Probable Cause

The Deanonymizing Effects of Approximate DRAM

Amir Rahmati, Matthew Hicks, Dan Holcomb, Kevin Fu

Approximate Computing

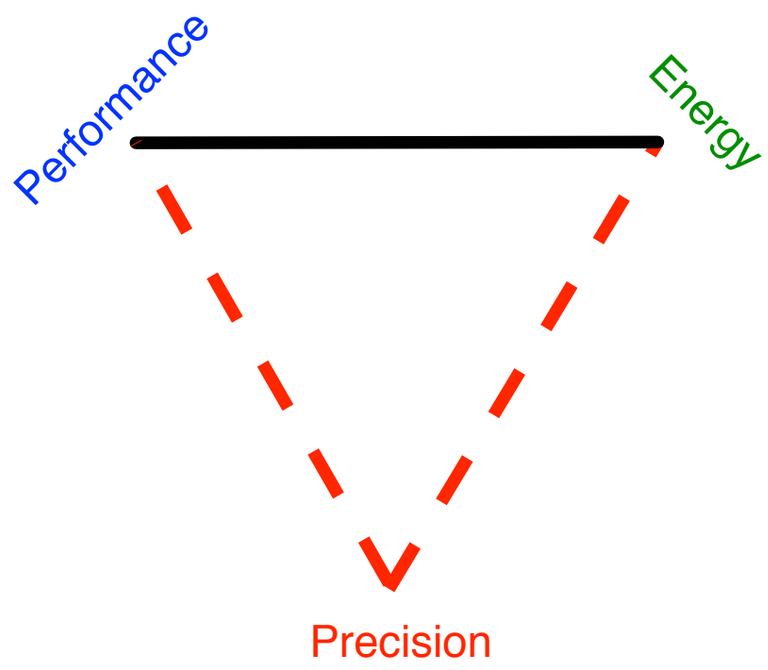

Approximate Computing

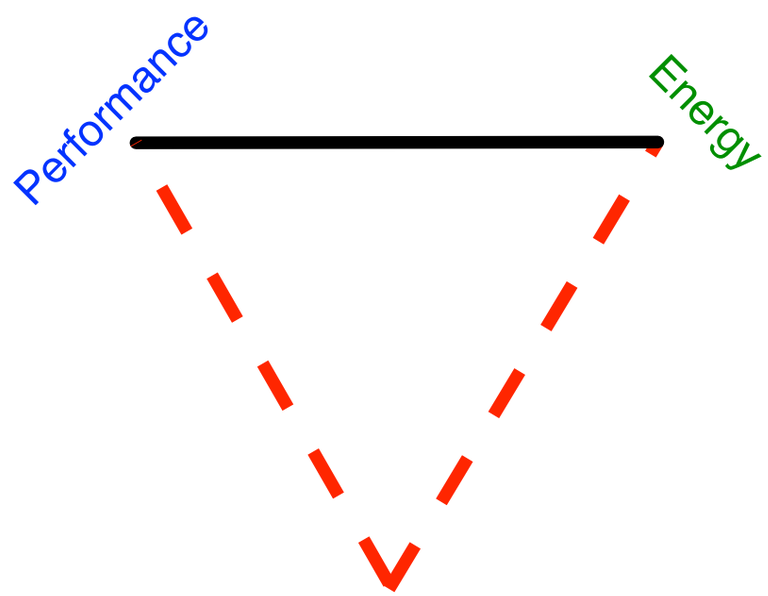

Precise computation
is **not** required
in many applications:

Approximate Computing

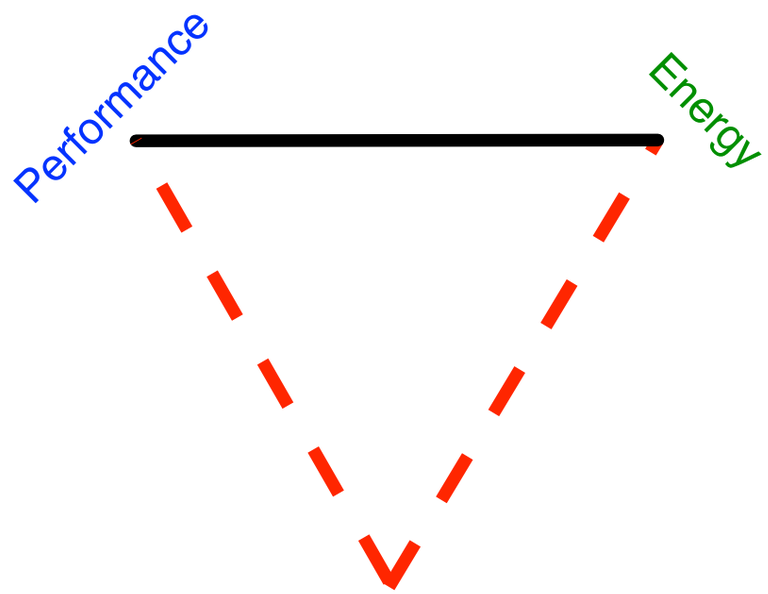

Precise computation is **not** required in many applications:

Machine learning, sensory data, information retrieval, physical simulation, computer vision…

Approximate Computing

Approximate Computing

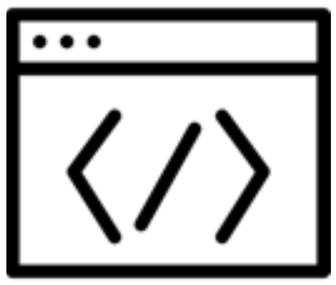

Programming Language

UncertainT (ASPLOS'14)
Enerj (PLDI'11)

Approximate Computing

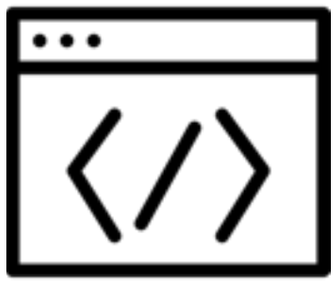

UncertainT (ASPLOS'14)
Enerj (PLDI'11)

Programming Language

Flikker (ASPLOS'11)
Approximate storage in solid state memory (Micro'13)

Storage

Approximate Computing

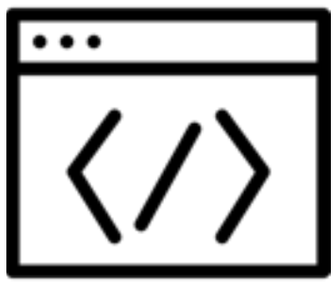

Programming Language

UncertainT (ASPLOS'14)
Enerj (PLDI'11)

Storage

Flikker (ASPLOS'11)
Approximate storage in solid state memory (Micro'13)

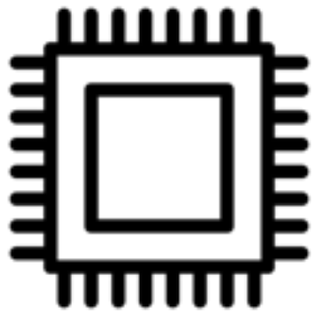

Architecture

Truffle (ASPLOS'12)
Relax (ISCA'10)
ERSA (DATE'10)

Approximate Computing

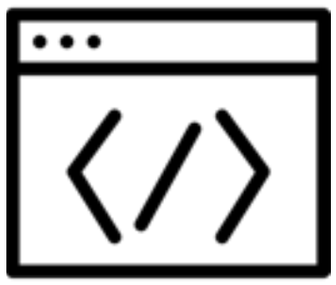

Programming Language

UncertainT (ASPLOS'14)
Enerj (PLDI'11)

Storage

Flikker (ASPLOS'11)
Approximate storage in solid state memory (Micro'13)

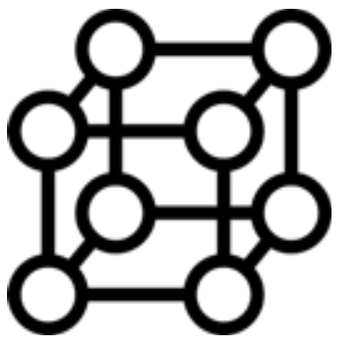

Algorithms

Green (PLDI'10)

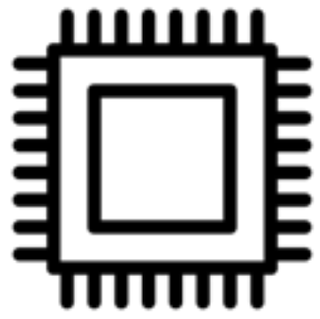

Architecture

Truffle (ASPLOS'12)
Relax (ISCA'10)
ERSA (DATE'10)

Approximate Computing

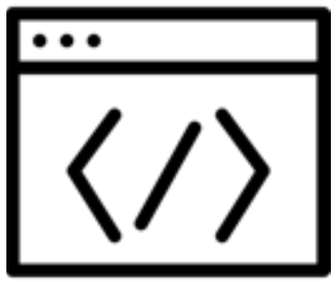

UncertainT (ASPLOS'14)
Enerj (PLDI'11)

Programming Language

Flikker (ASPLOS'11)
Approximate storage in solid state memory (Micro'13)

Storage

Security

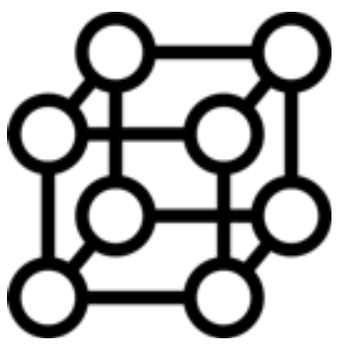

Green (PLDI'10)

Algorithms

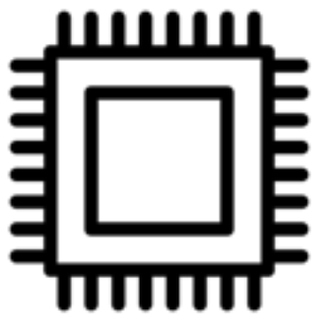

Truffle (ASPLOS'12)
Relax (ISCA'10)
ERSA (DATE'10)

Architecture

Approximate Computing

Security

How does Approximate Computing affect the end-user?

Privacy Implications of Approximate DRAM

Privacy Implications of Approximate DRAM

Identify the origin of data by looking at the error pattern

Overview

Overview

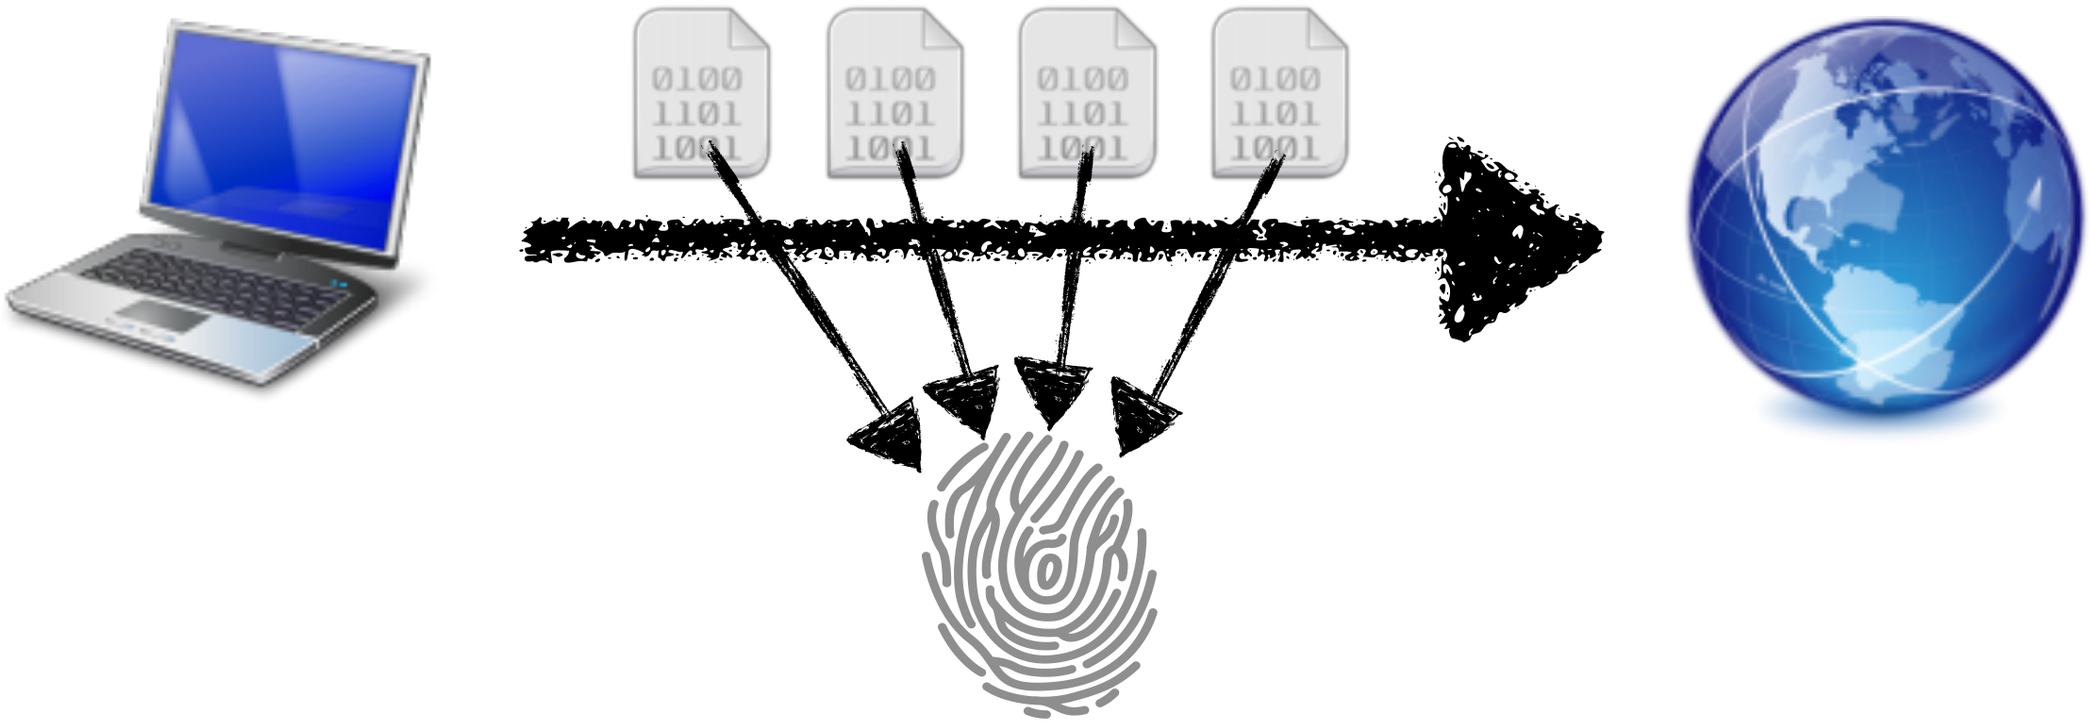

Overview

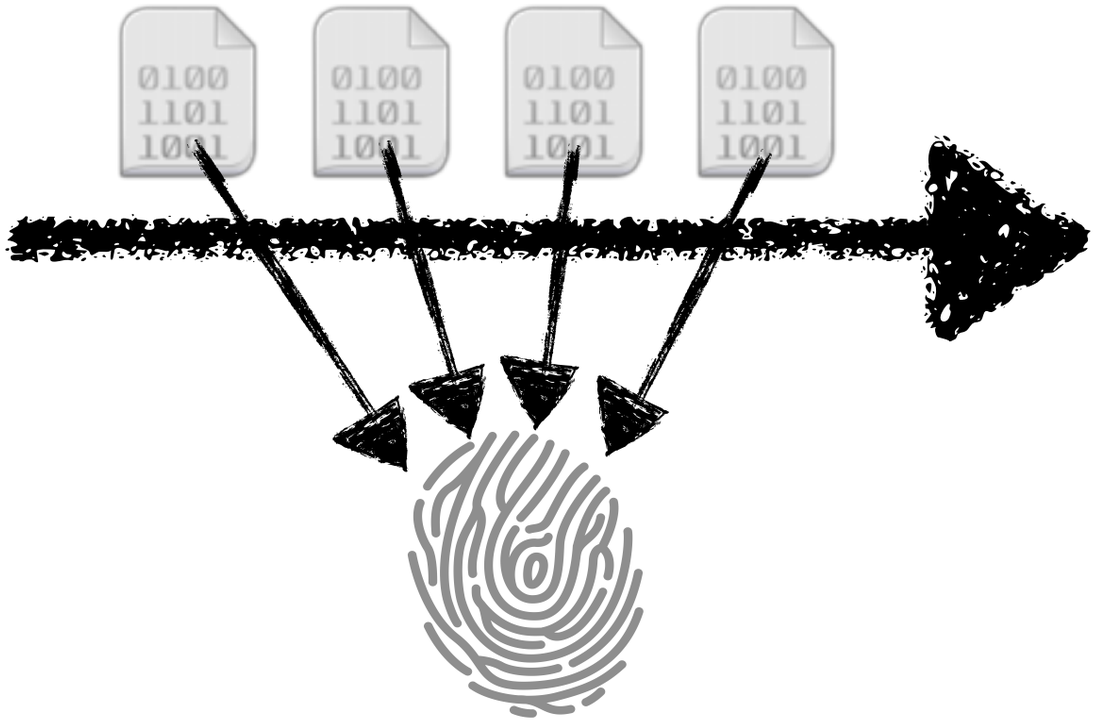

?

Overview

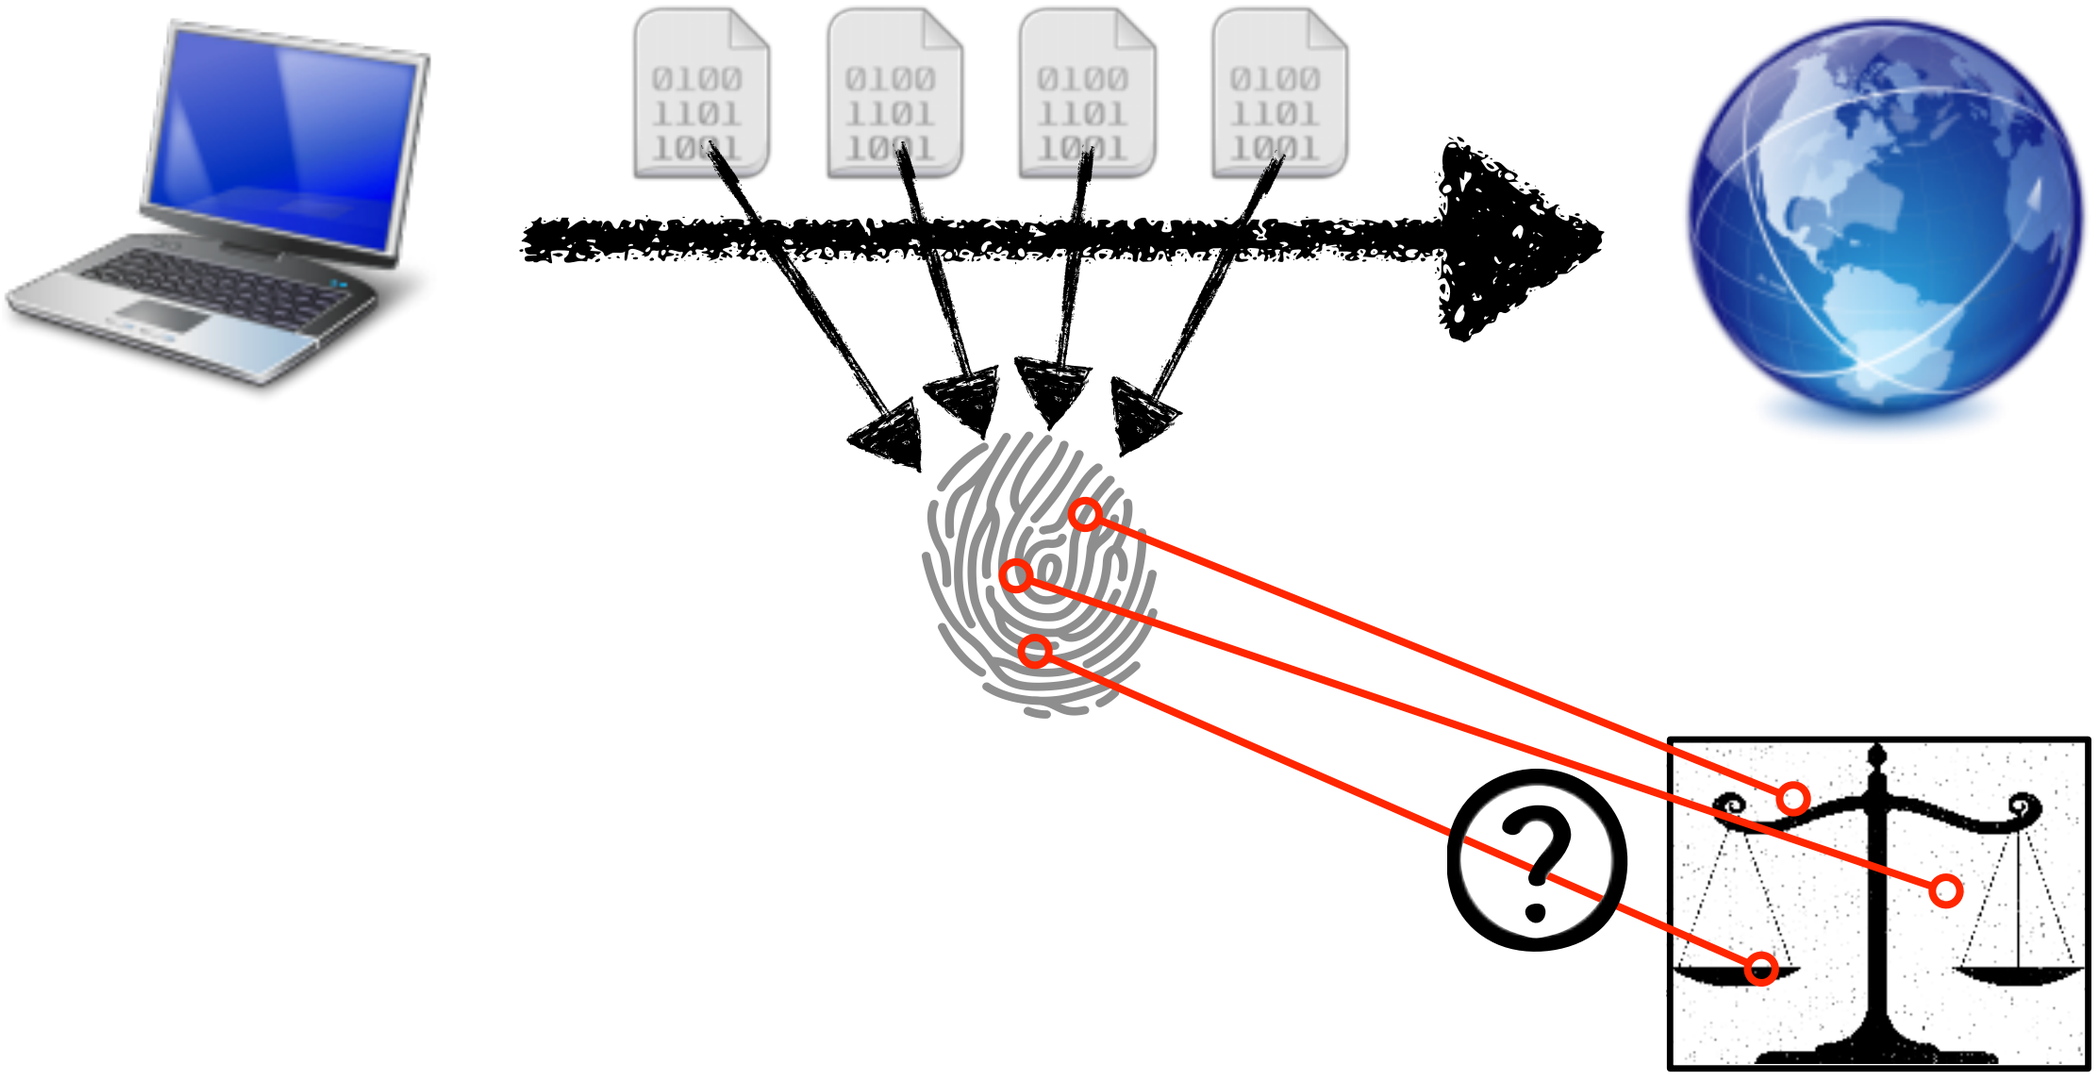

Overview

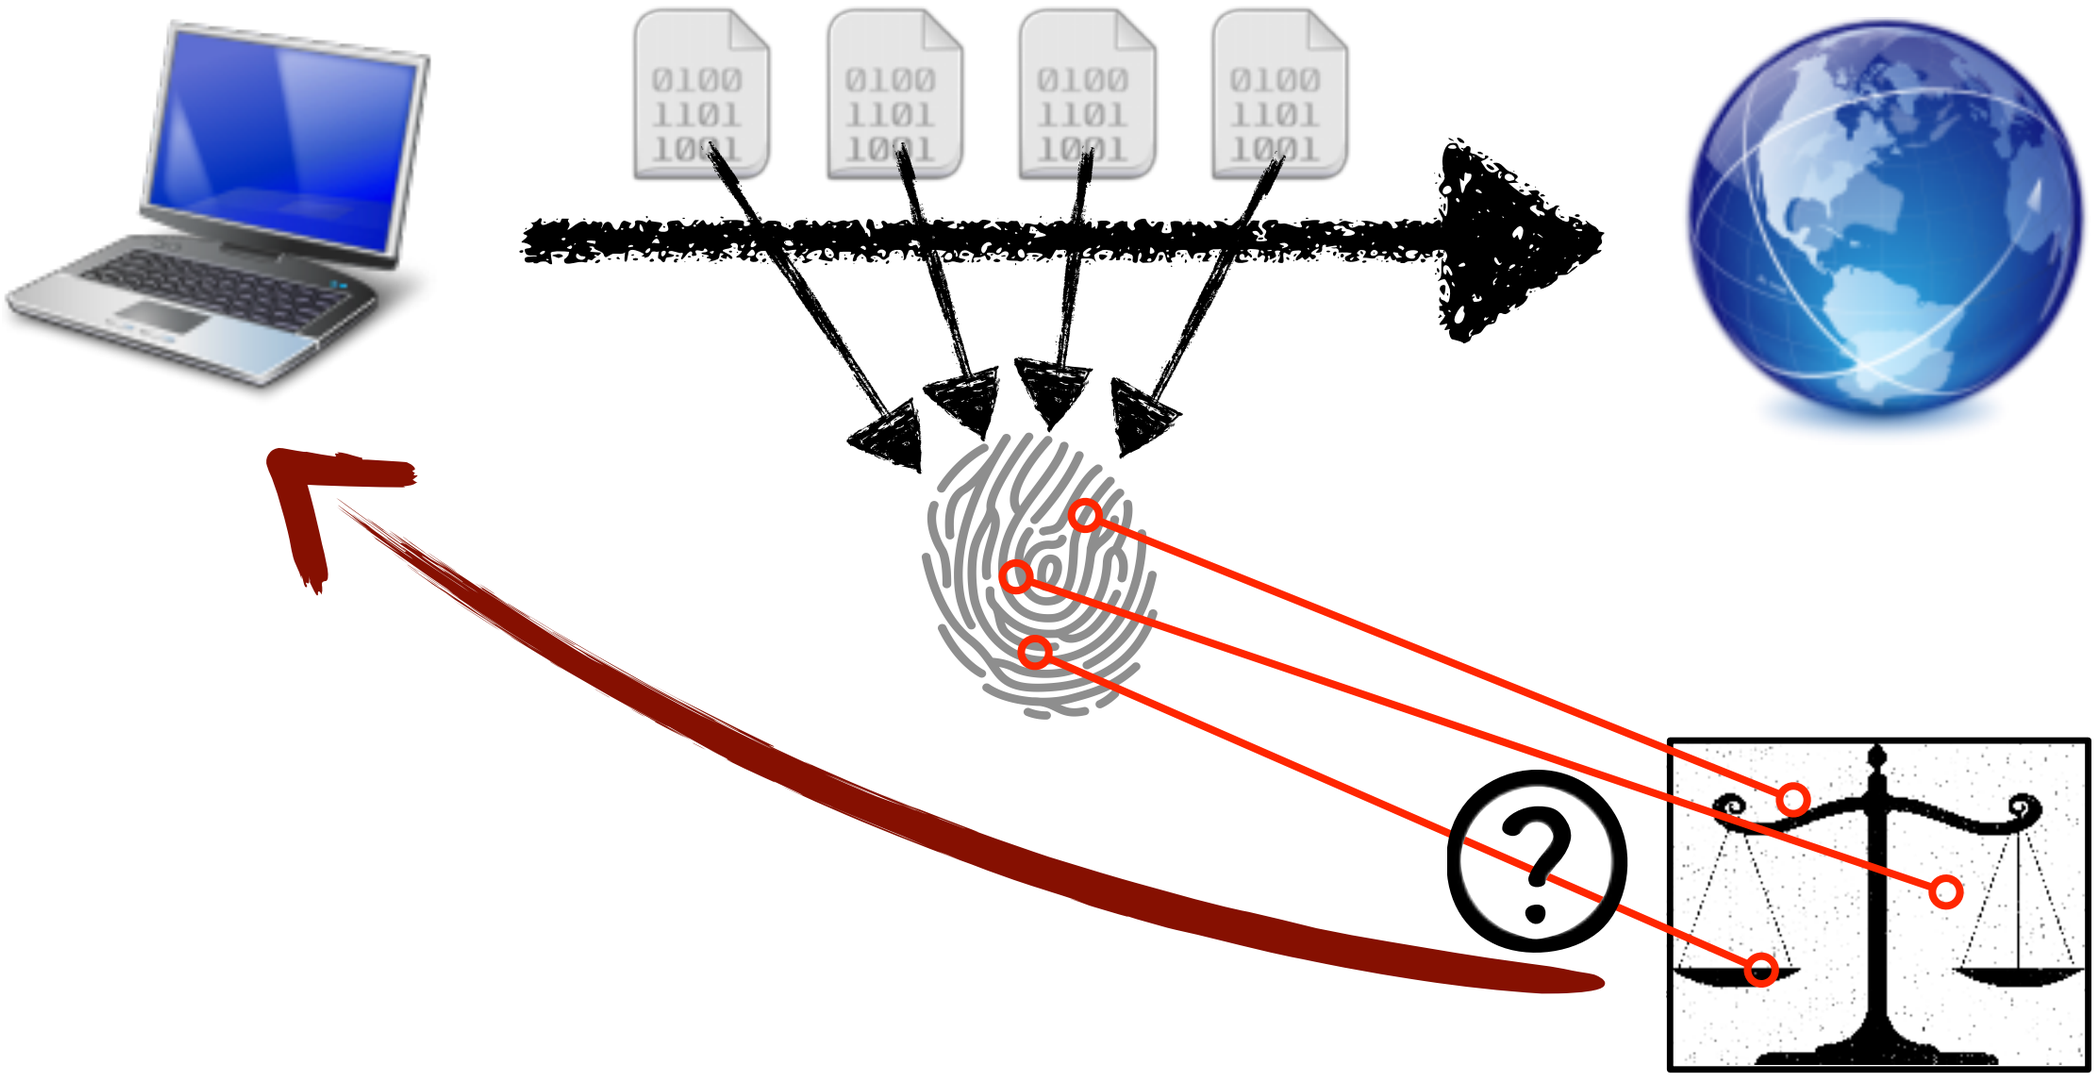

Background on DRAM

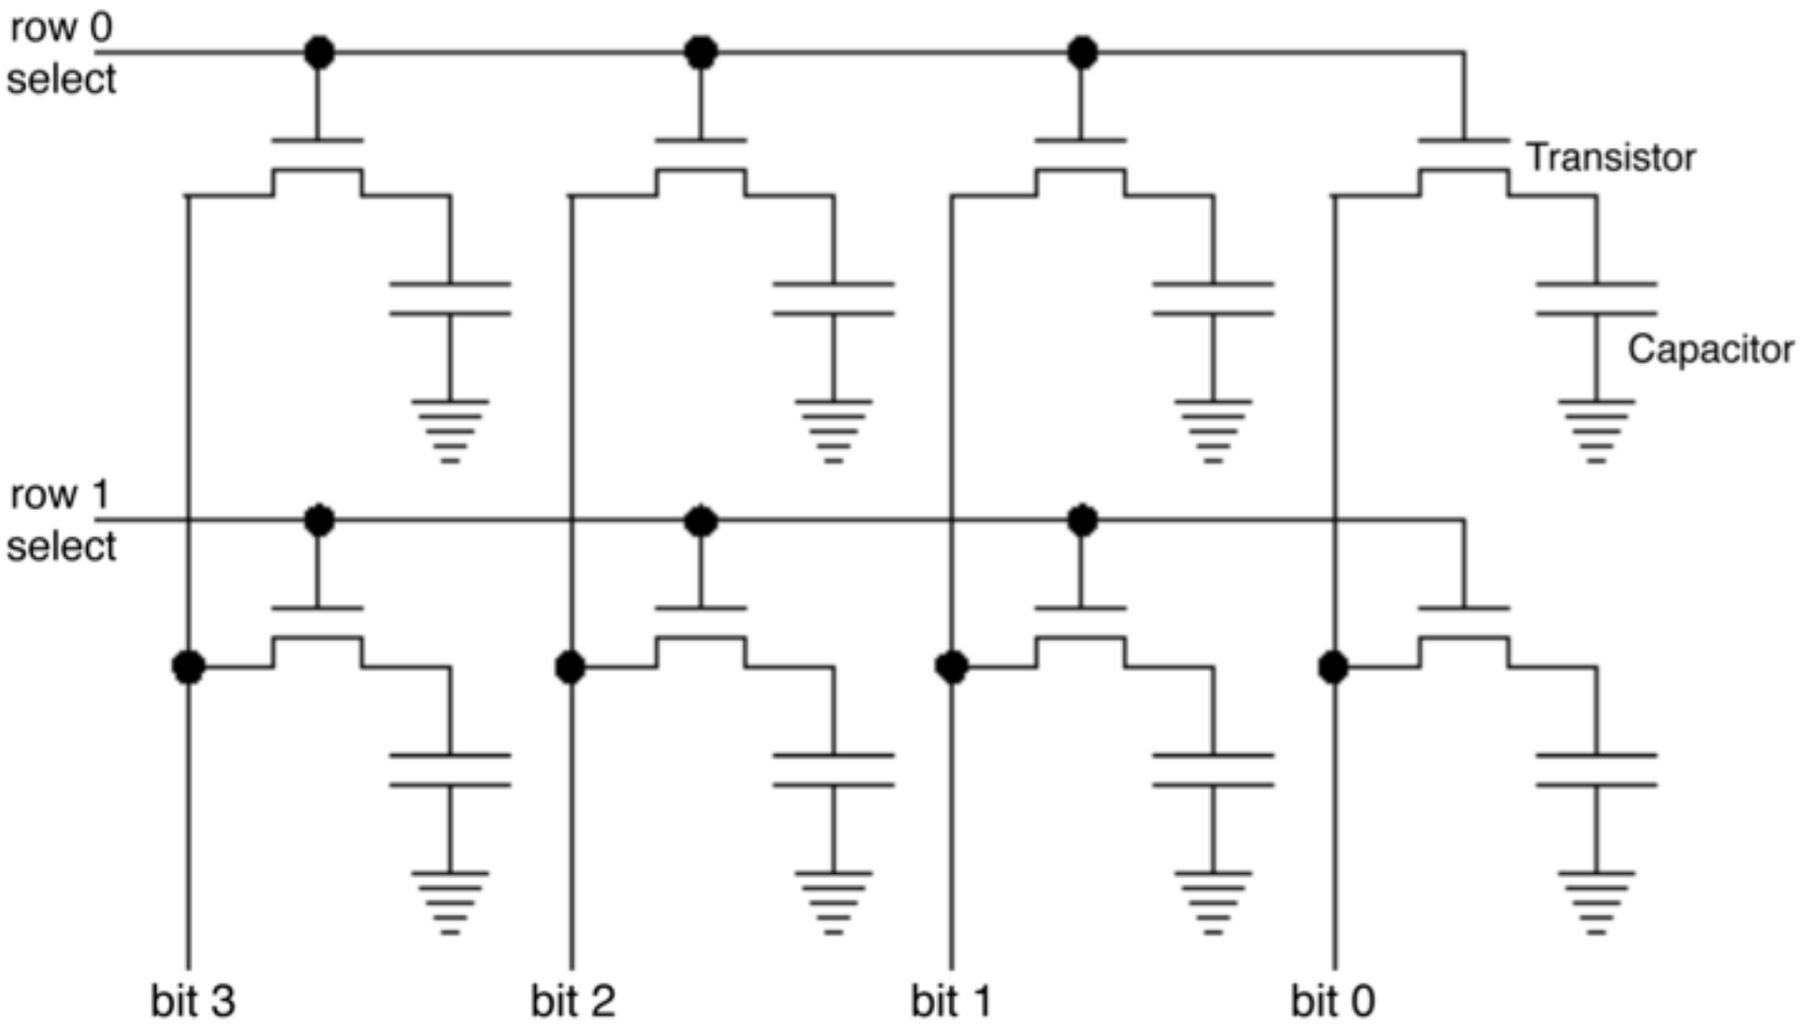

Background on DRAM

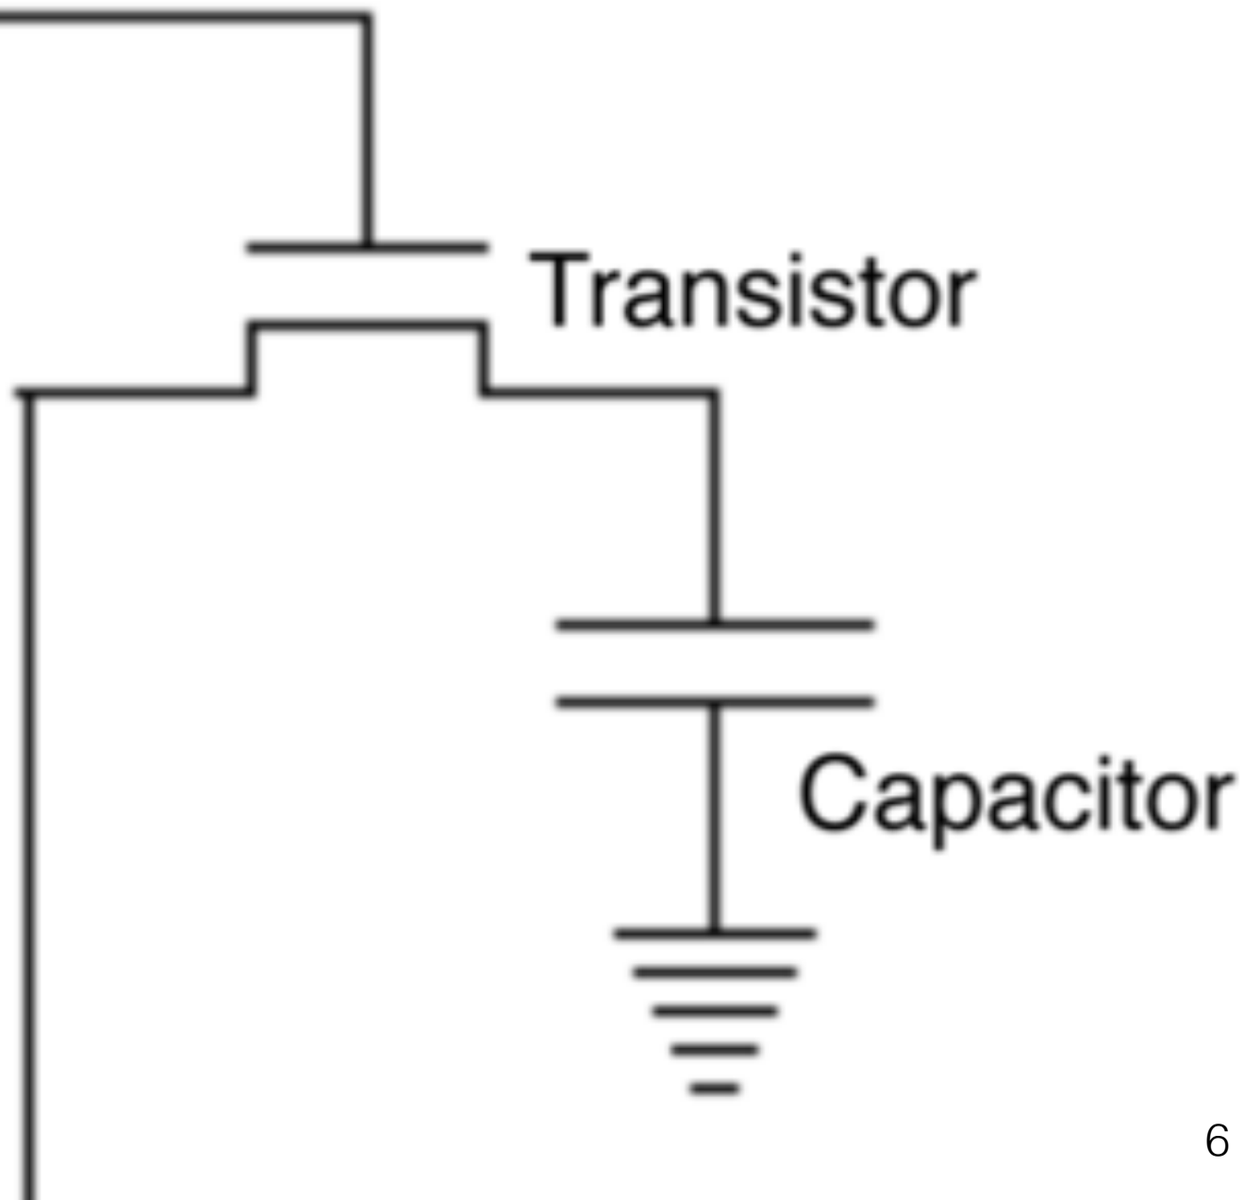

Background on DRAM

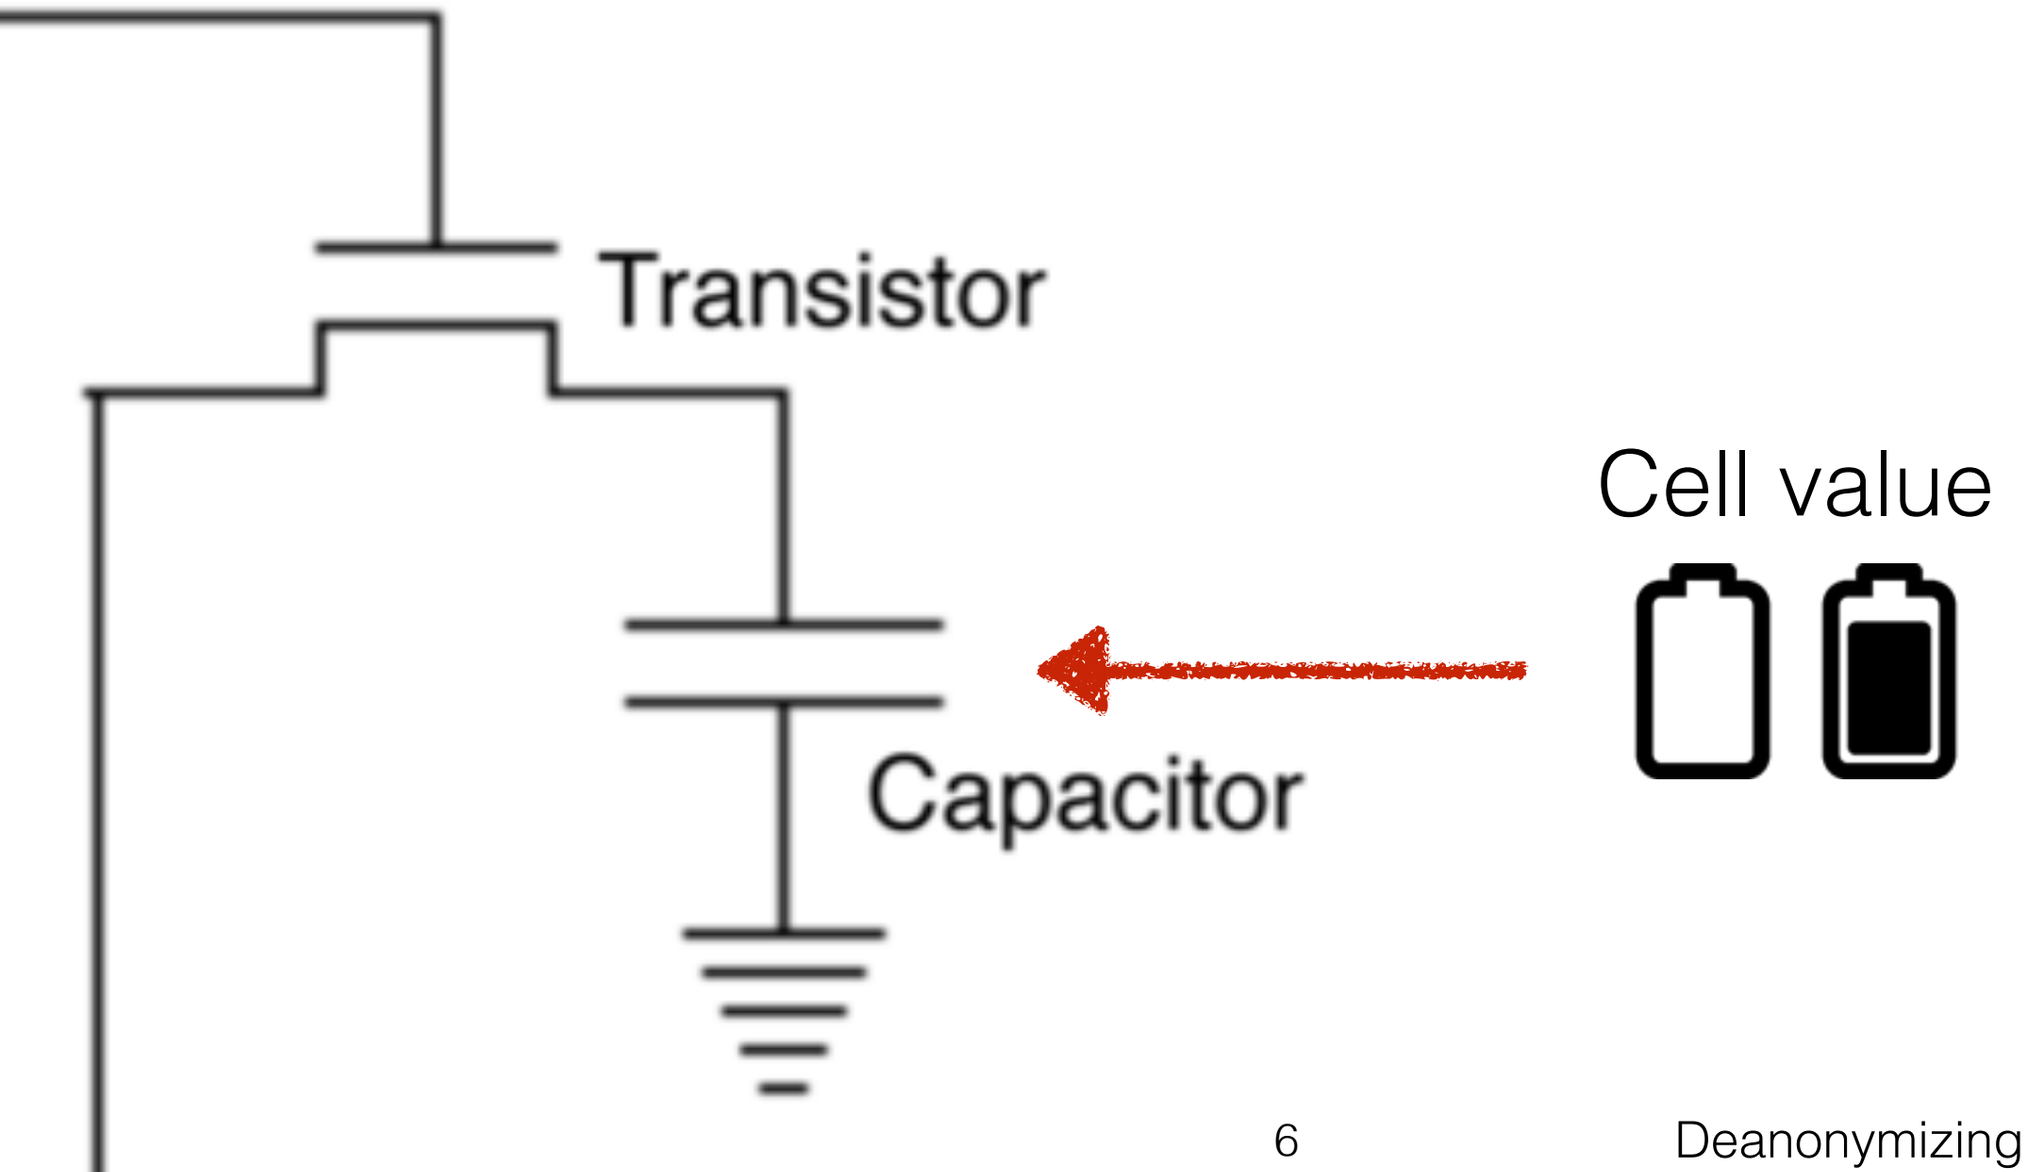

Background on DRAM

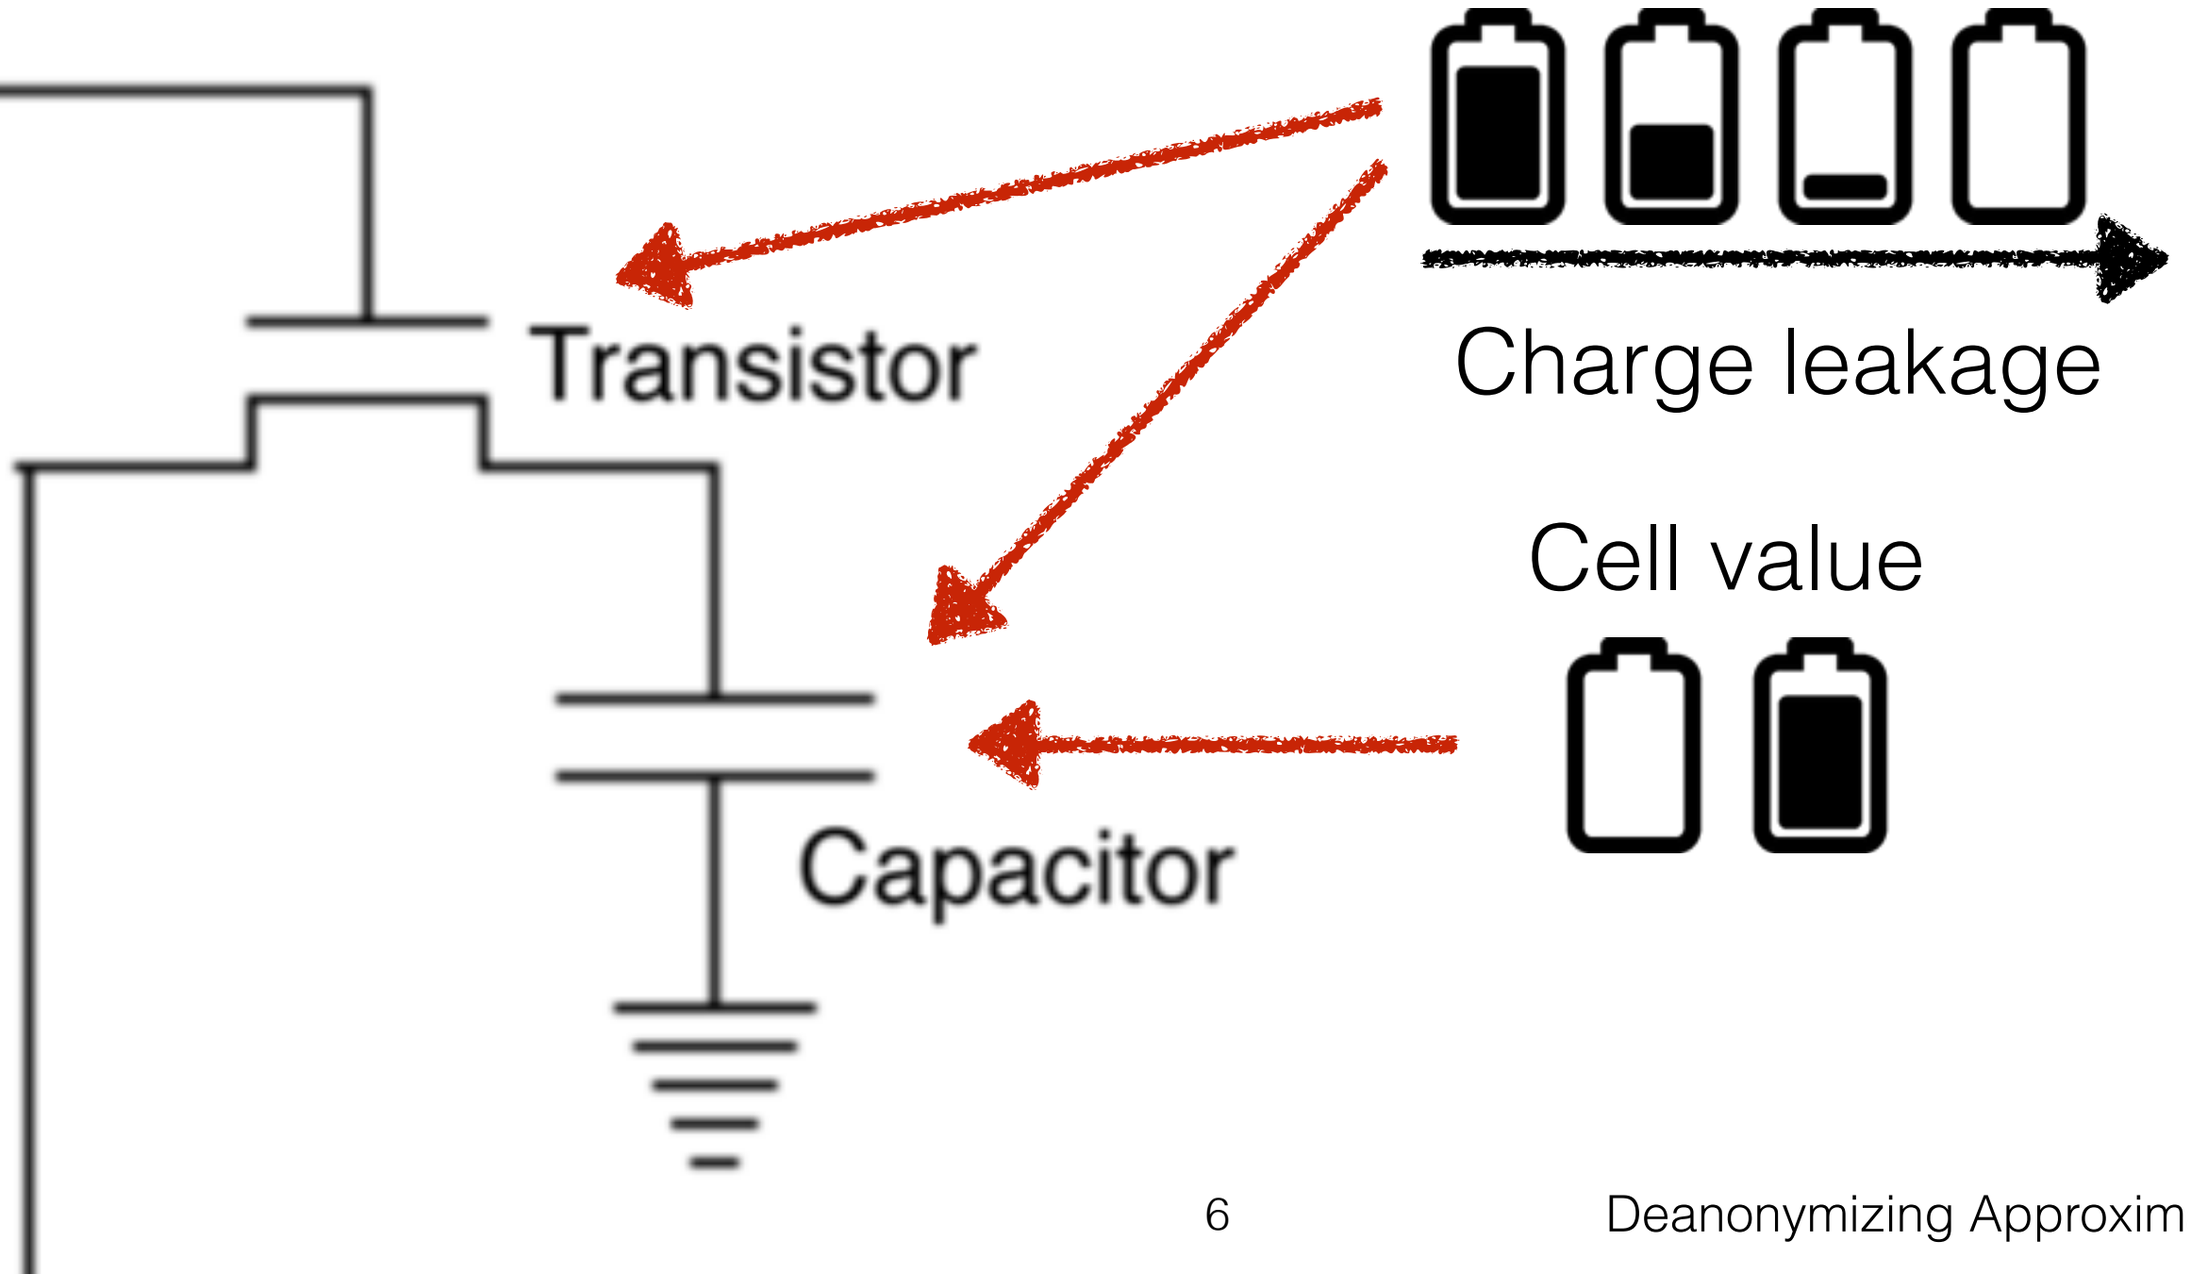

Background on DRAM

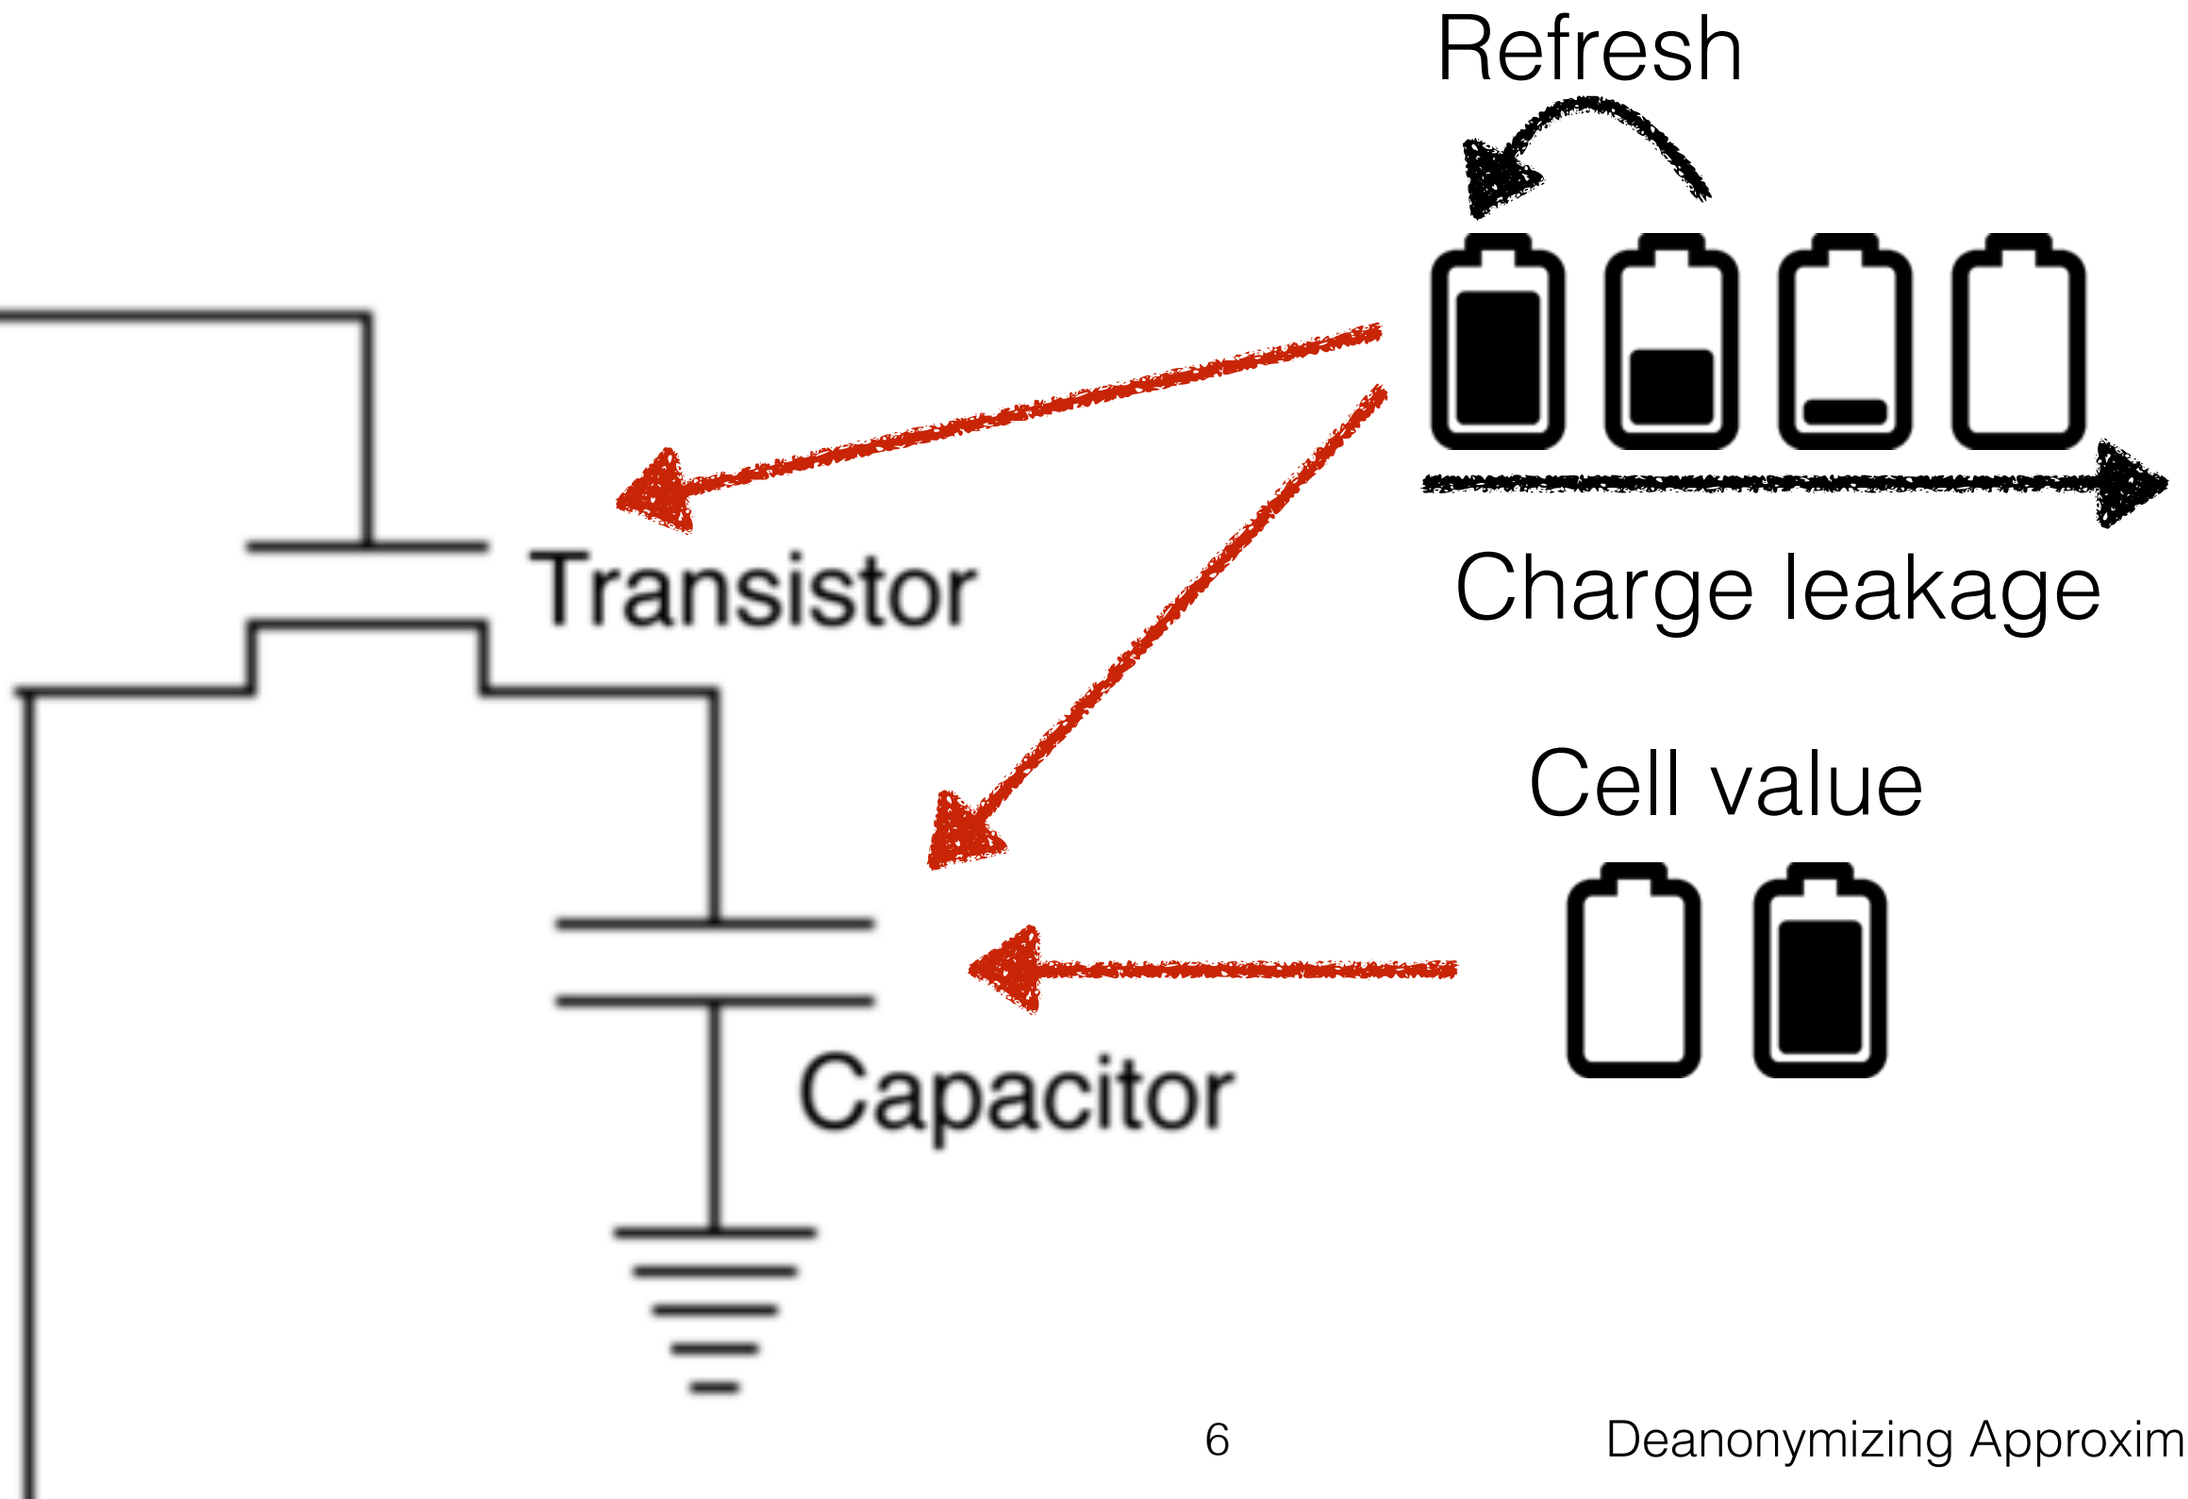

Background on DRAM

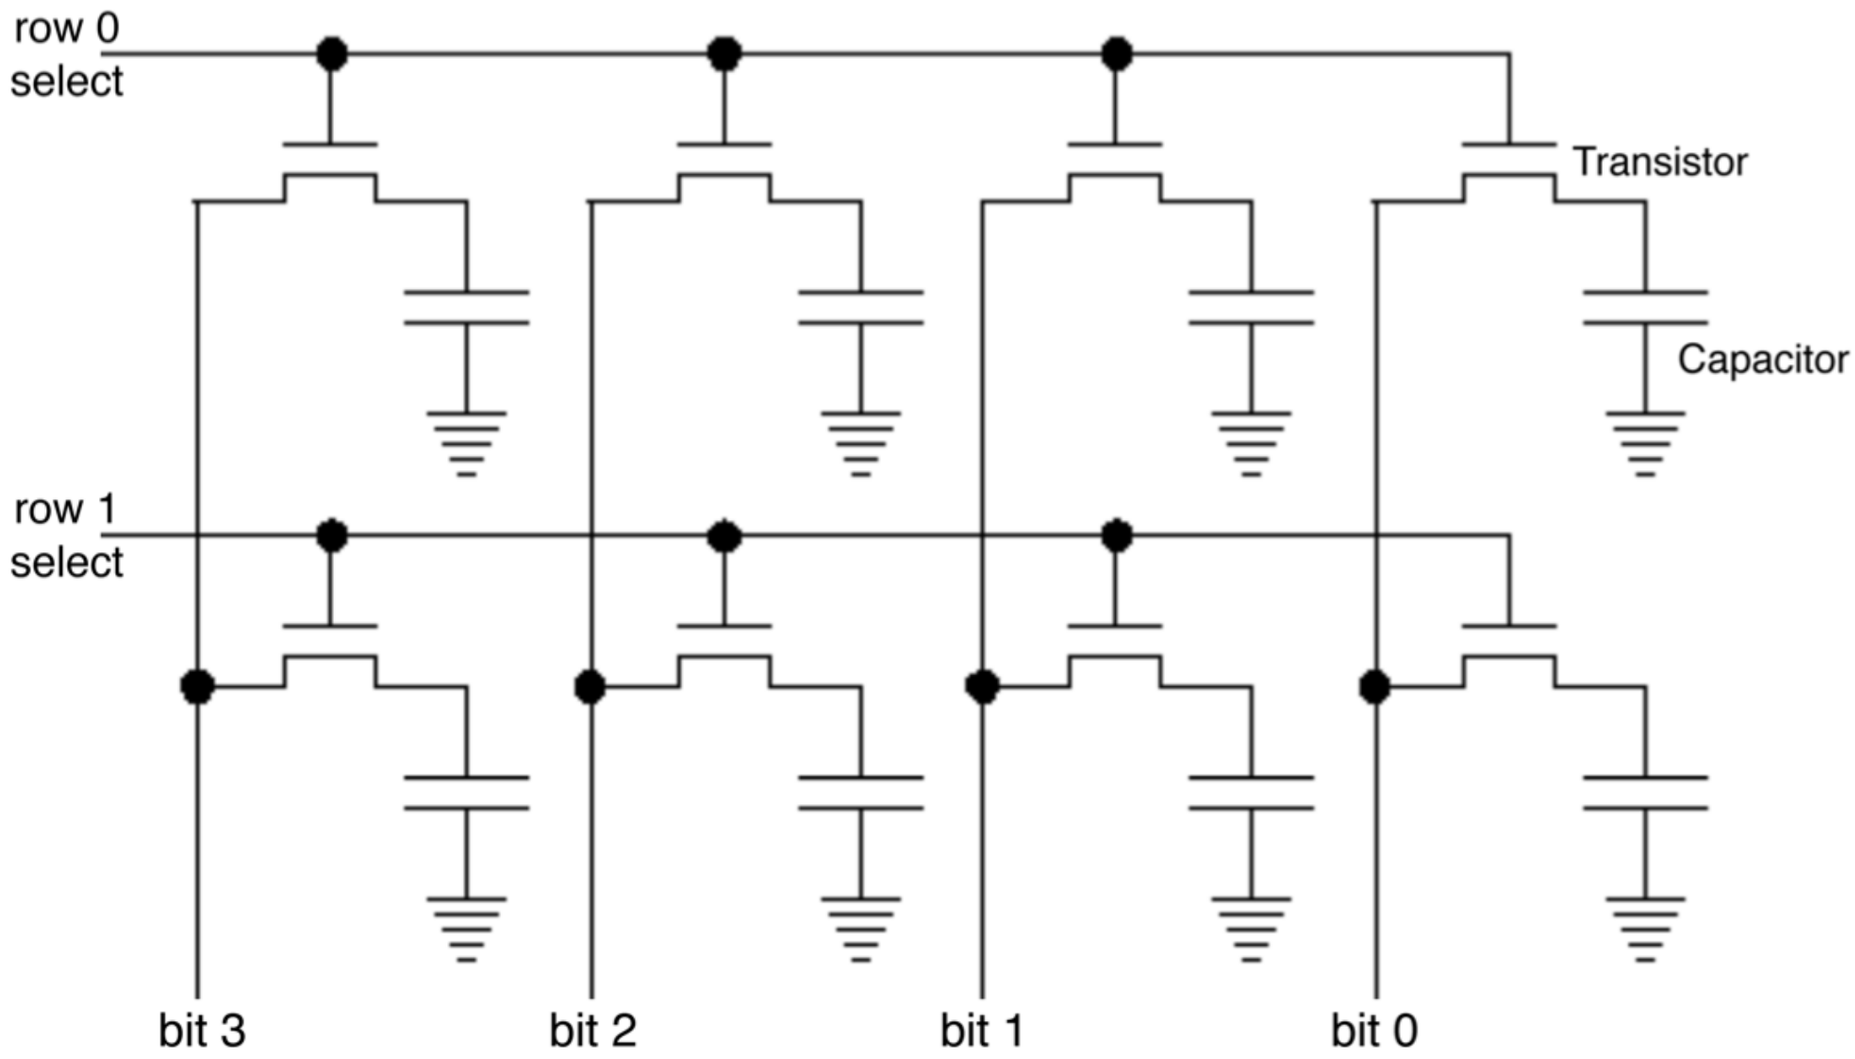

Background on DRAM

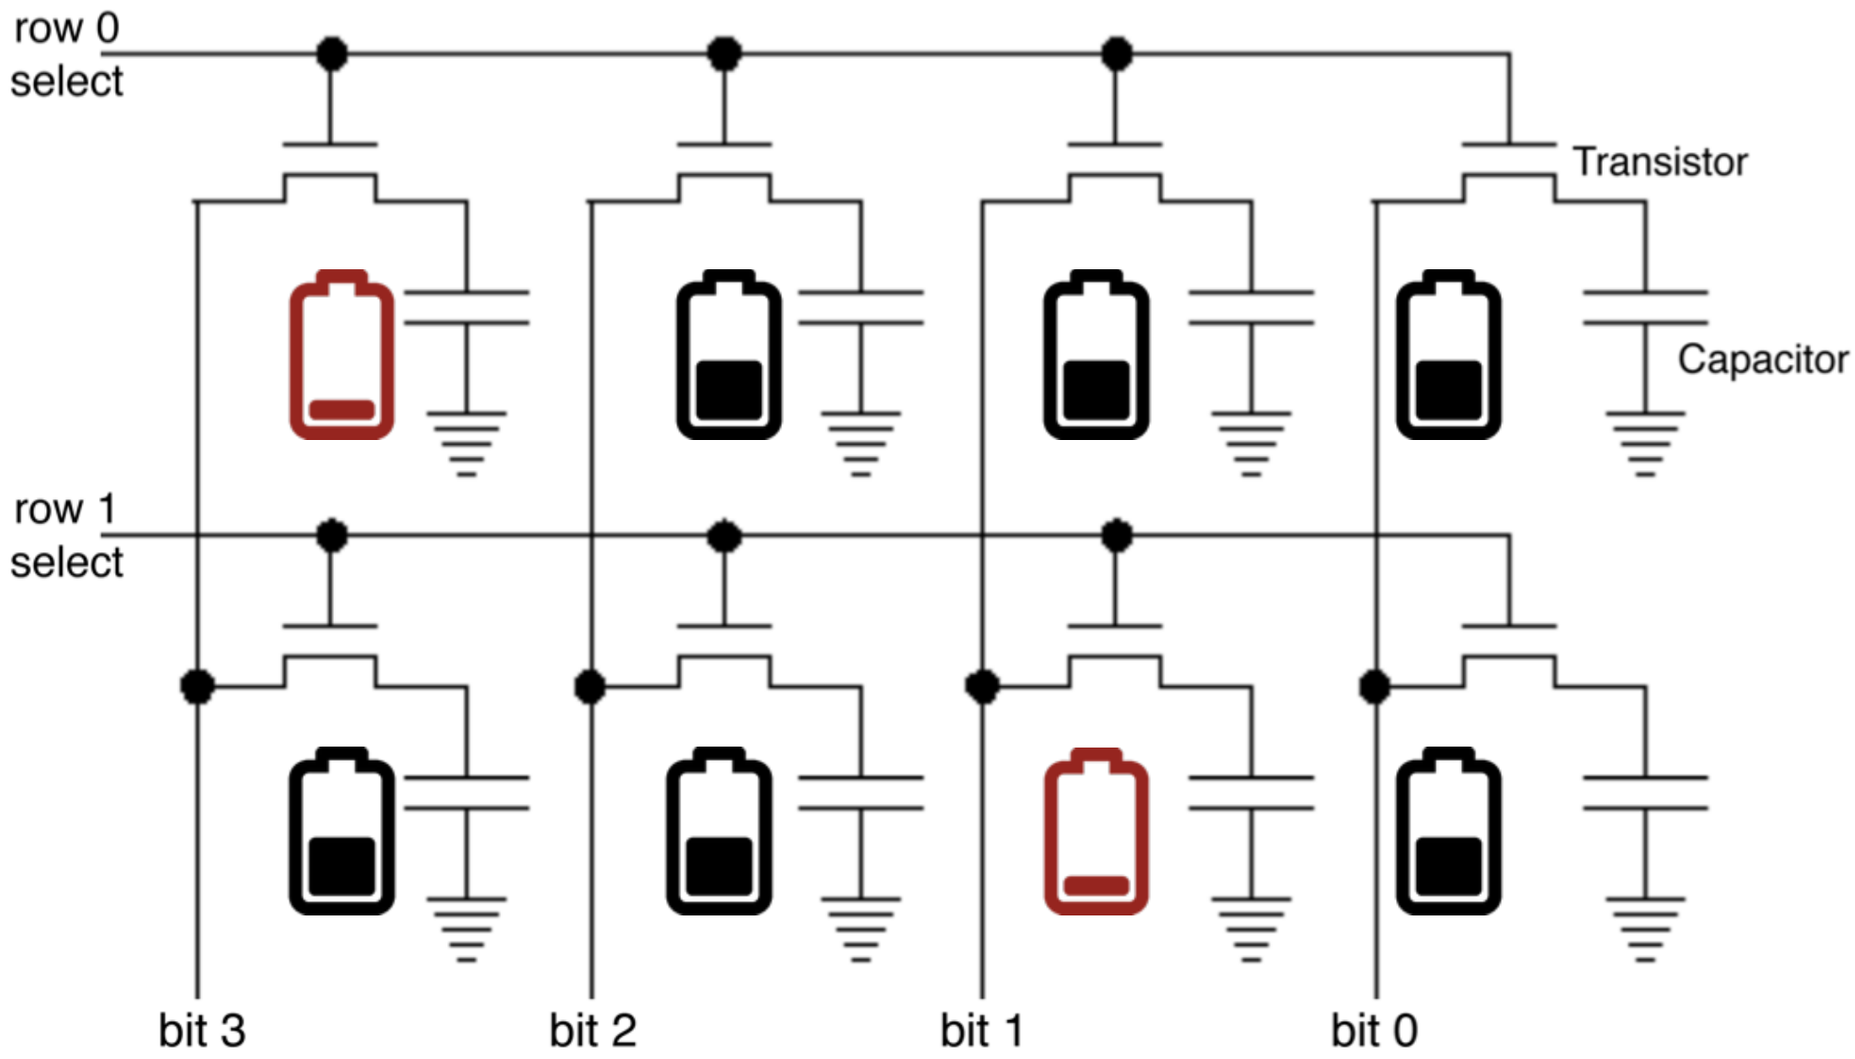

Toy Example

Device A

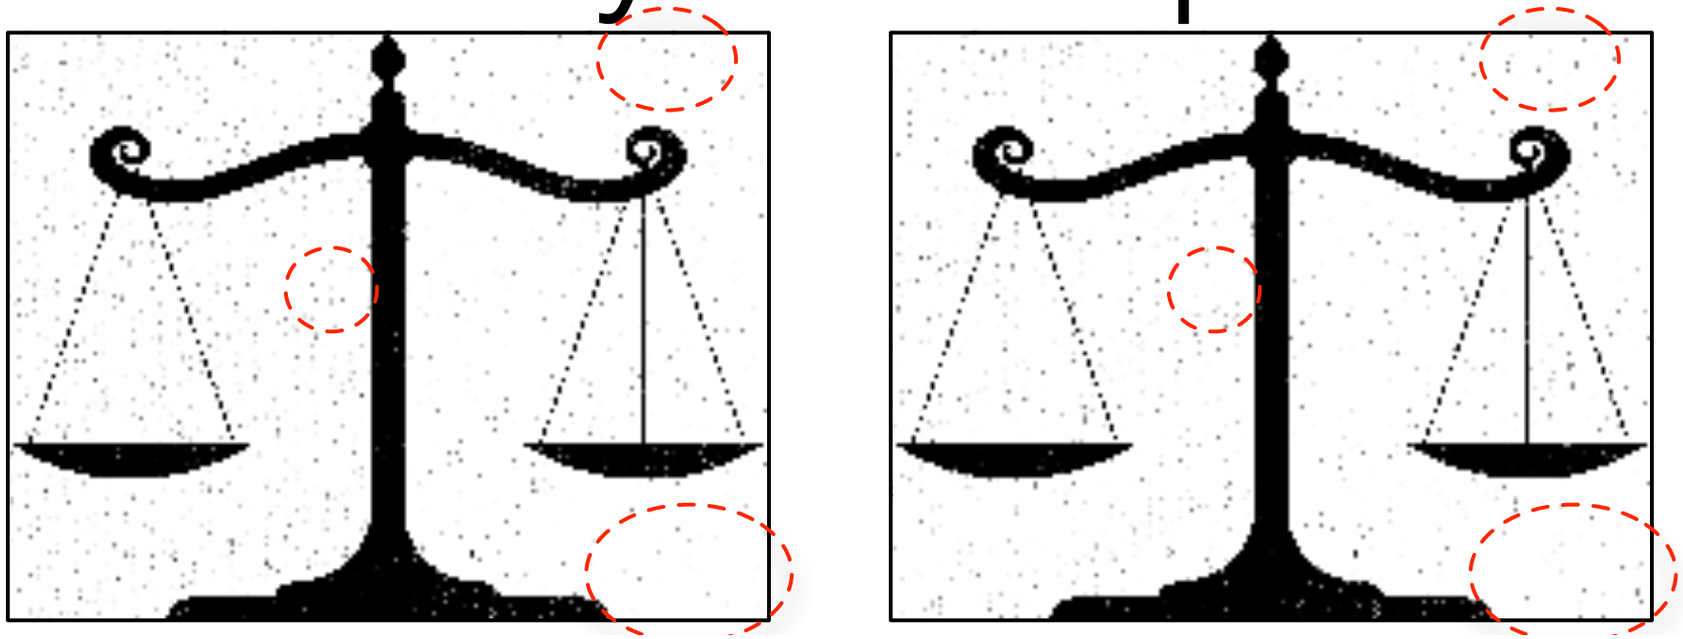

Device B

Toy Example

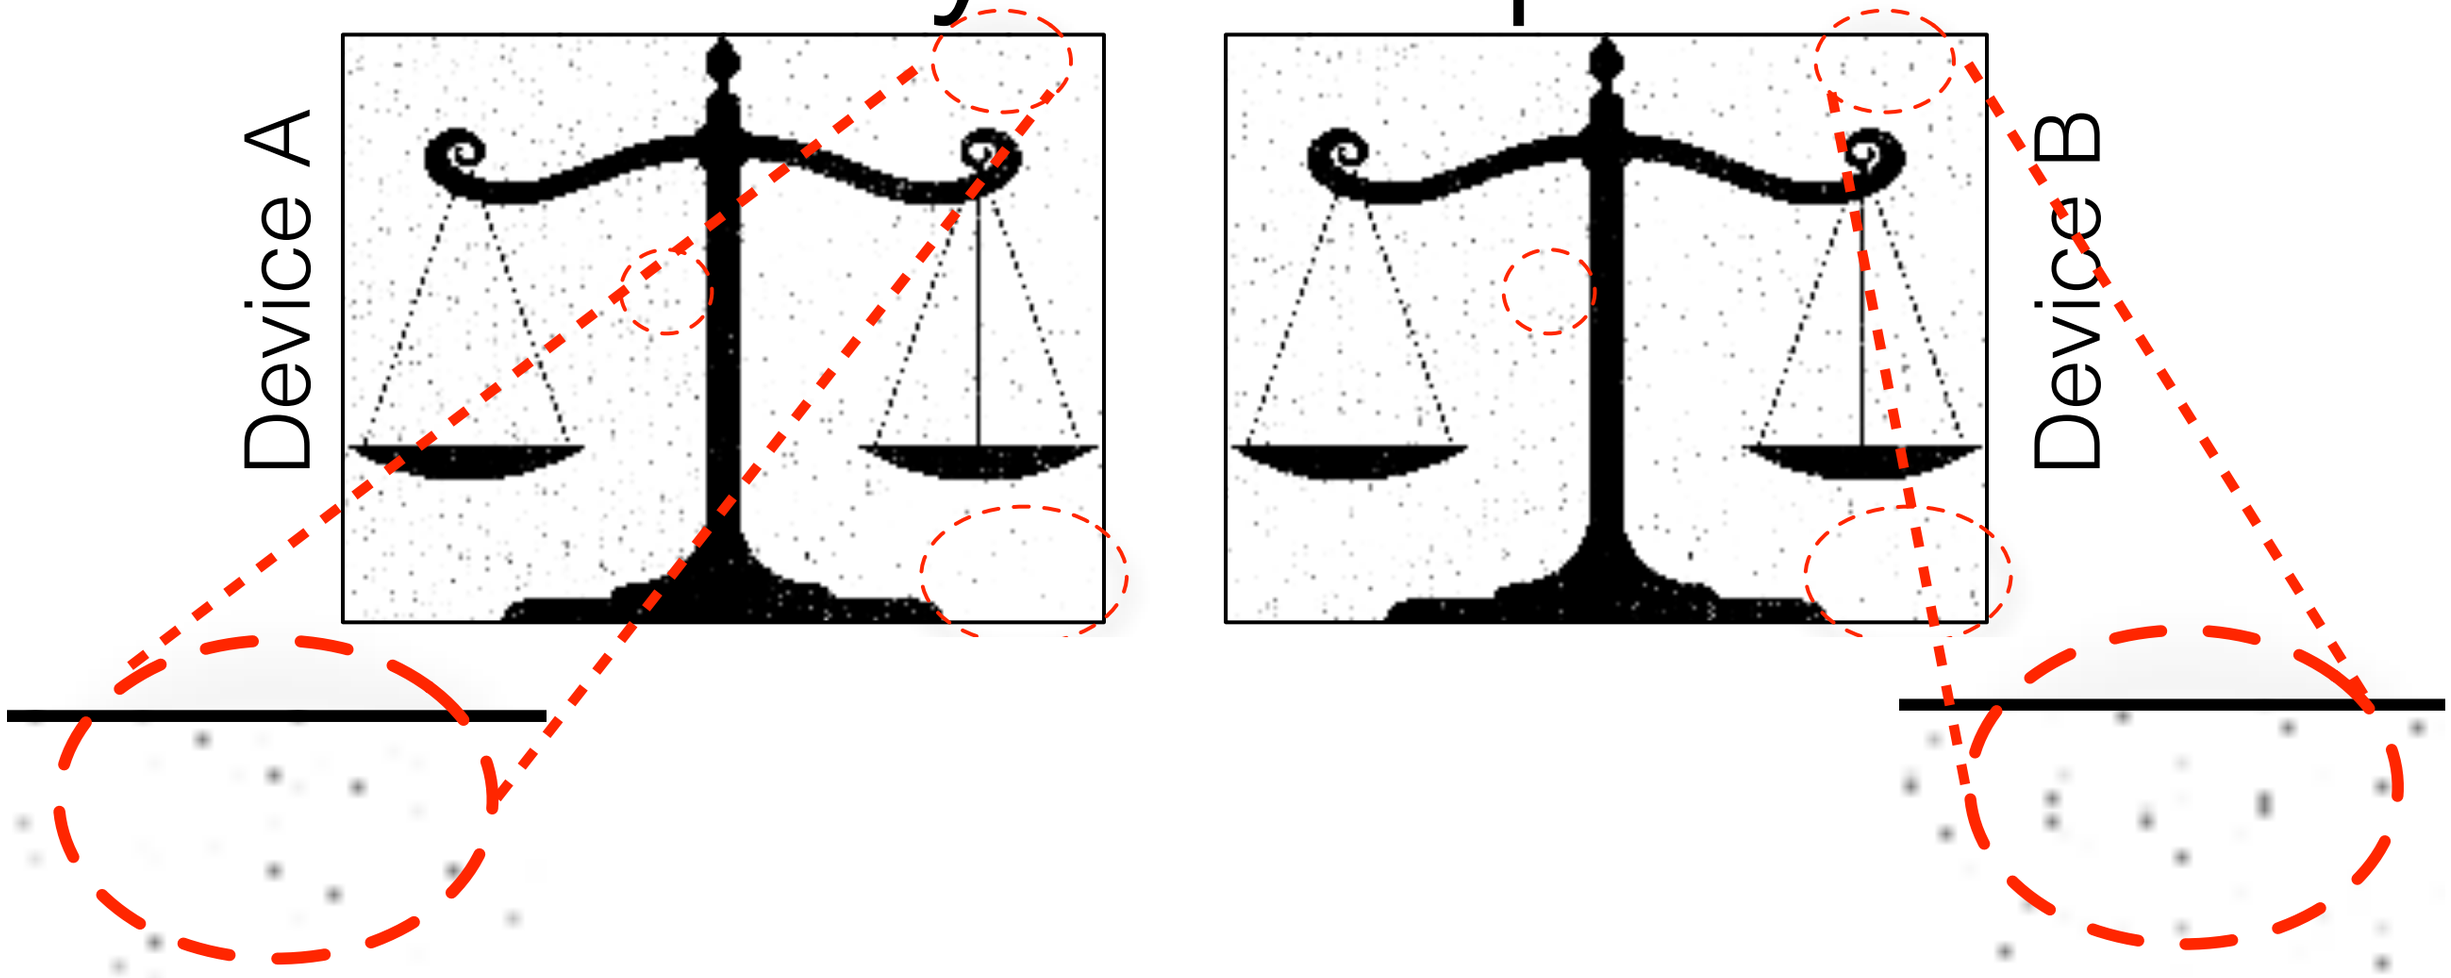

Toy Example

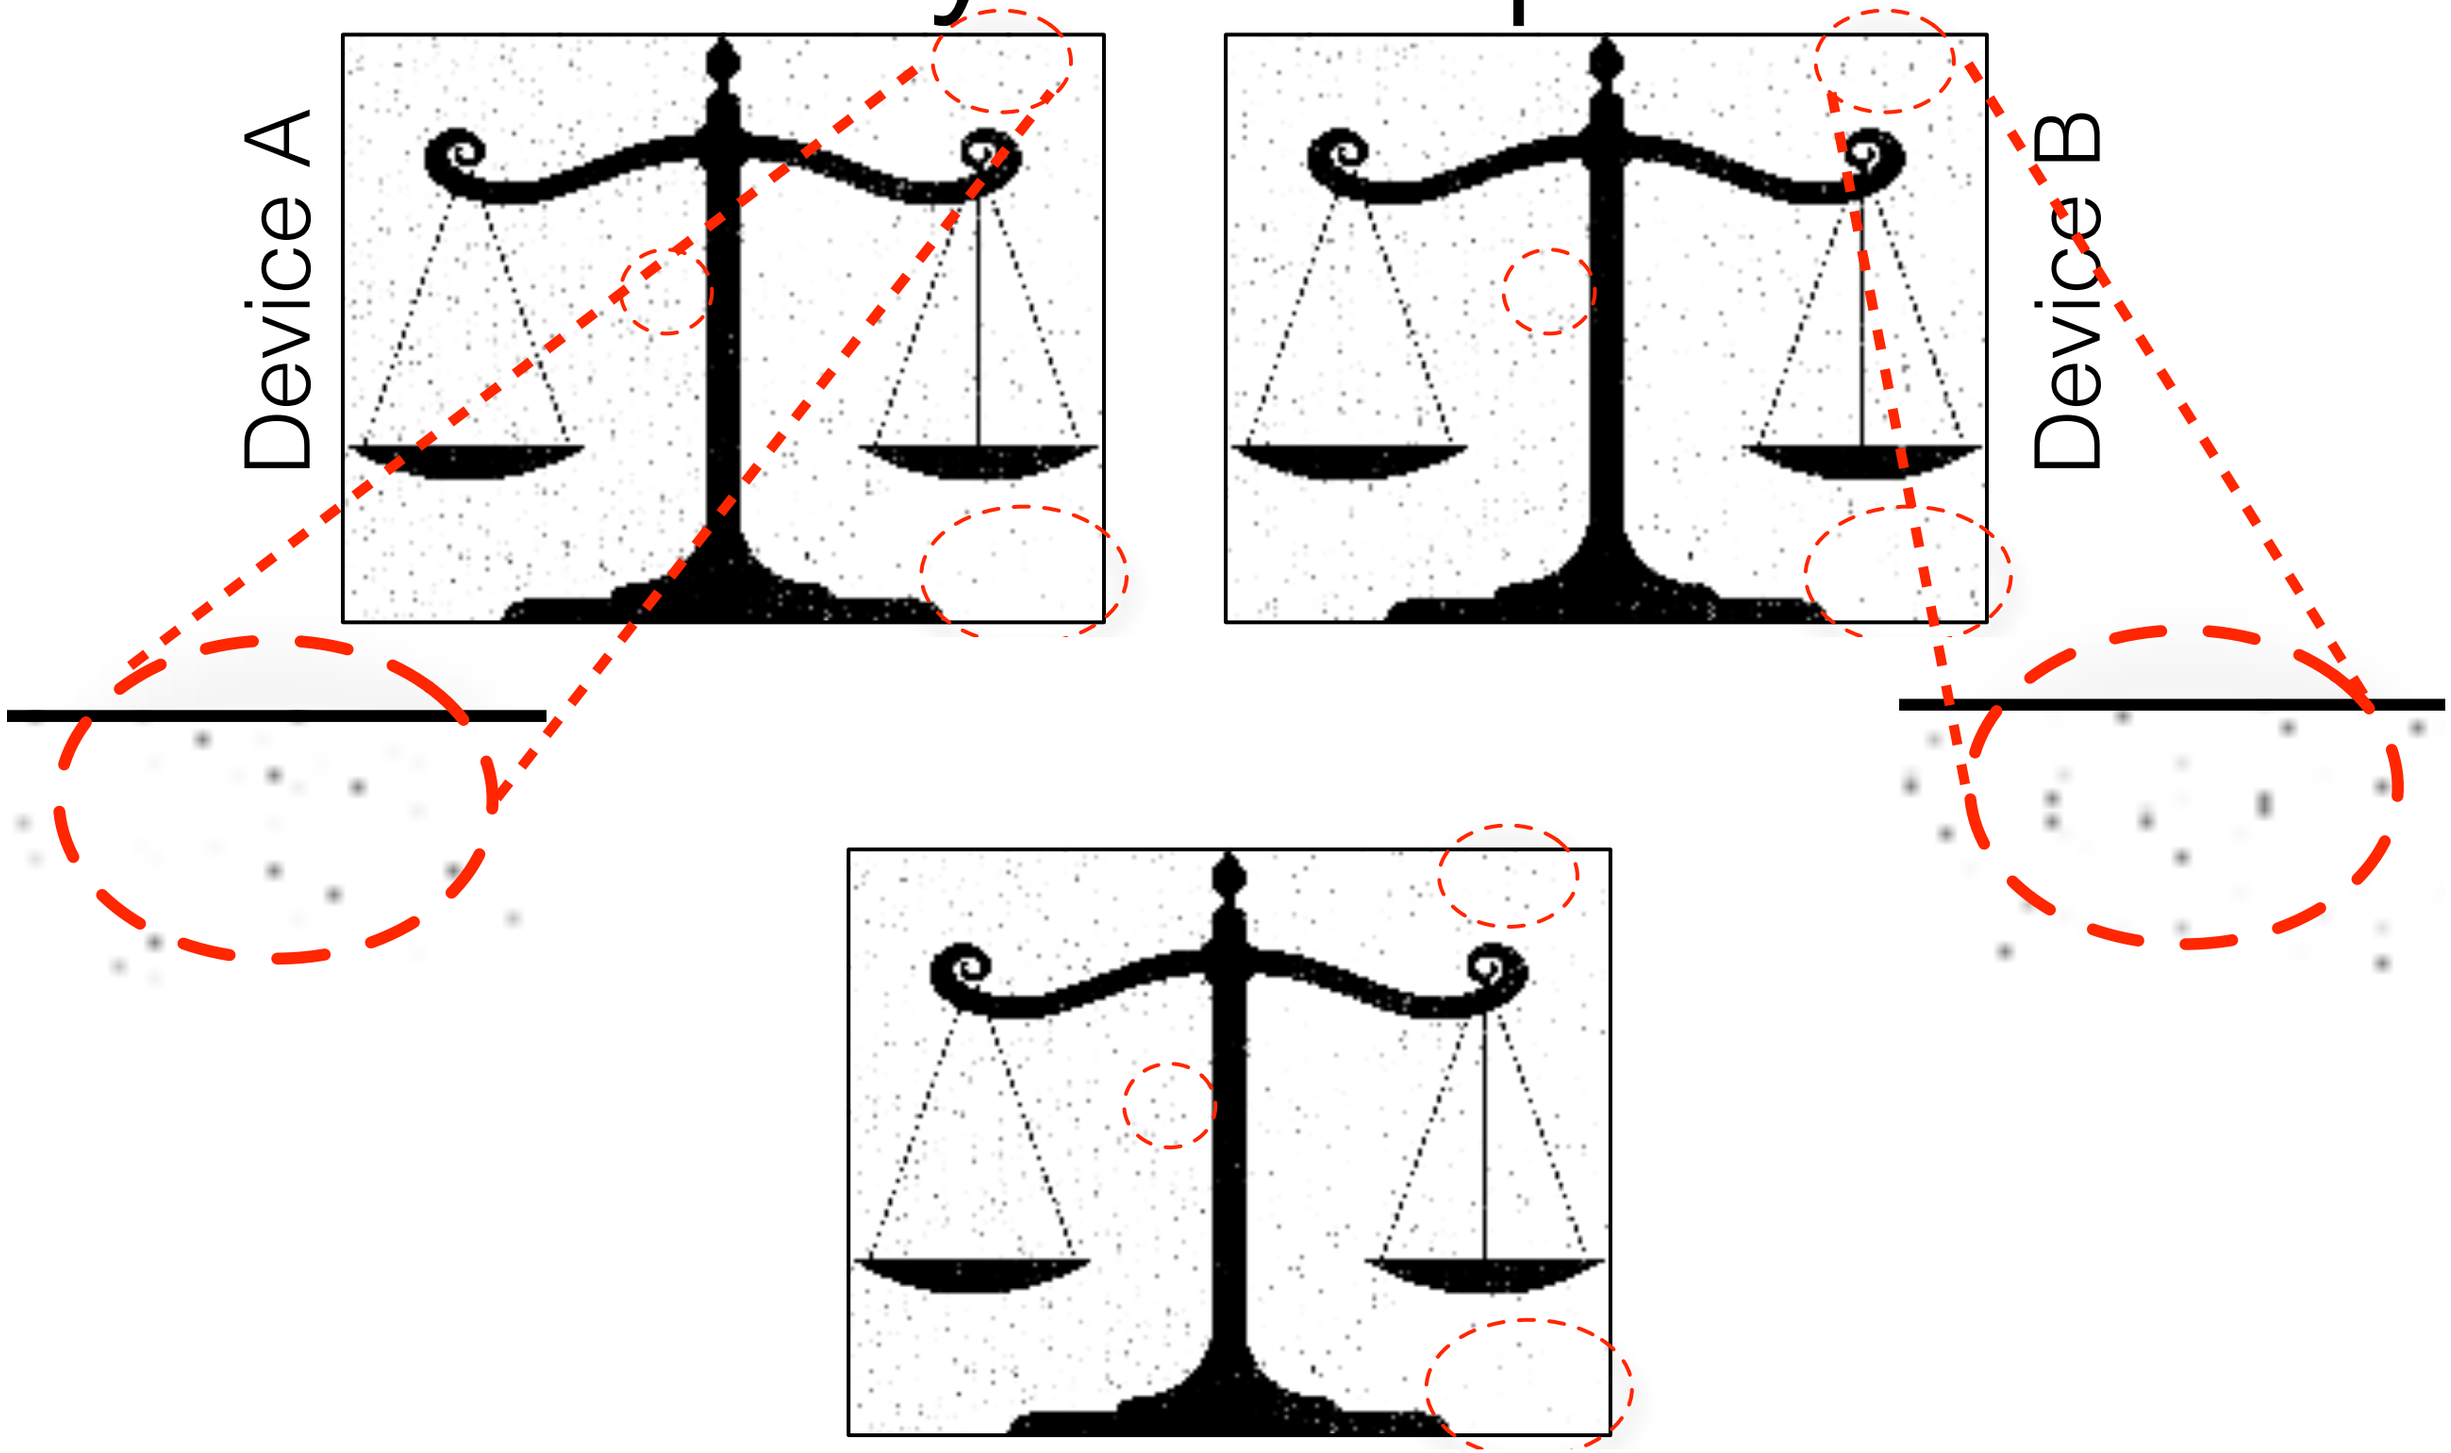

Toy Example

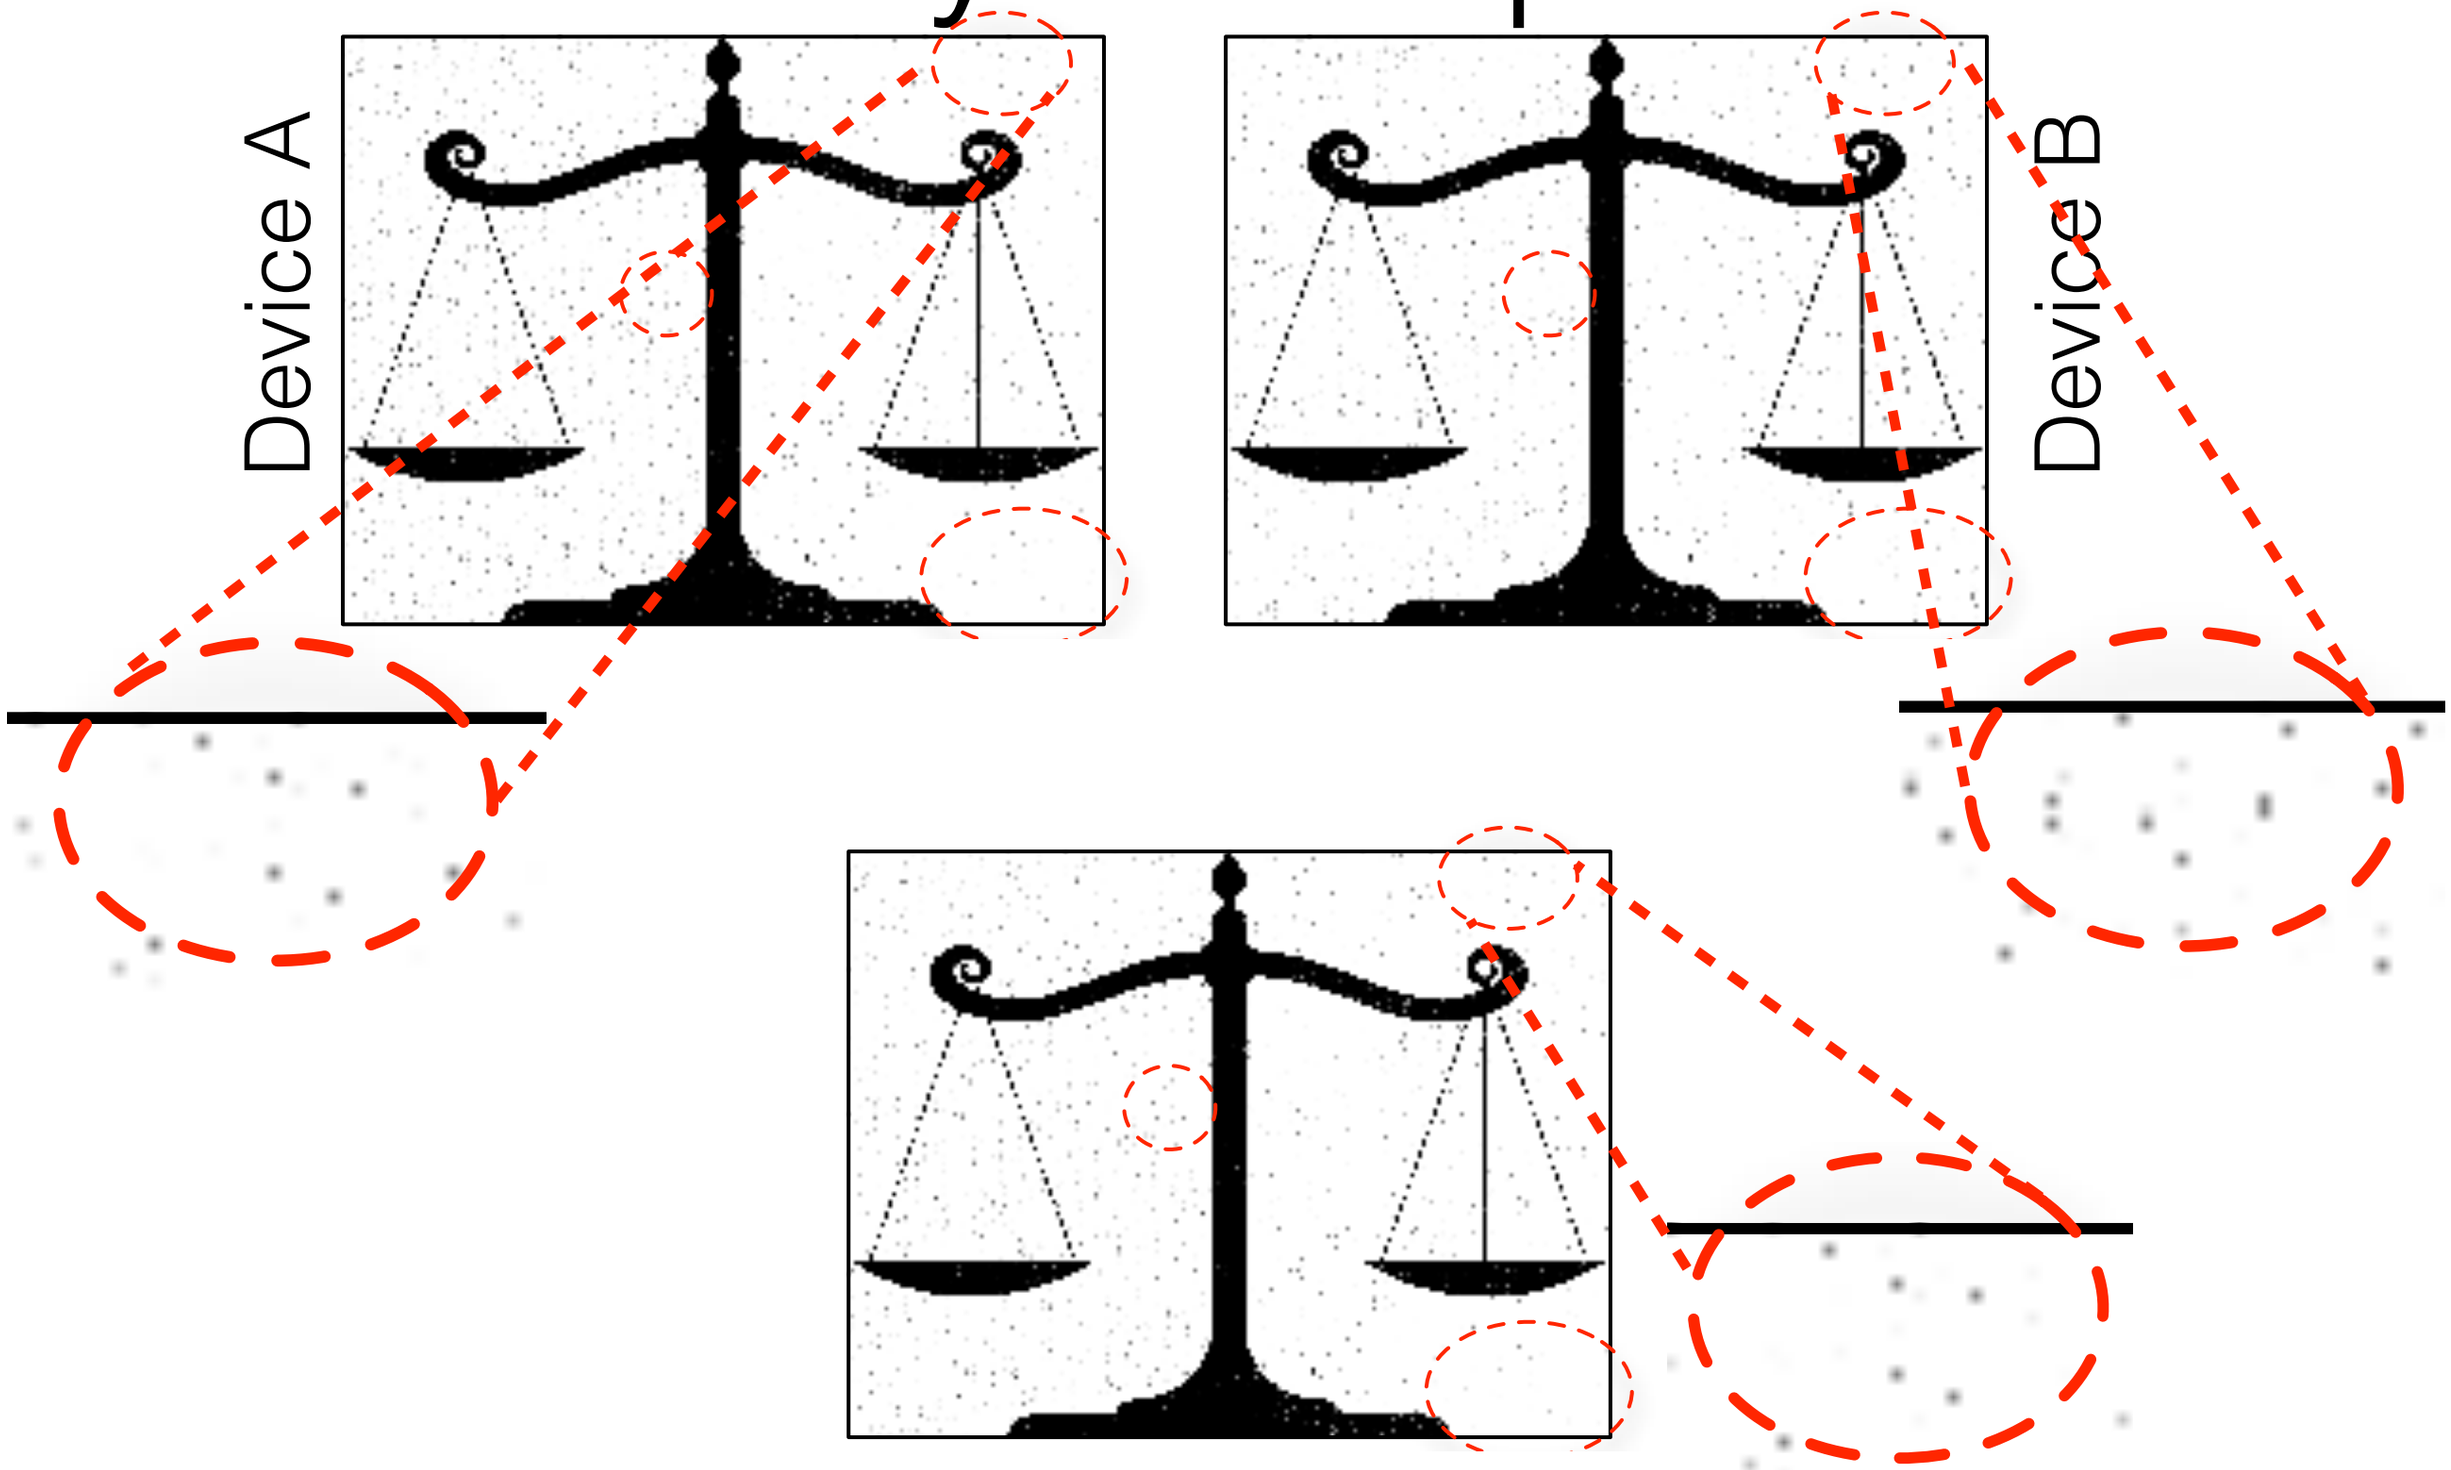

Toy Example

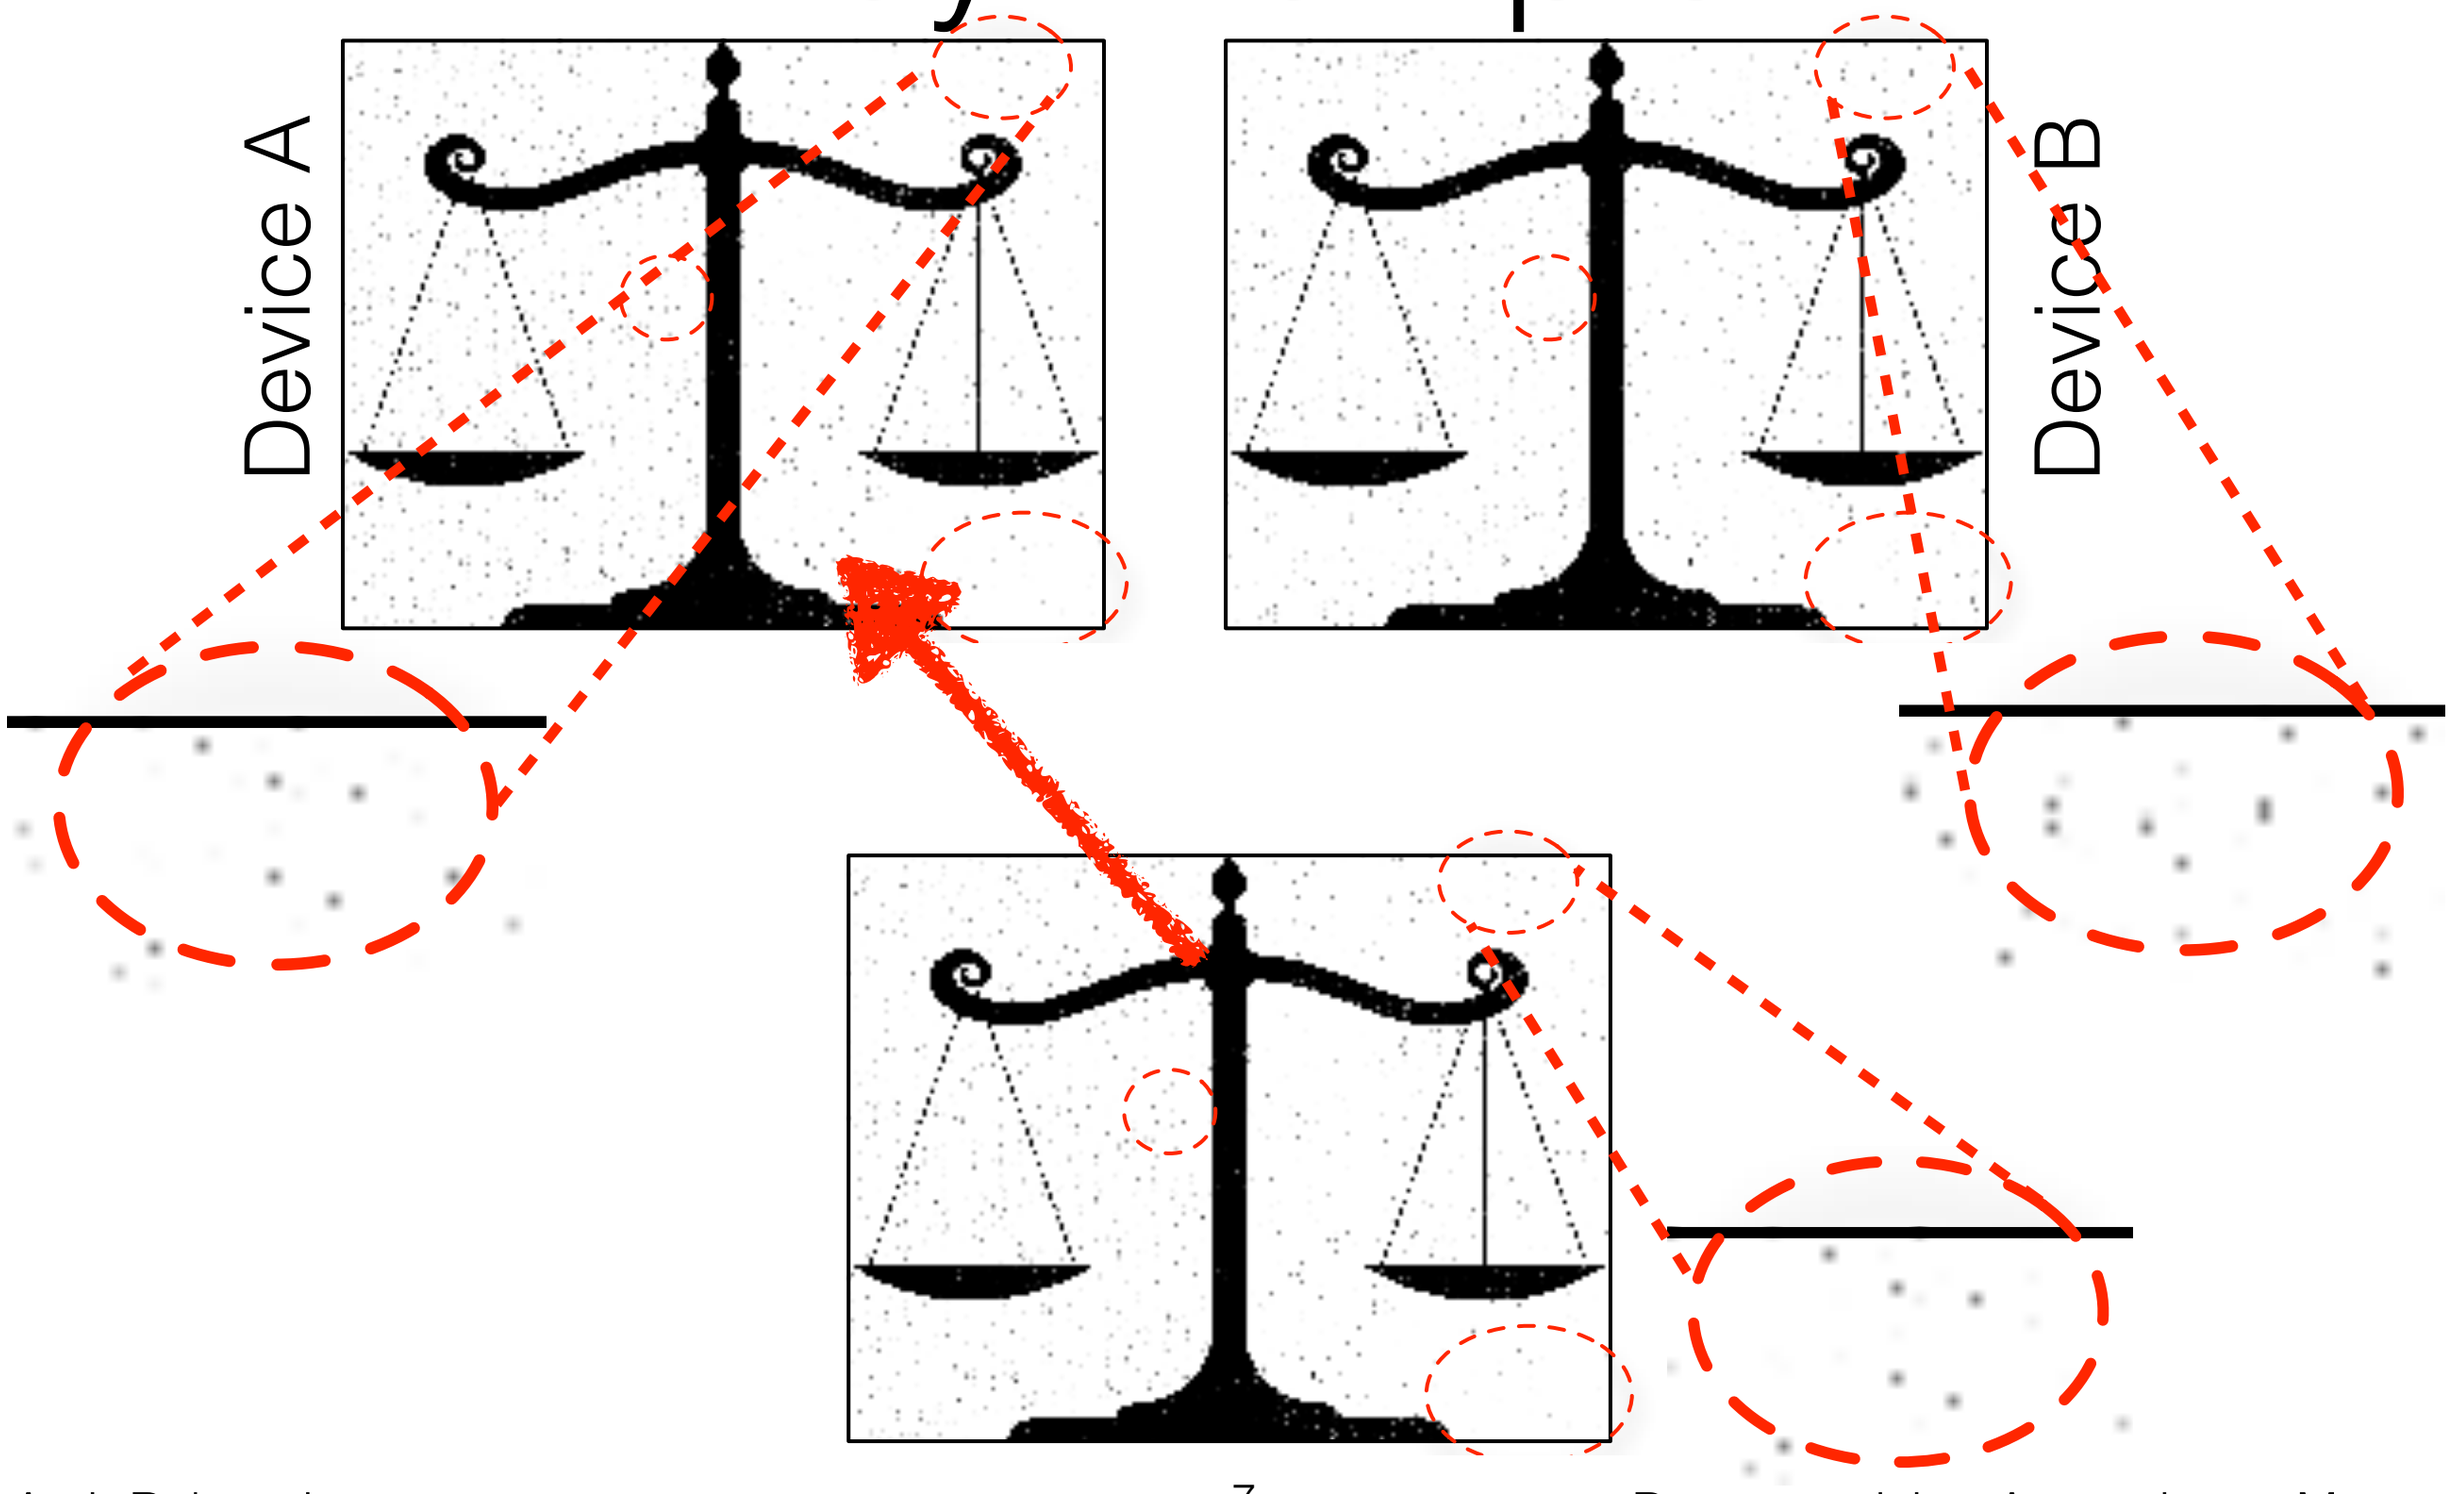

Distance Metric

Distance Metric

Hamming
Distance

Distance Metric

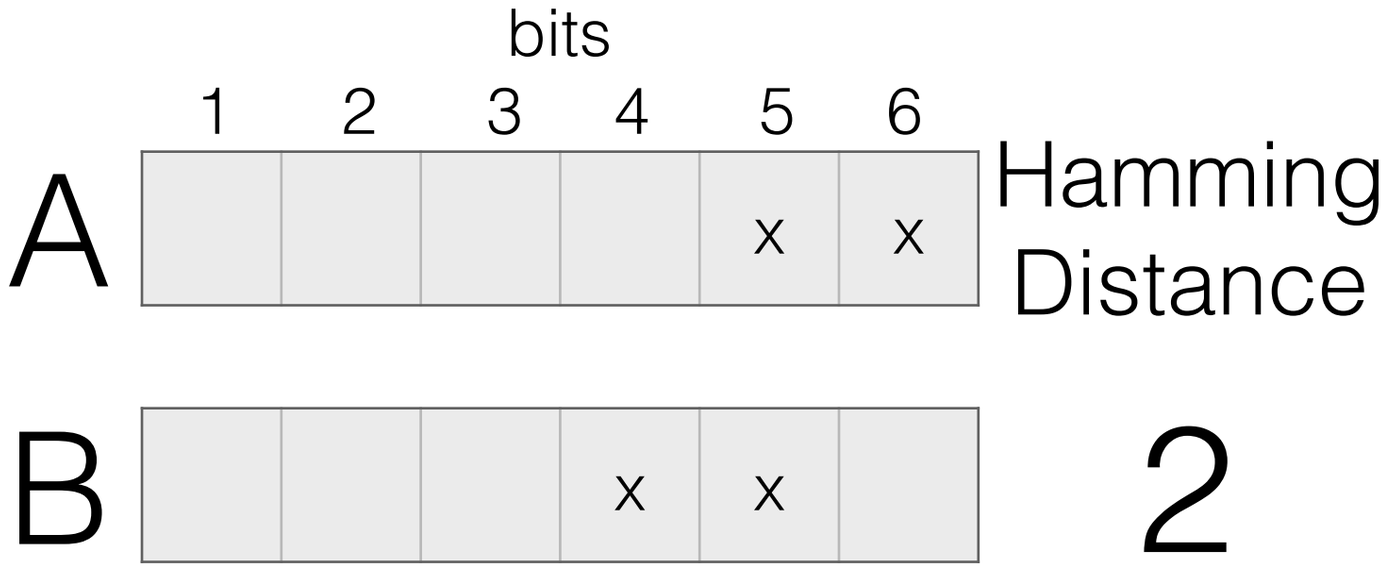

Hamming Distance

2

Distance Metric

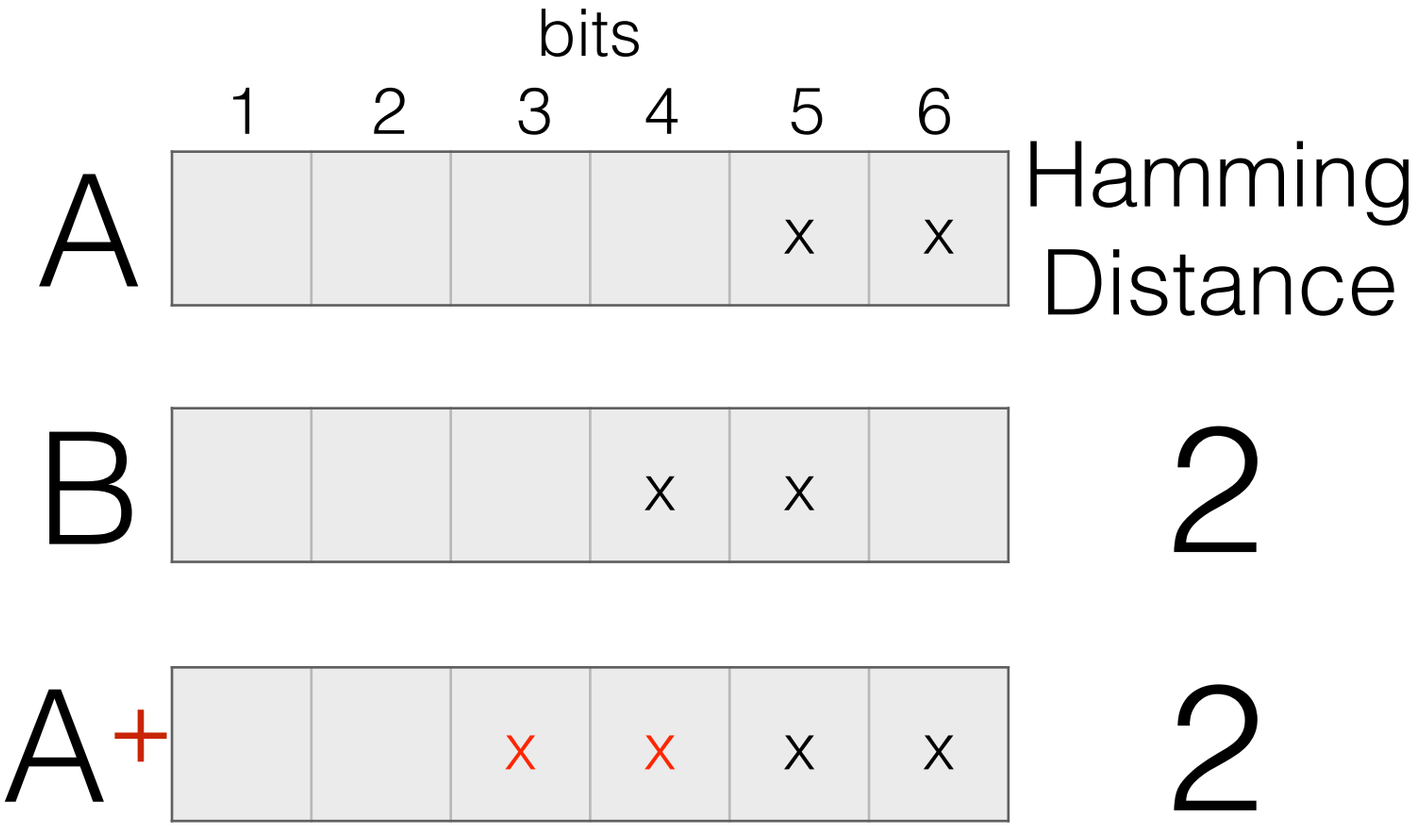

Distance Metric

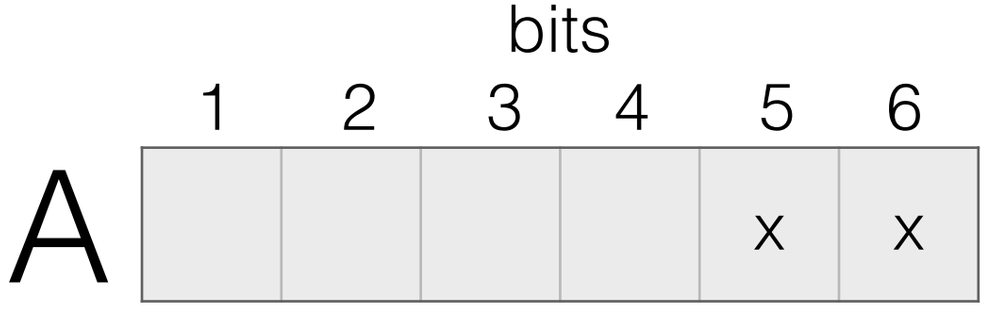

	Hamming Distance	Jaccard Distance
B	2	1
A+	2	0

Distance Metric

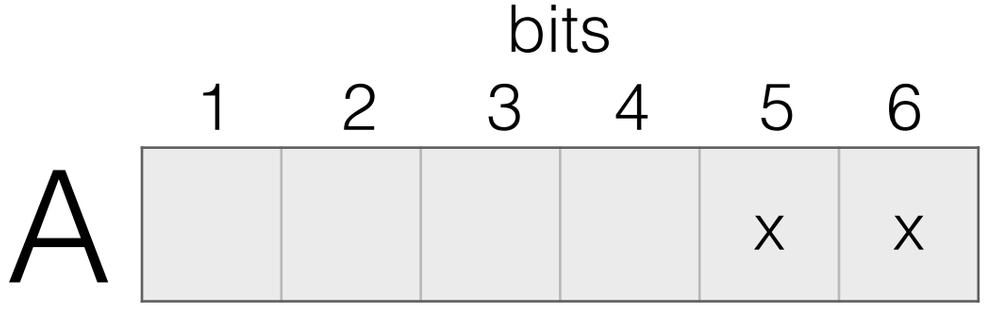

	Hamming Distance	Jaccard Distance
B	2	1
A+	2	0
C	3	1

Distance Metric

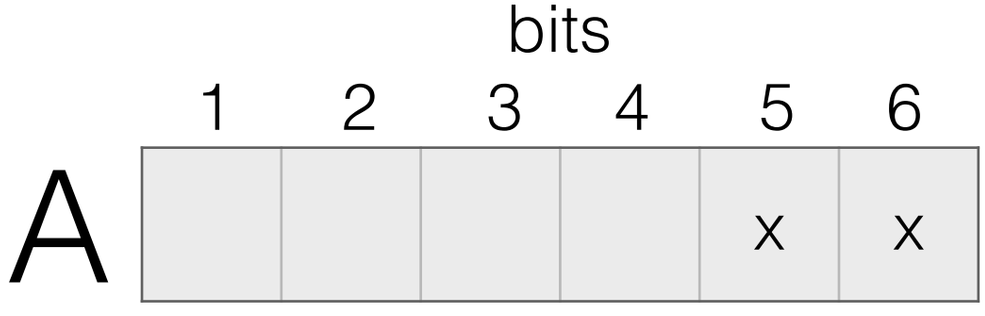

	Hamming Distance	Jaccard Distance	Jaccard Distance / Hamming Weight
B	2	1	.5
A+	2	0	0
C	3	1	.33

Distance Metric

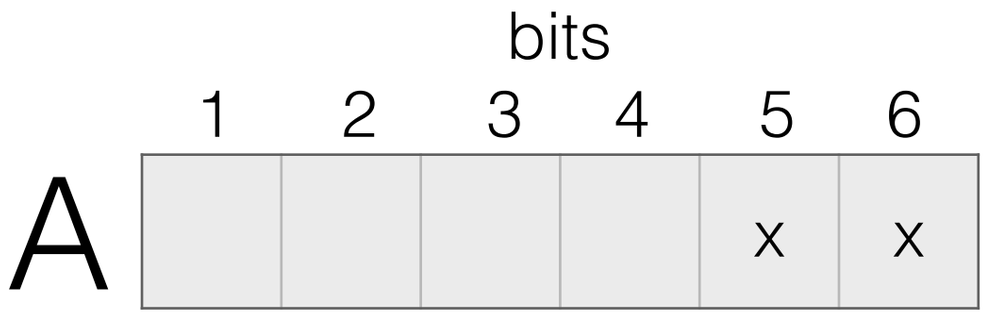

	Hamming Distance	Jaccard Distance	Jaccard Distance / Hamming Weight
B	2	1	.5
A+	2	0	0
C	3	1	.33

Putting Memory Fingerprint Together

- [ ] Page granularity

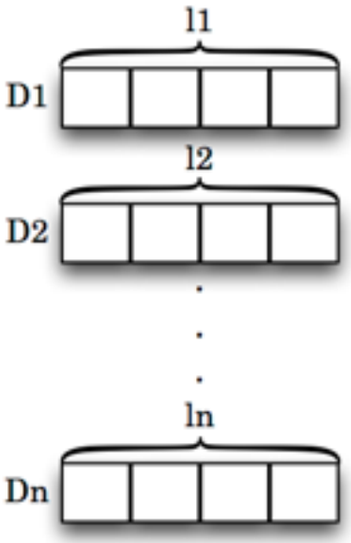

Putting Memory Fingerprint Together

Page granularity

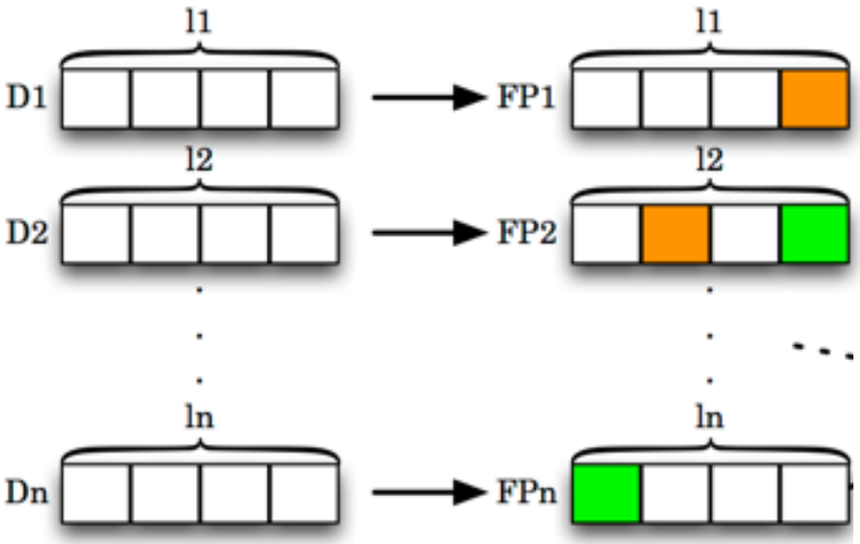

Putting Memory Fingerprint Together

Page granularity

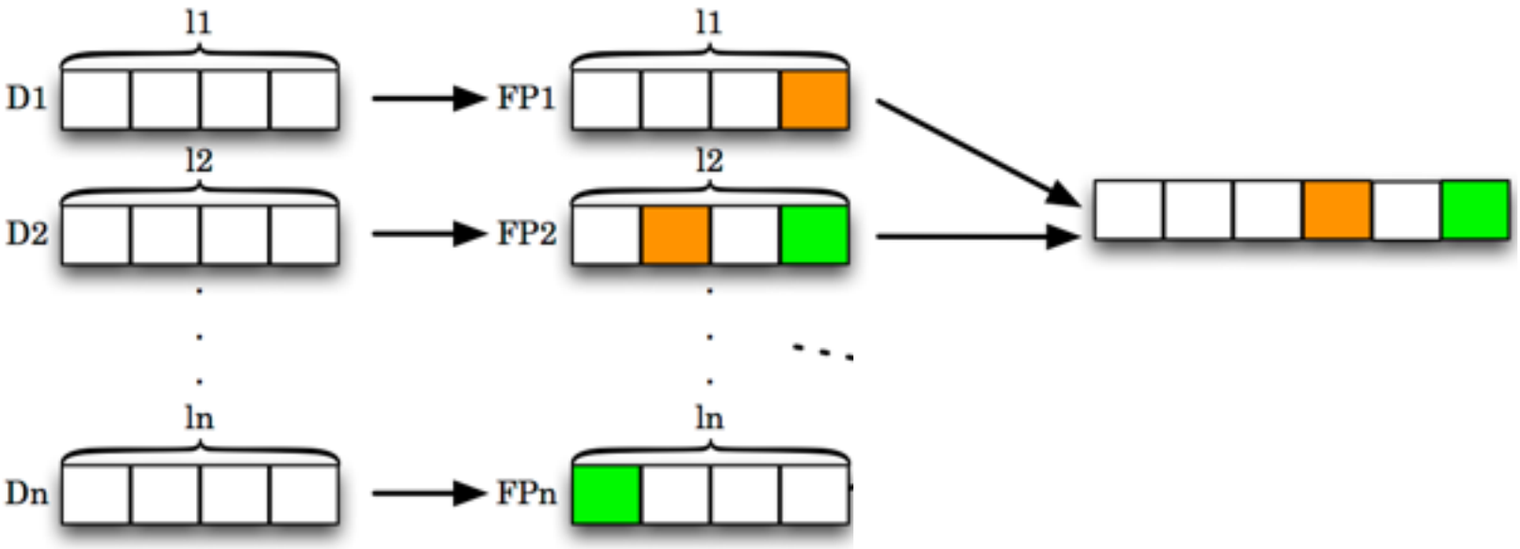

Putting Memory Fingerprint Together

Page granularity

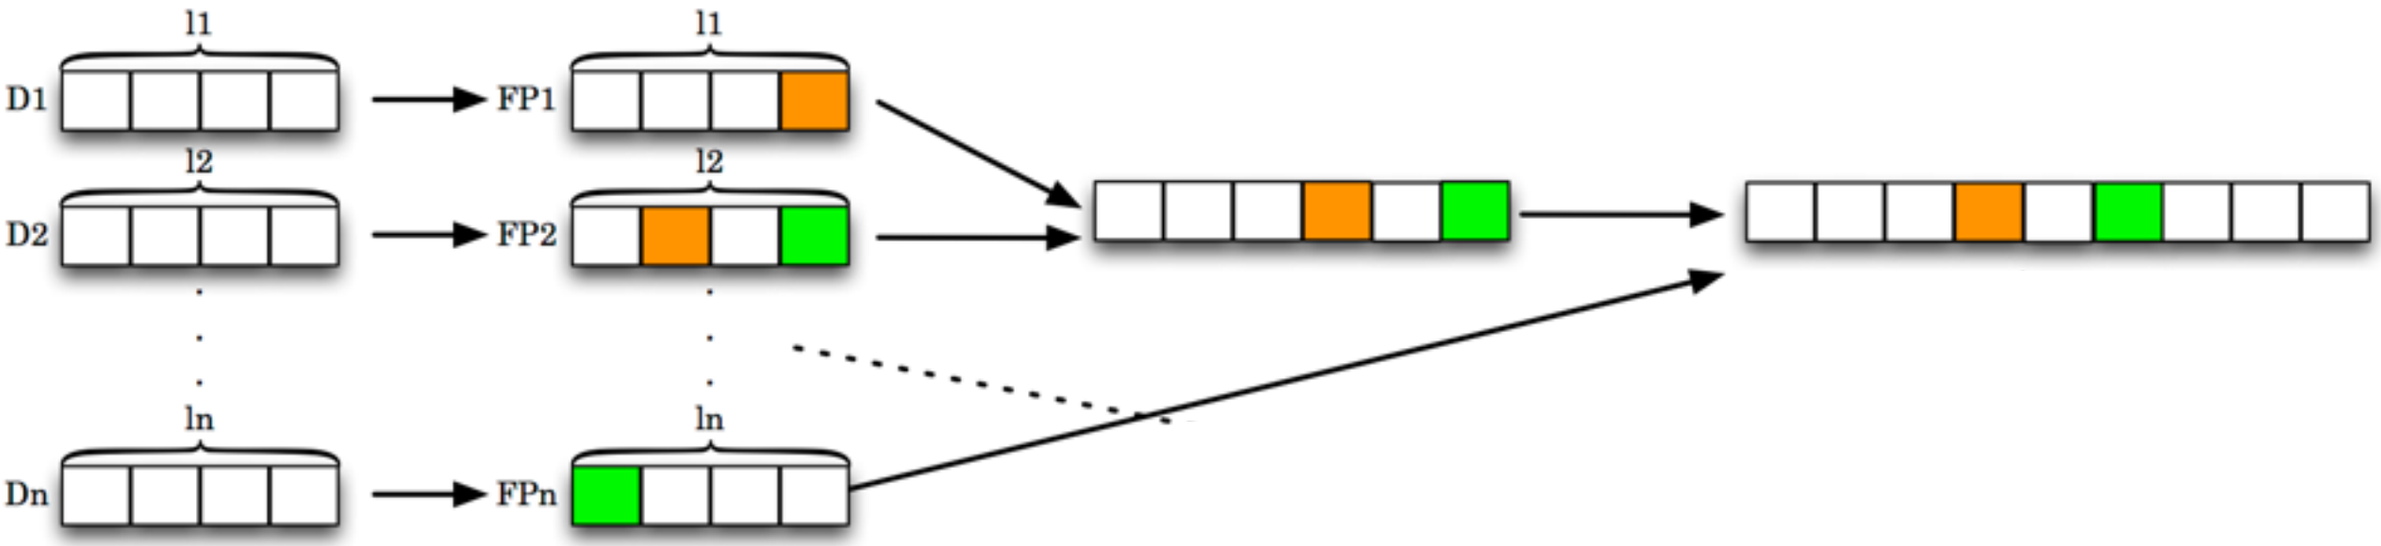

Putting Memory Fingerprint Together

Page granularity

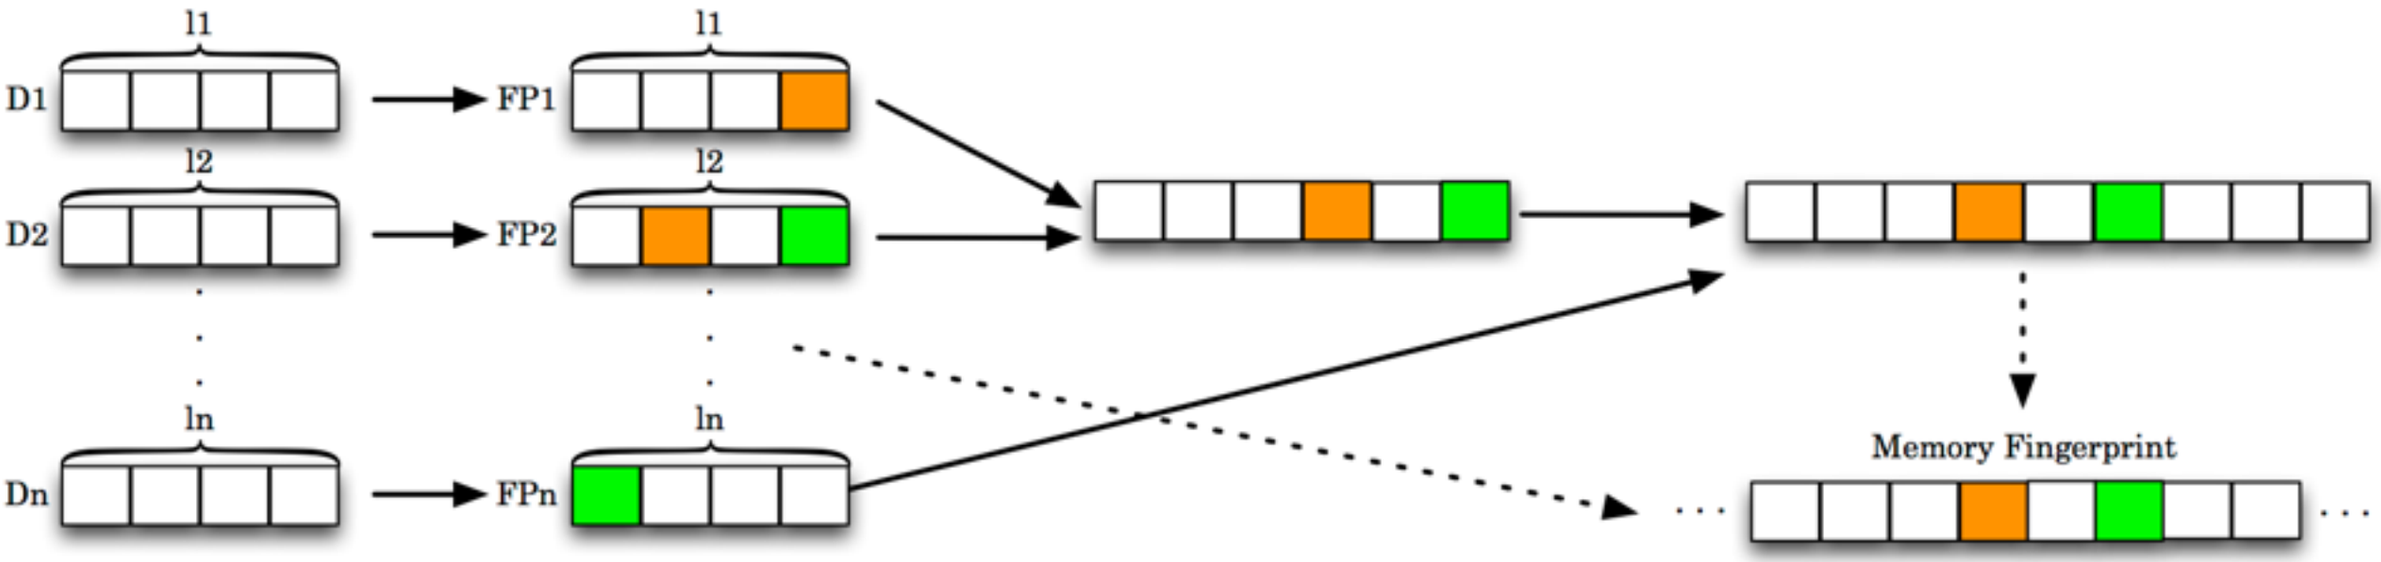

Experimental Setup

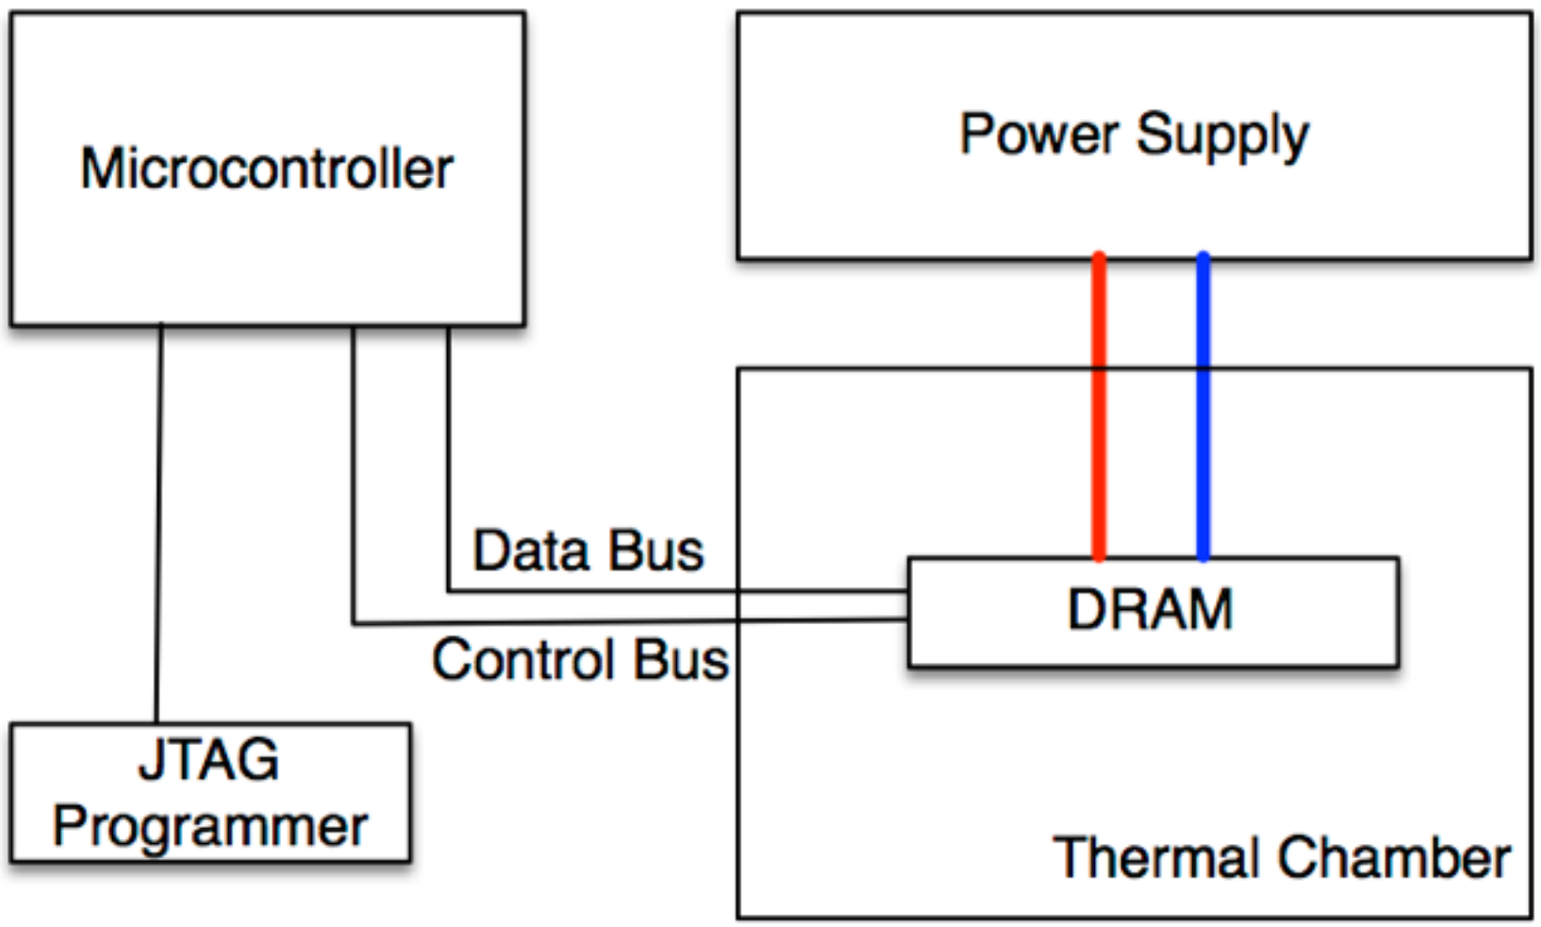

Uniqueness

How unique are the fingerprints?

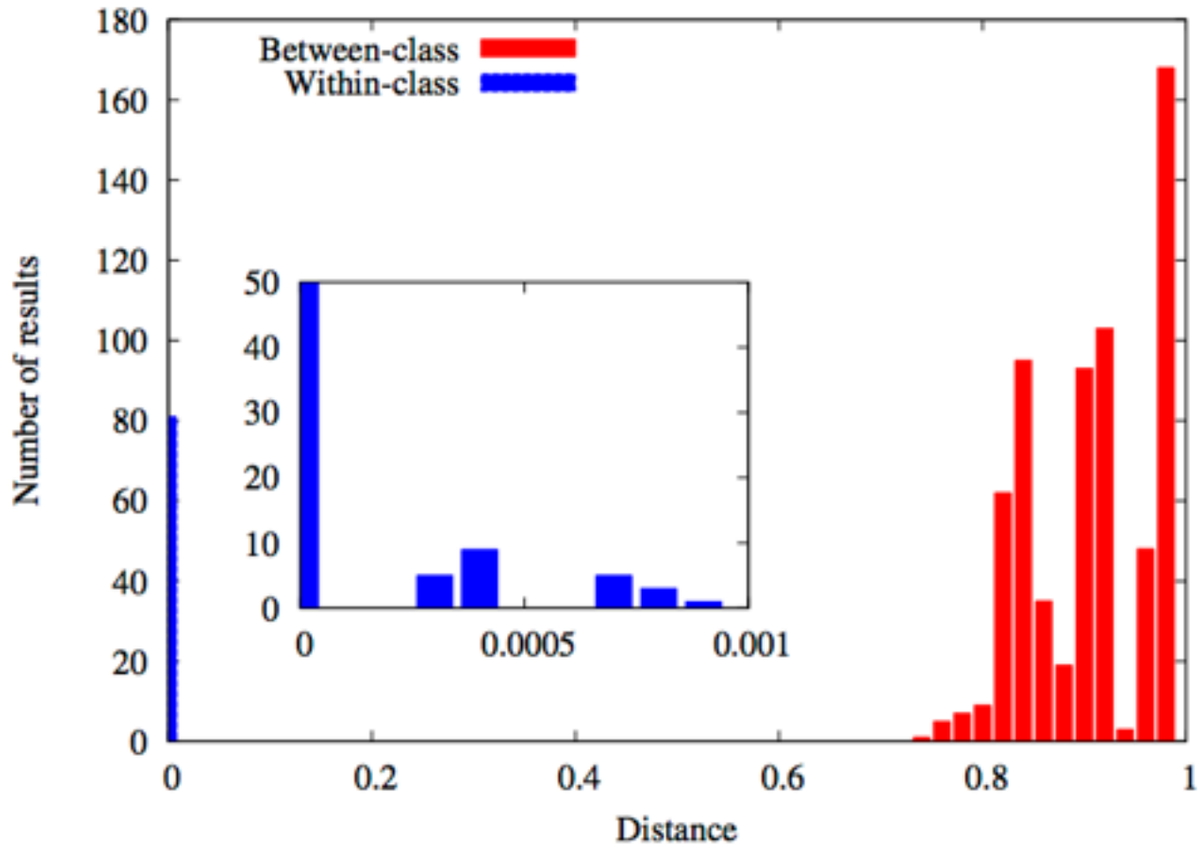

Uniqueness

How unique are the fingerprints?

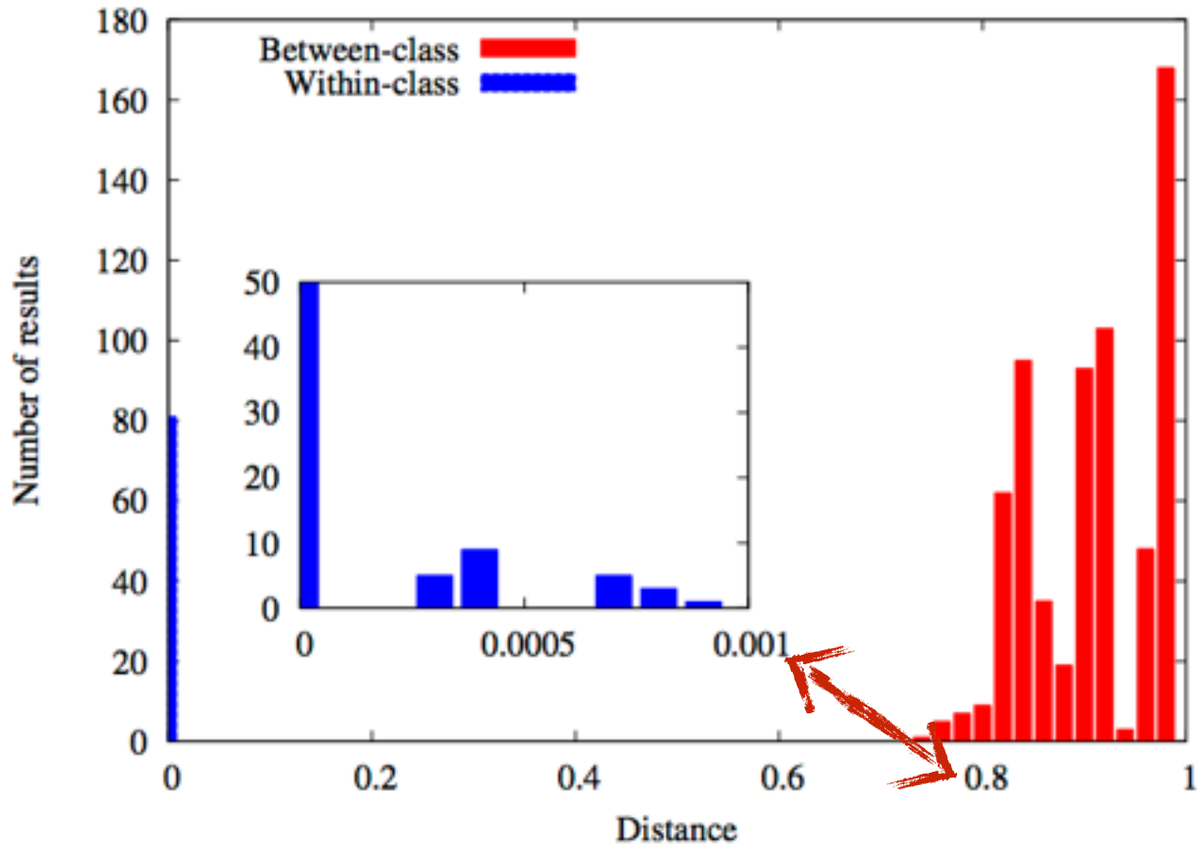

Two order of magnitude difference

Order of Failure

Does the fingerprint hold across different levels of approximation?

Order of Failure

Does the fingerprint hold across different levels of approximation?

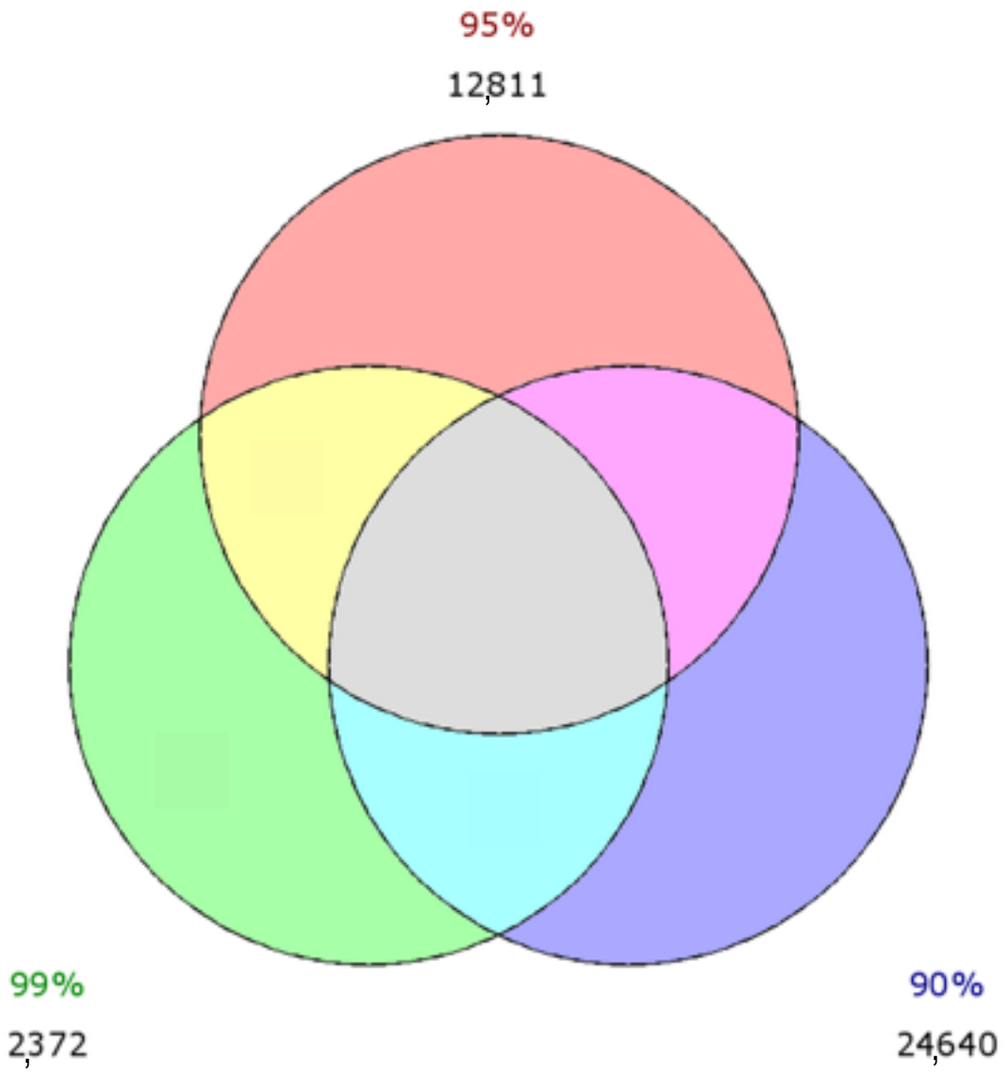

Order of Failure

Does the fingerprint hold across different levels of approximation?

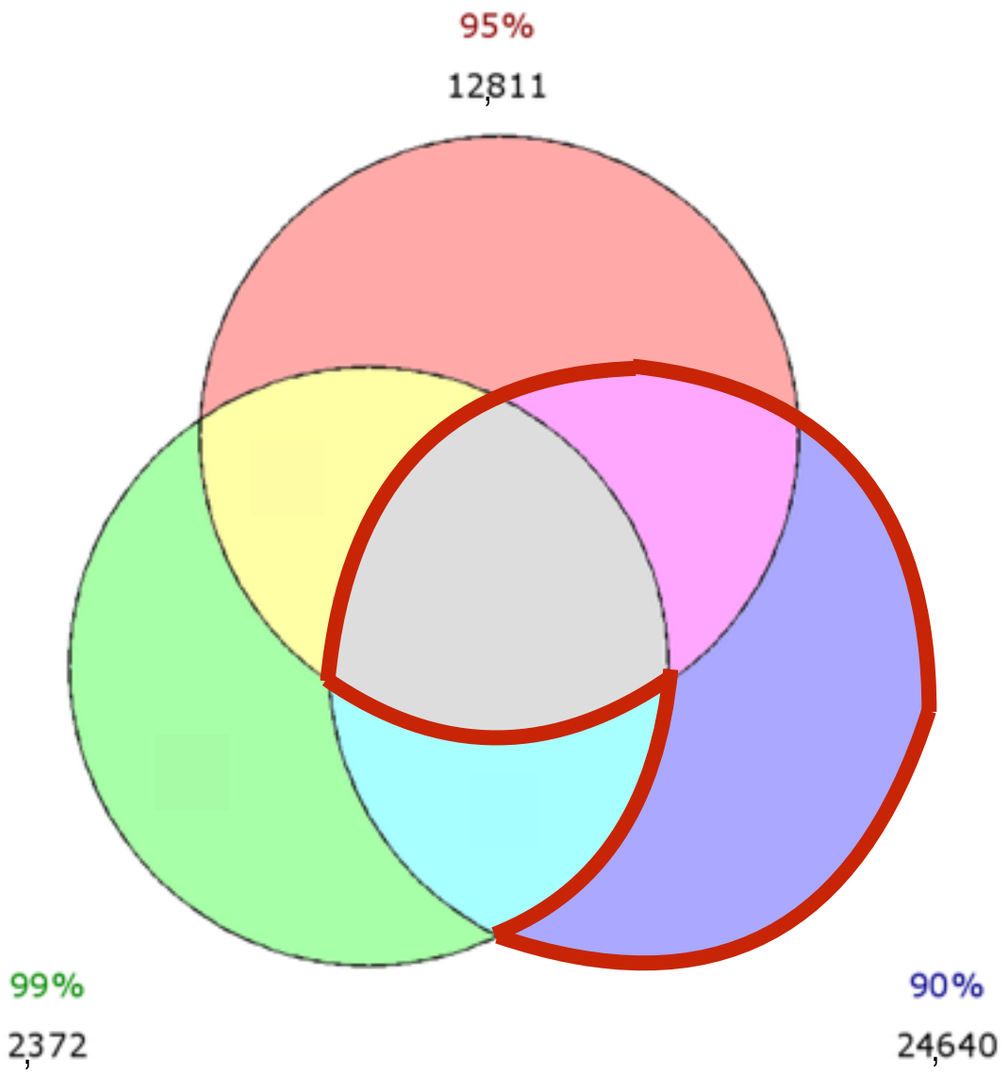

Order of Failure

Does the fingerprint hold across different levels of approximation?

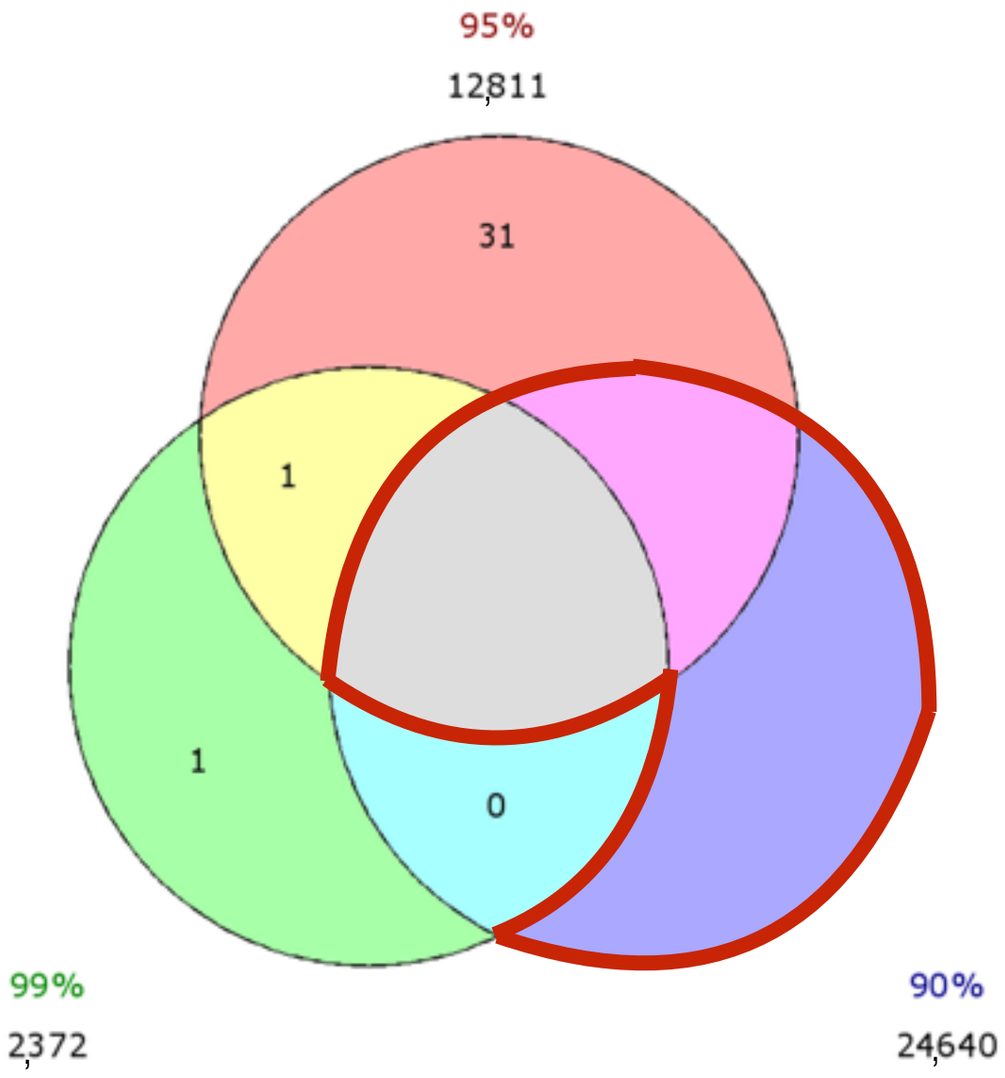

Order of Failure

Does the fingerprint hold across different levels of approximation?

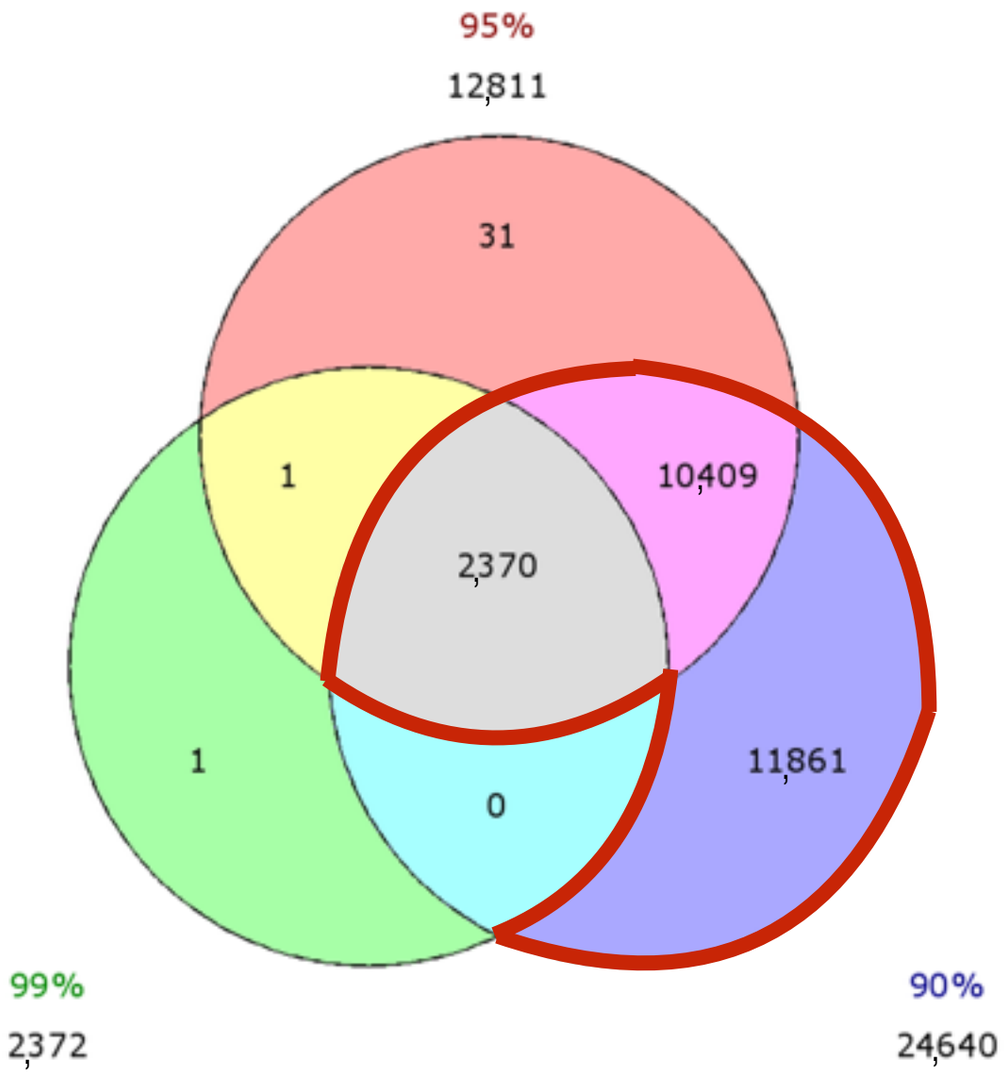

Level of Approximation

How do different levels of approximation affect identification?

Level of Approximation

How do different levels of approximation affect identification?

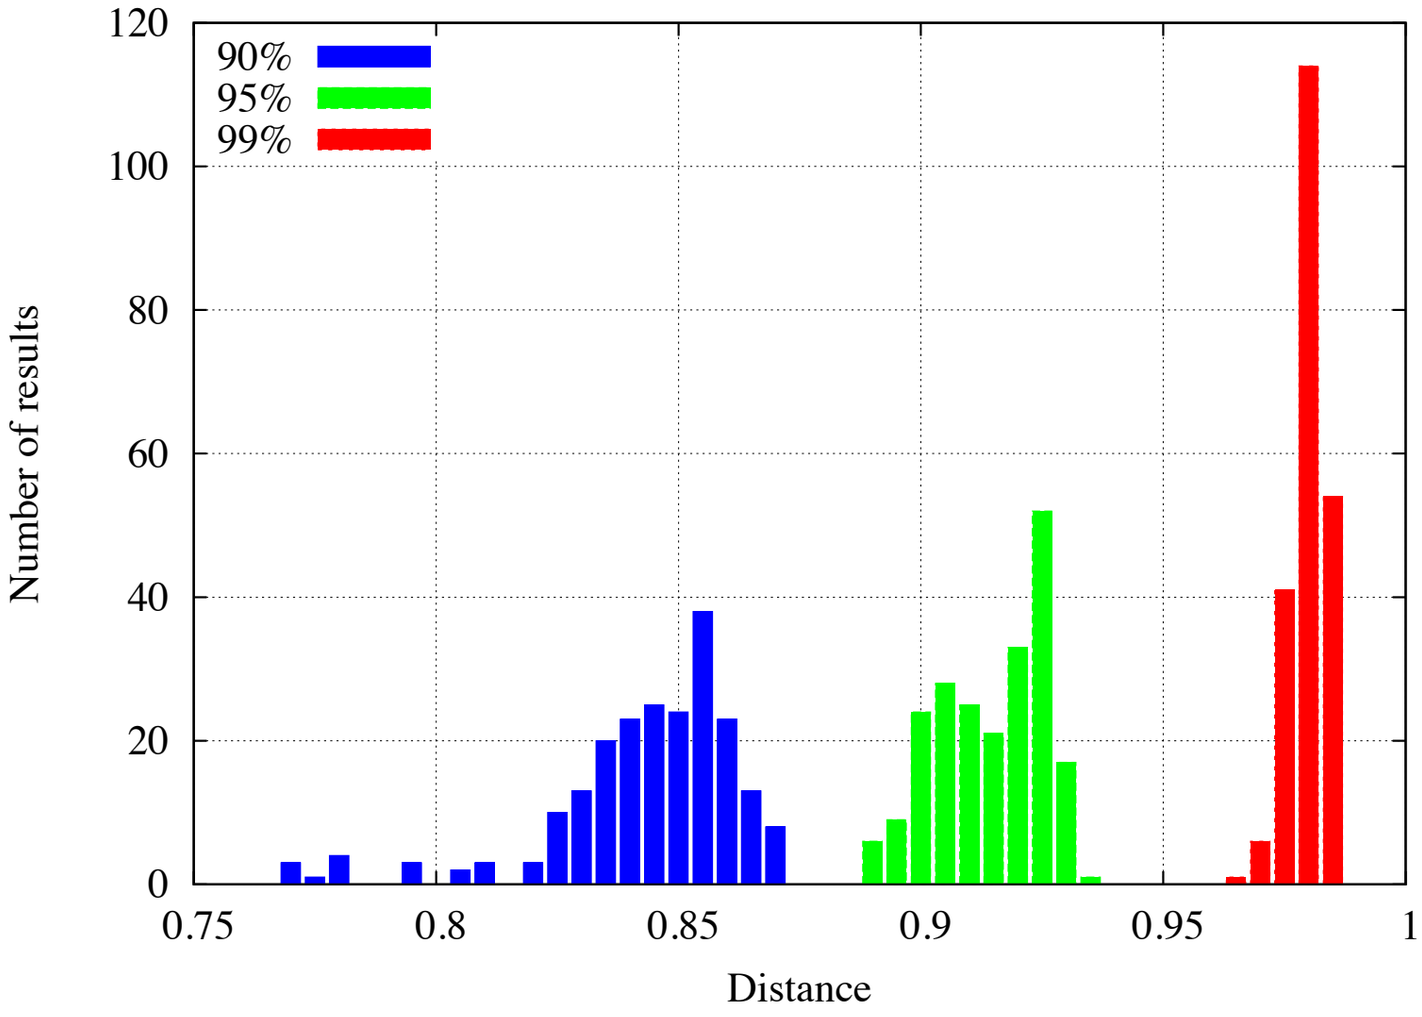

Thermal Effect

How does change in temperature affect identification?

Thermal Effect

How does change in temperature affect identification?

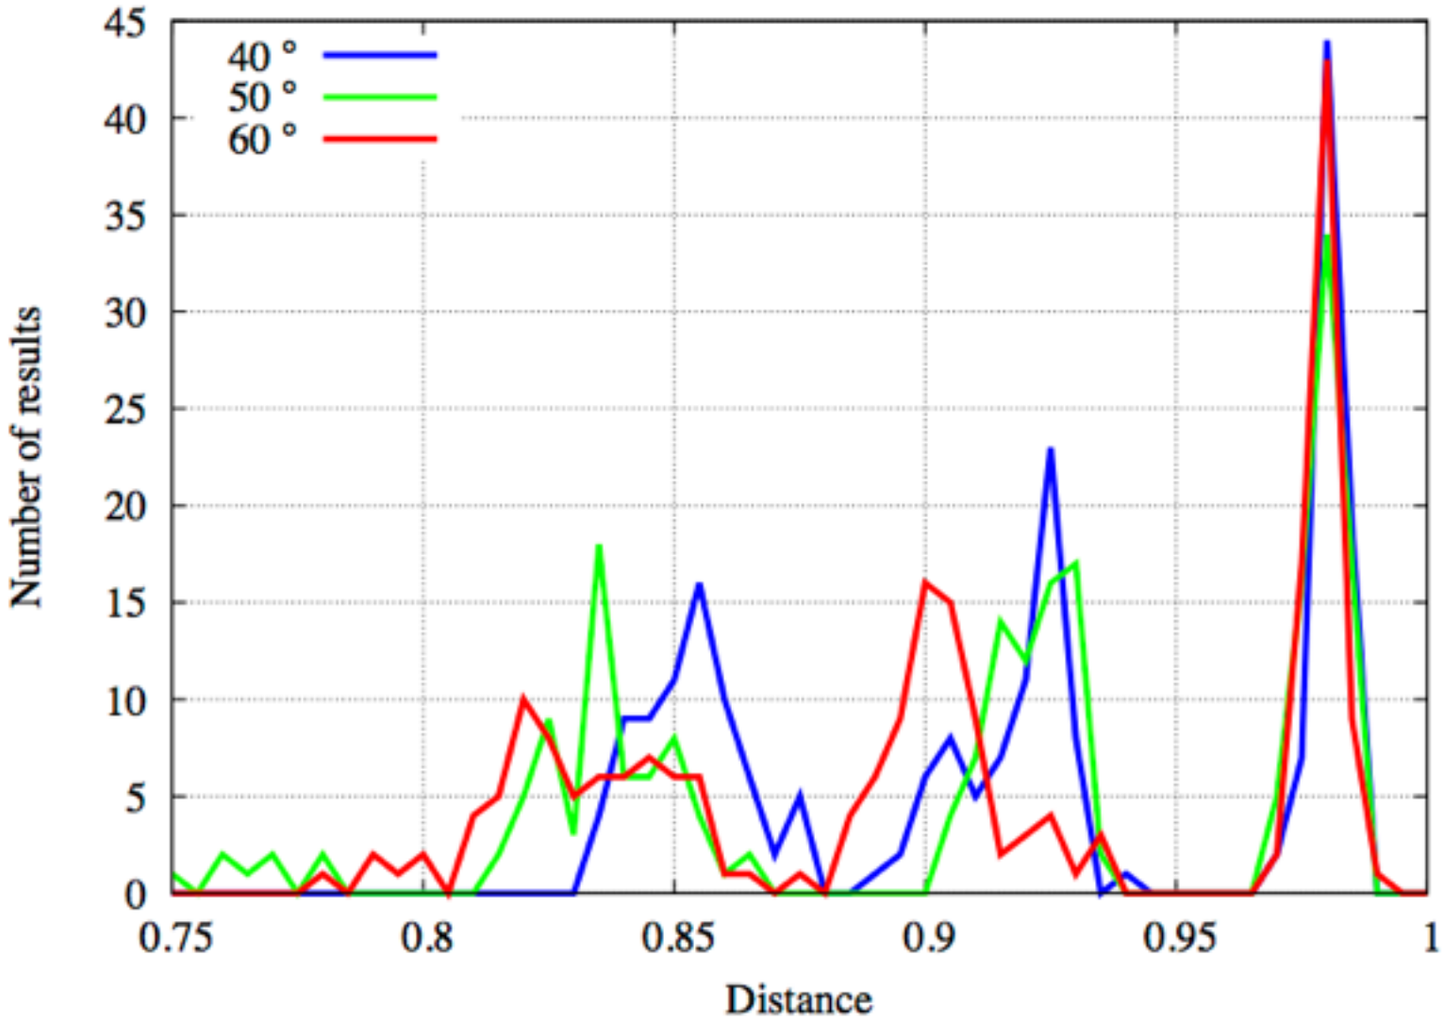

Thermal Effect

How does change in temperature affect identification?

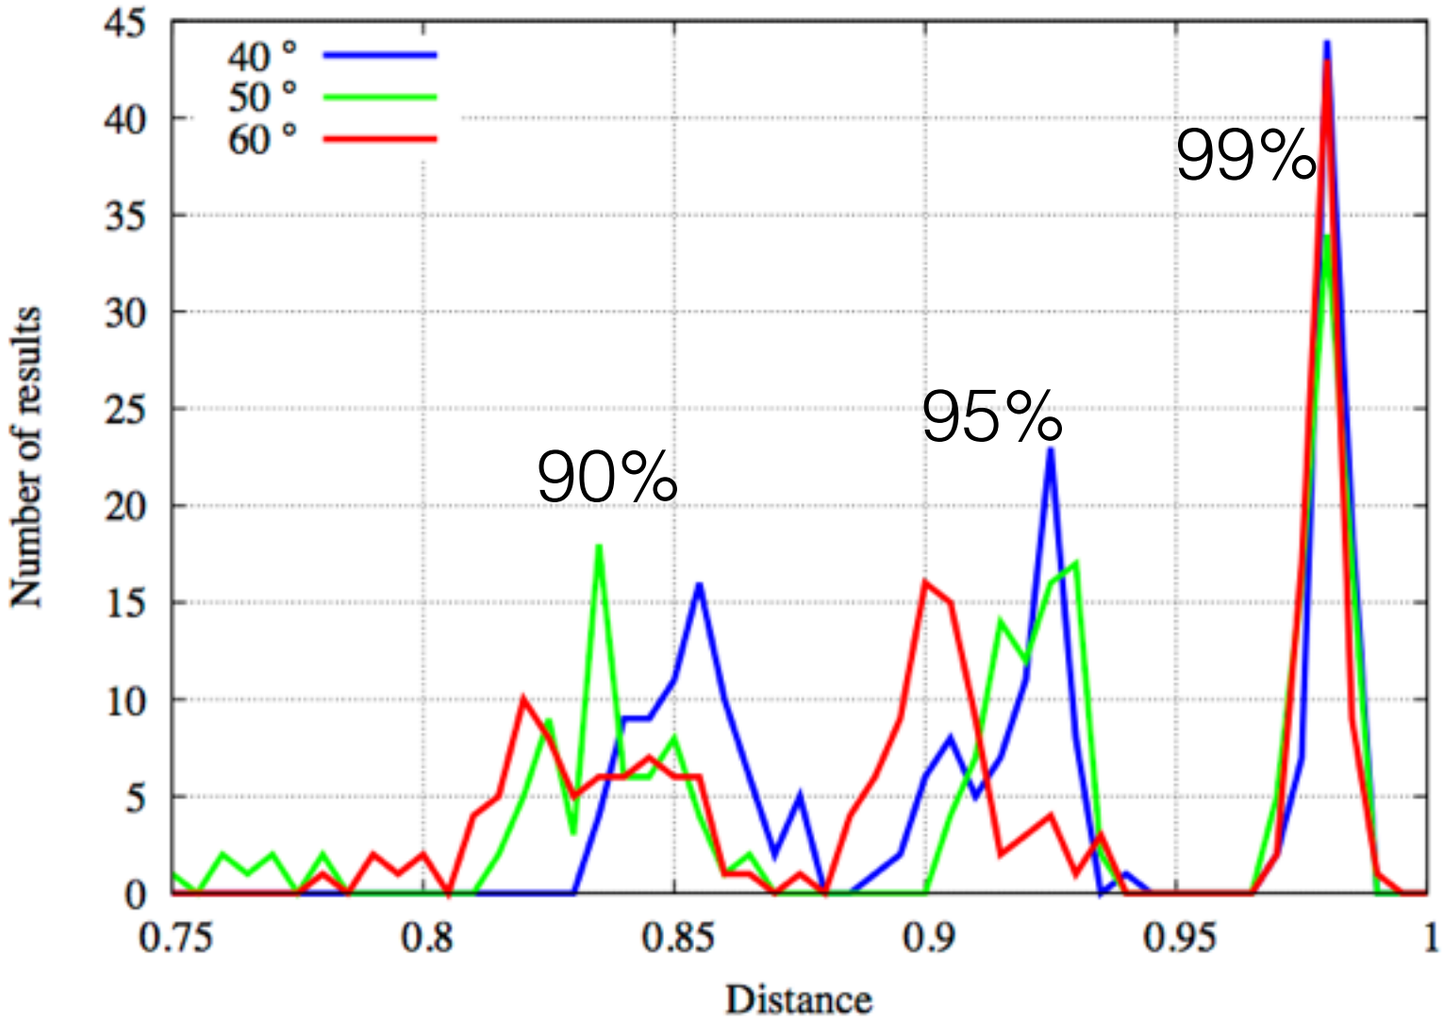

Consistency

How consistent are the fingerprints?

Consistency

How consistent are the fingerprints?

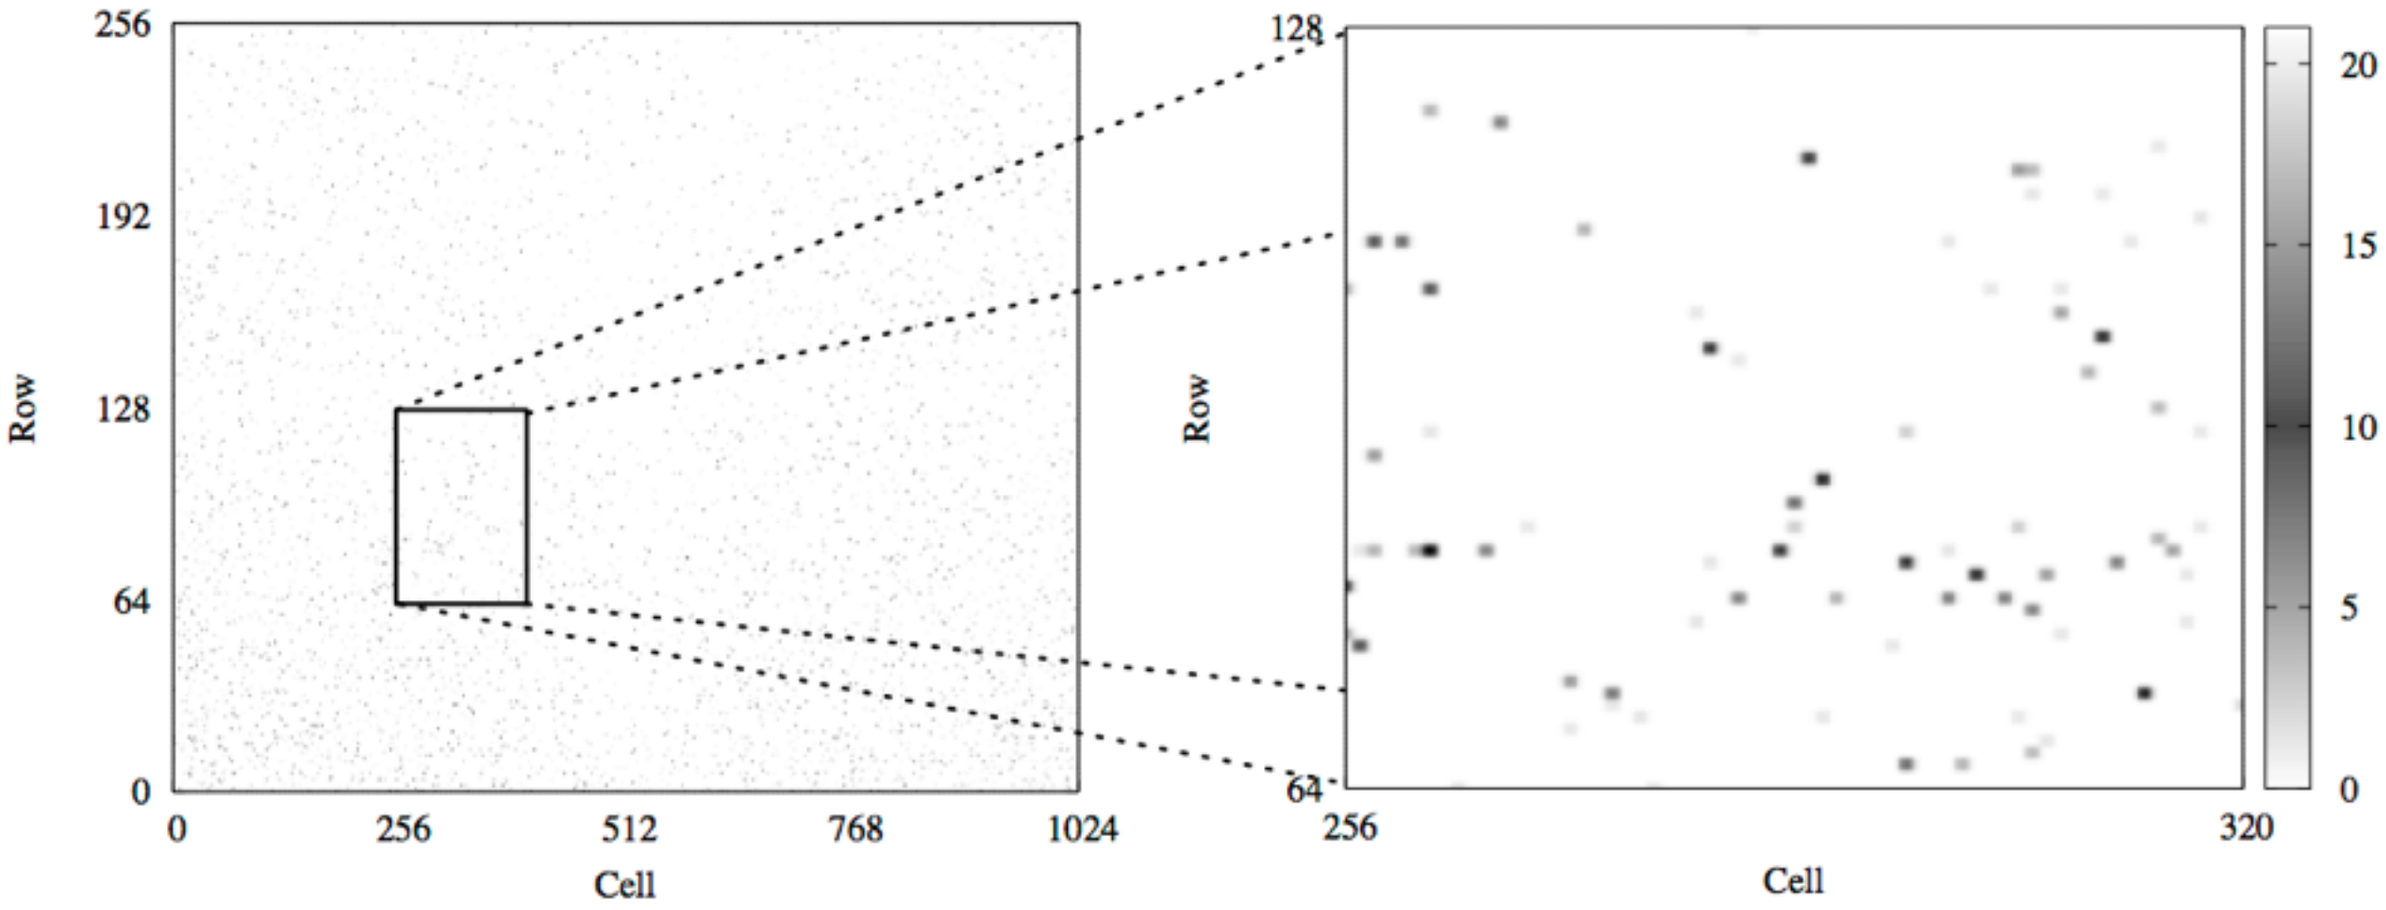

Types of Attack

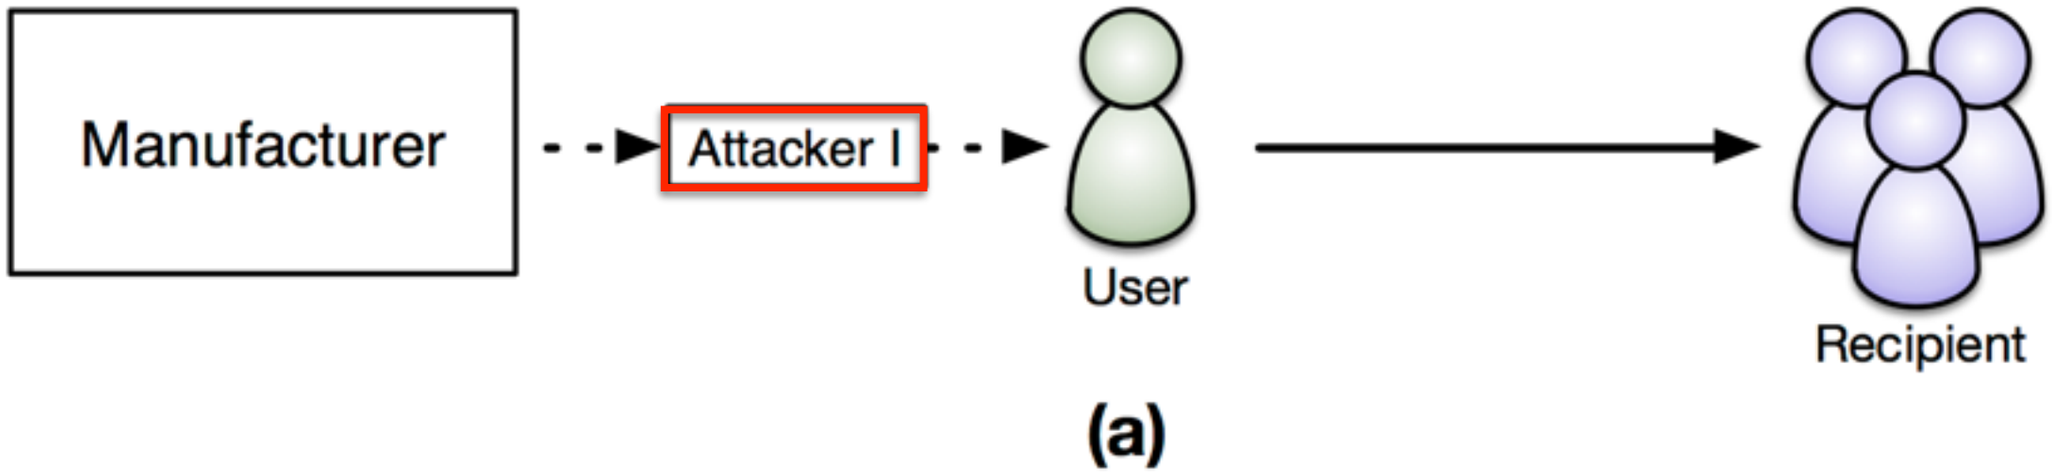

(a)

Types of Attack

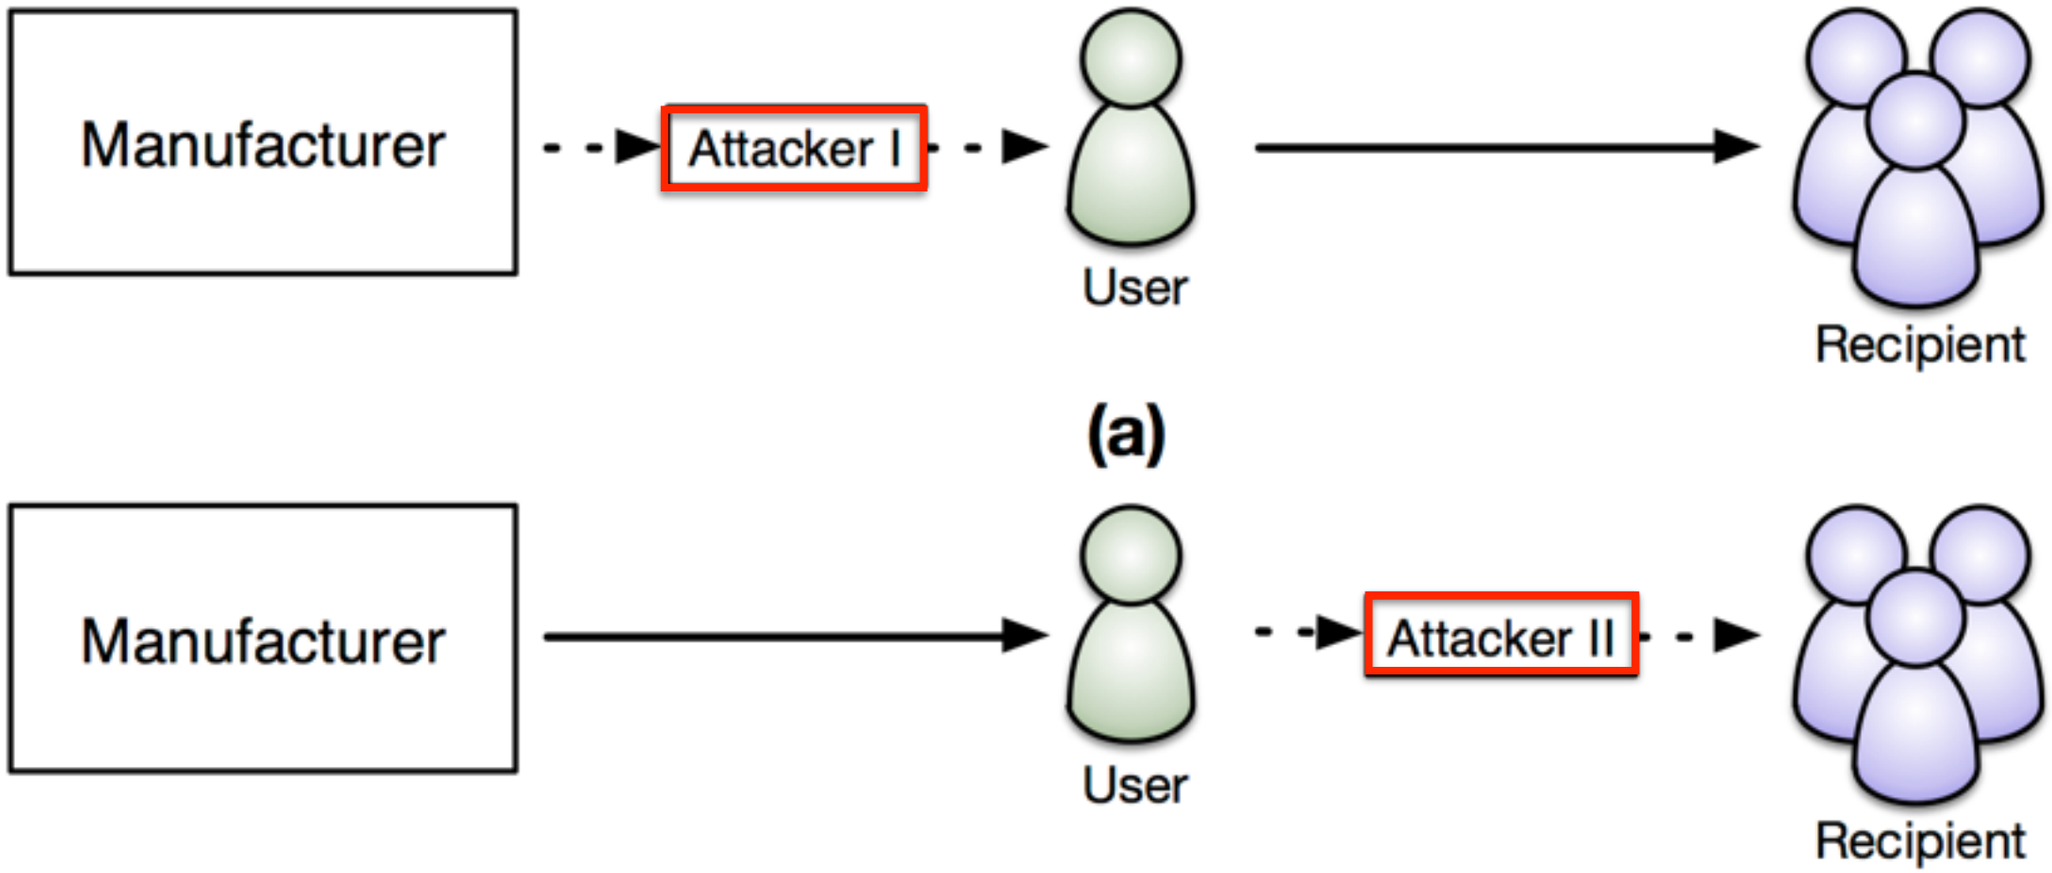

End to End Feasibility

End to End Feasibility

- Commodity system

End to End Feasibility

- Commodity system
- Edge detection tool

End to End Feasibility

- Commodity system
- Edge detection tool

- 1000X10MB traces

End to End Feasibility

- Commodity system
- Edge detection tool

- 1000X10MB traces

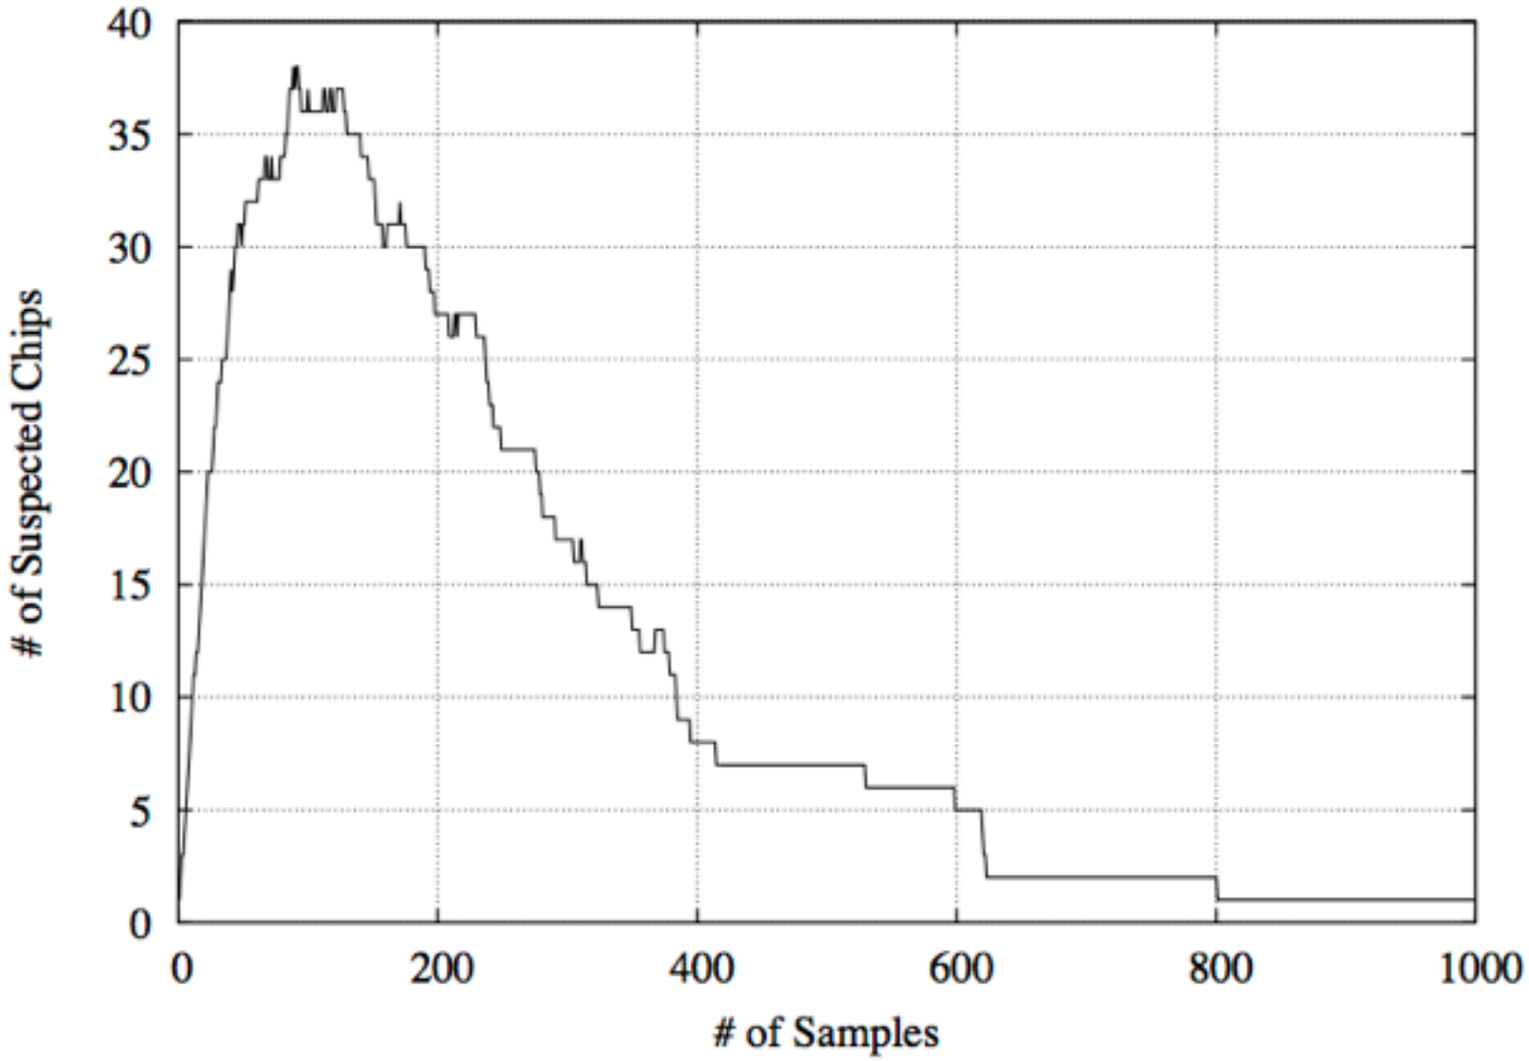

End to End Feasibility

- Commodity system
- Edge detection tool

- 1000X10MB traces

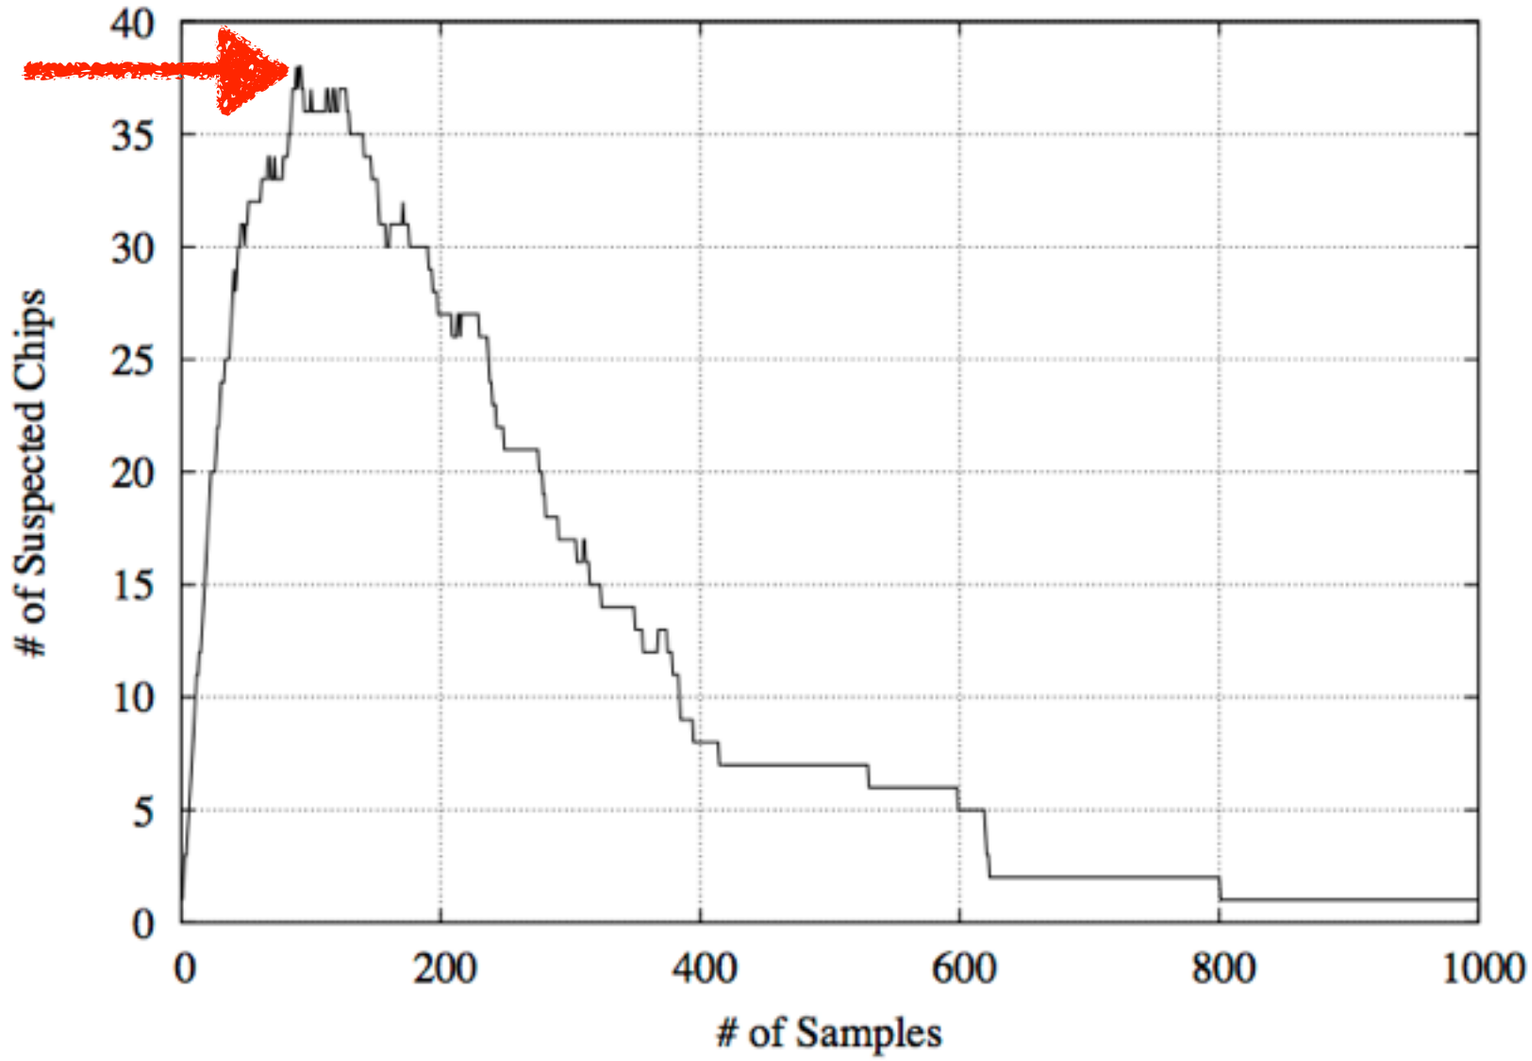

Chance of Mismatch

How much entropy does a page of memory provide?

Chance of Mismatch

How much entropy does a page of memory provide?

For memory of size **M** bits where
A bits of errors are tolerated:

$$Max\ unique\ fingerprints = \binom{M}{A}$$

Chance of Mismatch

How much entropy does a page of memory provide?

For memory of size **M** bits where **A** bits of errors are tolerated:

$$Max\ unique\ fingerprints = \binom{M}{A}$$

Given noise threshold of **T** bits using Hamming bound:

$$\frac{\sum_{i=1}^{T}\binom{M}{i}}{\binom{M}{A}} \leq Chance\ of\ mismatching \leq \frac{\sum_{i=1}^{2T}\binom{M}{i}}{\binom{M}{A}}$$

Chance of Mismatch

How much entropy does a page of memory provide?

For memory of size **M** bits where **A** bits of errors are tolerated:

$$Max\ unique\ fingerprints = \binom{M}{A}$$

Given noise threshold of **T** bits using Hamming bound:

$$\frac{\sum_{i=1}^{T}\binom{M}{i}}{\binom{M}{A}} \leq Chance\ of\ mismatching \leq \frac{\sum_{i=1}^{2T}\binom{M}{i}}{\binom{M}{A}}$$

One page of memory	
$M = 32768$ bits, $A = 1\%$, $T = 32$ bits	
Max possible fingerprints	8.70×10^{795}
Max unique fingerprints	$\geq 1.07 \times 10^{590}$
Chance of mismatching	$\leq 9.29 \times 10^{-591}$
Total Entropy	2423 bits

Conclusion

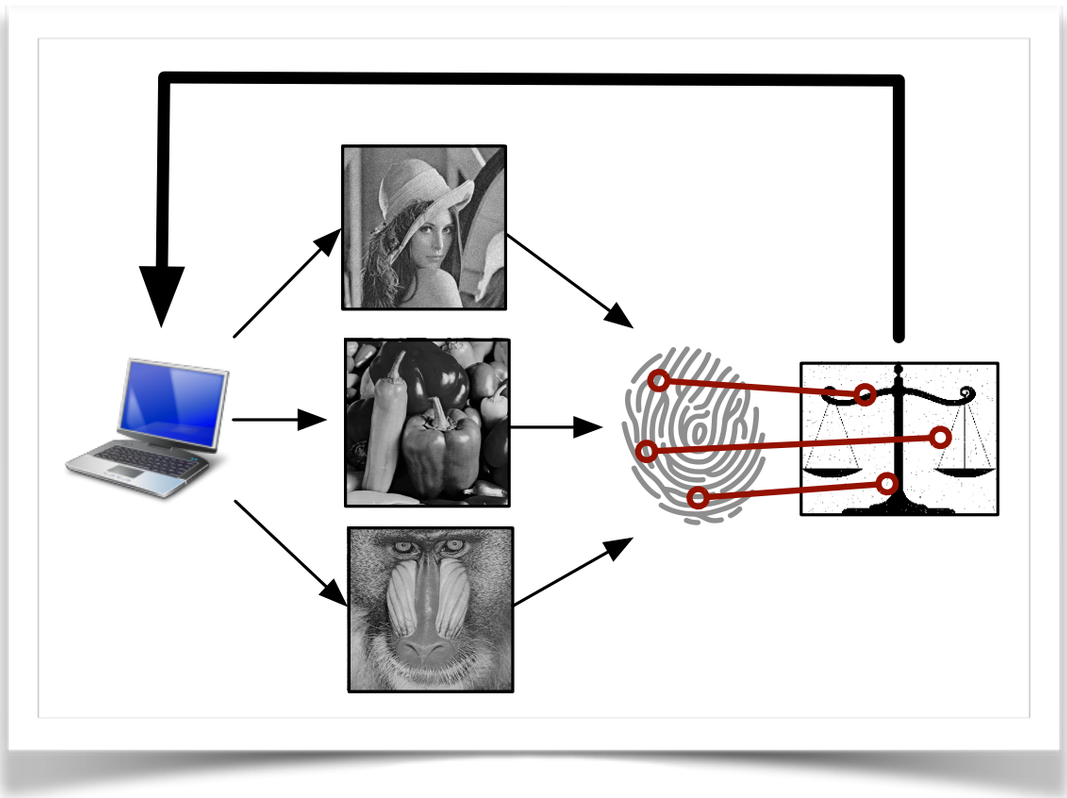

https://github.com/impedimentToProgress/ProbableCause

Conclusion

Consider **Security & Privacy** as a primary design criteria in emerging systems

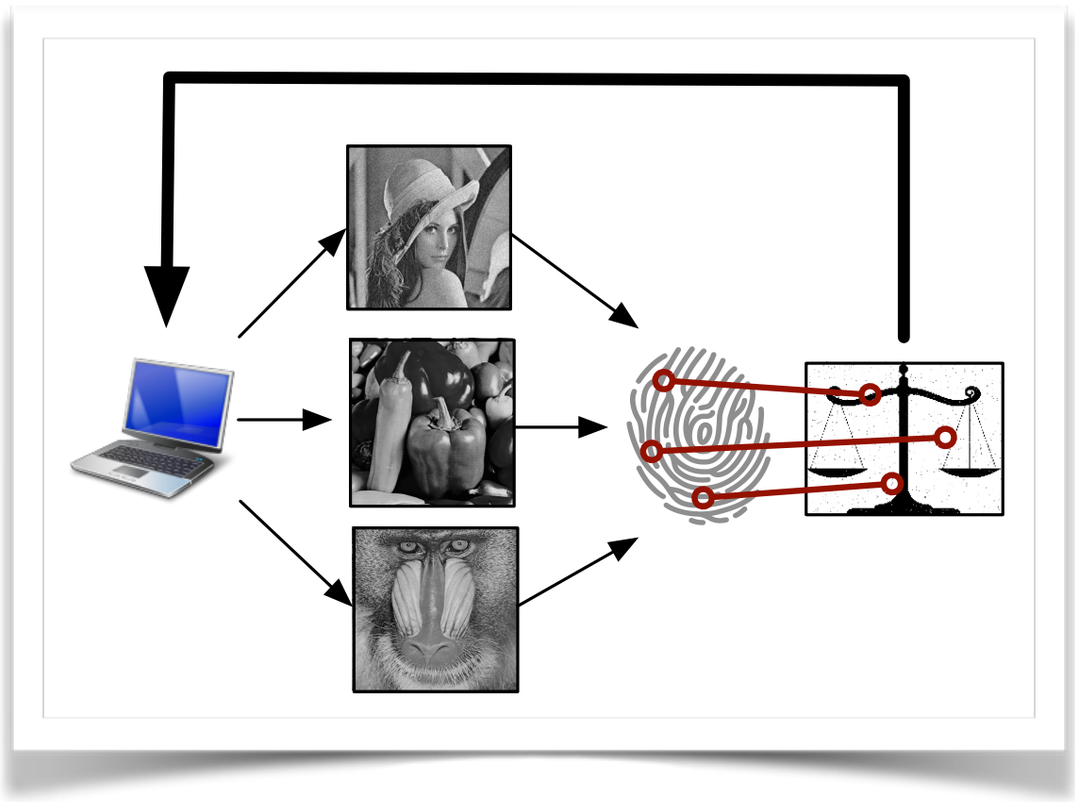

https://github.com/impedimentToProgress/ProbableCause

Backup Slides

Defenses

- Data Segregation
- Noise
- Data Scrambling

Error Localization

- Recalculate from known inputs
- Noise detection algorithms
- Speculative distance calculation

پایان