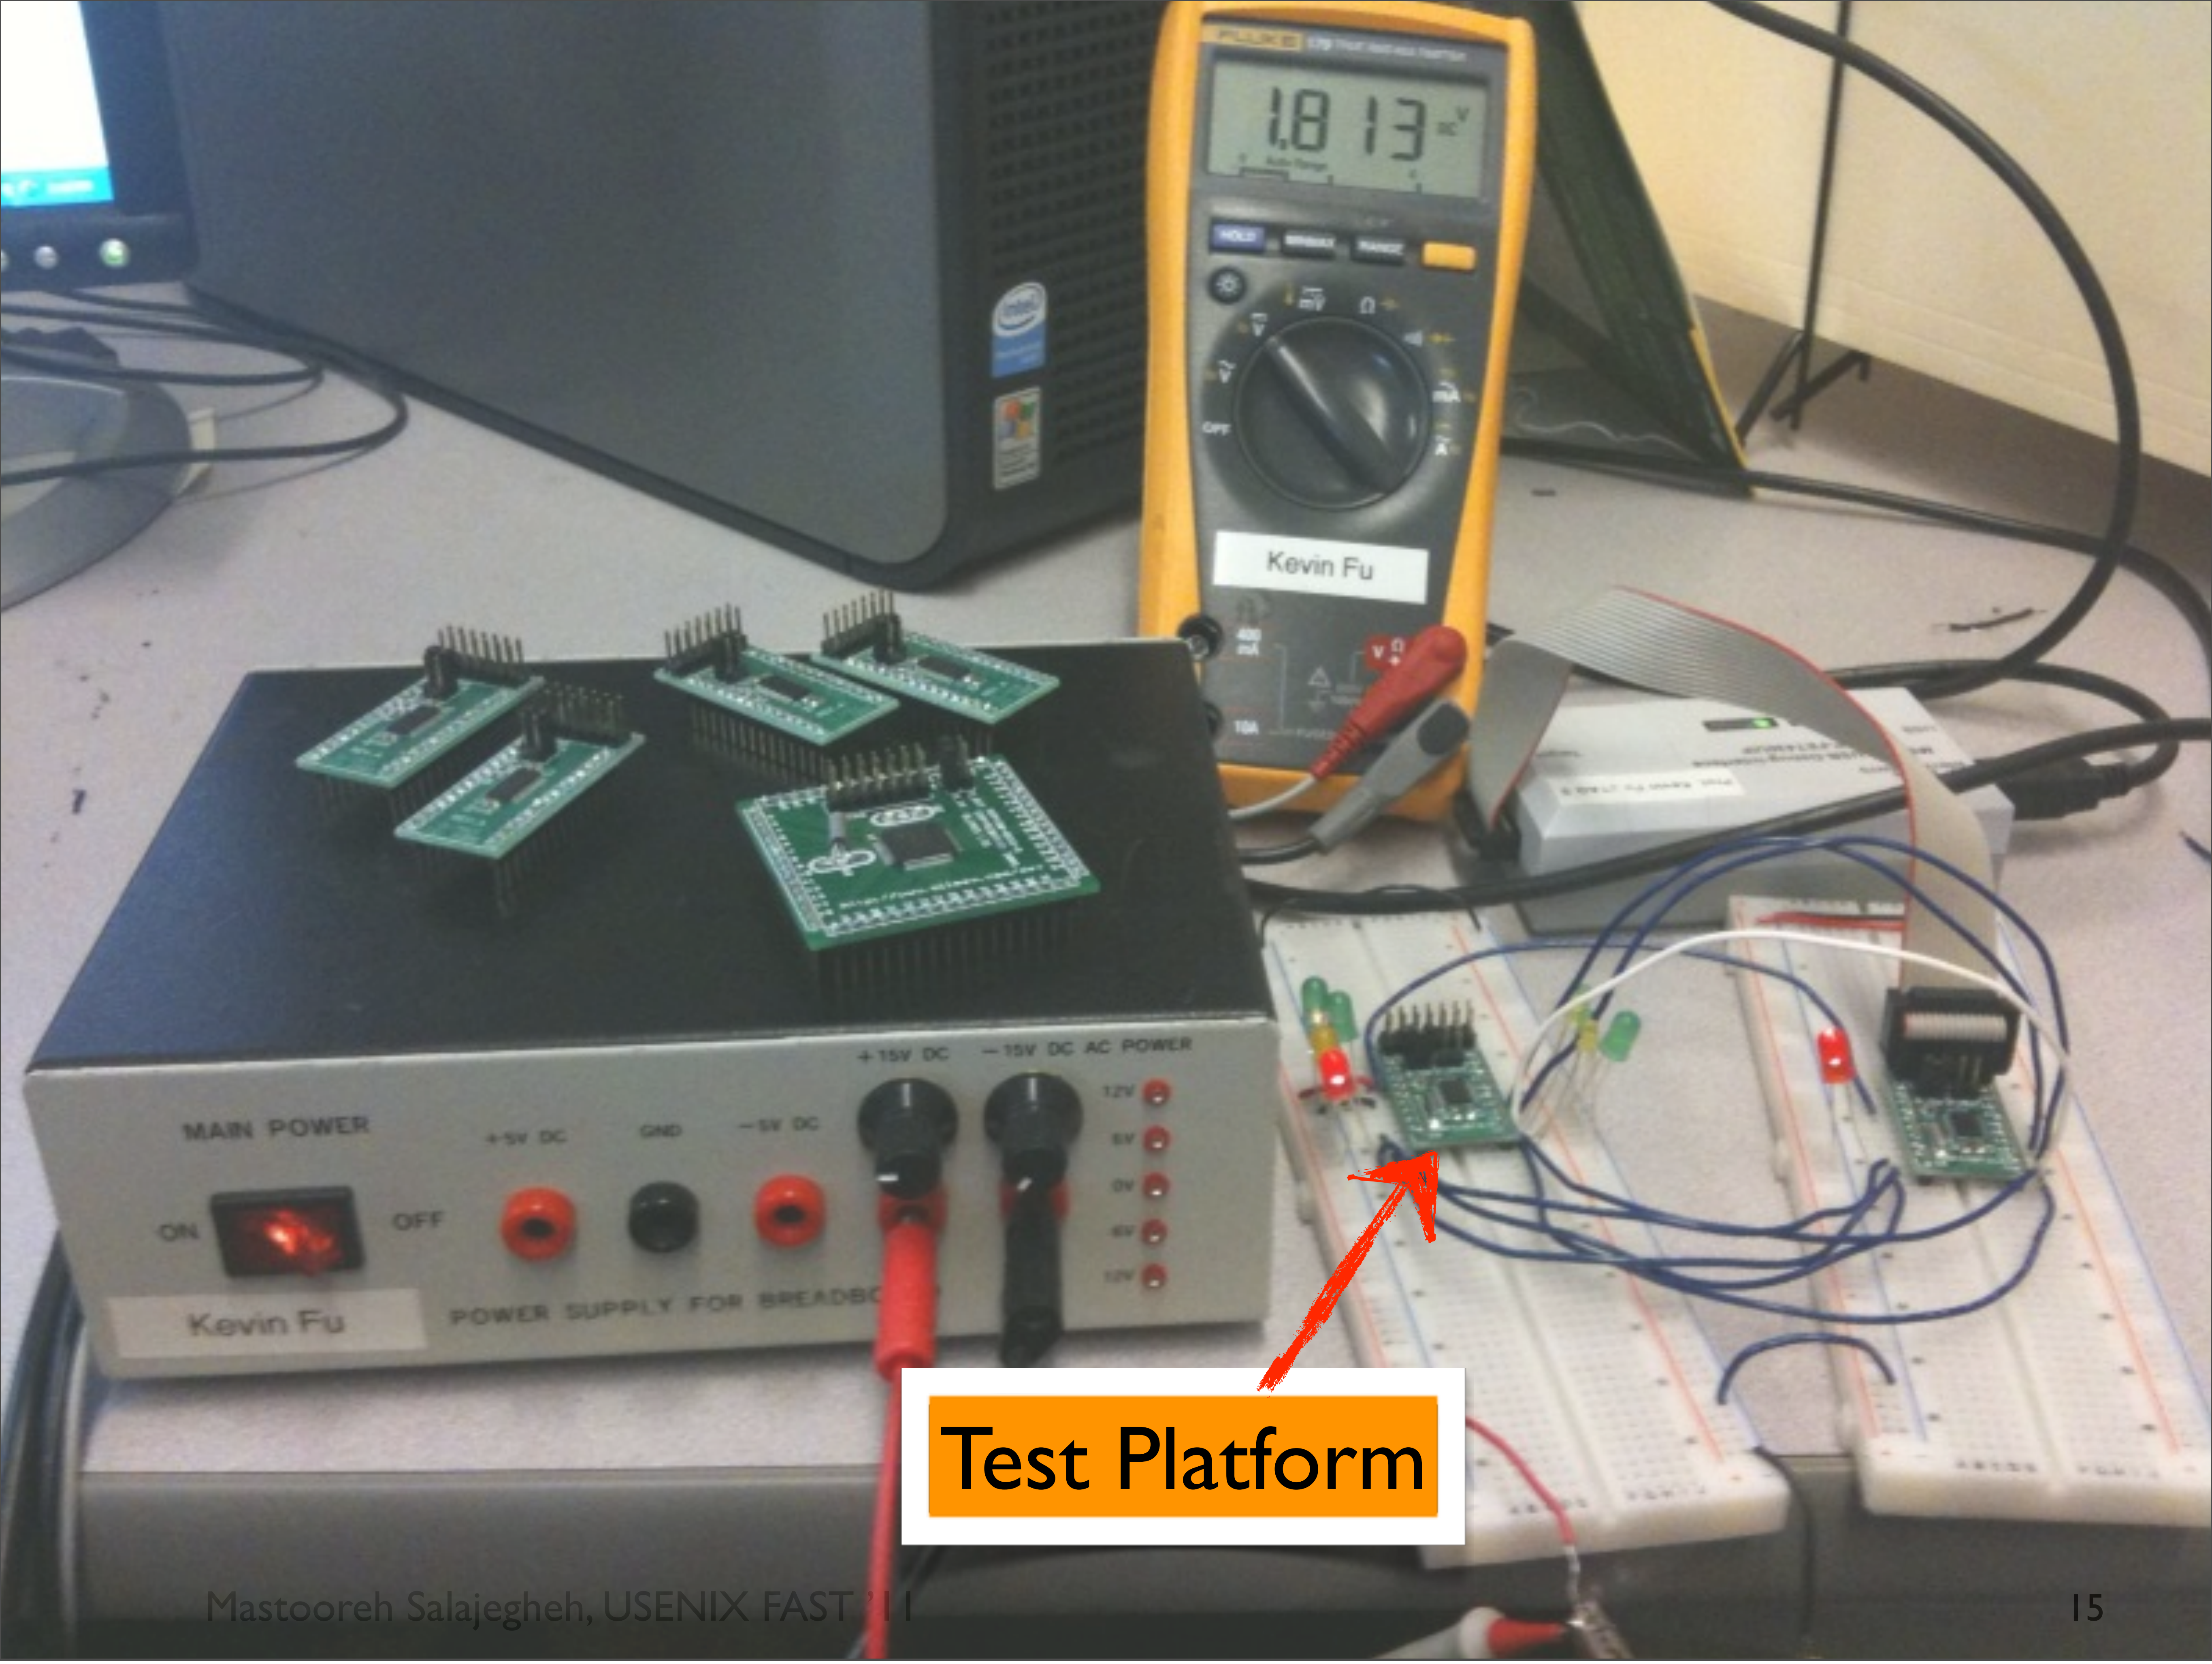
1.813
Kevin Fu
MAIN POWER
ON
OFF
+5V DC
GND
−5V DC
+15V DC
−15V DC AC POWER
Kevin Fu
Test Platform

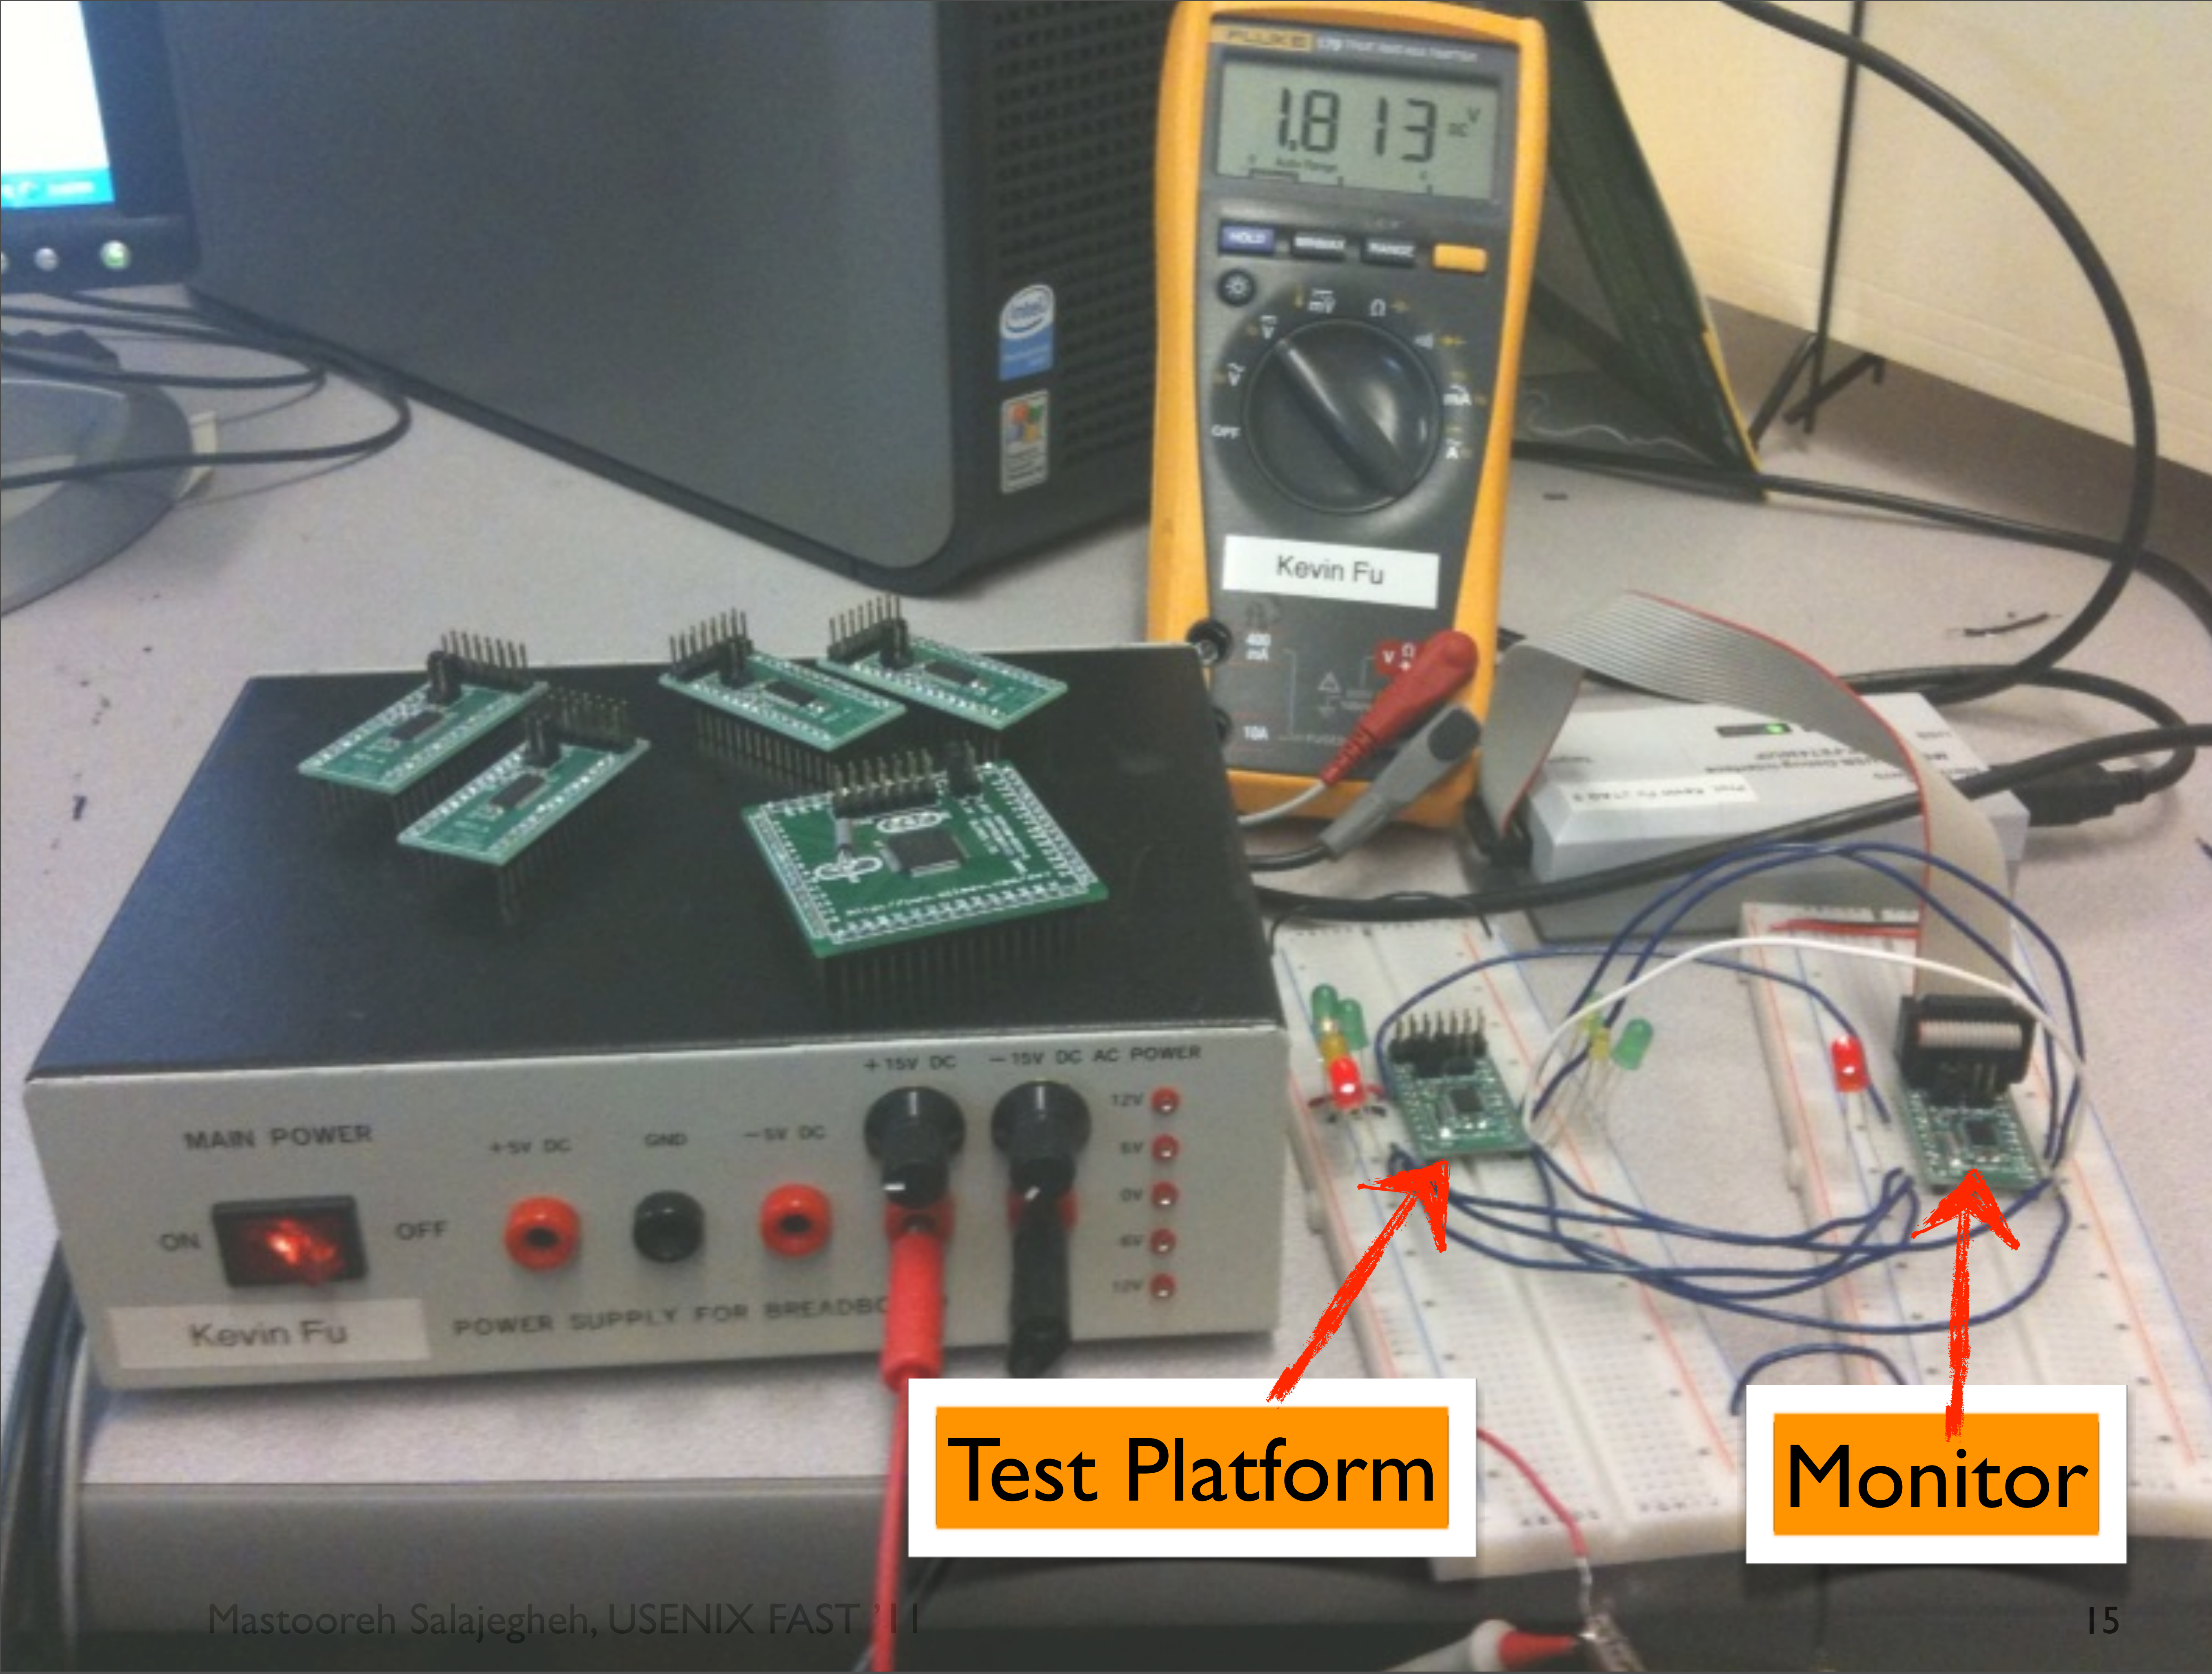
Test Platform
Monitor
Kevin Fu
1.813
MAIN POWER
ON
OFF
+5V DC
GND
-5V DC
+15V DC
-15V DC
AC POWER
12V
5V
0V
-5V
-12V
Kevin Fu
POWER SUPPLY FOR BREADBO

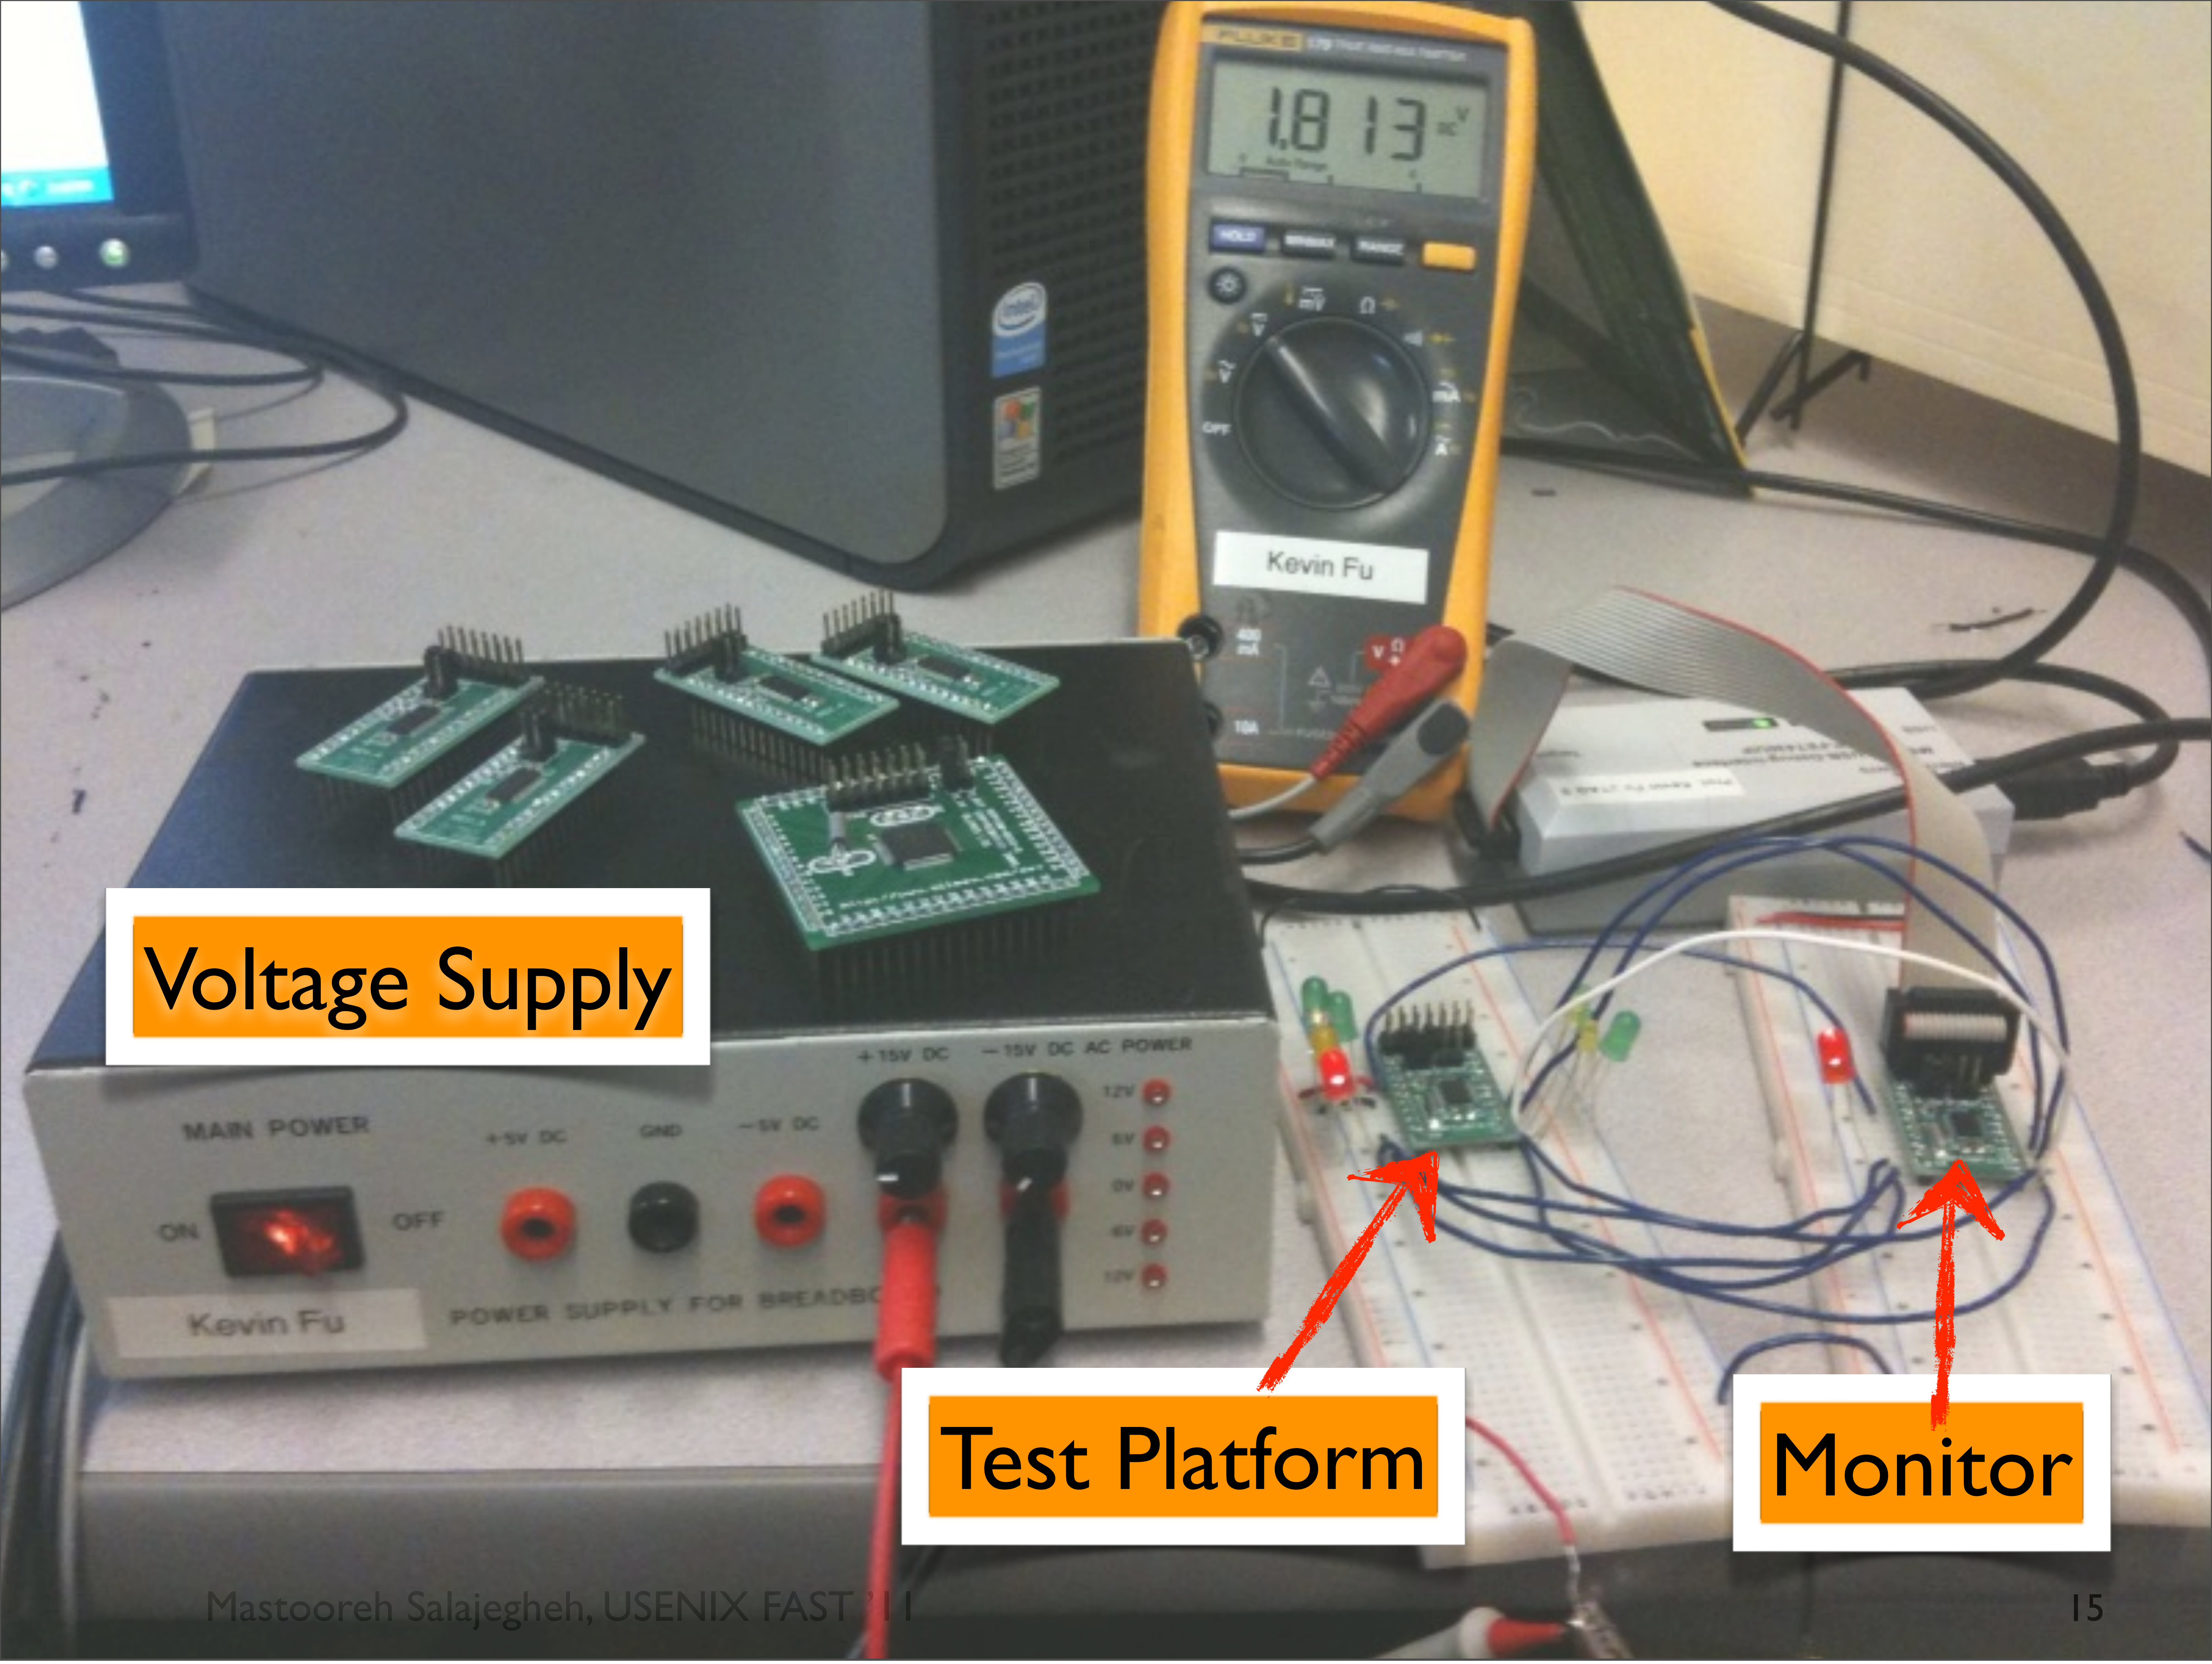
Voltage Supply
Test Platform
Monitor
1.813
Kevin Fu
MAIN POWER
ON
OFF
+5V DC
GND
−5V DC
+15V DC
−15V DC
AC POWER
Kevin Fu

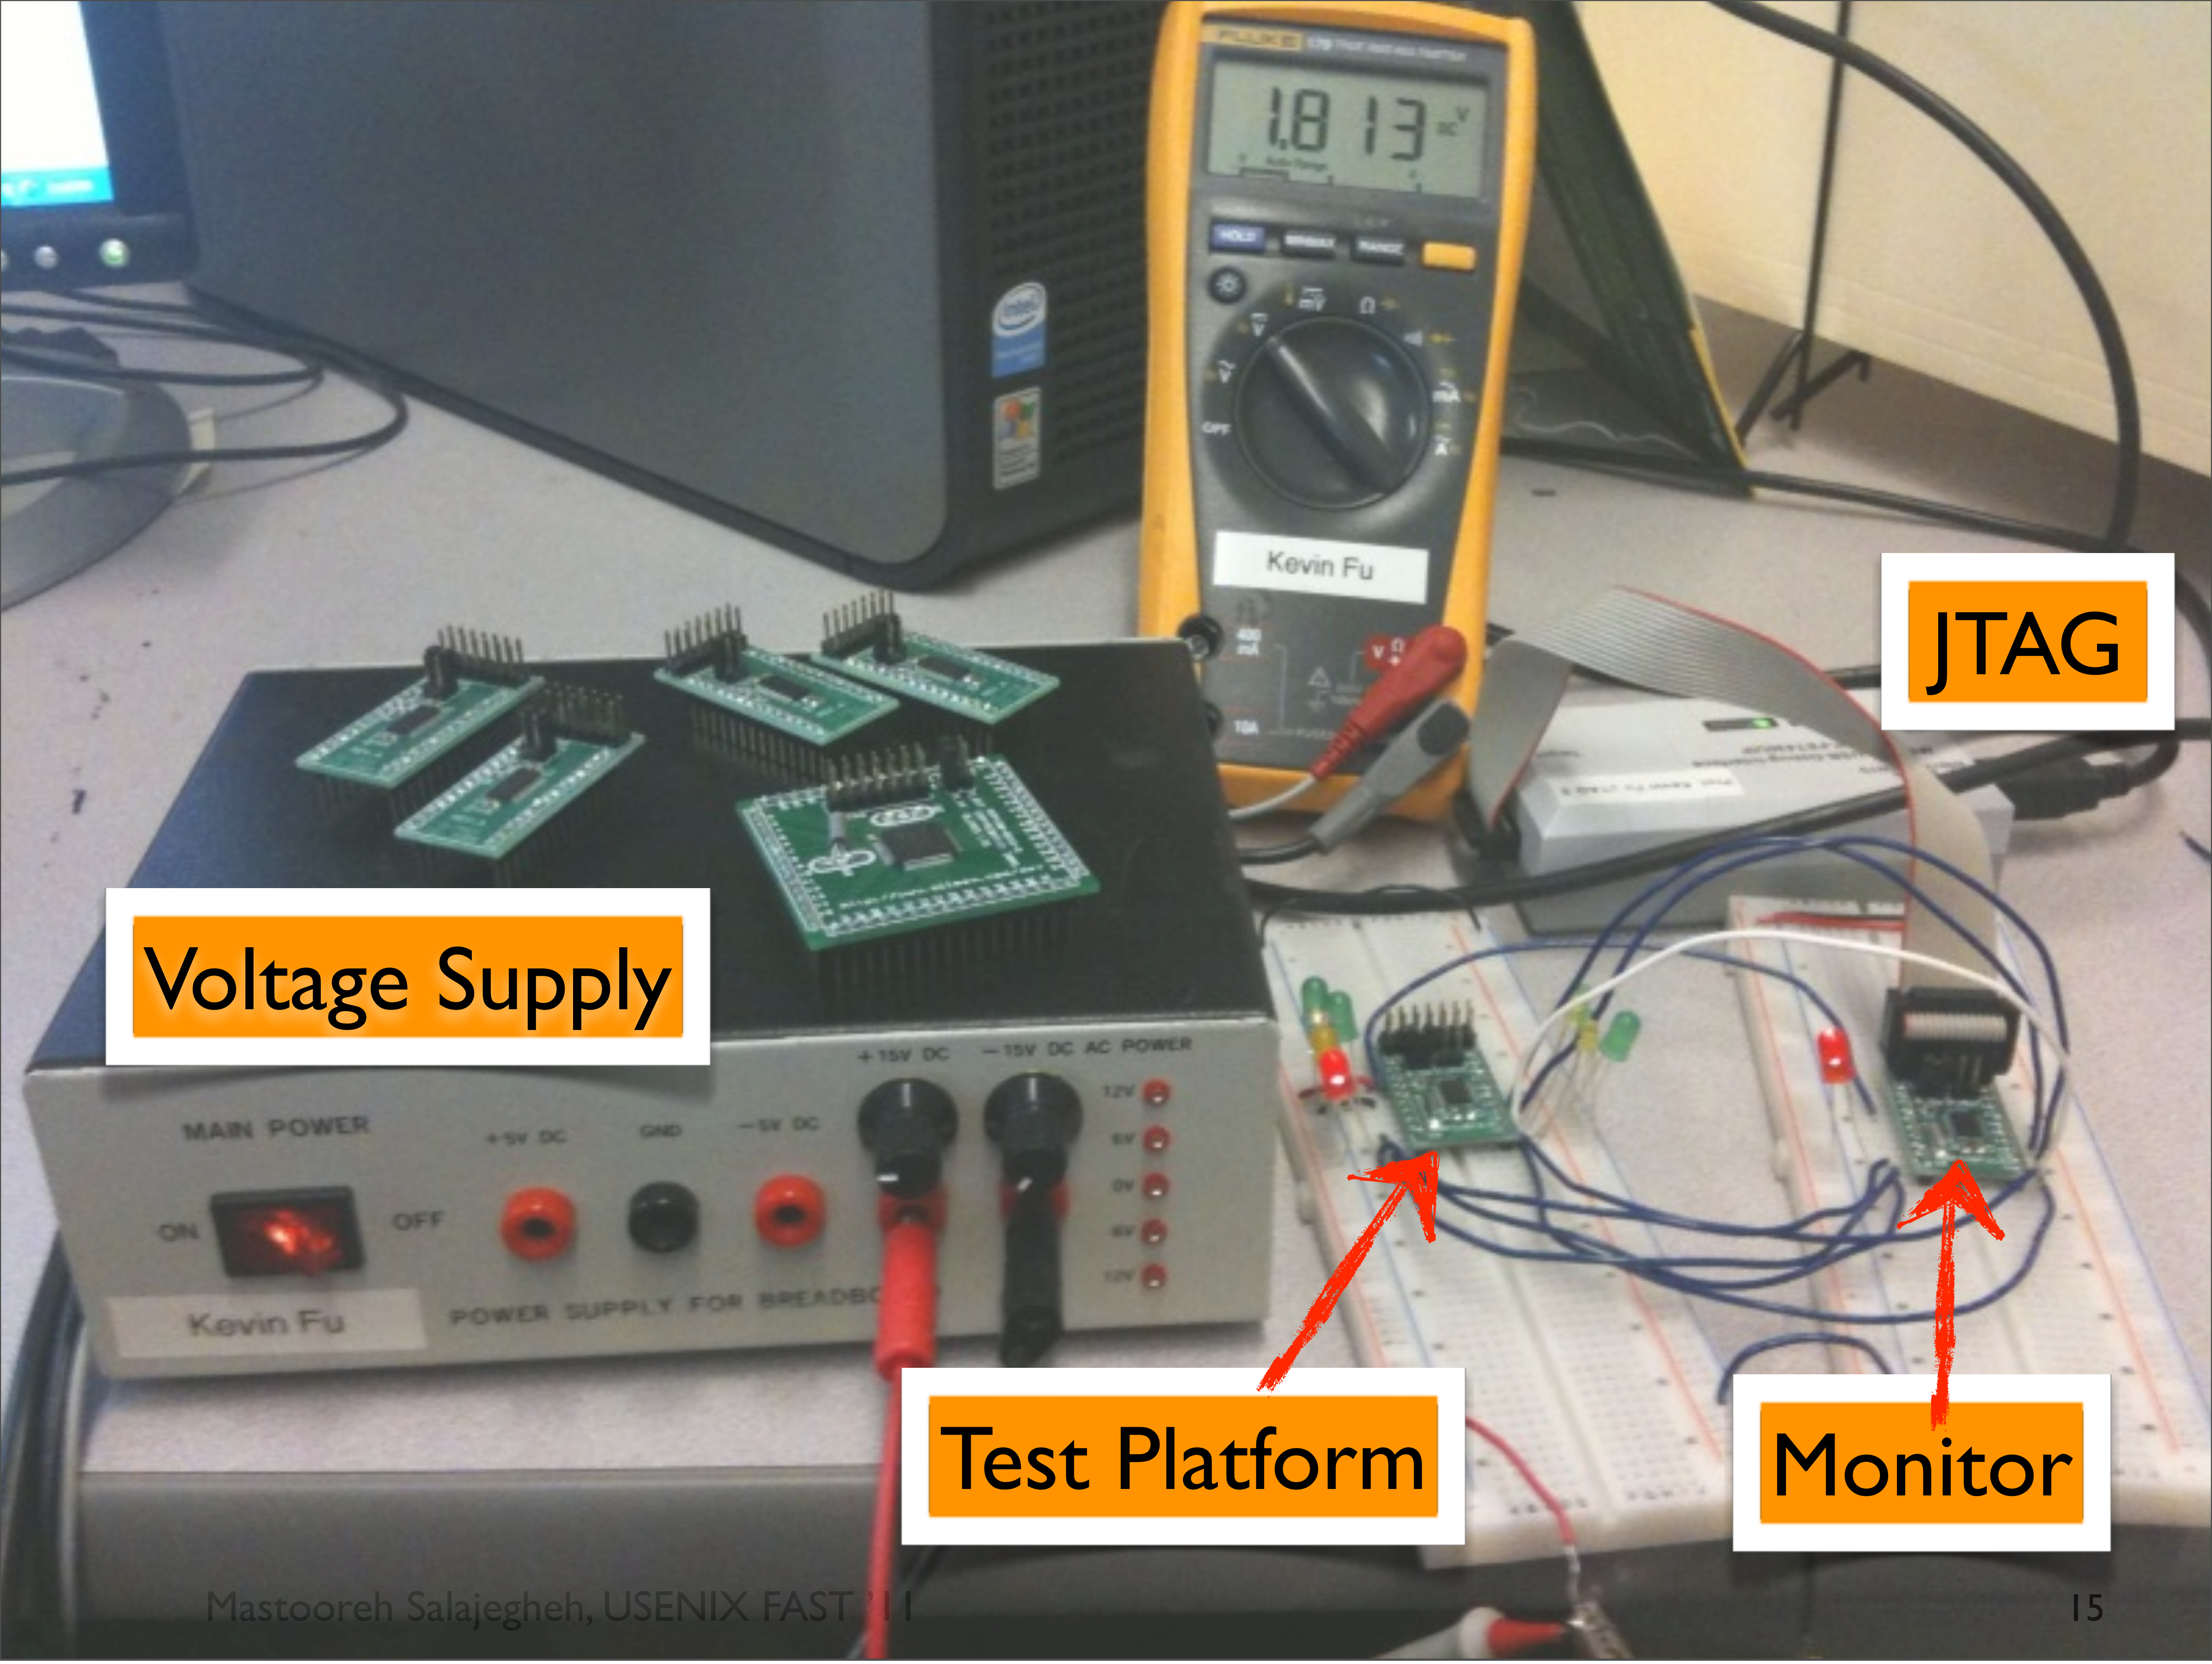
JTAG
Voltage Supply
Test Platform
Monitor
Kevin Fu
MAIN POWER
ON
OFF
+5V DC
GND
-5V DC
+15V DC
-15V DC
AC POWER
Kevin Fu

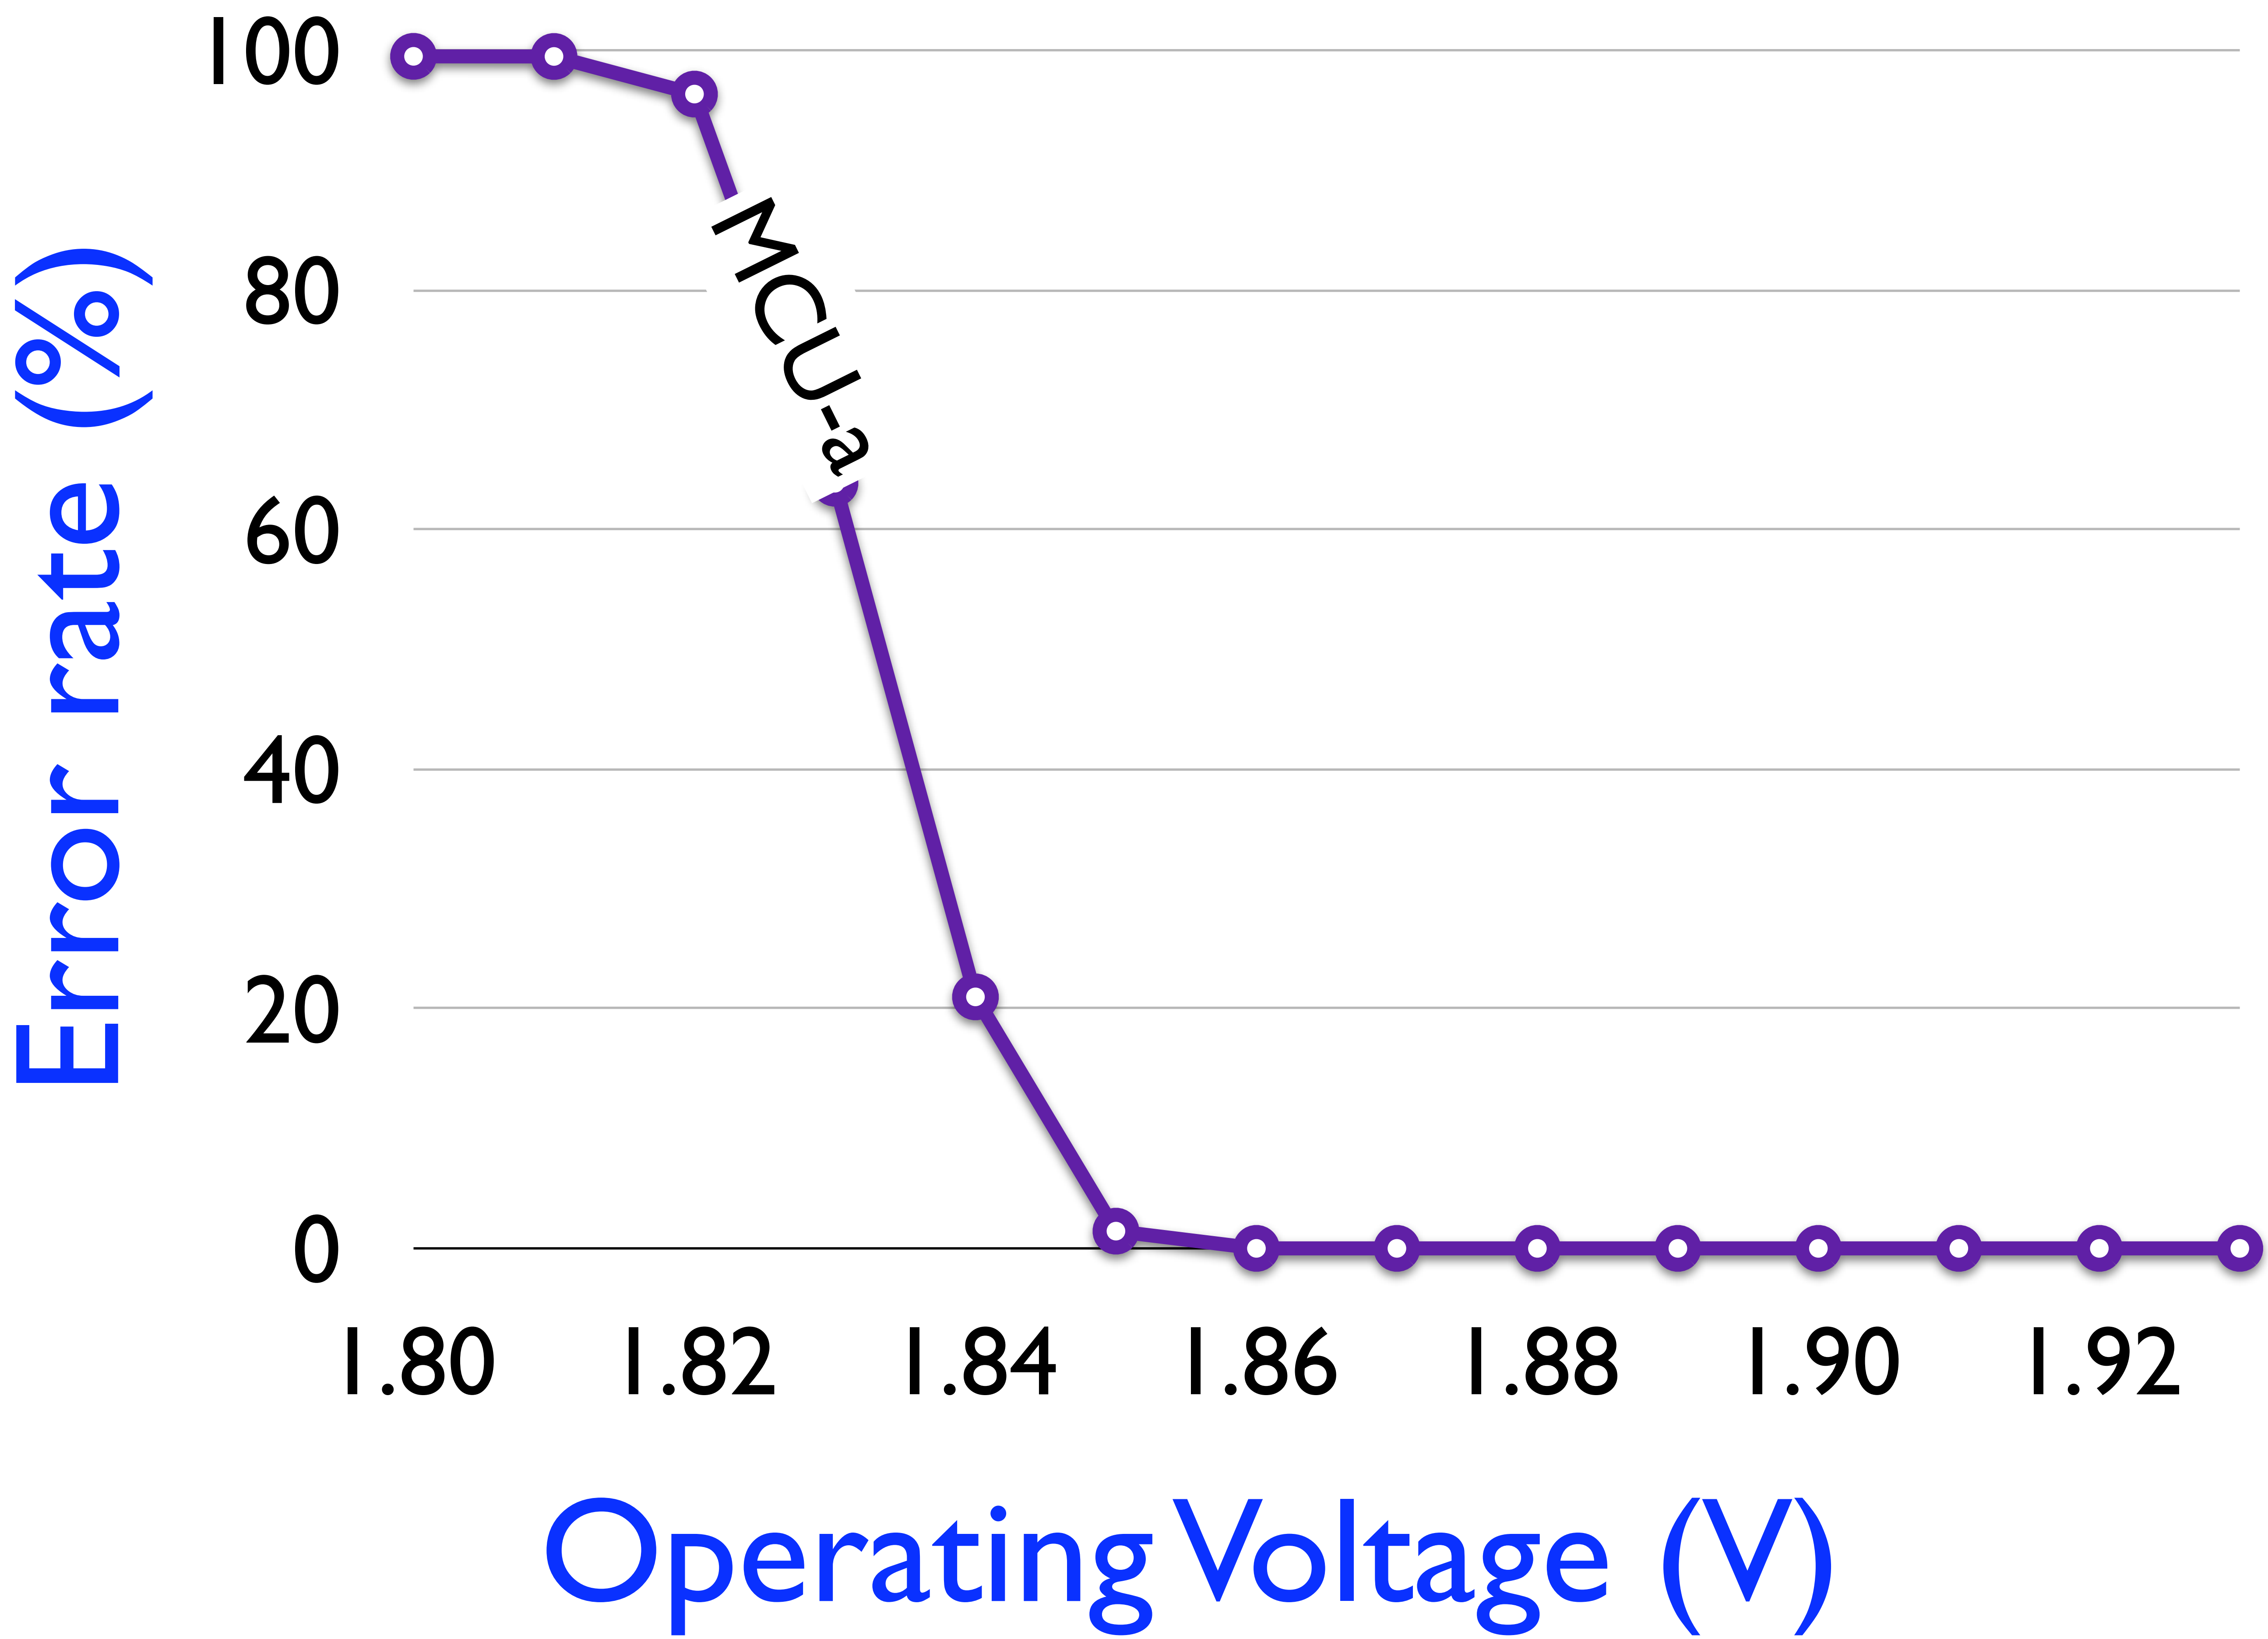
100
80
60
40
20
0
Error rate (%)
MCU-a
1.80
1.82
1.84
1.86
1.88
1.90
1.92
Operating Voltage (V)

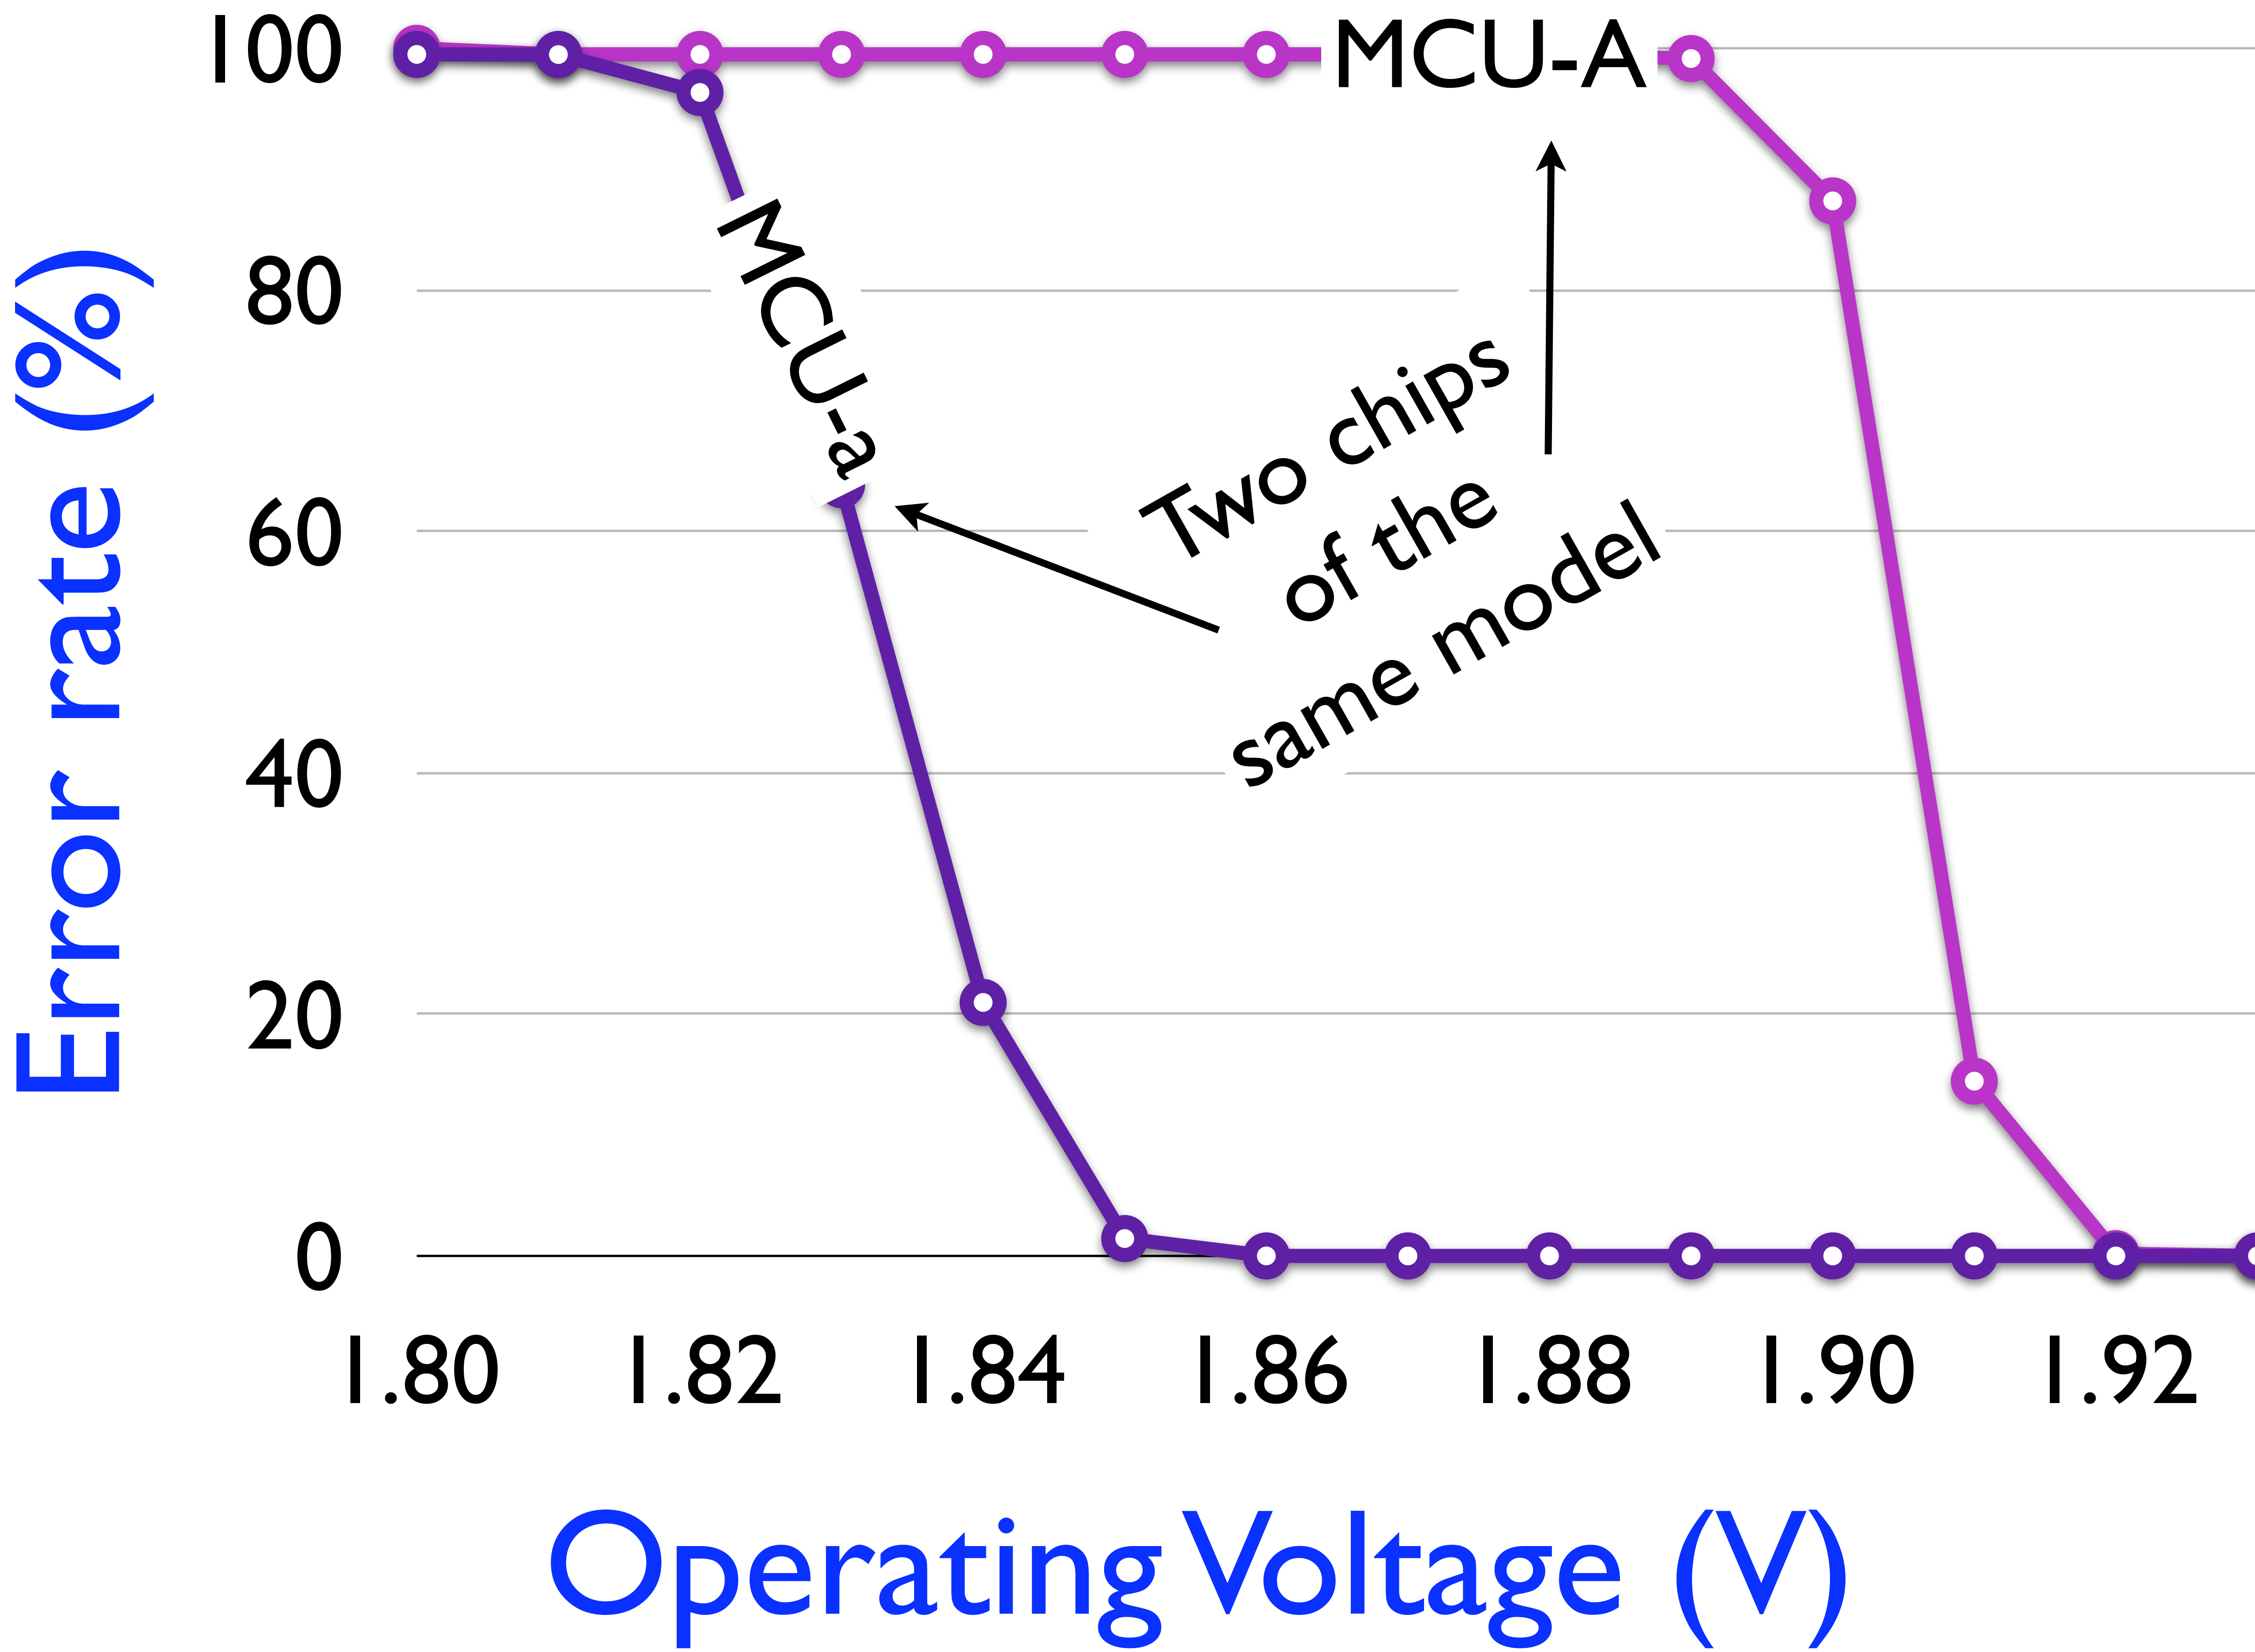

Error rate (%)
100
80
60
40
20
0
1.80
1.82
1.84
1.86
1.88
1.90
1.92
Operating Voltage (V)
MCU-A
MCU-a
Two chips of the same model

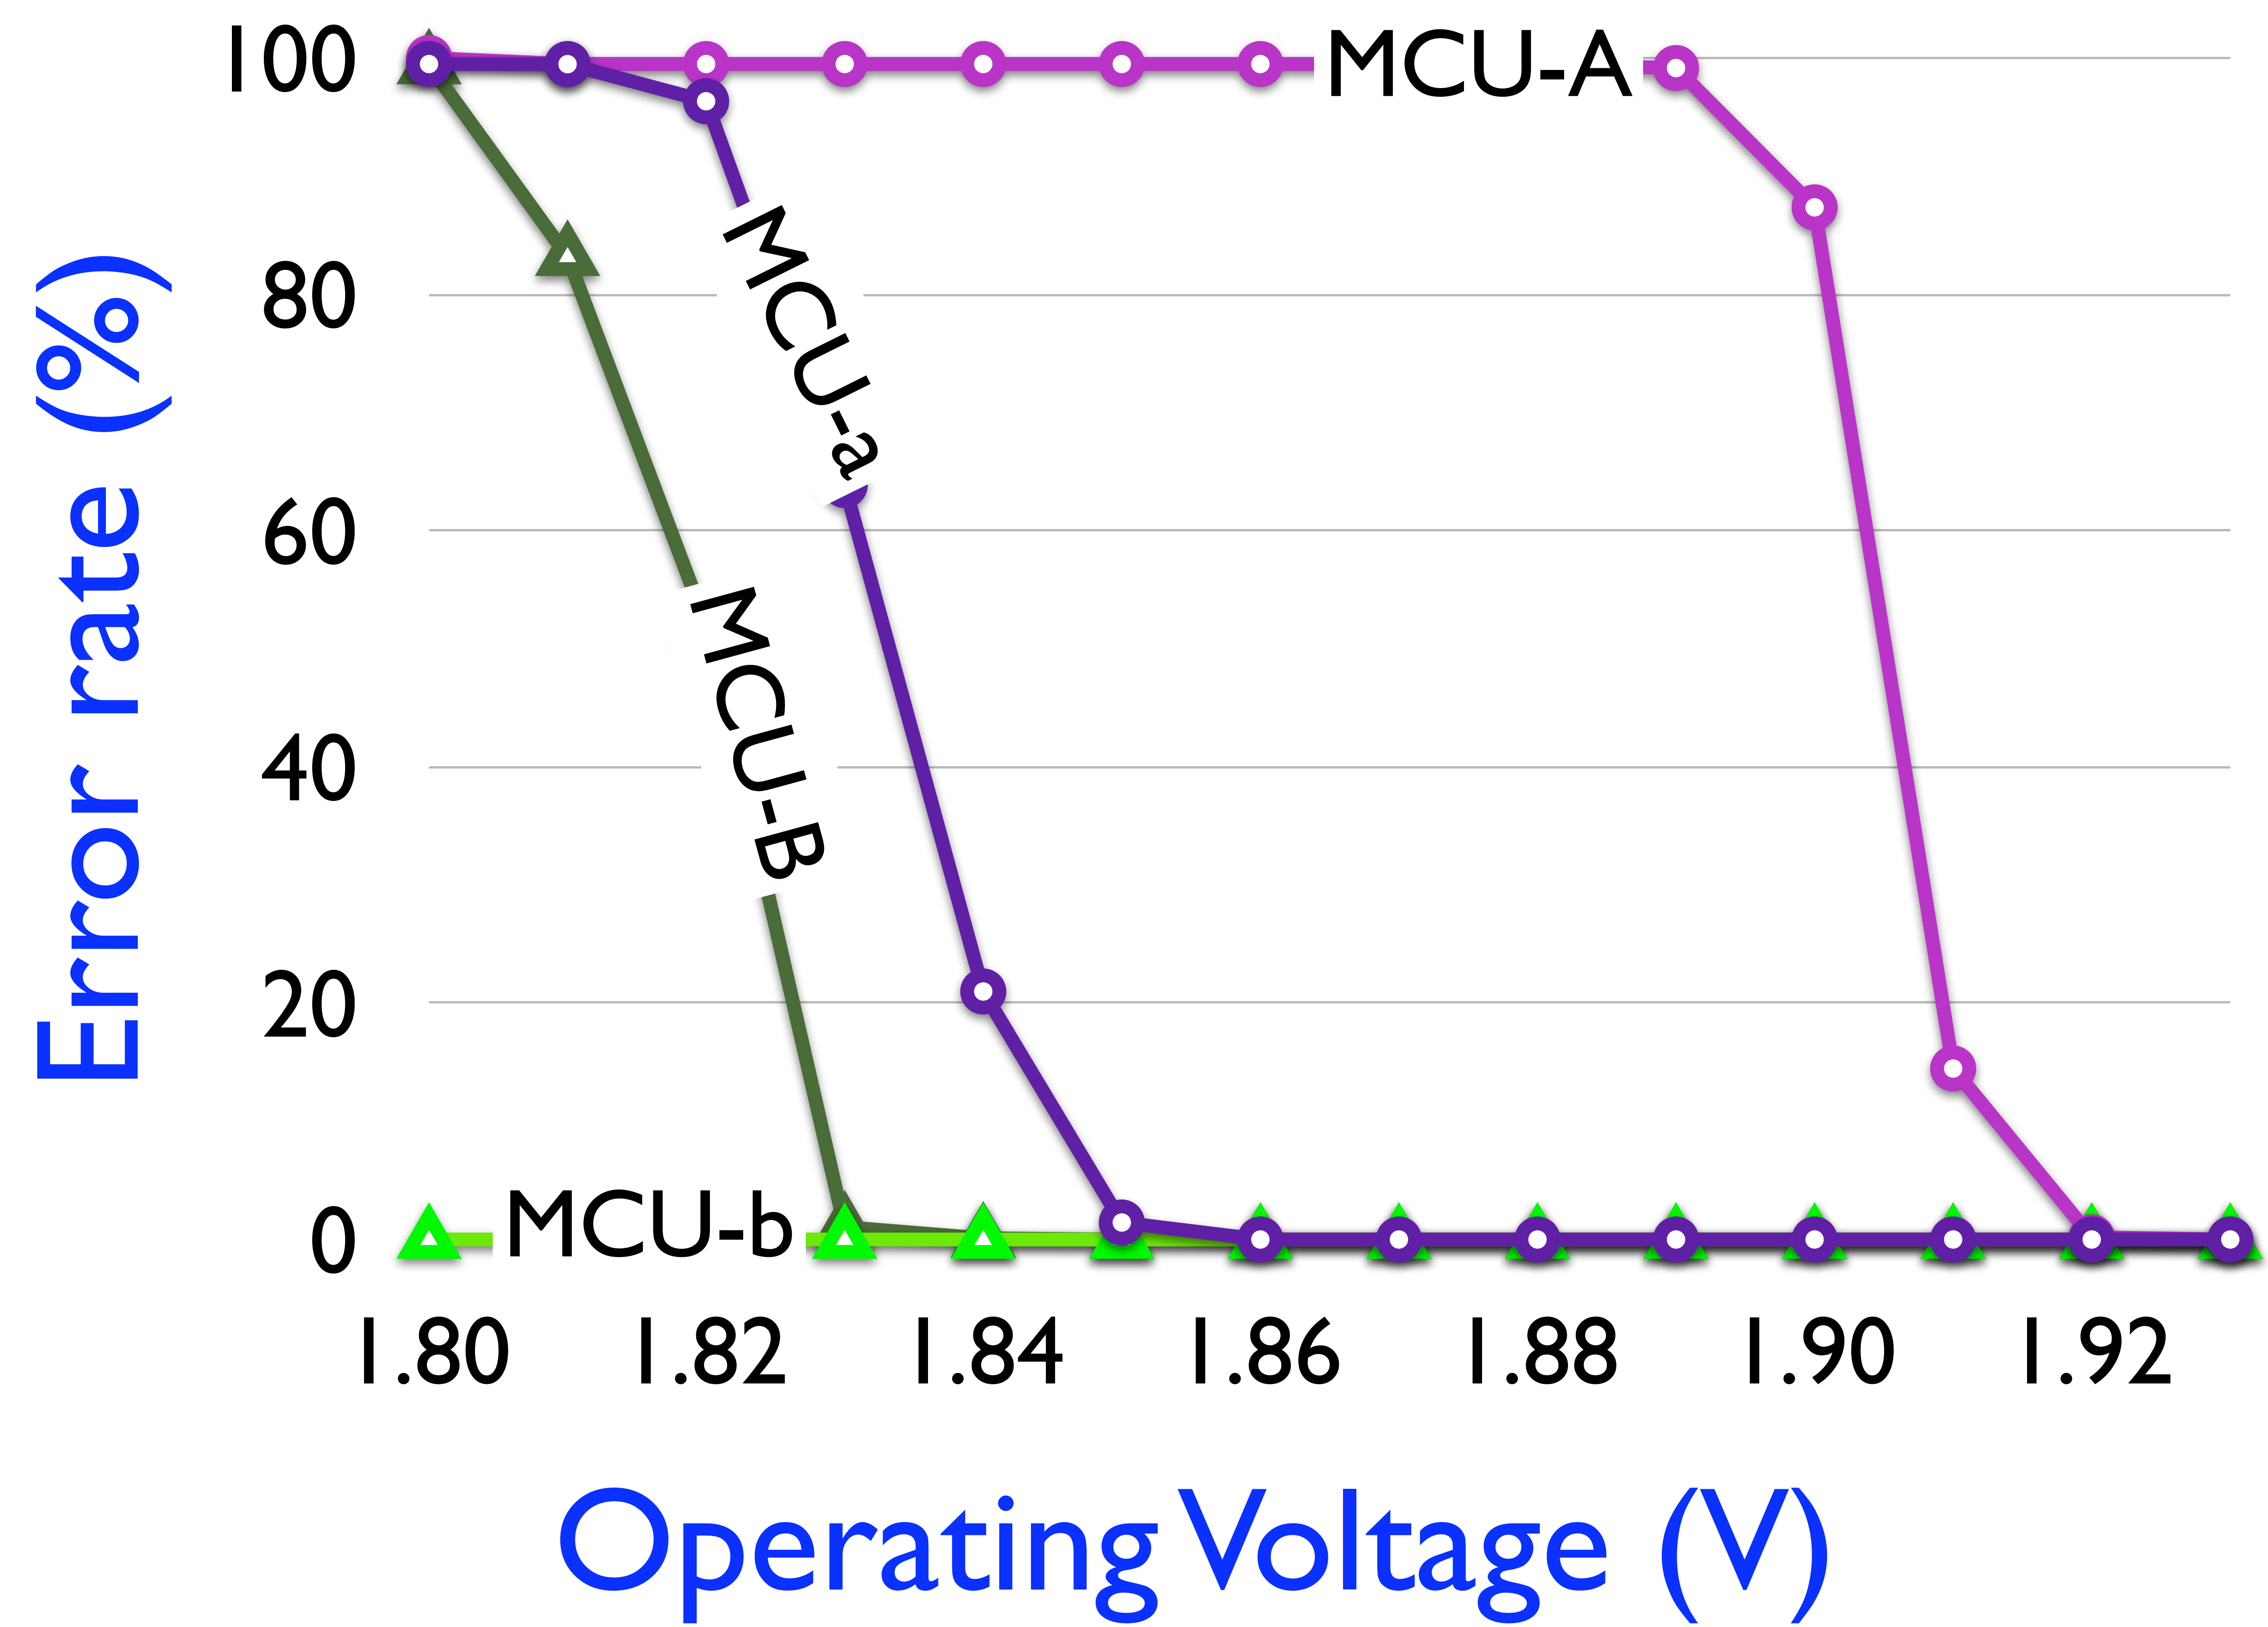

100
80
60
40
20
0
Error rate (%)
MCU-A
MCU-a
MCU-B
MCU-b
1.80
1.82
1.84
1.86
1.88
1.90
1.92
Operating Voltage (V)

Data Hamming Weight

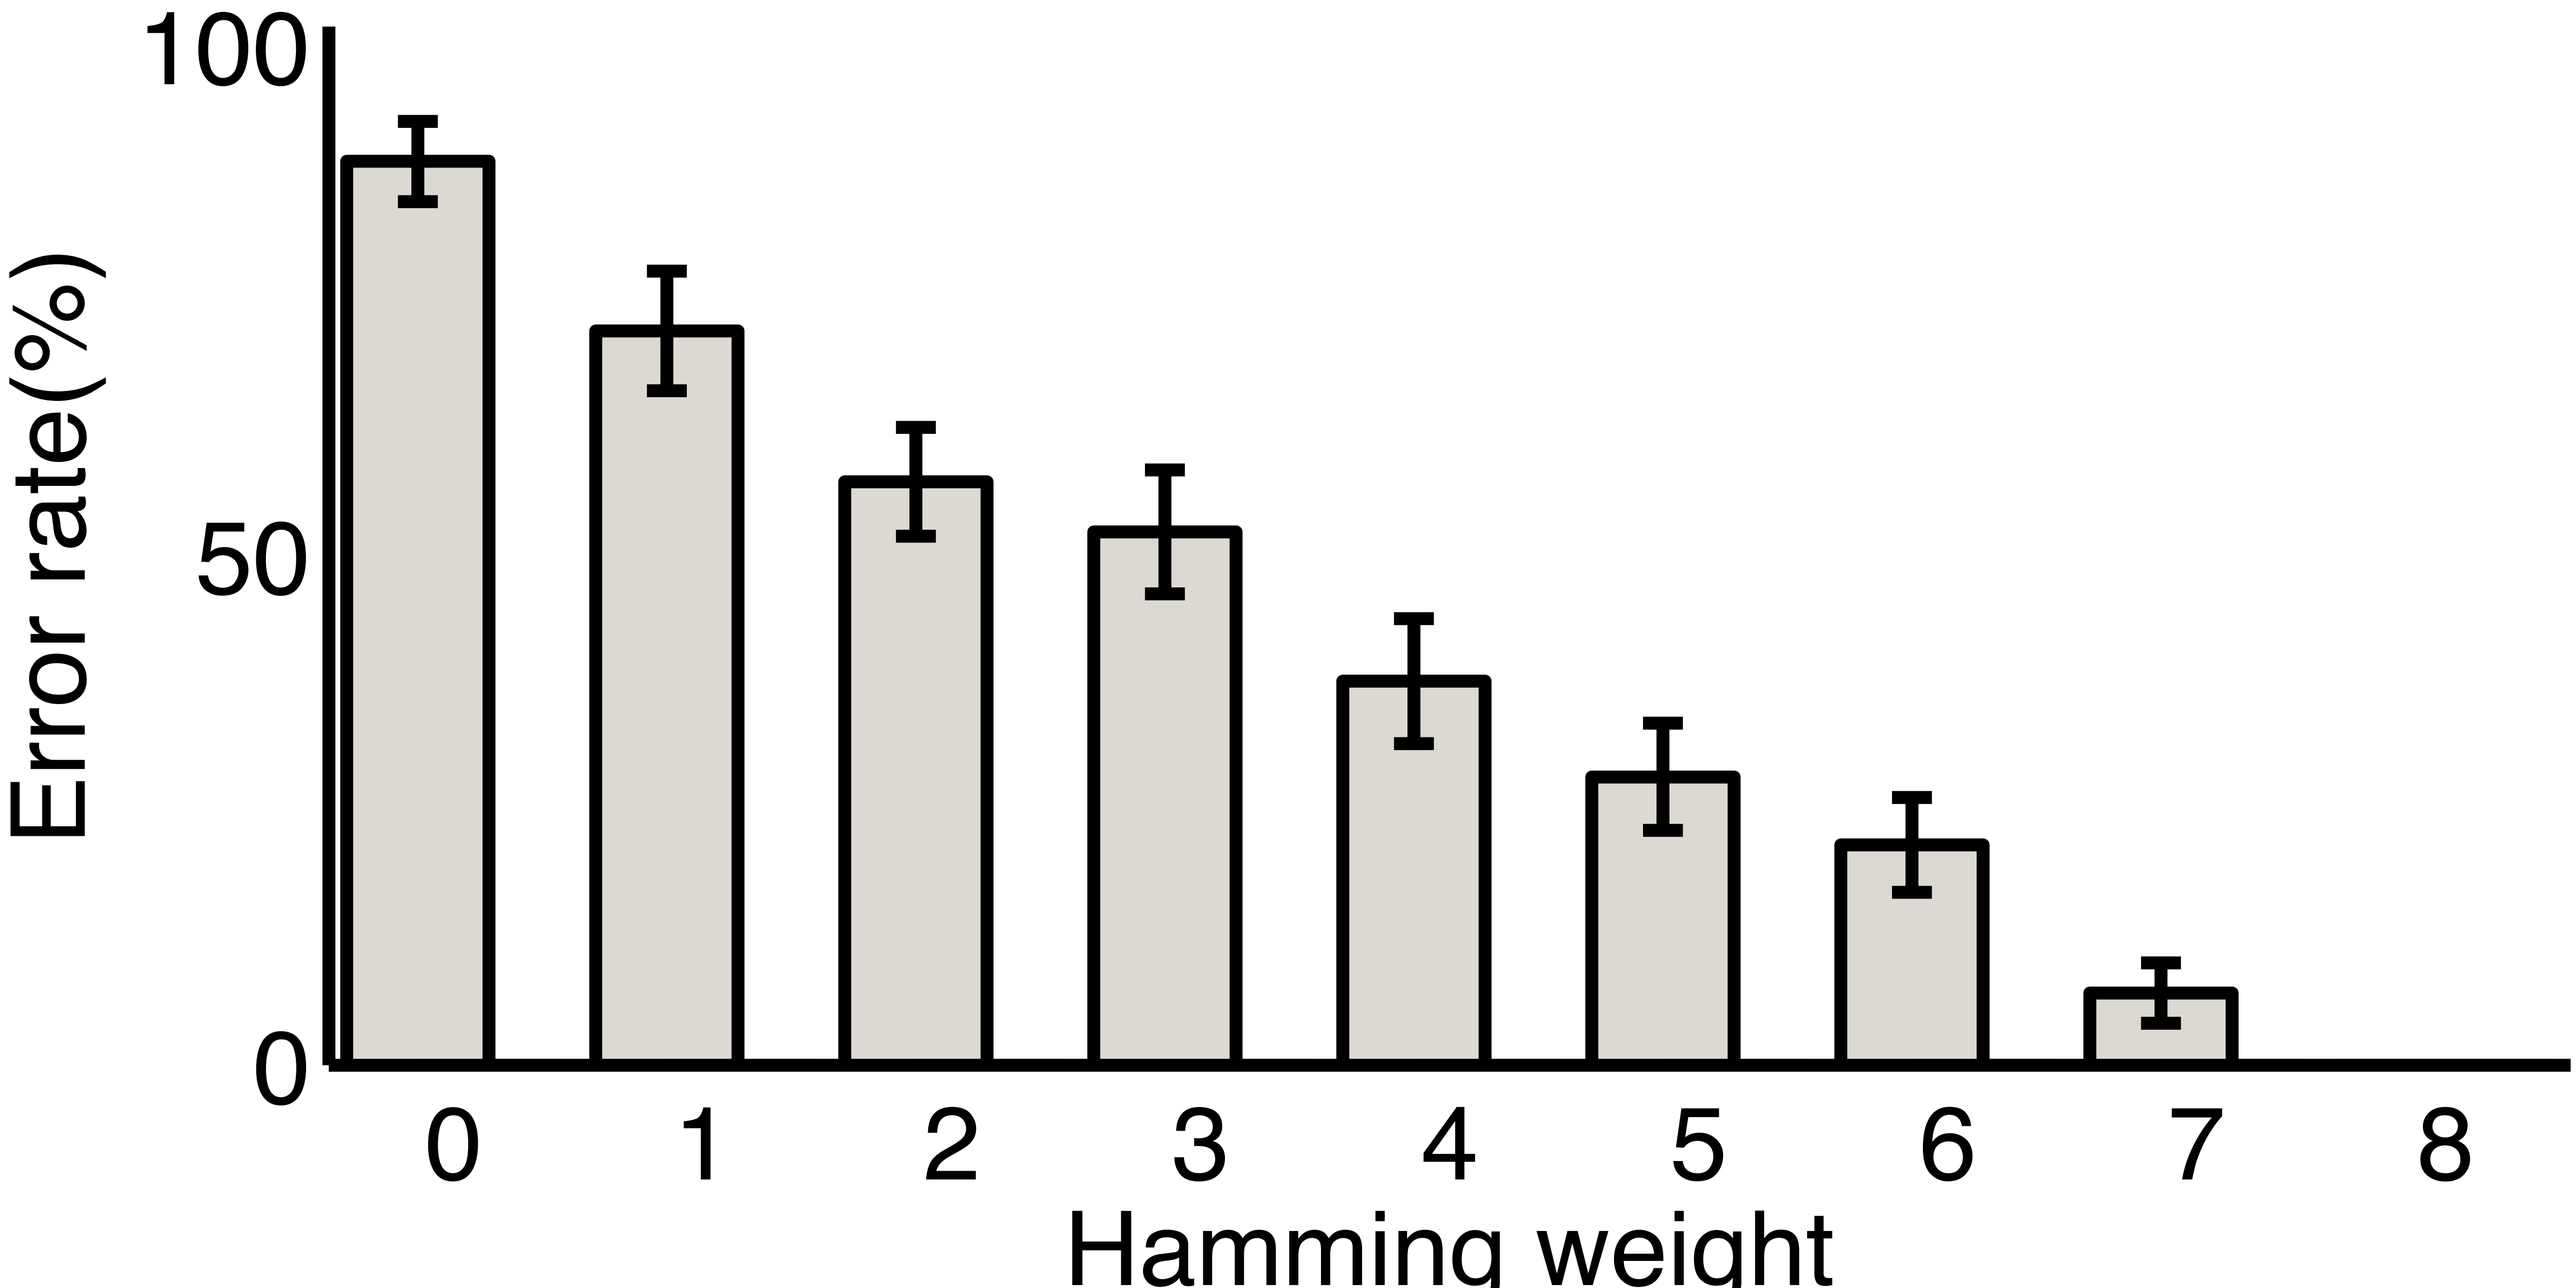

Data Hamming Weight

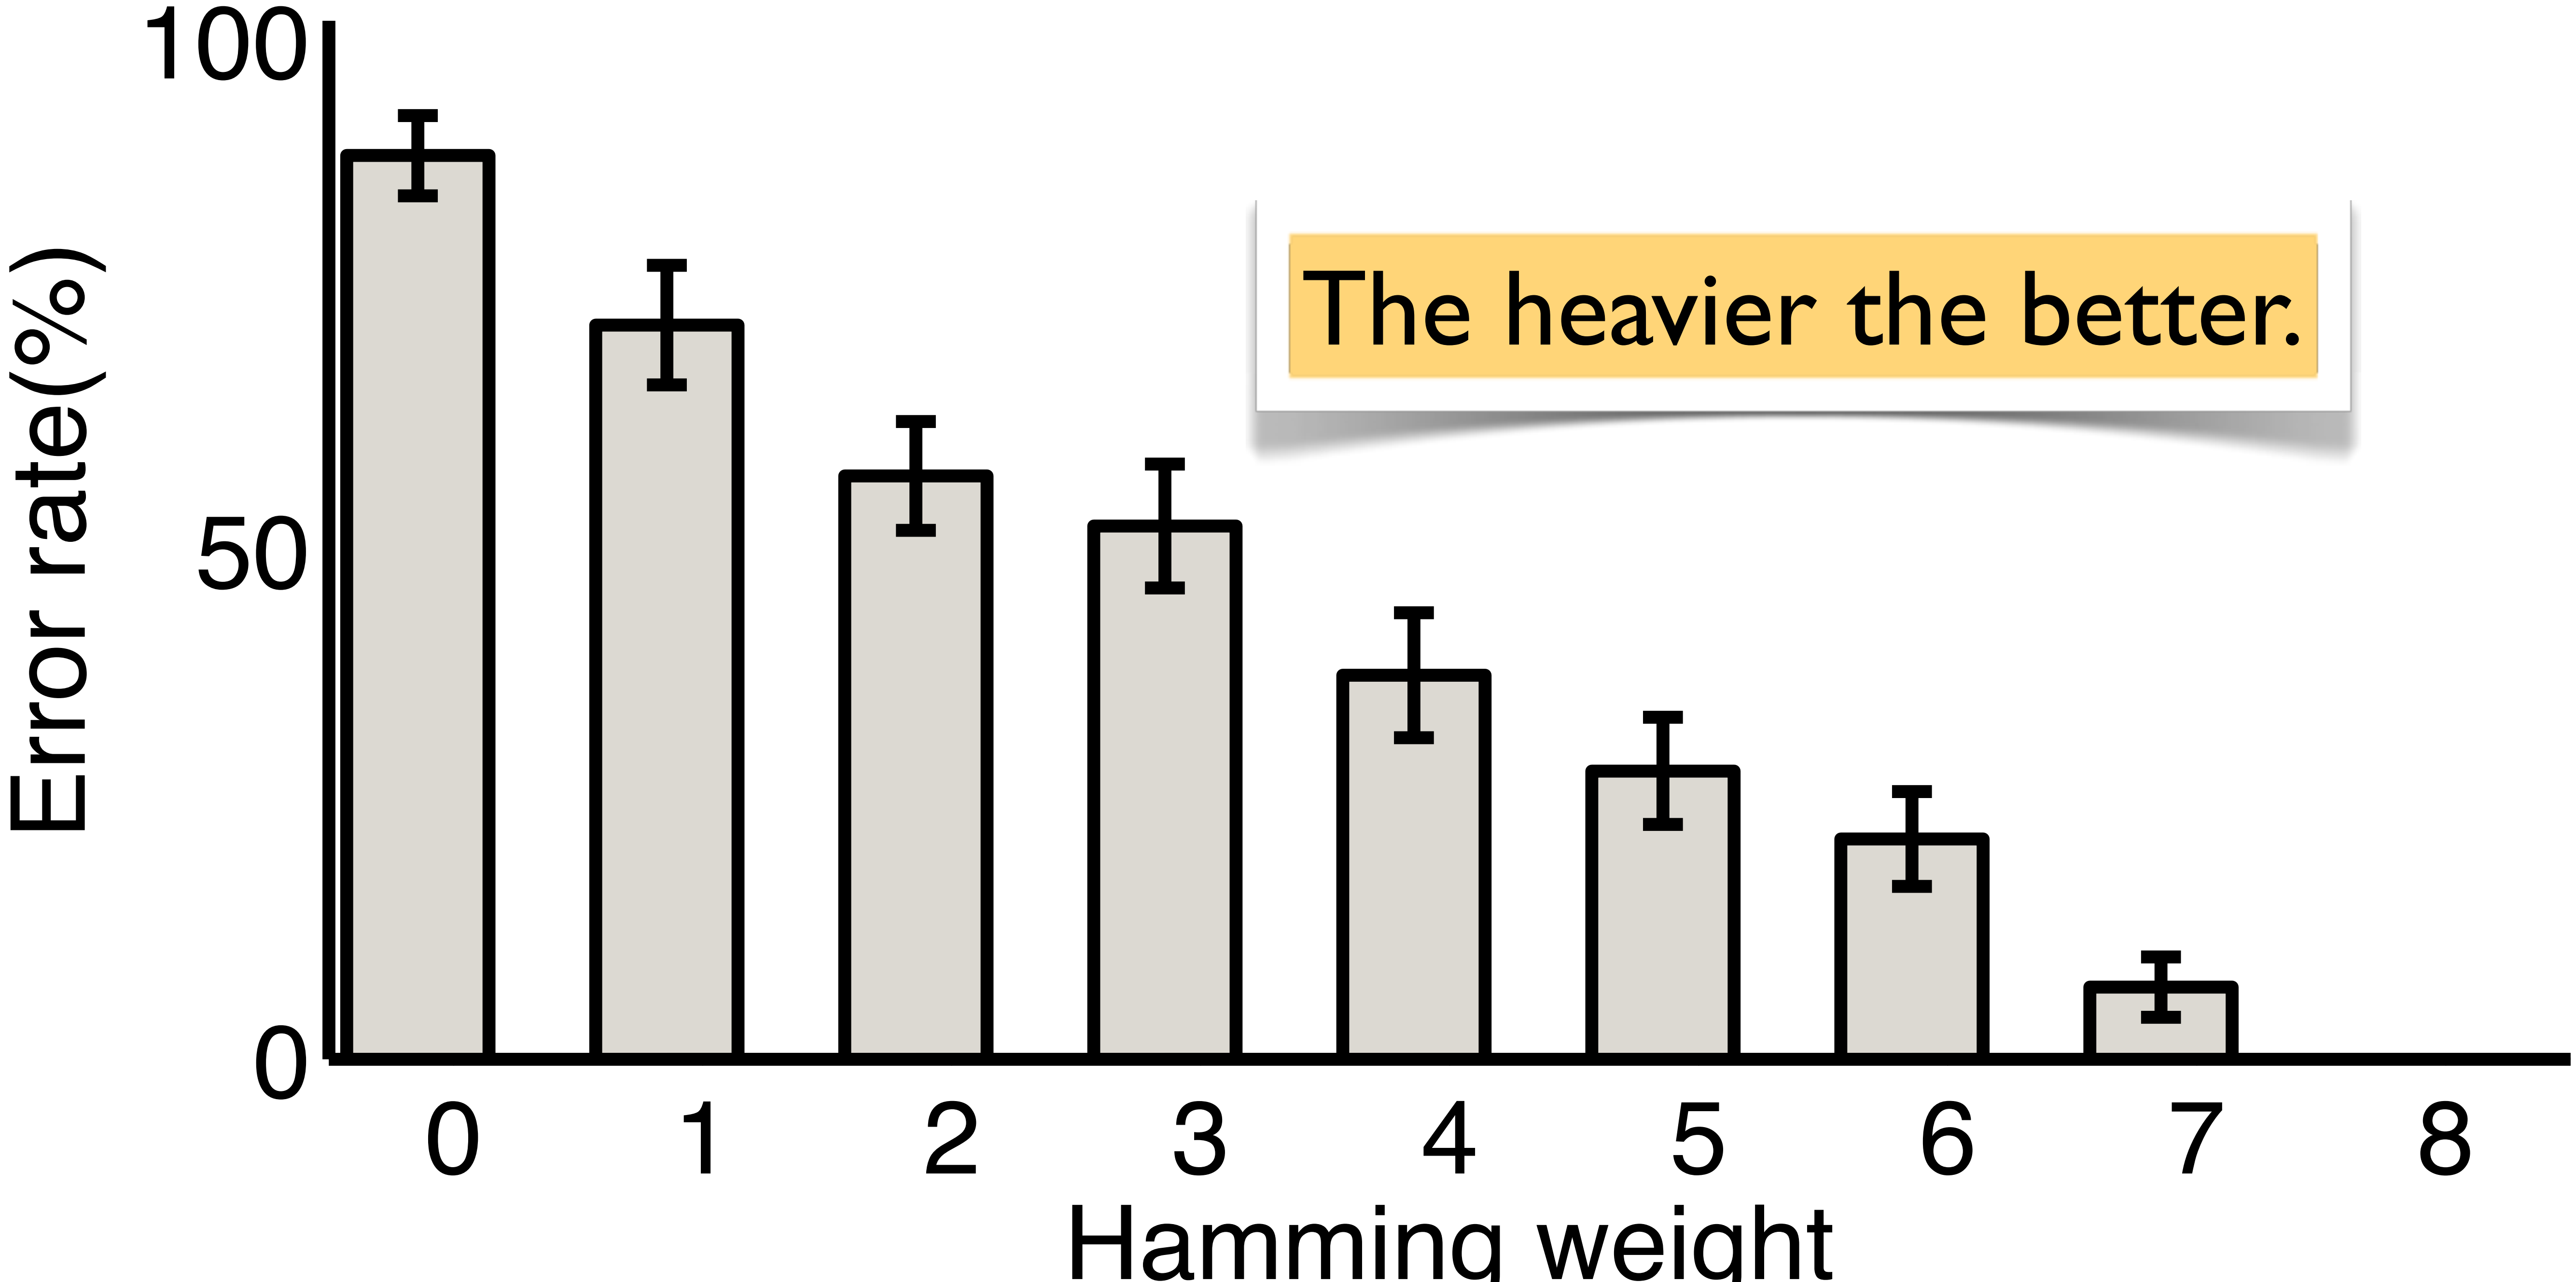

The heavier the better.

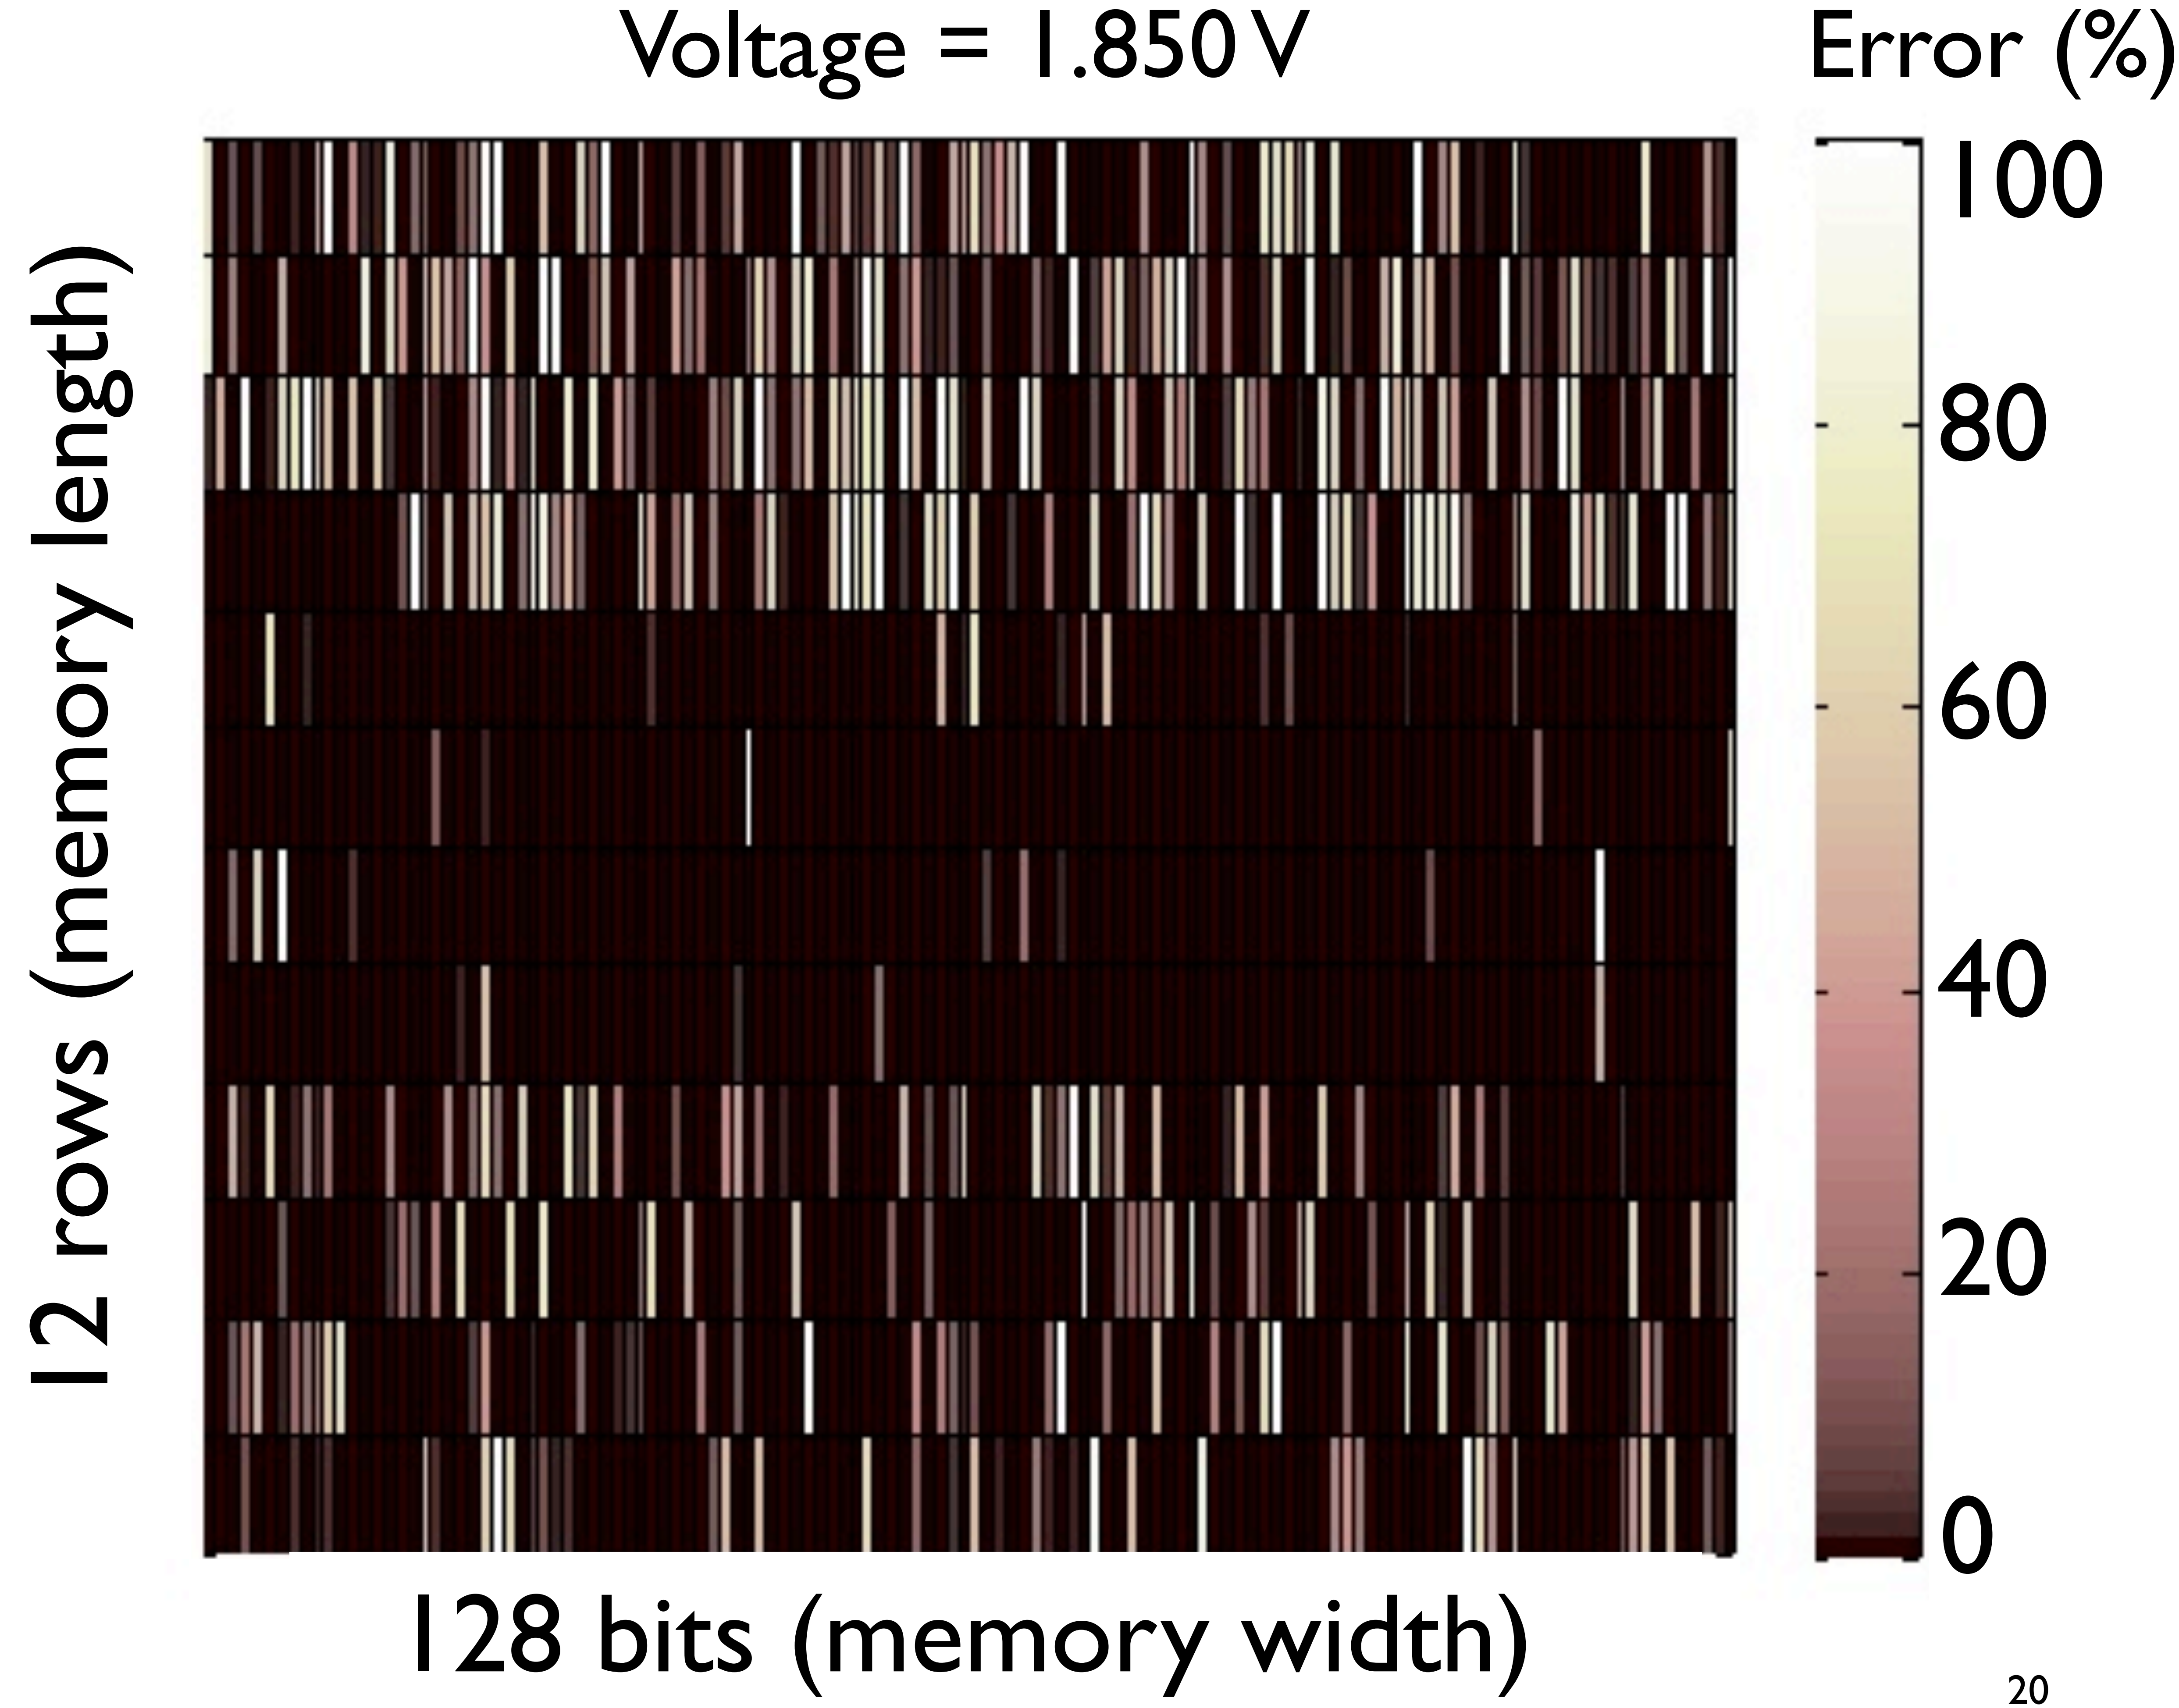
Voltage = 1.850 V
Error (%)
100
80
60
40
20
0
12 rows (memory length)
128 bits (memory width)

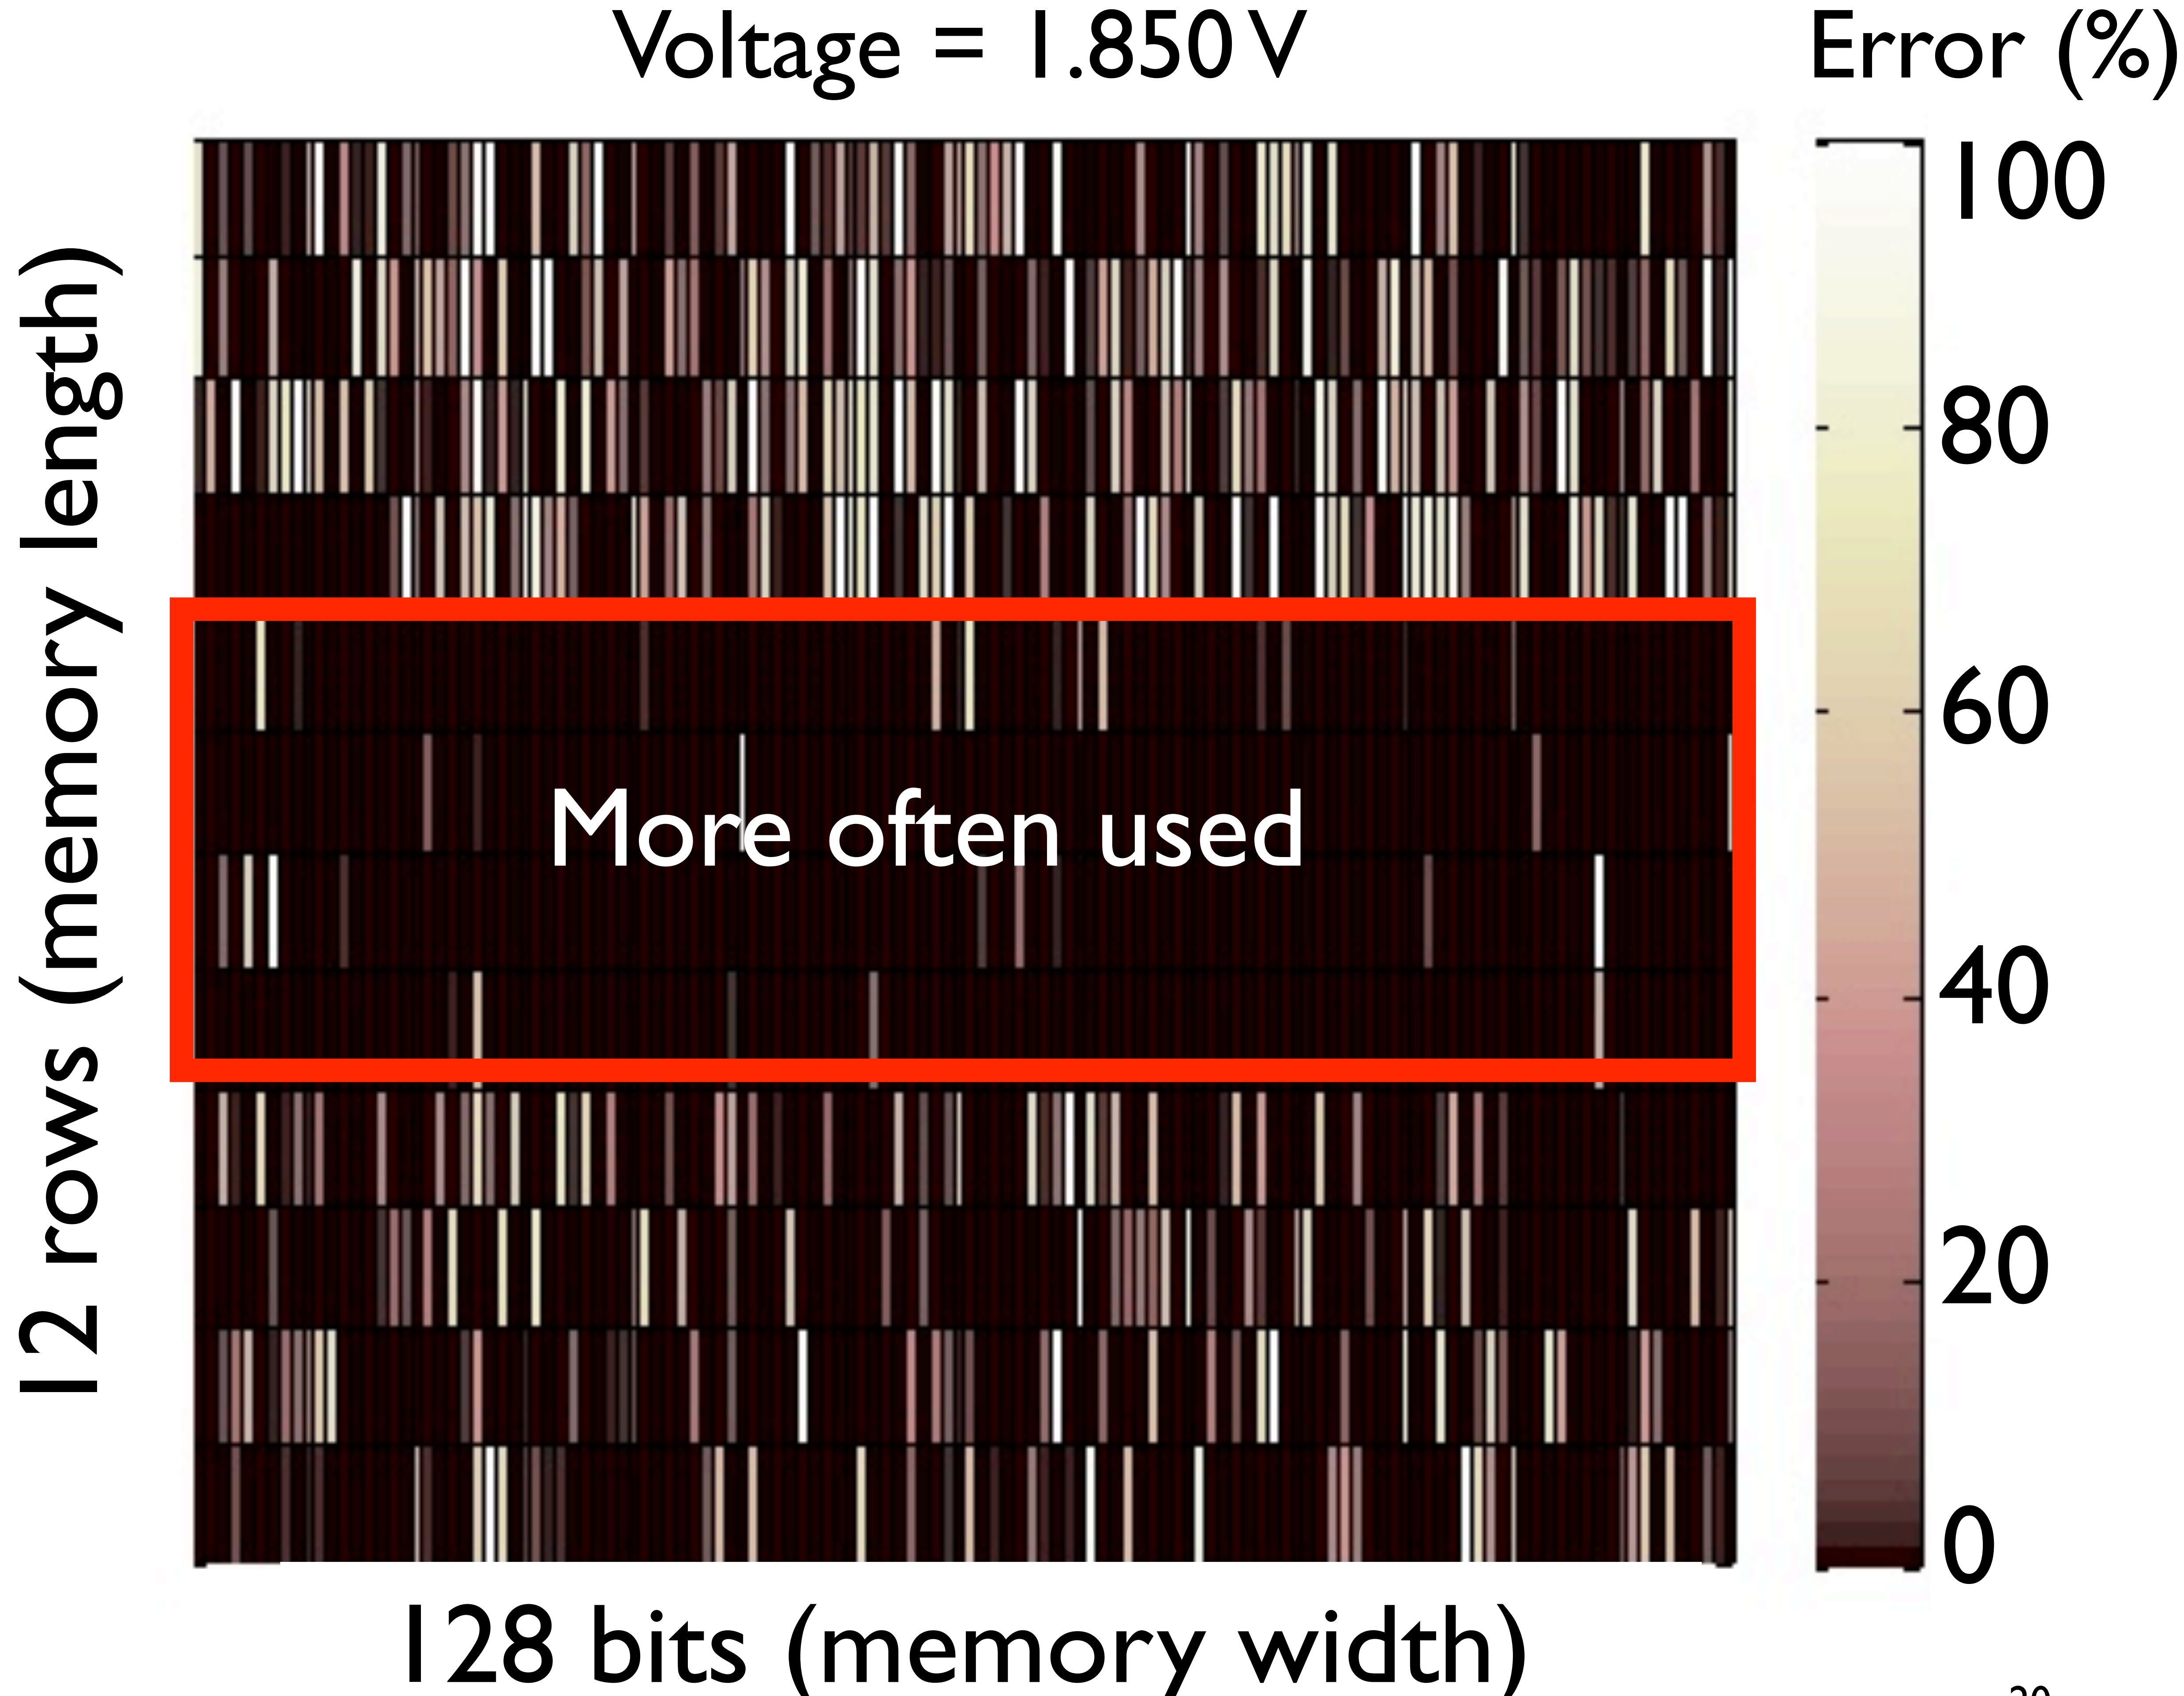

Voltage = 1.850 V
Error (%)
100
80
60
40
20
0
12 rows (memory length)
More often used
128 bits (memory width)

Accumulative Behavior

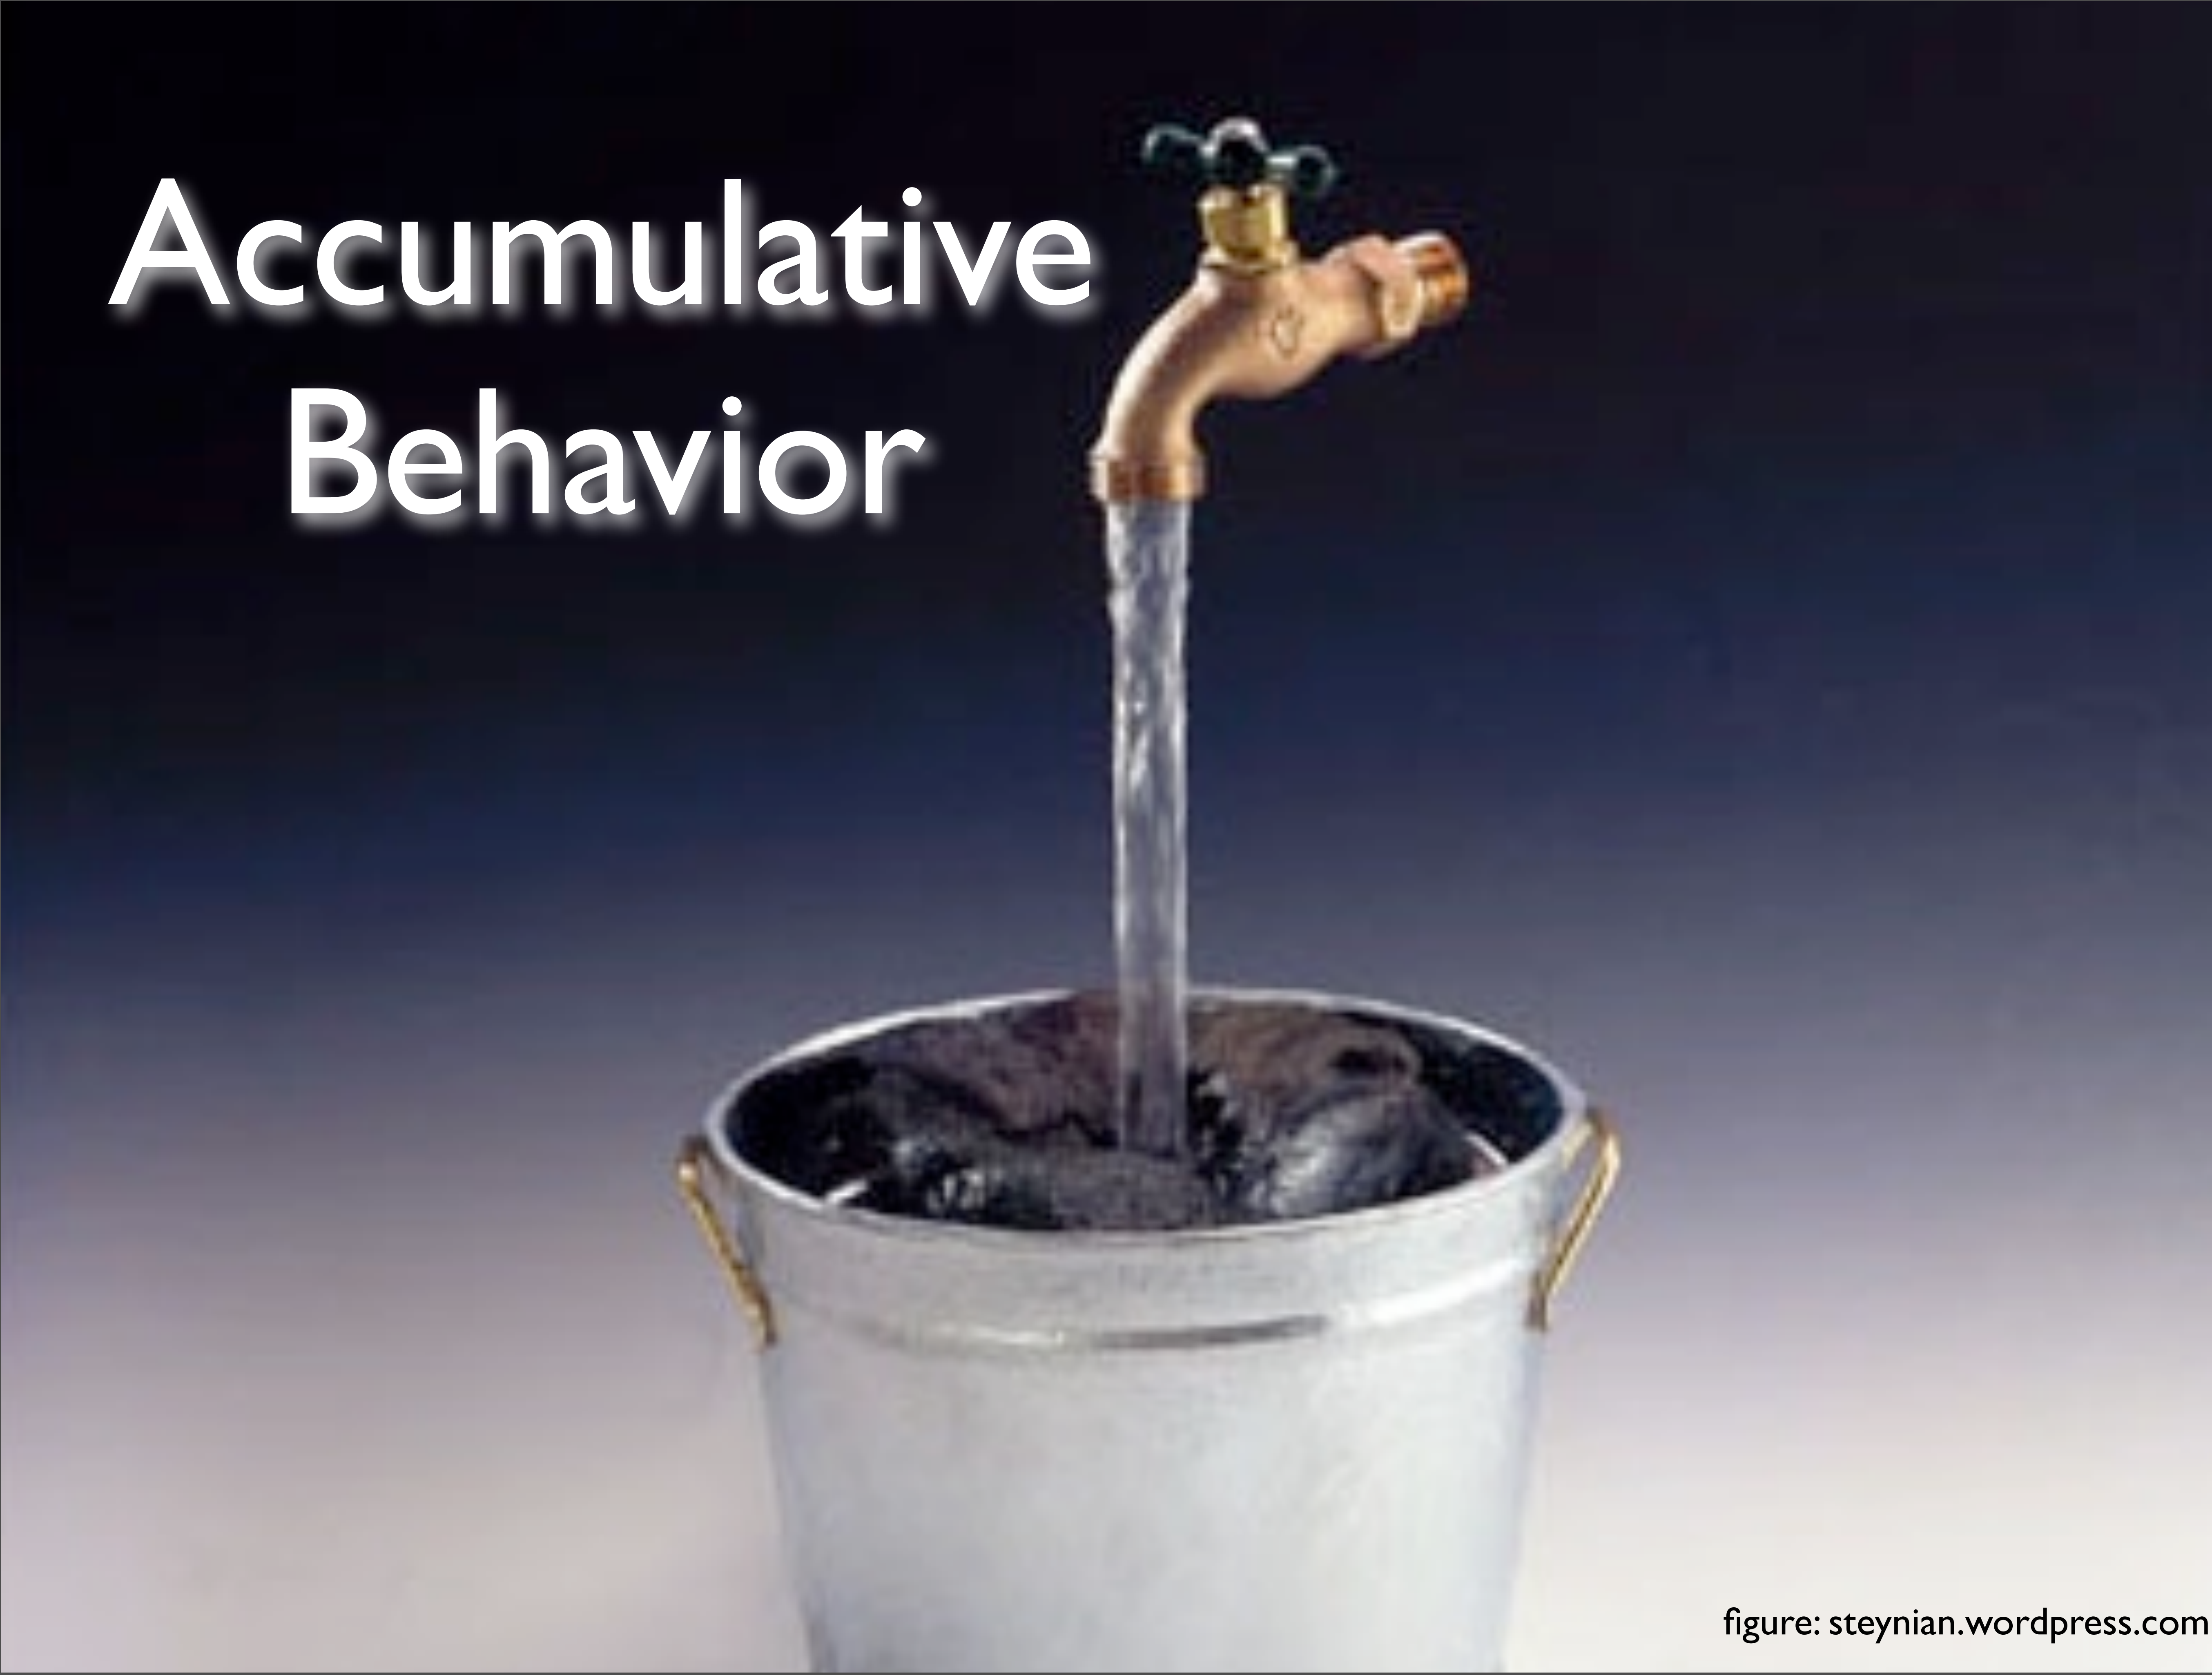

figure: steynian.wordpress.com

Design of a Low-voltage Storage

Modeling Flash Memory

- A set of n cells
- Cell state: $< c_1, ..., c_n >$
- $c_i \in \{0, 1\}$
- Initial state: $\forall i, c_i = 1$
- Update: set a subset of $< c_1, ..., c_n >$ to 0
  - Once $c_i = 0$ then a write cannot $c_i \leftarrow 1$
  - Write Once Bits: Wits

Modeling Flash Memory

At low voltage:

- $c_i \leftarrow 0$ might fail and c_i remains 1. [Pavan:97]
- There might not be enough charge stored in a cell to represent a 0: Half-Wits

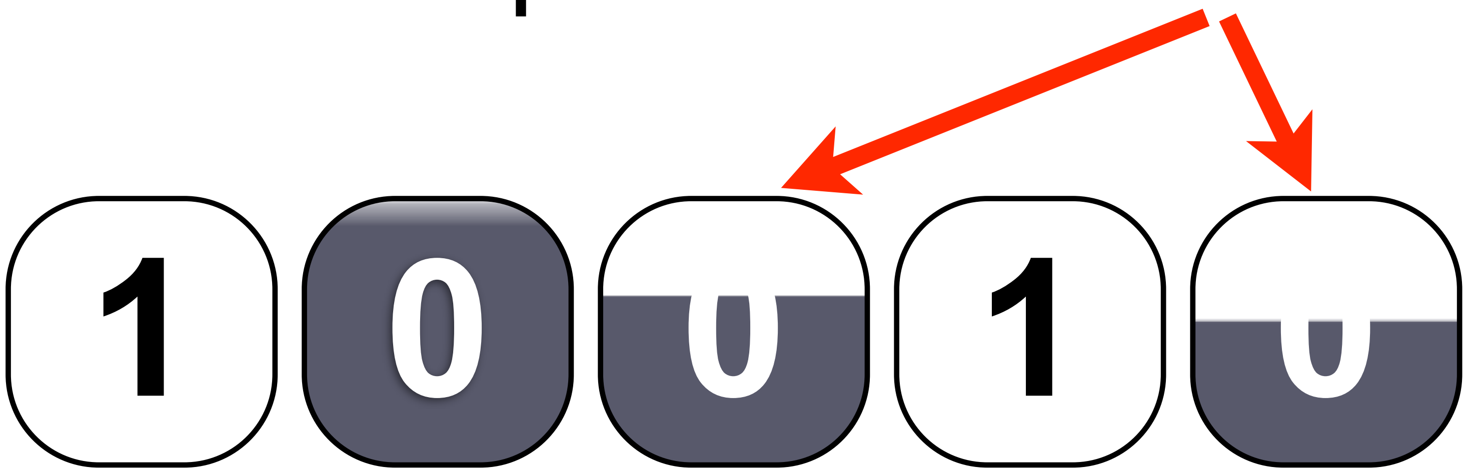

Design Goals

Minimize:

- Energy consumption
- Error rate
- Delay

Proposed Techniques

1. In-place writes
2. Multiple-place writes
3. ReedSolomon-Berger Codes

Negative Logic

Cell

Initialization: 1

Negative Logic

Cell

Initialization: 1

Charge

Written: 0

1. In-place writes

- Repeatedly attempt a write to the same location.

1. In-place writes

- Repeatedly attempt a write to the same location.
- Example: One bit over time

1

1. In-place writes

- Repeatedly attempt a write to the same location.
- Example: One bit over time

1

1→0

1. In-place writes

- Repeatedly attempt a write to the same location.
- Example: One bit over time

1 | 1→0 | 1→0

1. In-place writes

- Repeatedly attempt a write to the same location.
- Example: One bit over time

1 | 1→0 | 1→0 | 0

In-place writes

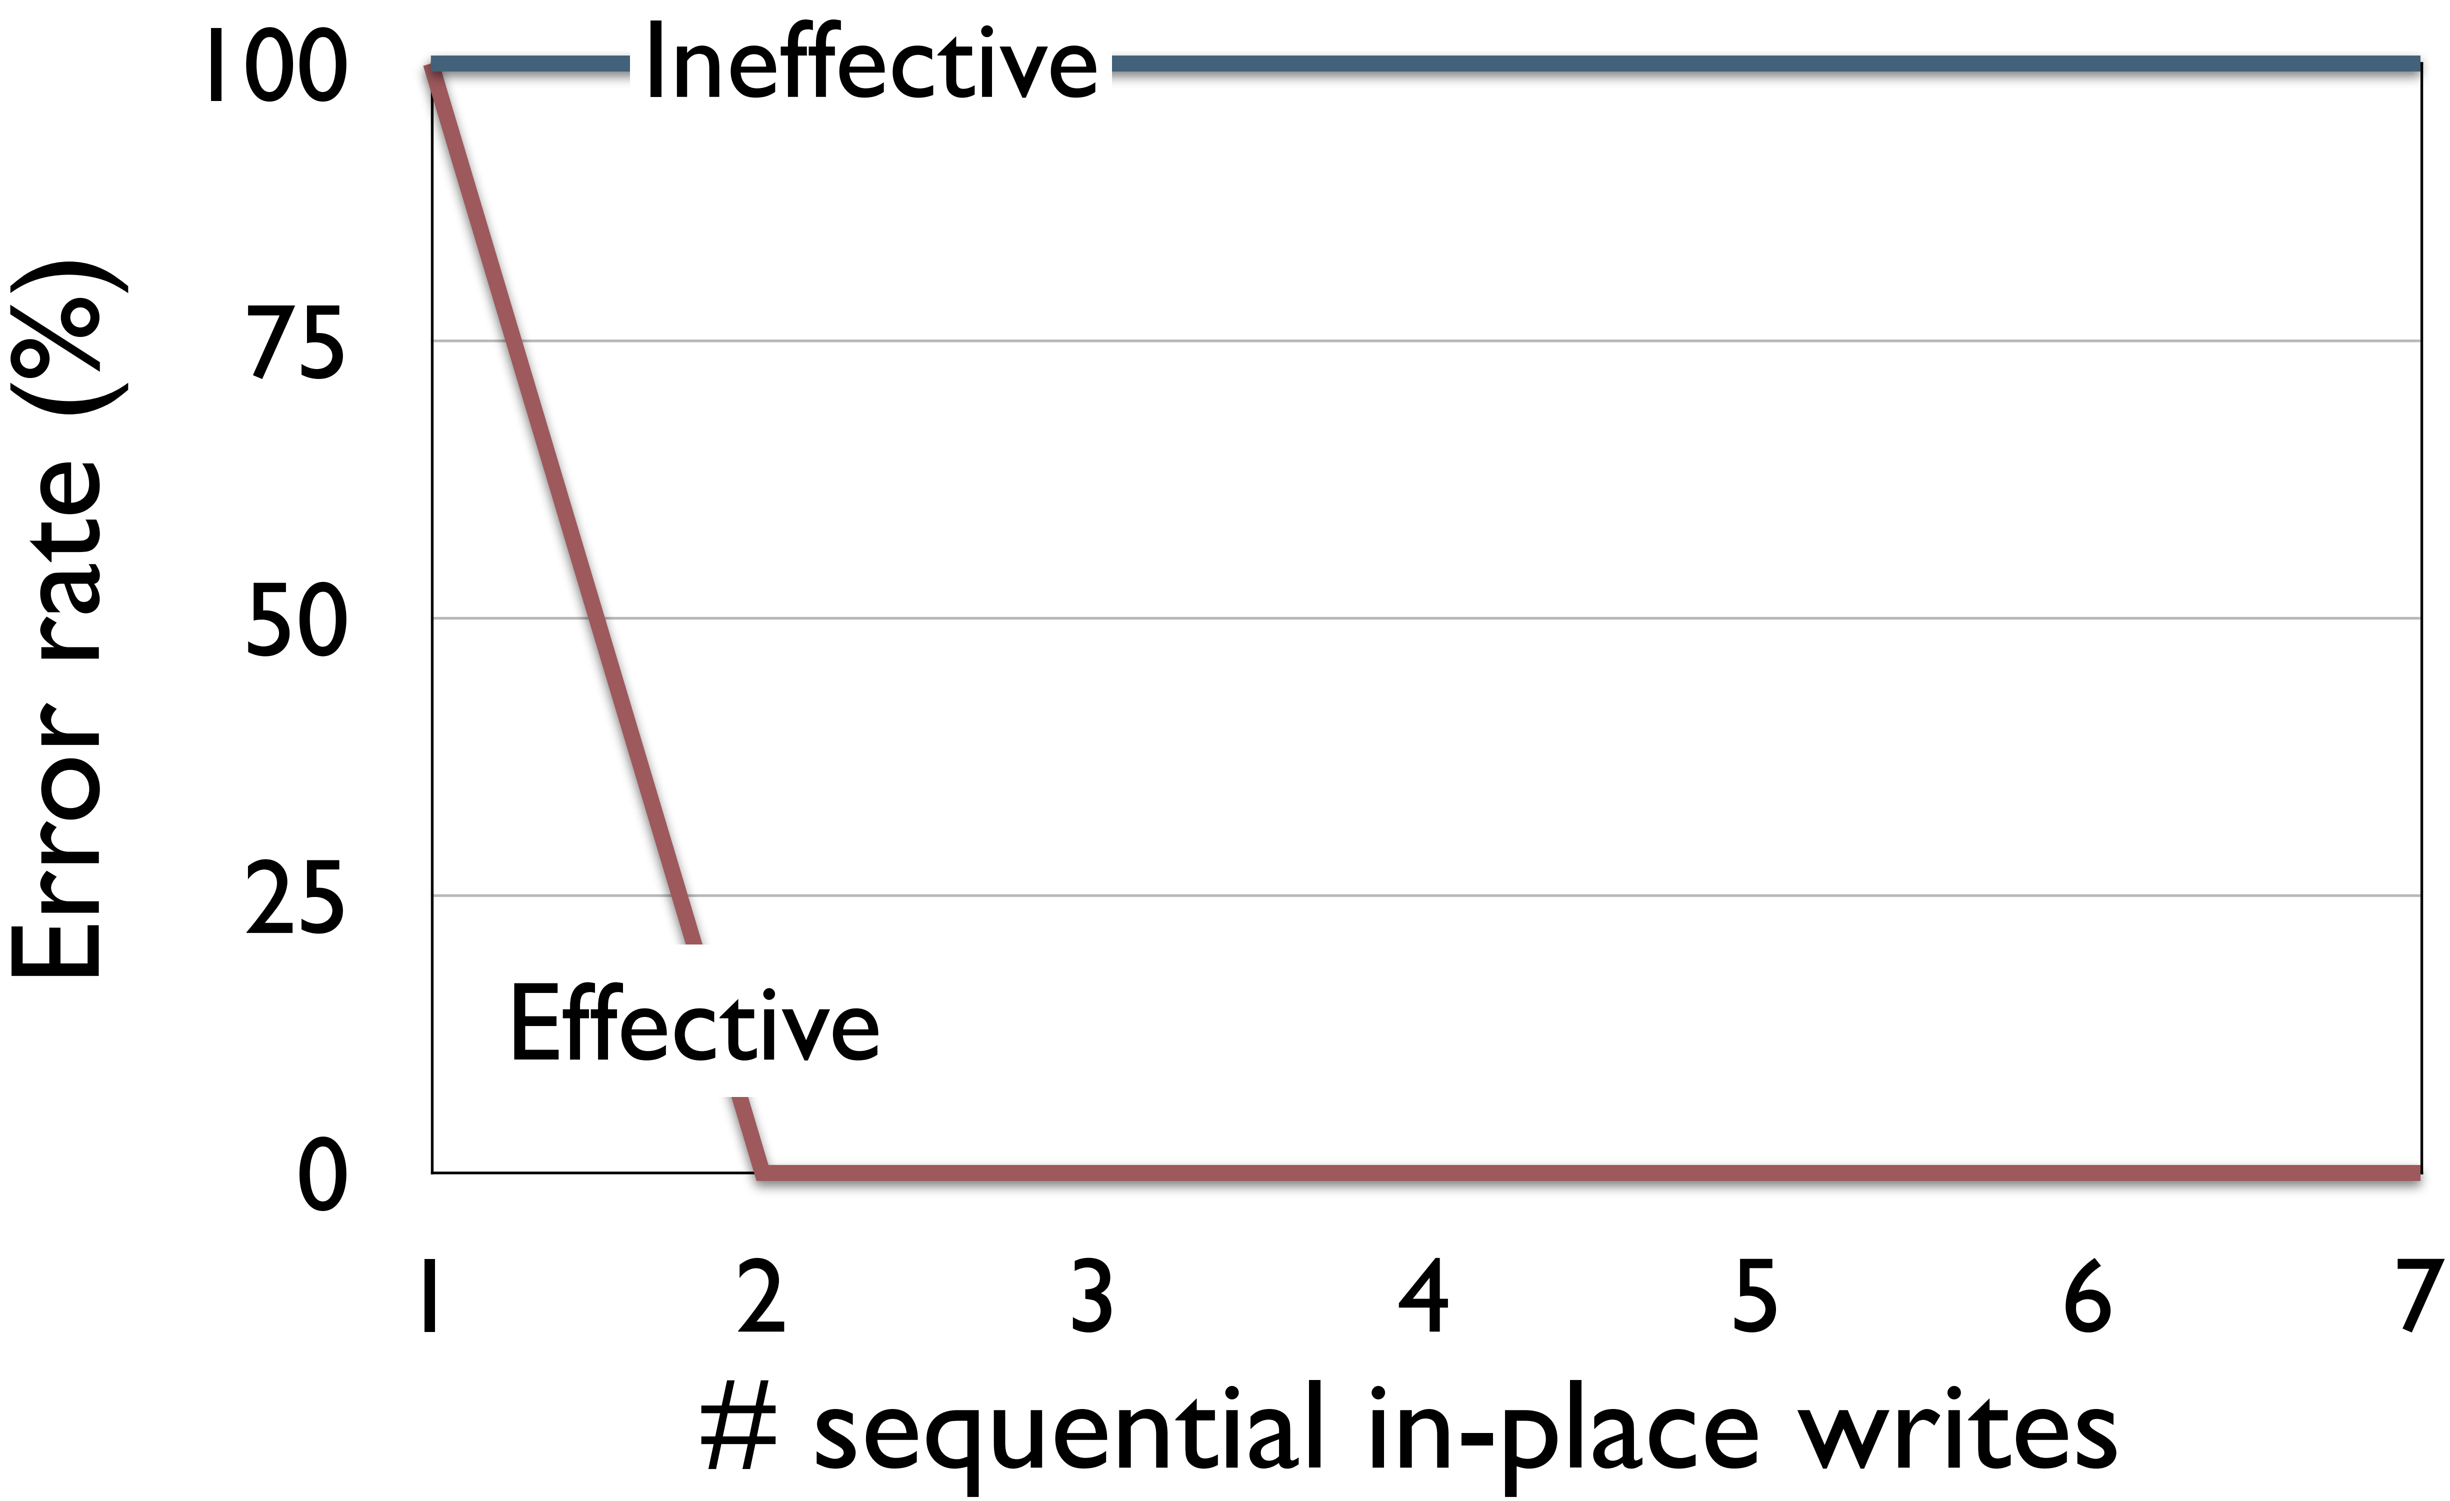

In-place writes

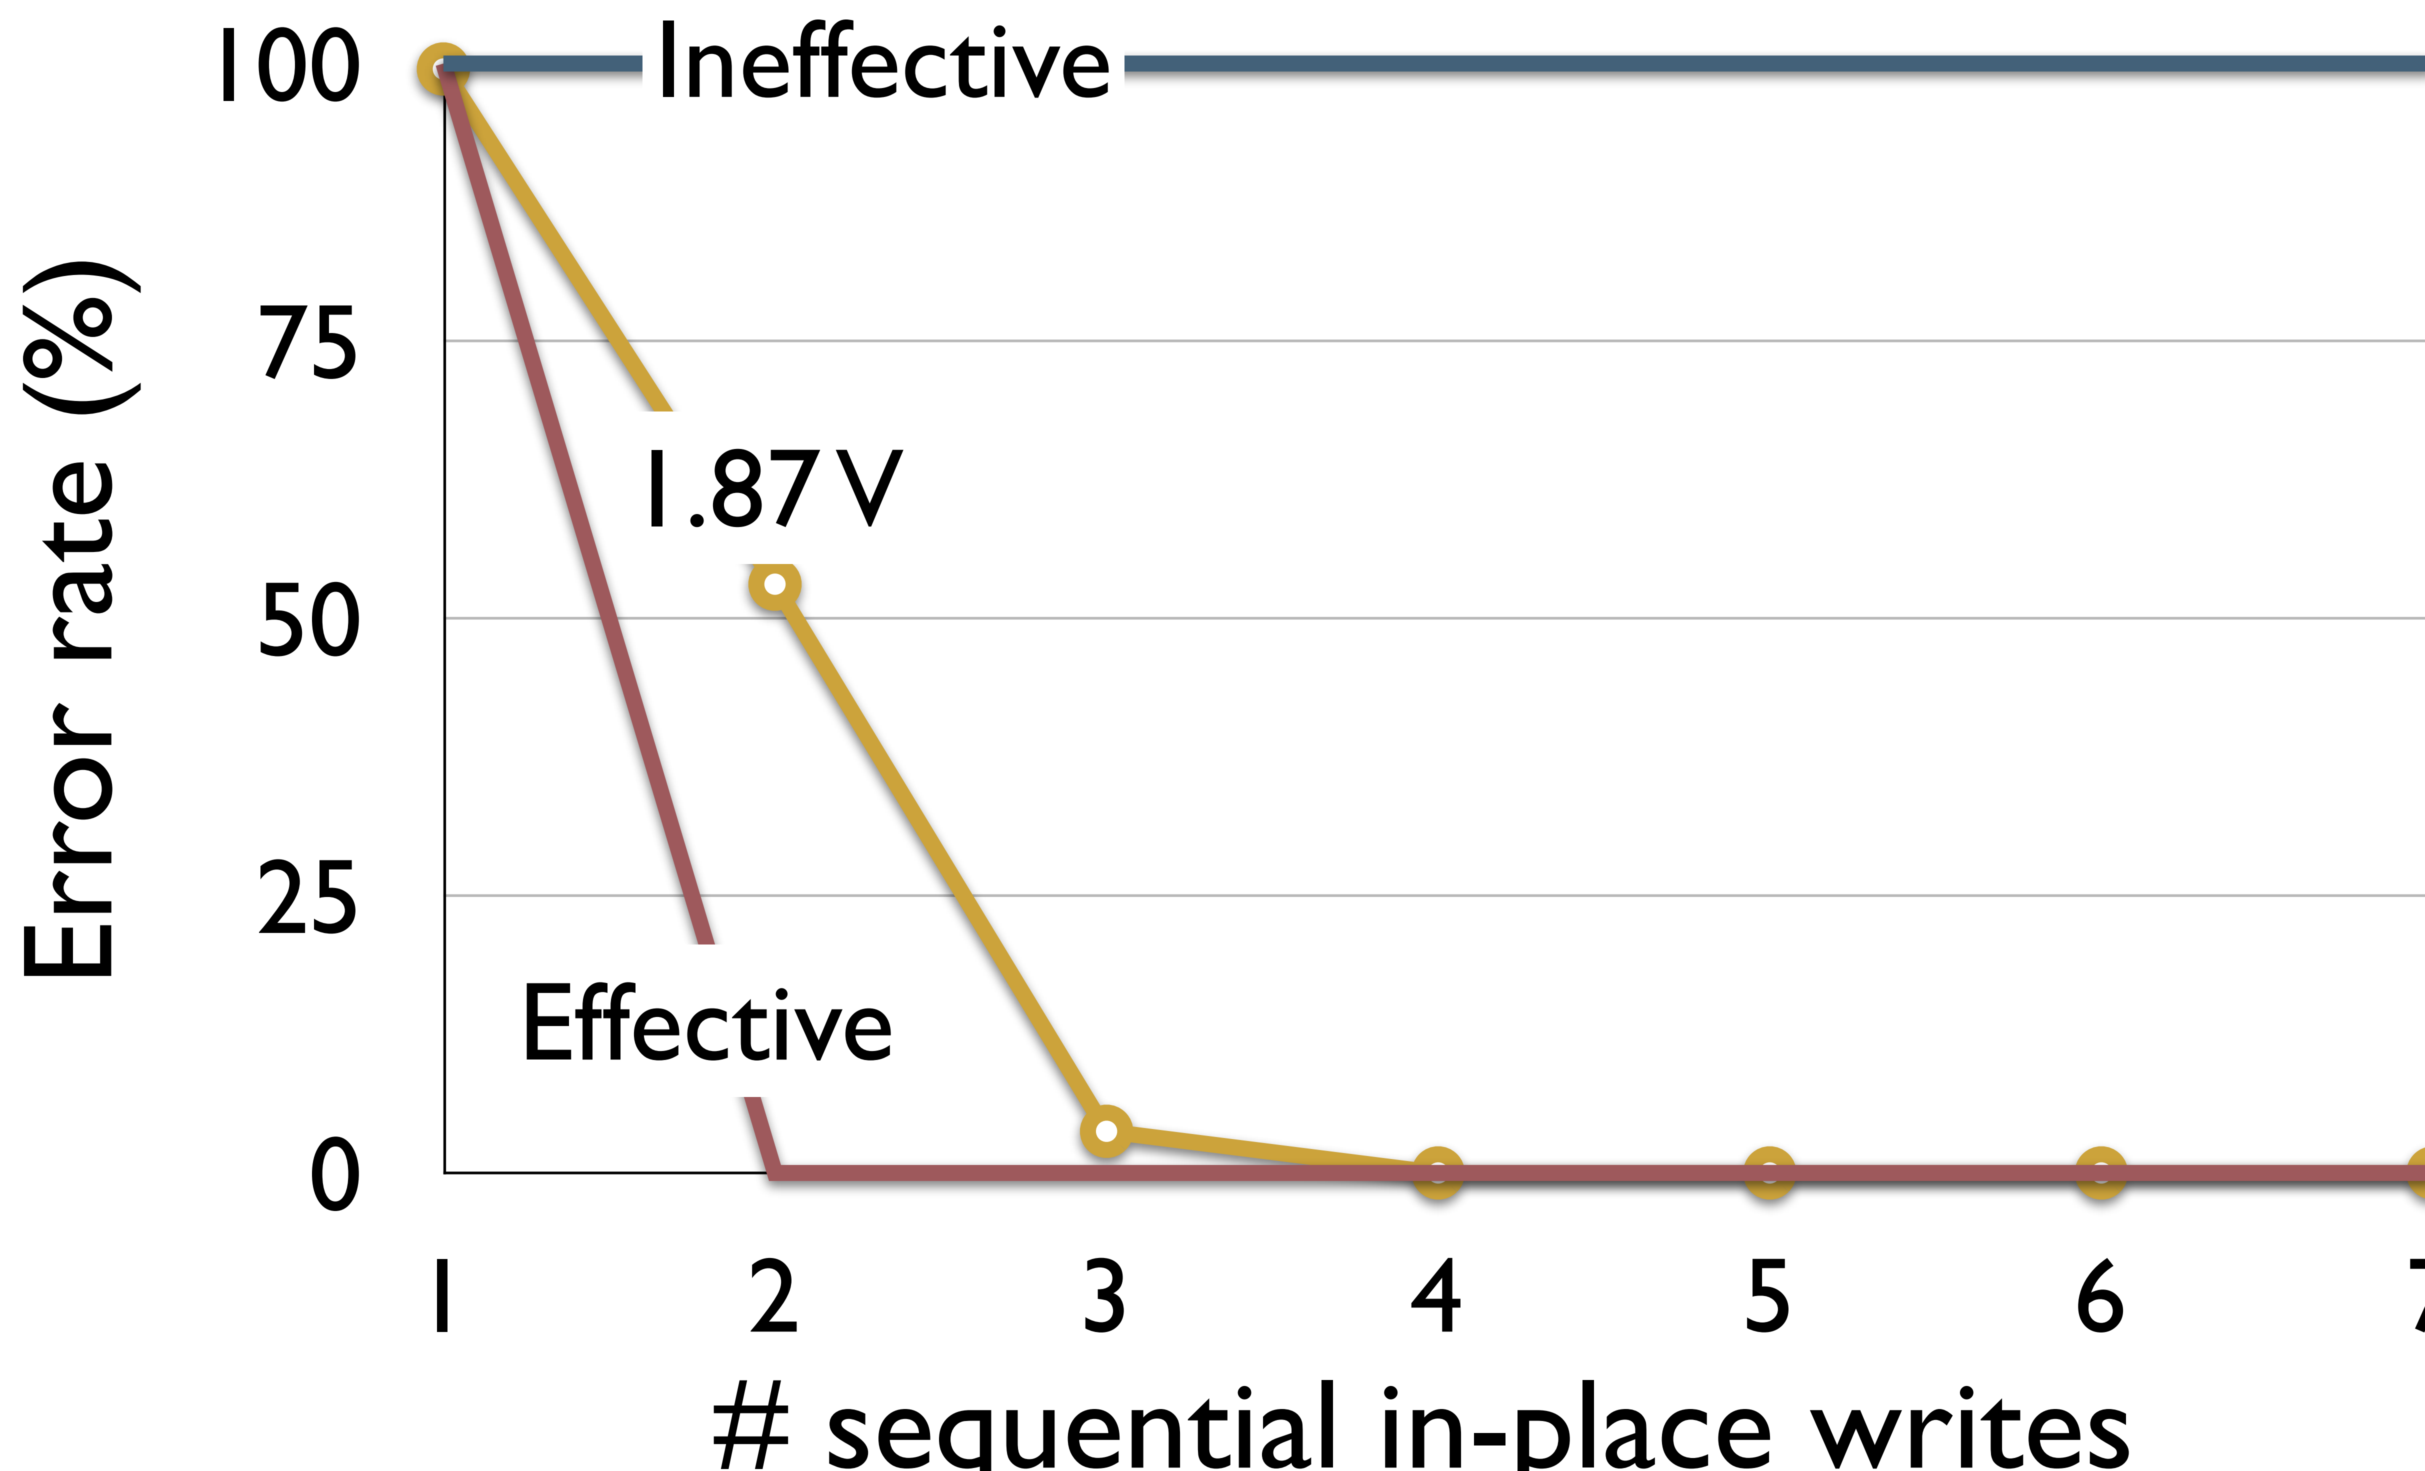

In-place writes

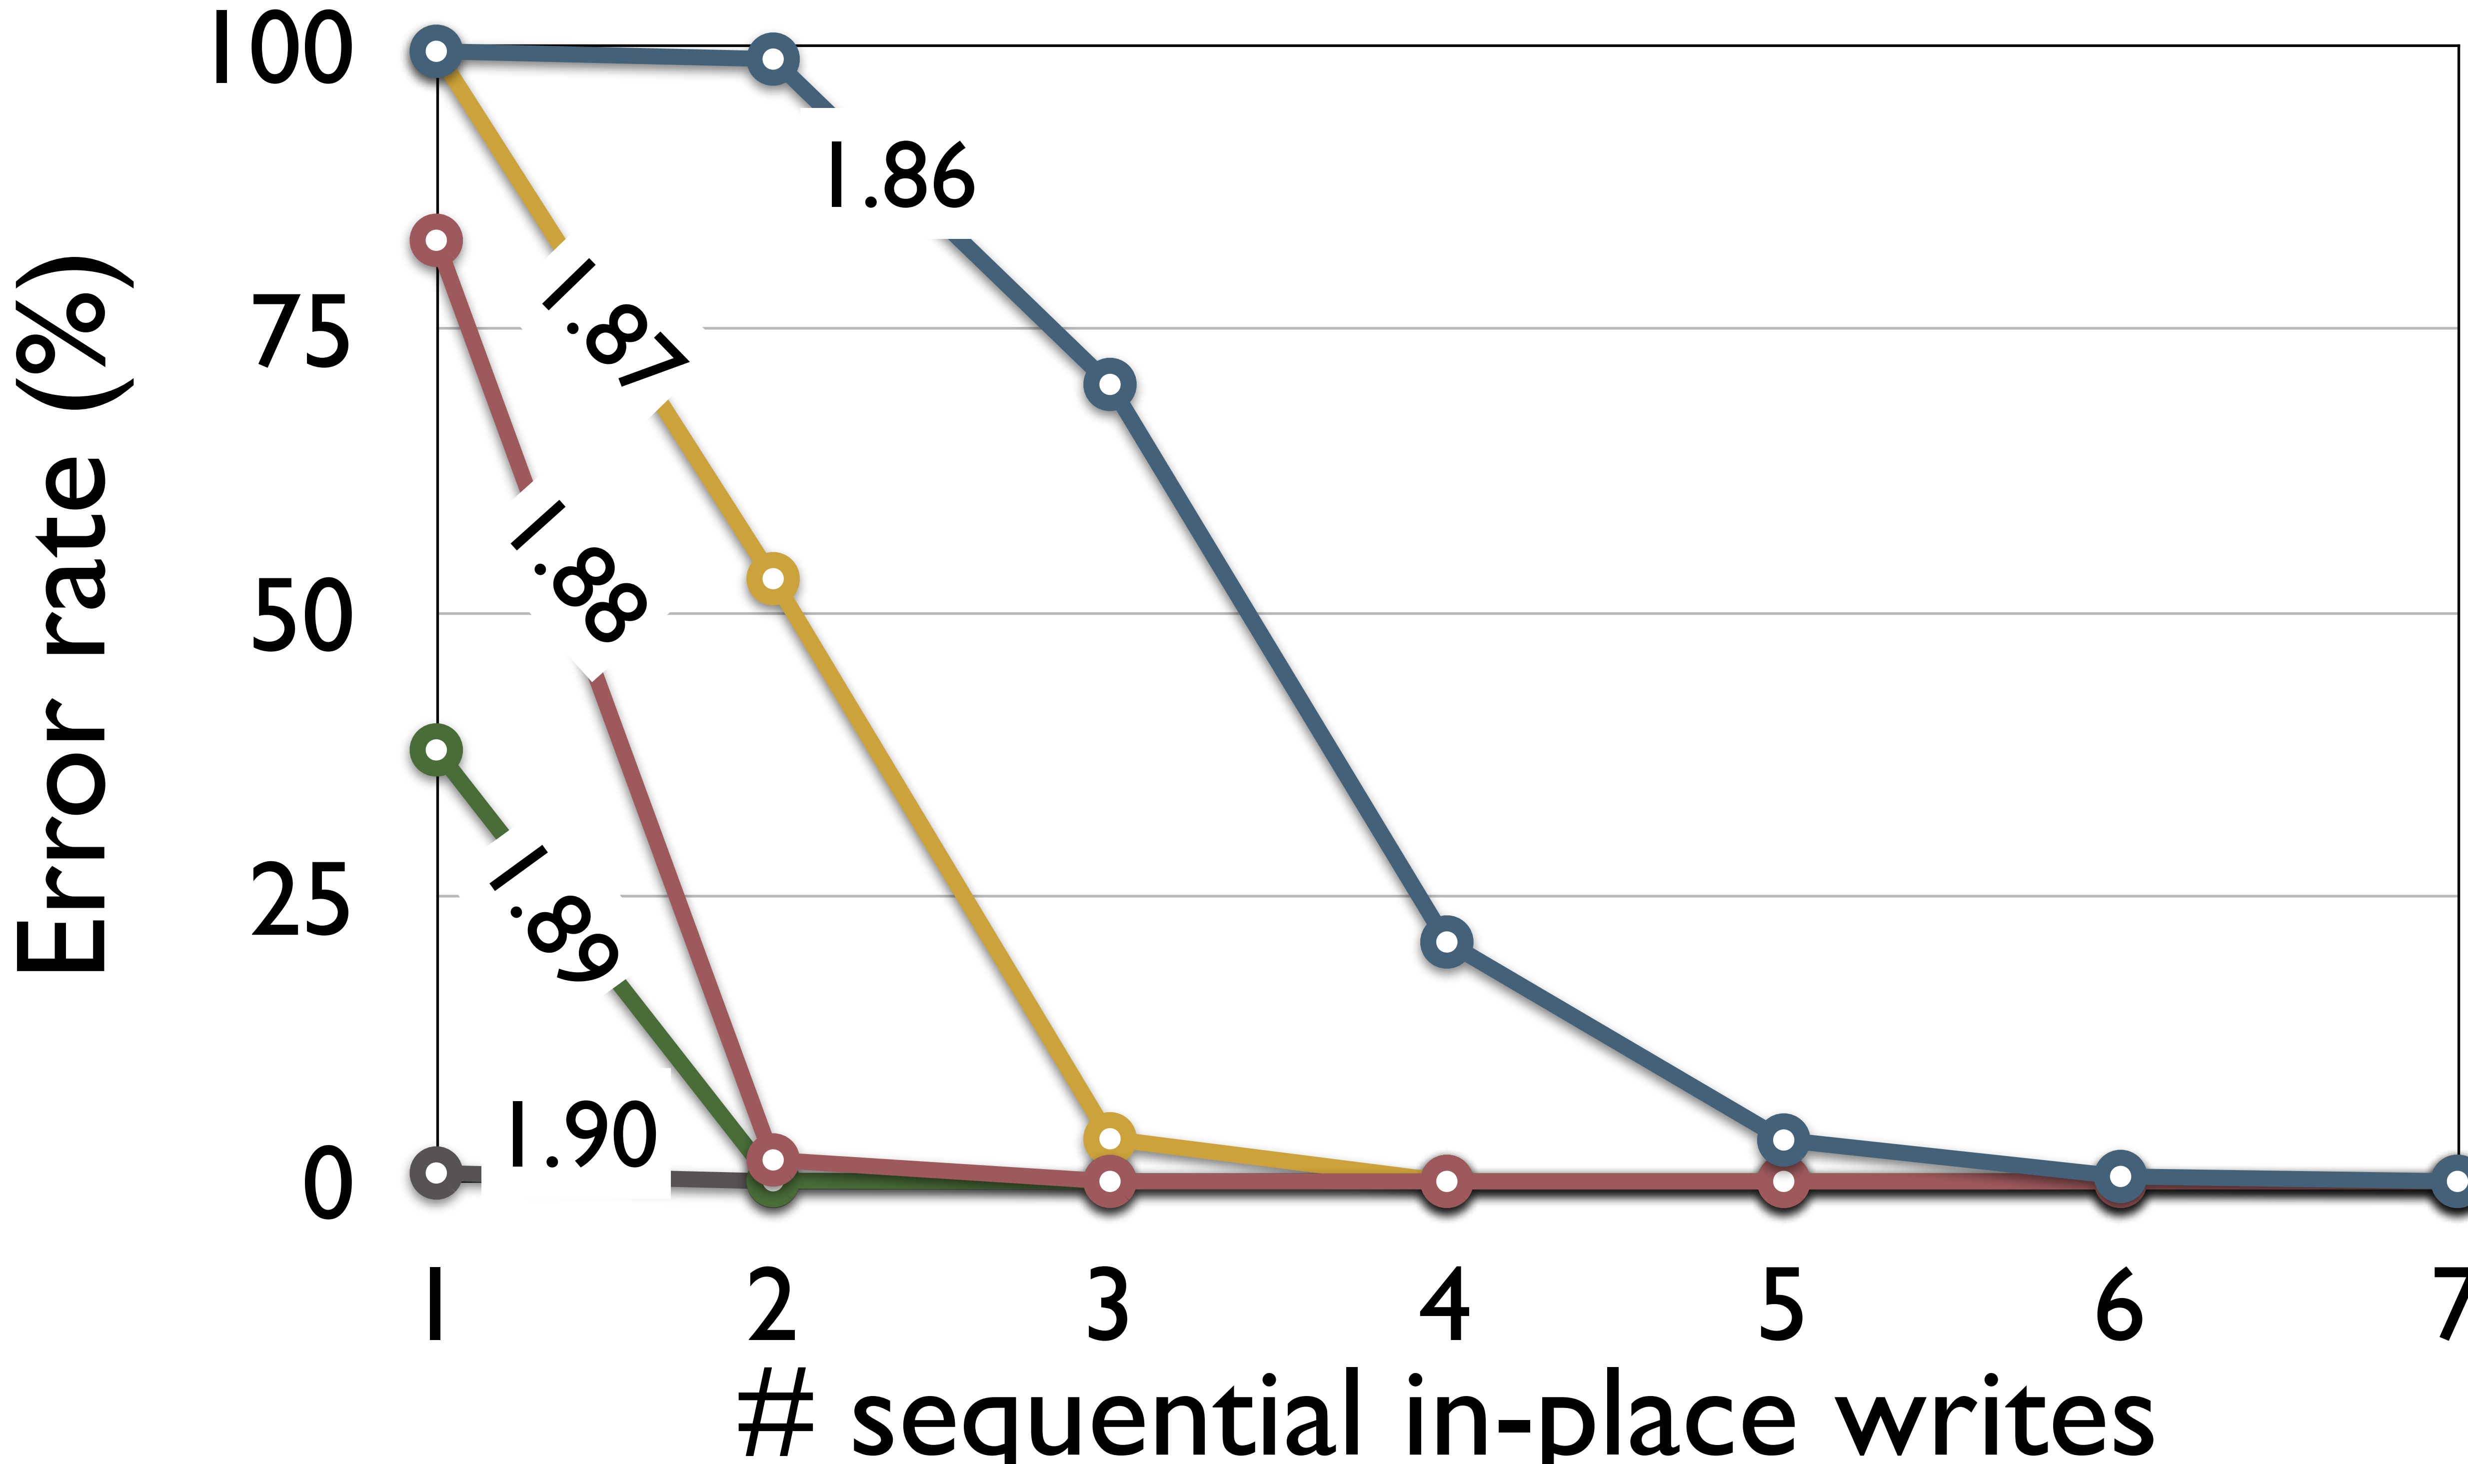

2. Multiple-place writes

- Encoding: Write to more than one location.

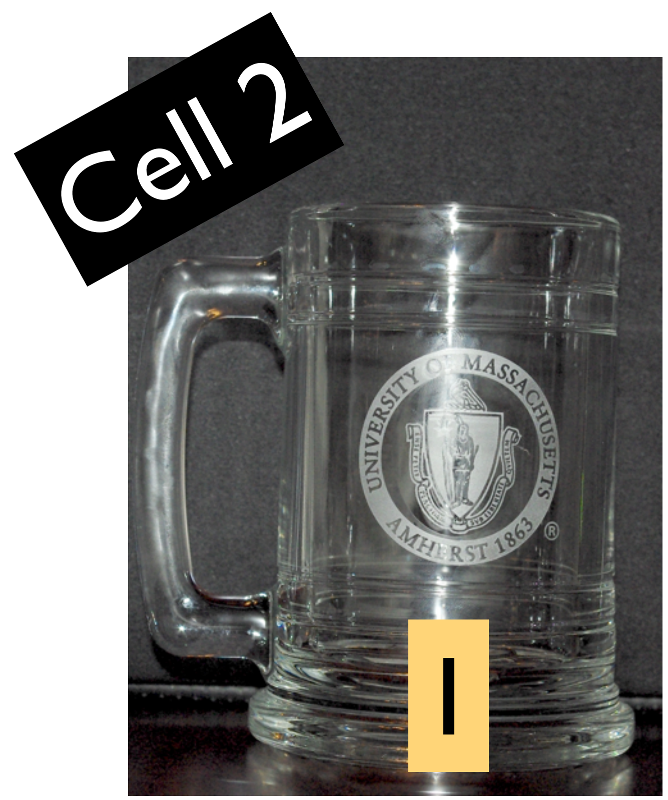

2. Multiple-place writes

- Encoding: Write to more than one location.

2. Multiple-place writes

- Encoding: Write to more than one location.

- Decoding: AND all of the values at read time.

2. Multiple-place writes

- Encoding: Write to more than one location.

- Decoding: AND all of the values at read time.

Multiple-place writes

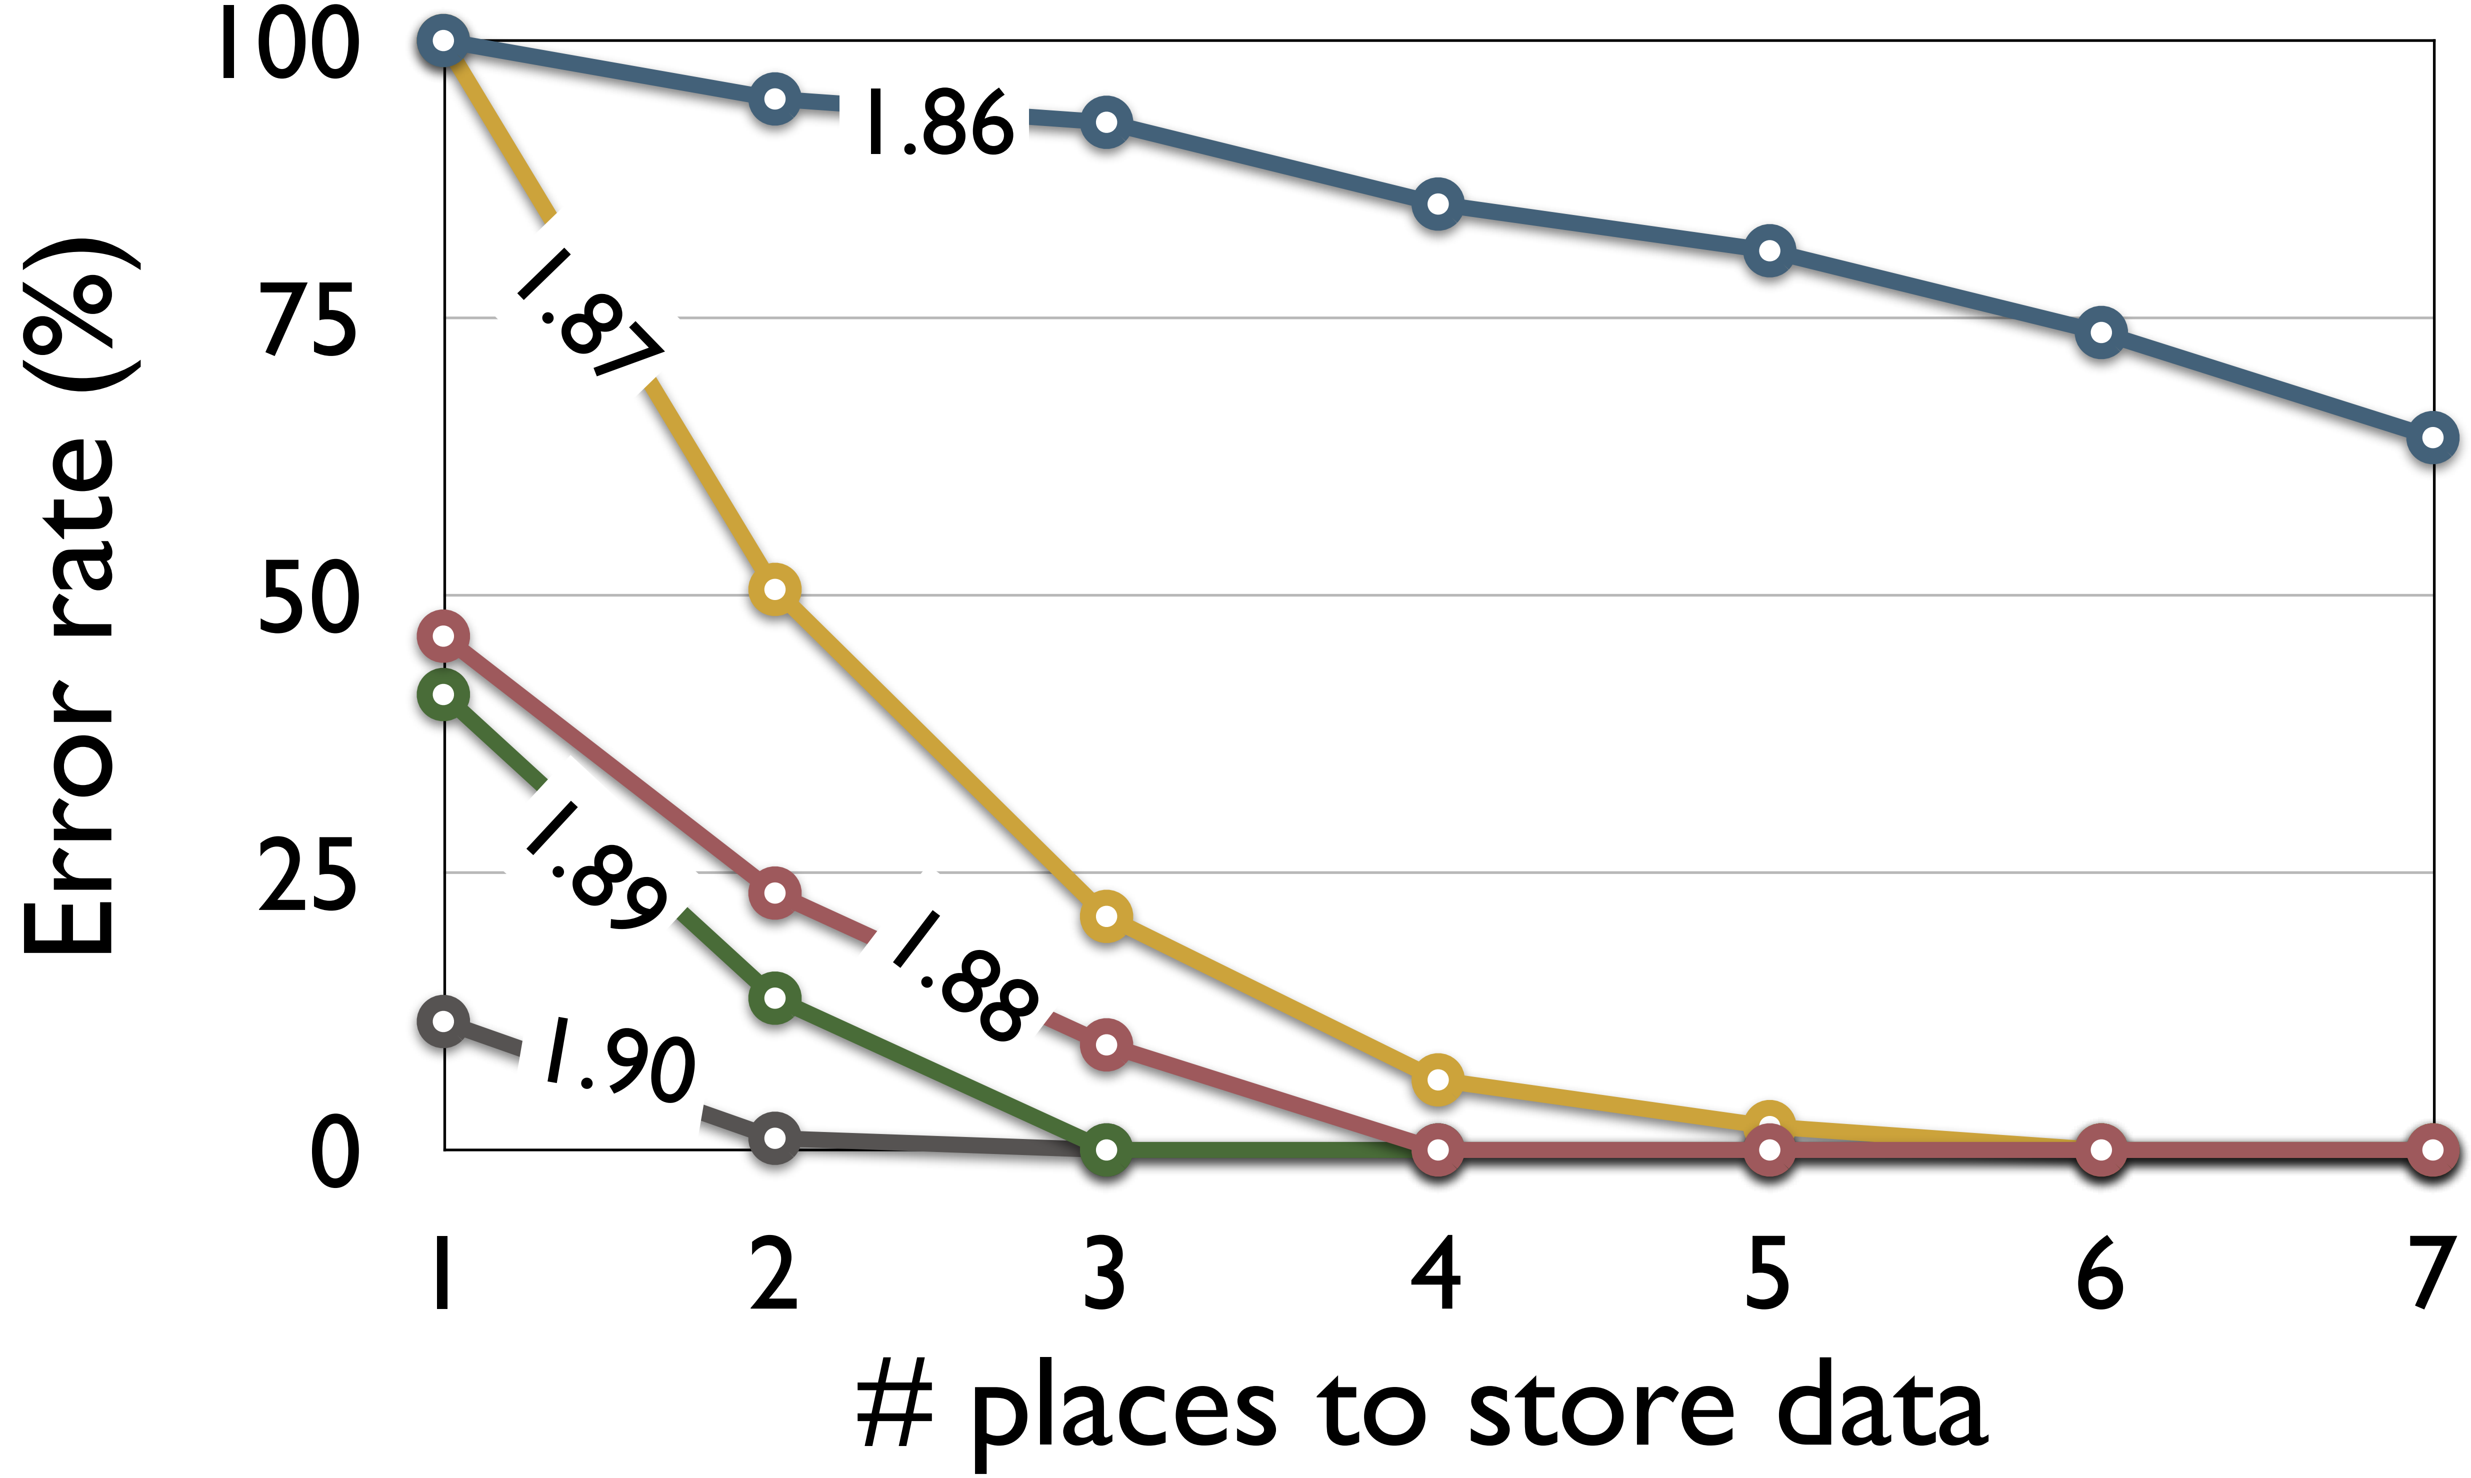

Design Goals

Minimize:

- Energy consumption
- Error rate ✔
- Delay

Comparison at 1.9V

Store: Accelerometer trace
Repeating the writes 2 times/ locations

Comparison at 1.9V

Store: Accelerometer trace
Repeating the writes 2 times/ locations

Method	Time (ms)	Energy (μJ)
In-place	15.43	38
Multiple-place	16.85	40

Comparison at 1.9V

Store: Accelerometer trace
Repeating the writes 2 times/ locations

Method	Time (ms)	Energy (μJ)
In-place	15.43	38
Multiple-place	16.85	40

Micro-benchmarks:

In-Place writes

at

low voltage

vs

Standard approach

at

high voltage

Half-wits Vs. Wits

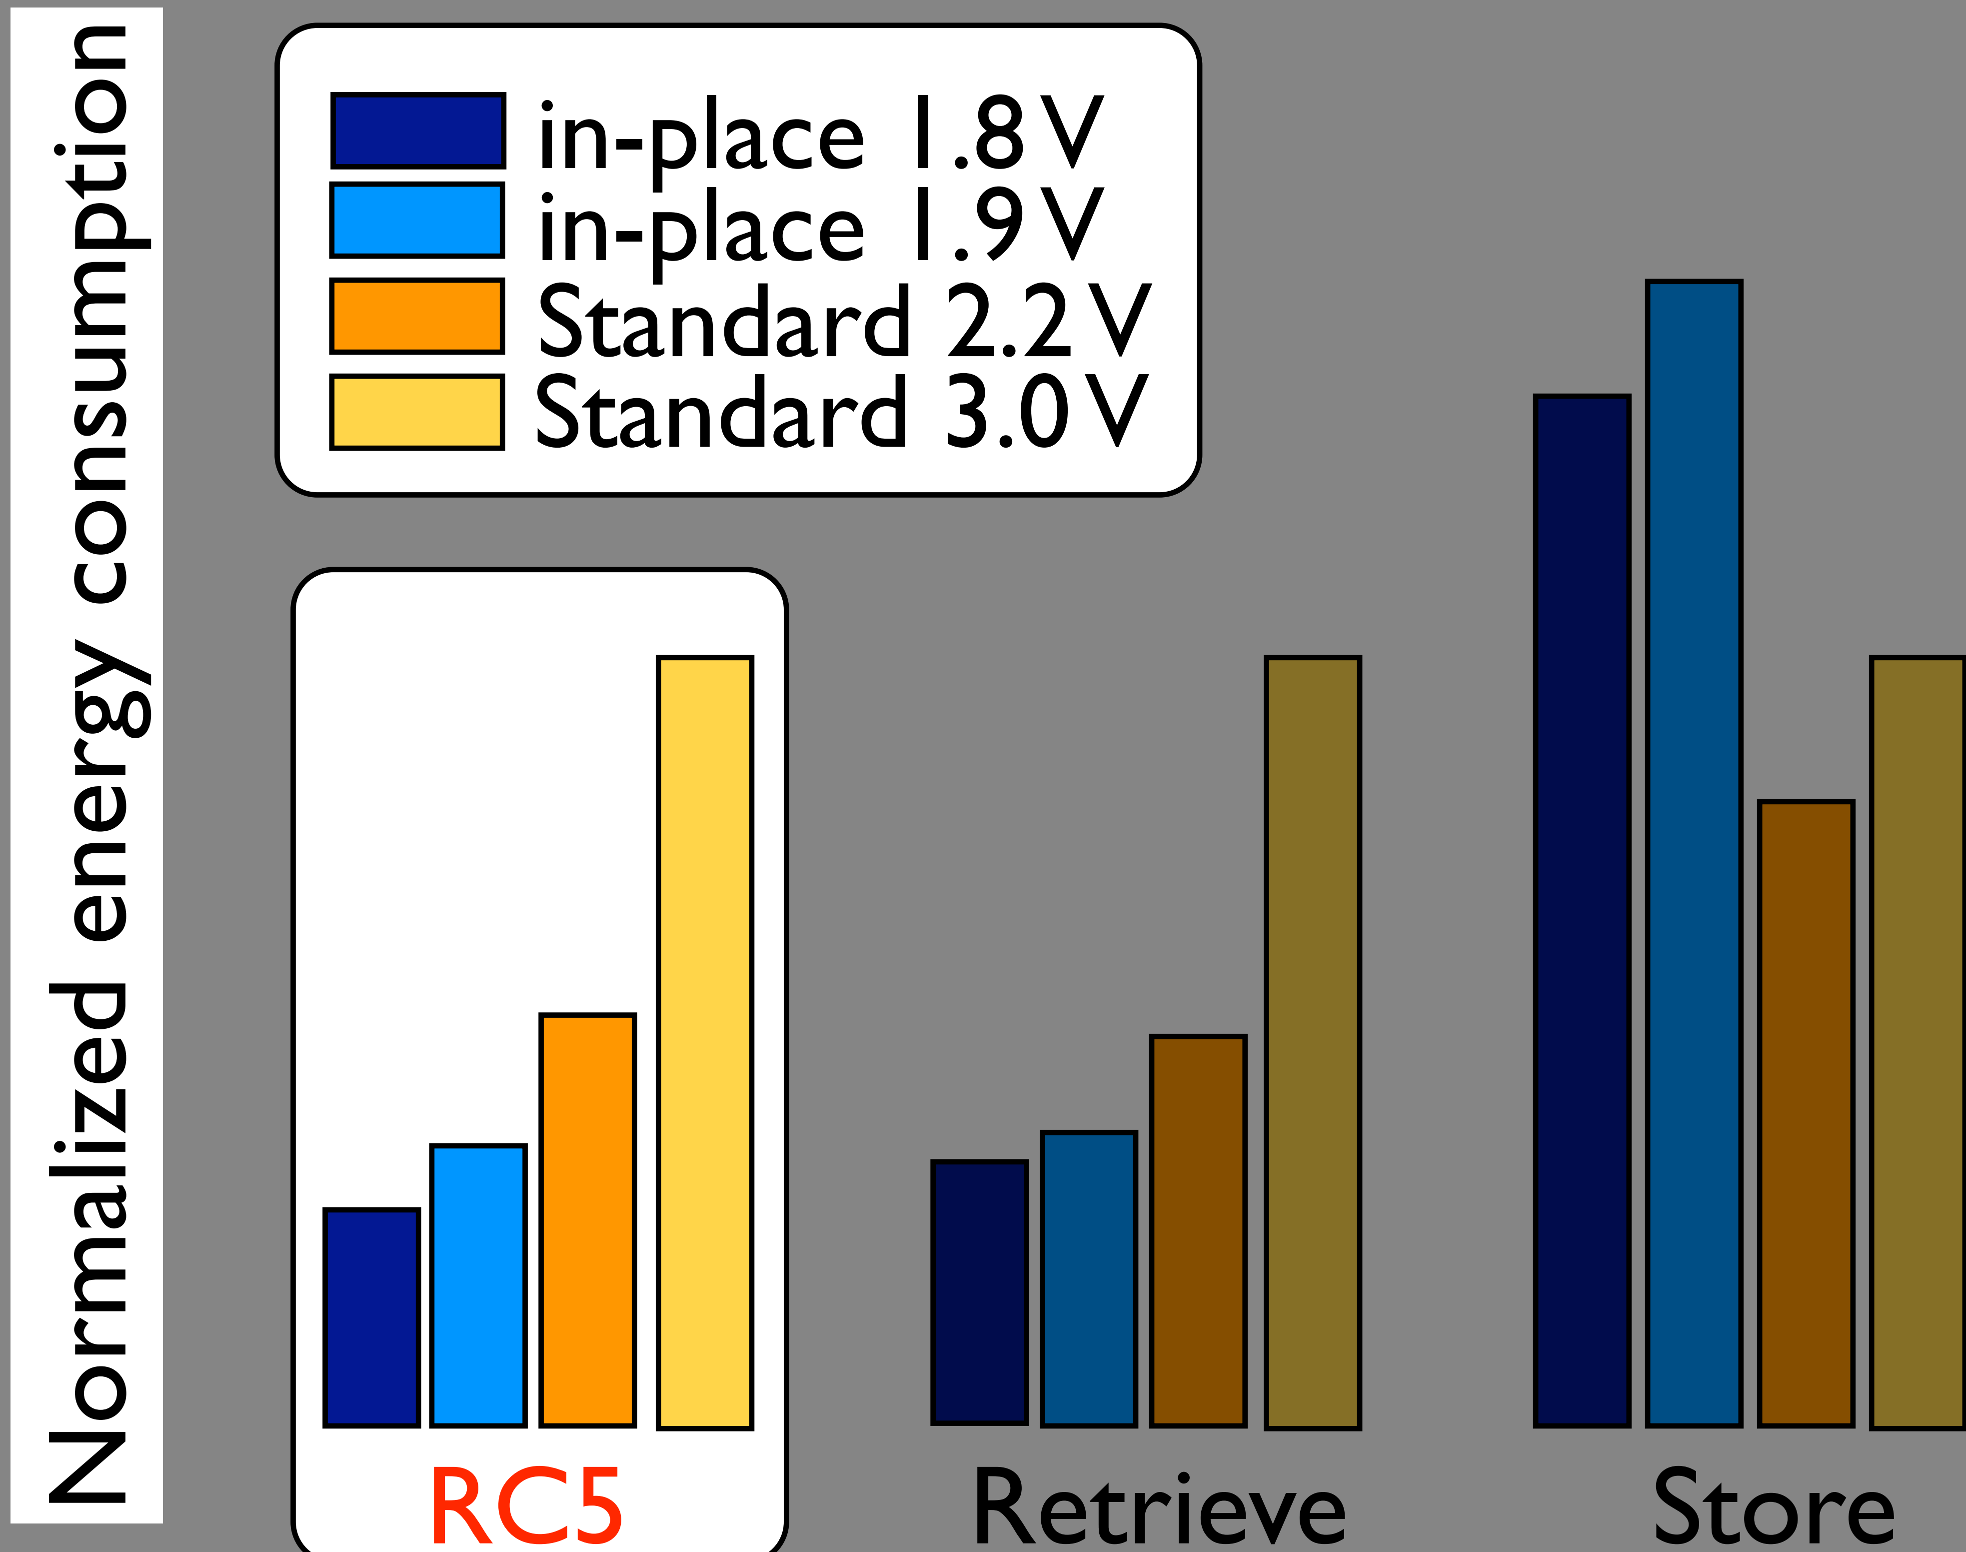

[Rivest:94]

Half-wits Vs. Wits

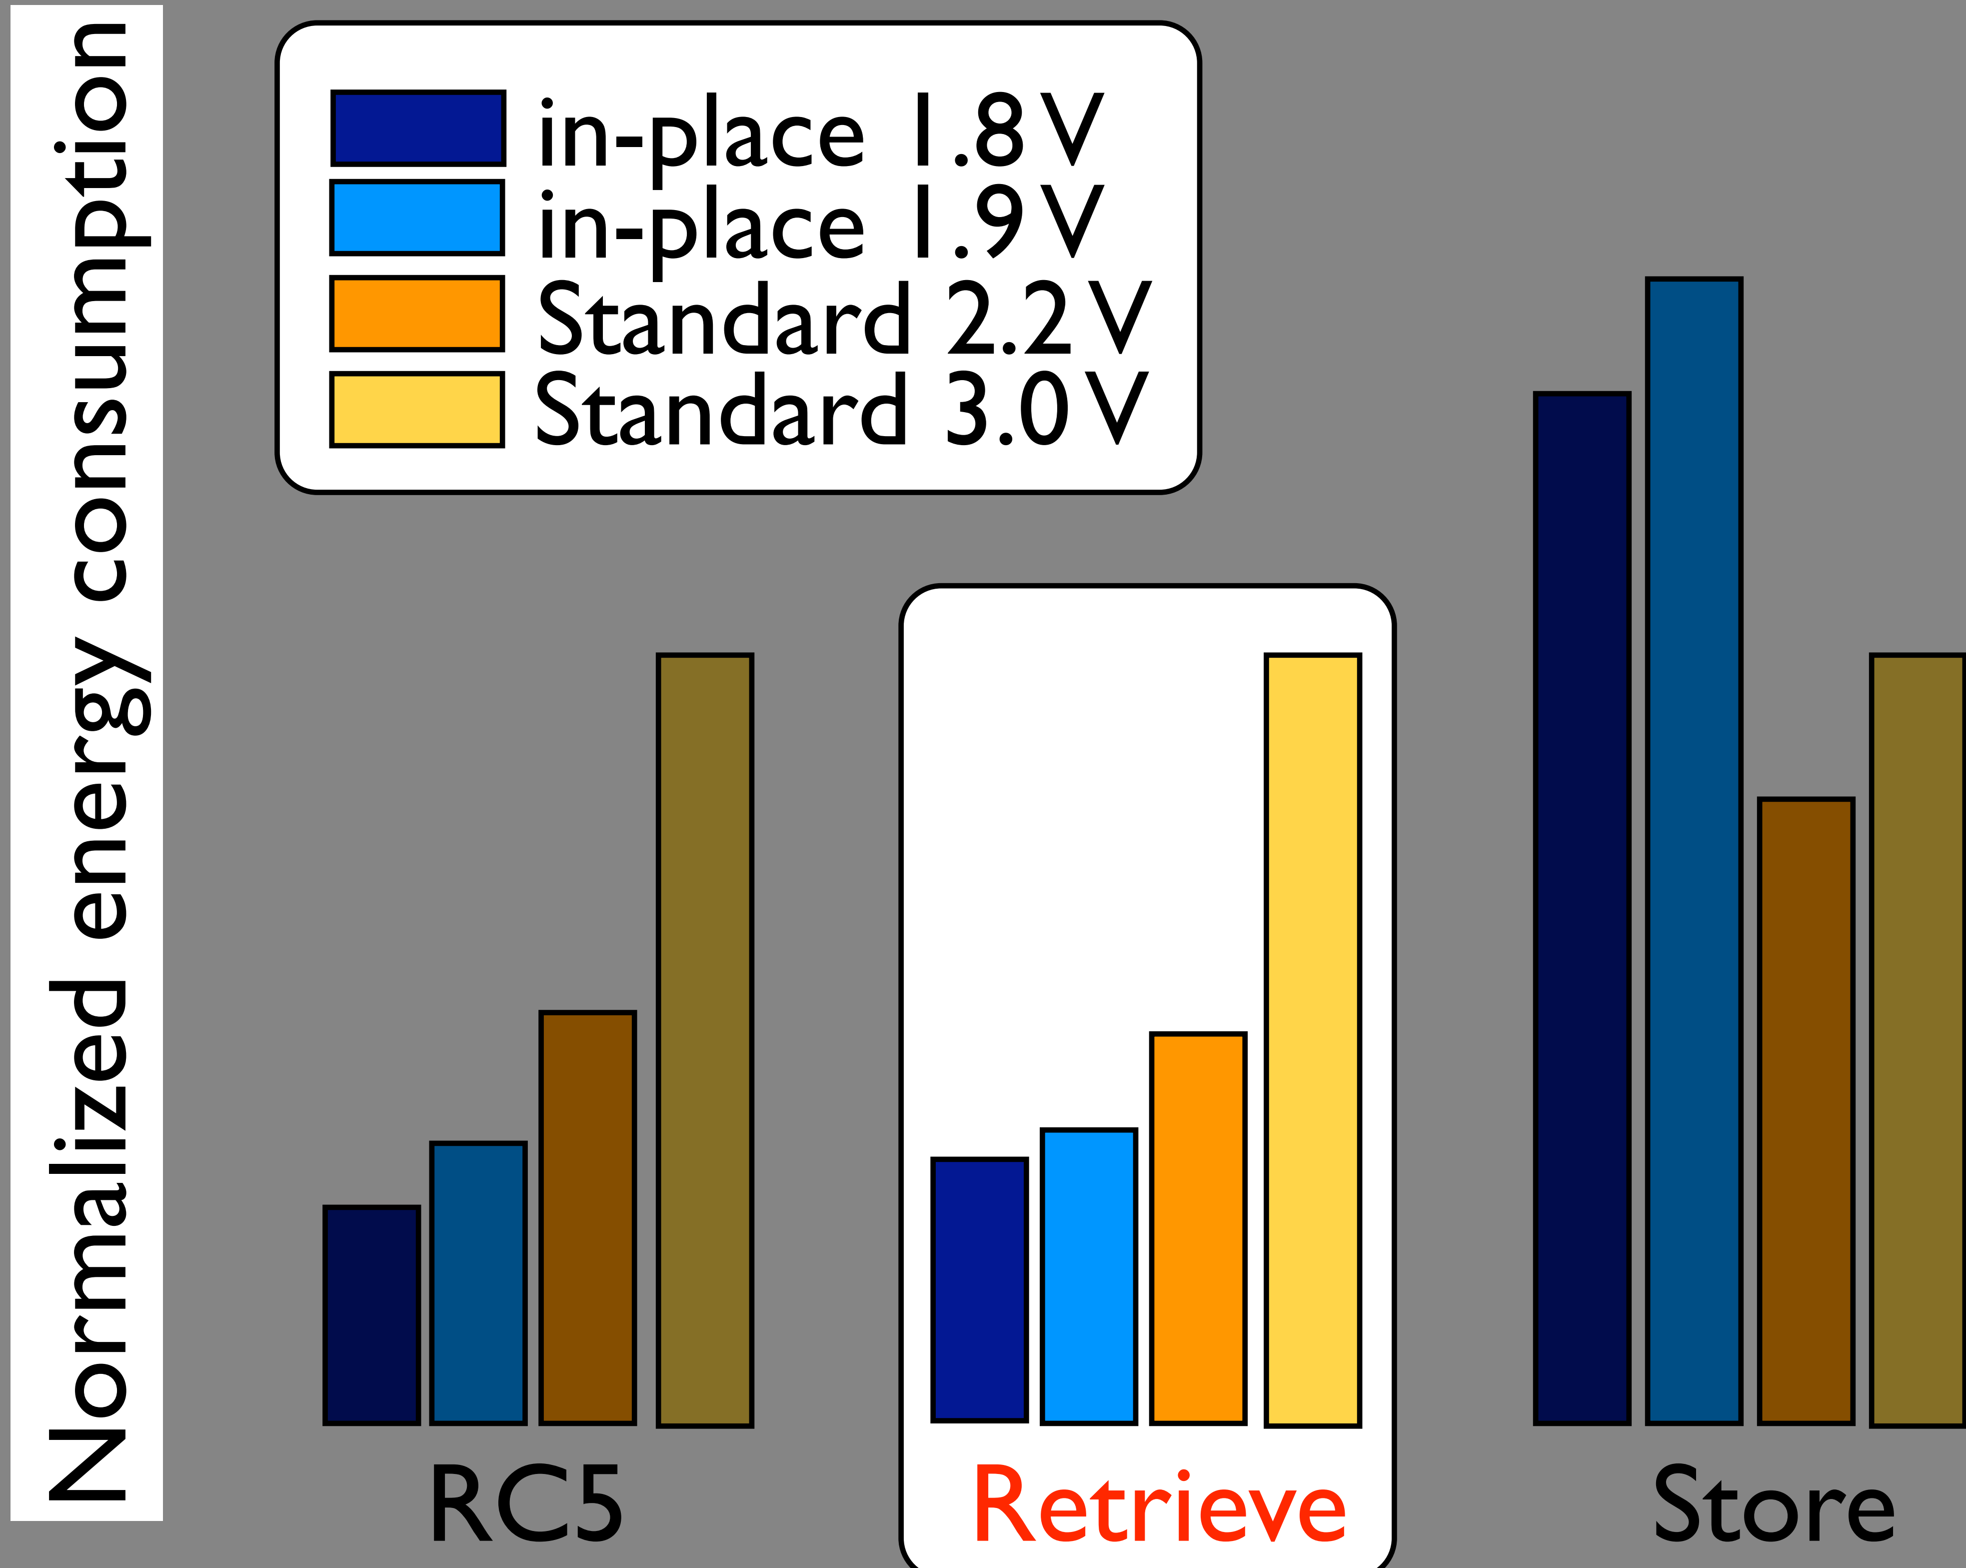

Half-wits Vs. Wits

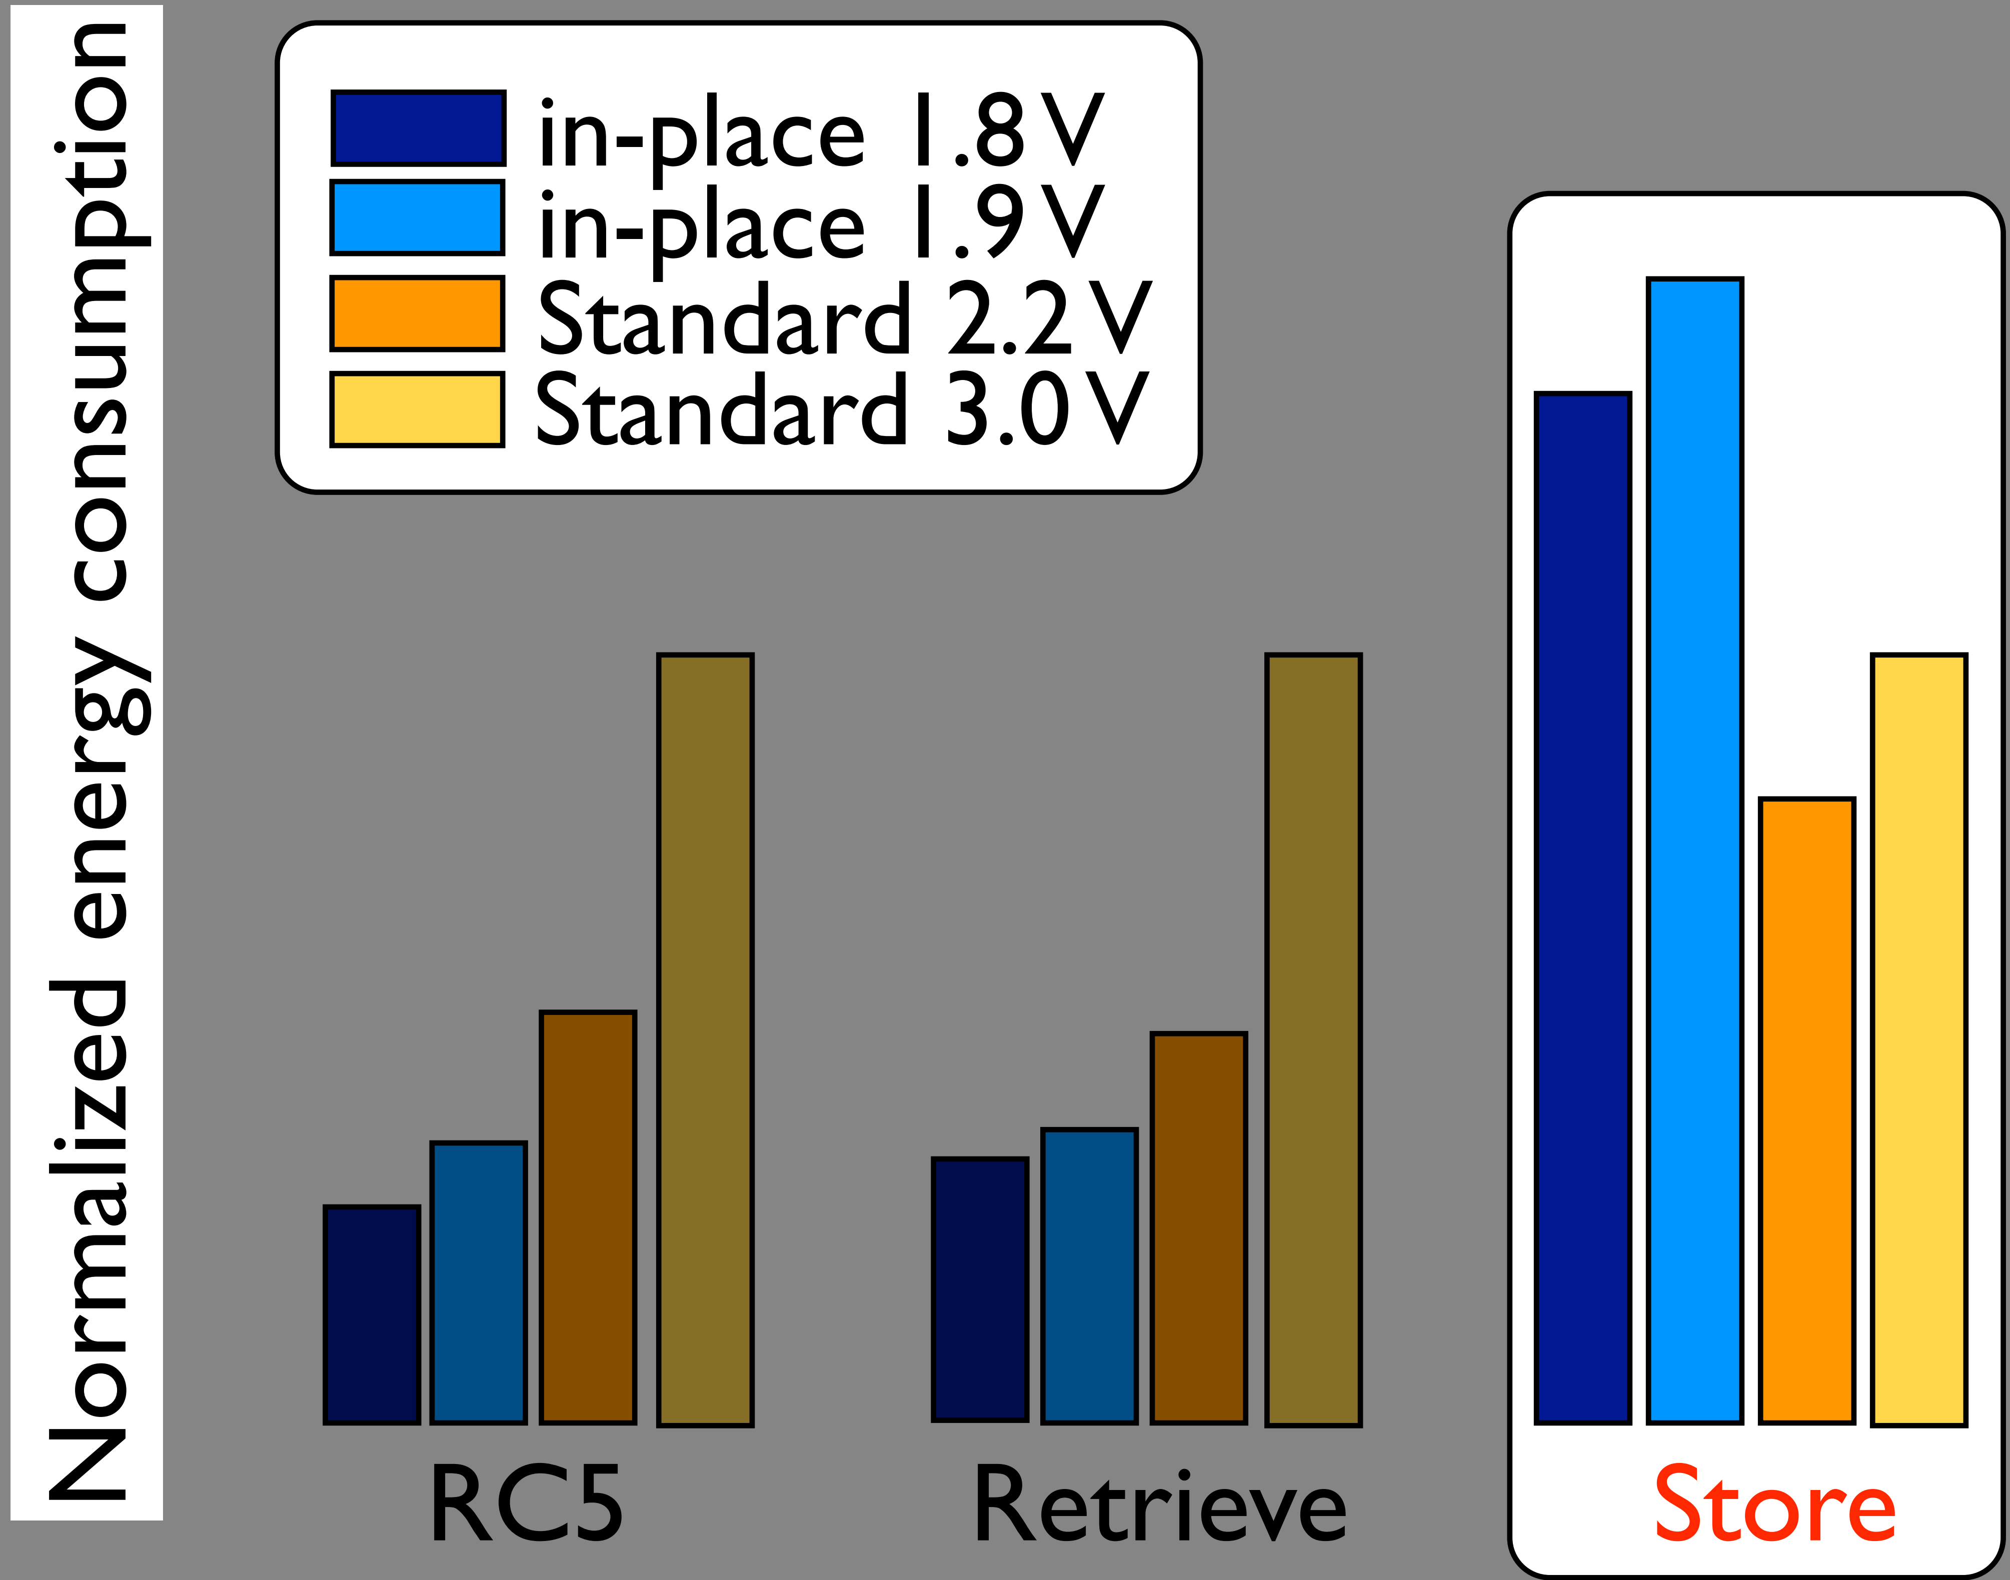

Half-wits Vs. Wits

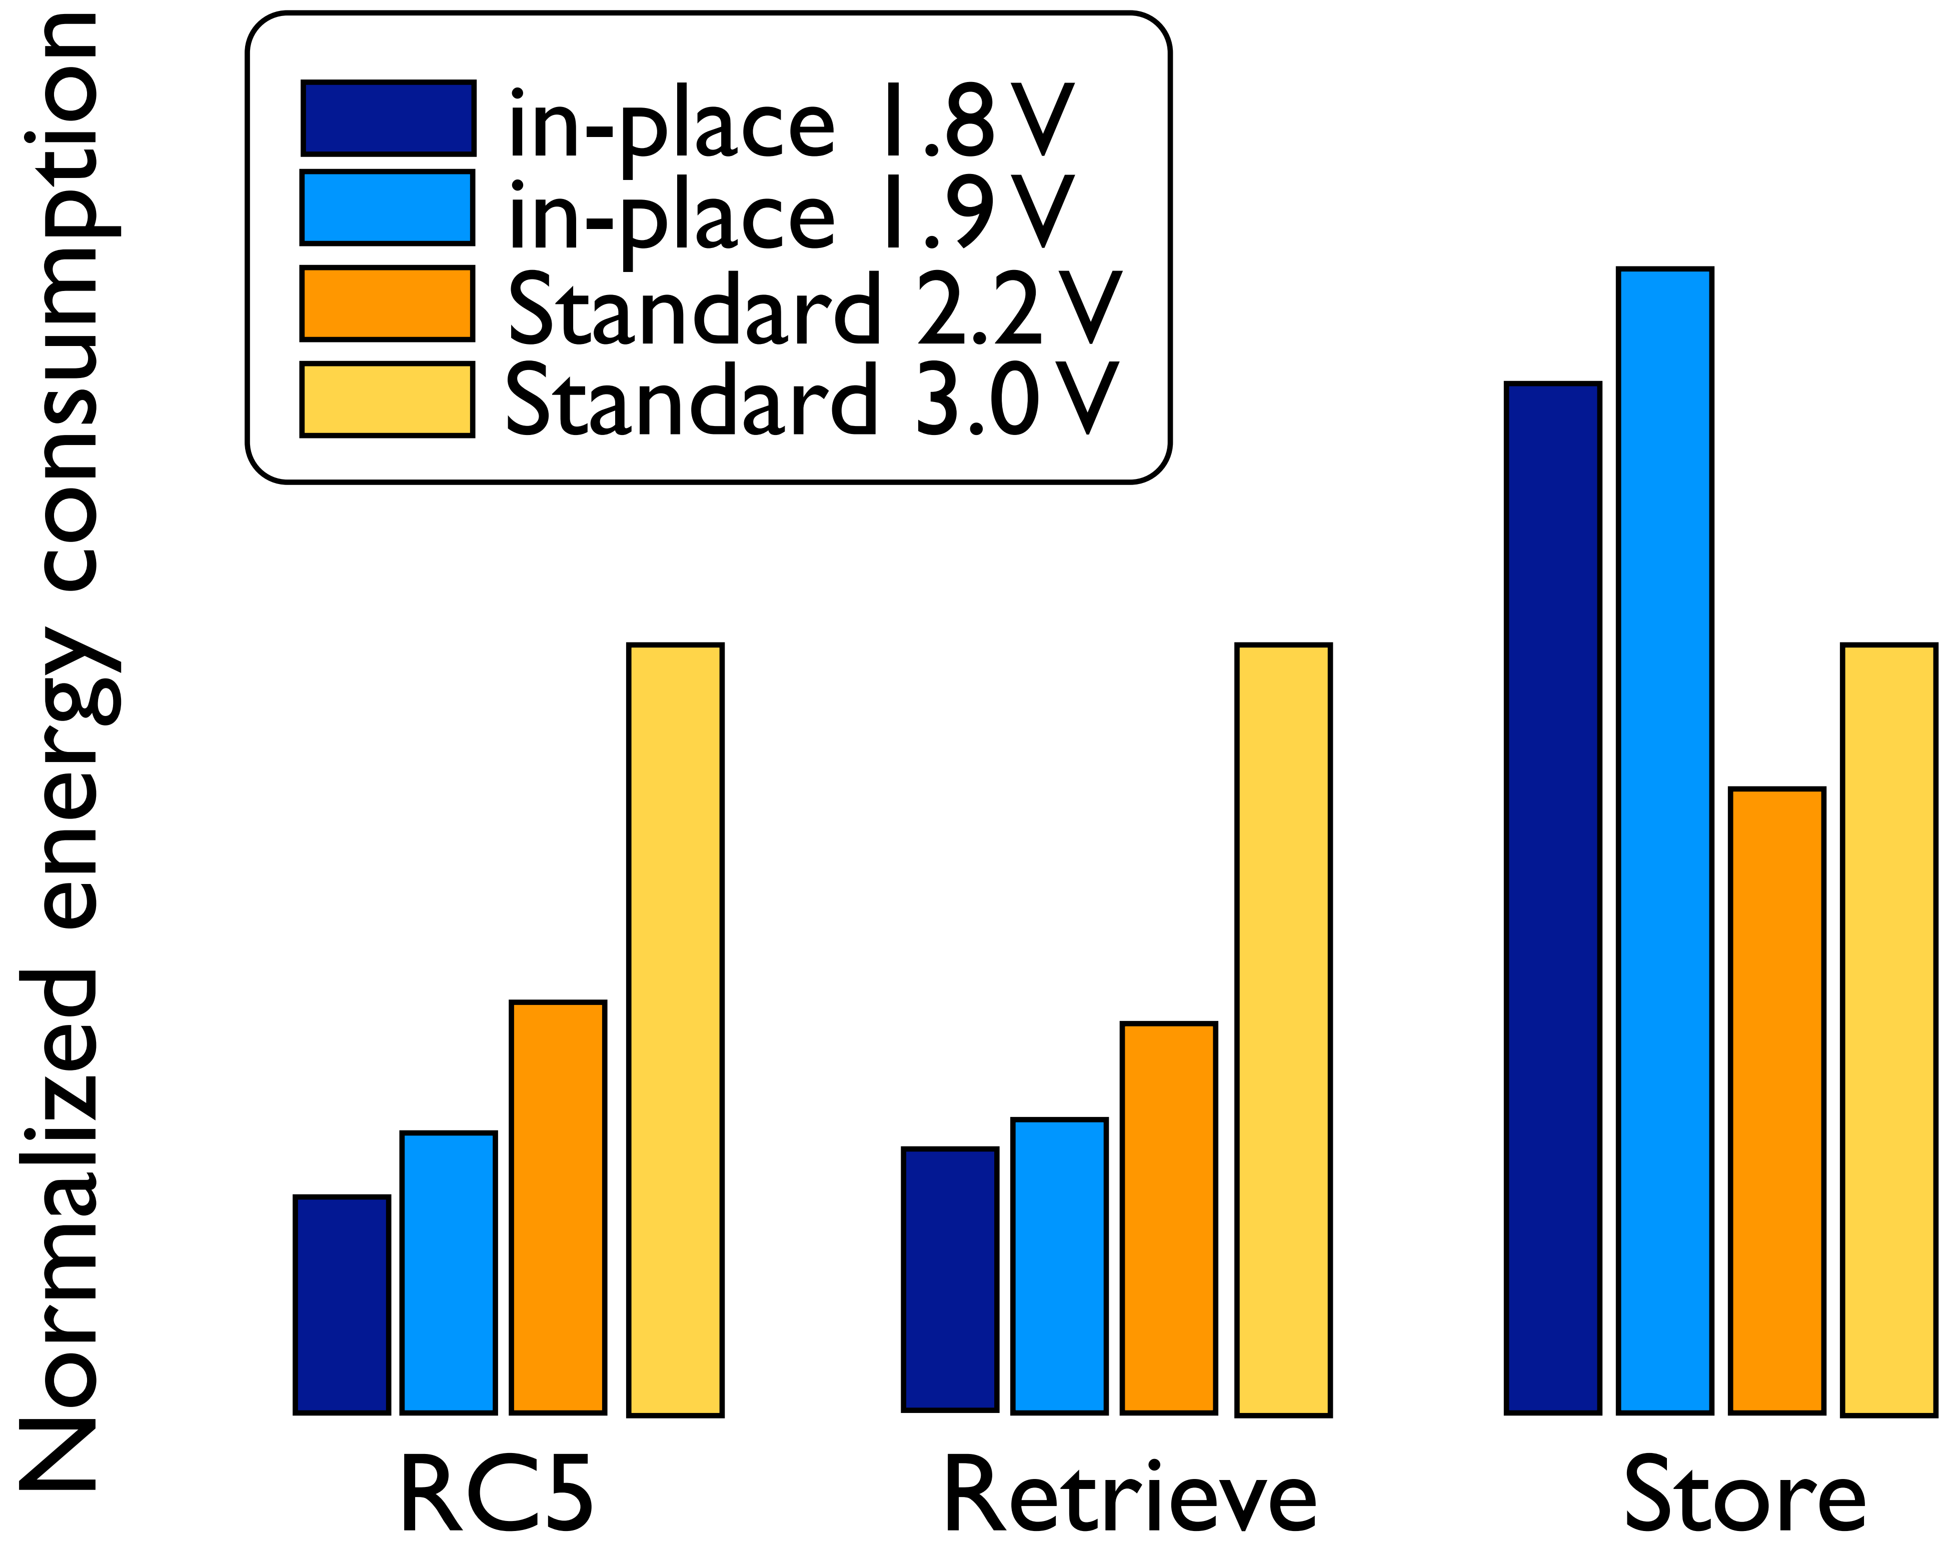

Hypothesis

For CPU-bound workloads:

Energy of low-voltage system < Energy of high-voltage system

Monitoring application

- Read 256 bytes of accelerometer data
- Aggregate data: Min, Max, Mean, Std. dev.
- Write the aggregation of 256 bytes of data

In-place Writes

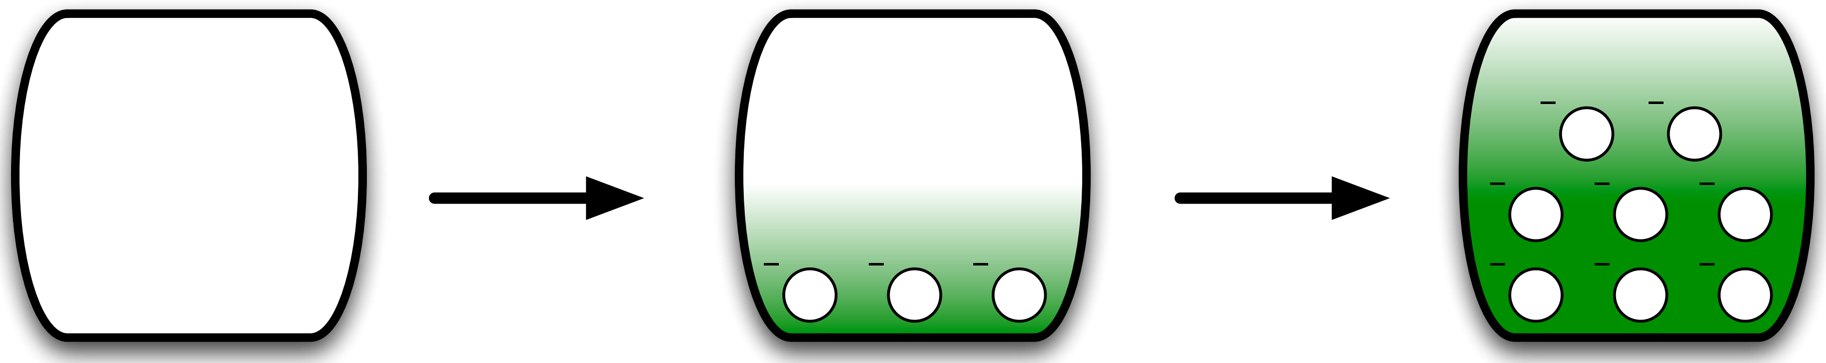

34%

Energy Savings

Monitoring application

Method	In-place 1.8 v	In-place 1.9 V	Standard 2.2 V	Standard 3.0 V
Energy (μJ)	270	300	410	760

Monitoring application

Method	In-place 1.8 v	In-place 1.9 V	Standard 2.2 V	Standard 3.0 V
Energy (μJ)	270	300	410	760

Monitoring application

Method	In-place 1.8 v	In-place 1.9 V	Standard 2.2 V	Standard 3.0 V
Energy (μJ)	270	300	410	760

Further improvements

1. Sign-bits and storing the complement [Papirla:09]

2. Memory mapping-table

Summary: Exploiting Half-Wits

- In-place writes on half-wits is an effective way to reduce wasted energy.
- Microcontrollers can work at a lower voltage and get more work done with the same amount of energy.
- The digital abstractions pay a higher price than necessary to provide reliability.