Efficient and Cost-Effective Testing for Side-Channel Vulnerabilities

Mark E. Marson
Technical Director

May 3, 2016

Rambus

Topics

- Side-Channel Overview
- Side-Channel Testing Methodology
- AES Example
- Demo
- Summary

Side-Channel Overview

- Discovered by Cryptography Research in mid-1990s (“DPA” and “SPA”)
- Low cost, non-invasive attacks on crypto HW
  - Key extraction
  - Reverse engineering
  - Device modification
- All cryptographic algorithms vulnerable
  - Symmetric crypto: DES, AES, HMAC,…
  - Asymmetric crypto: RSA, DH, EC variants,…
- Affects all types of hardware and software implementations, including:
  - ASICs, FPGAs, software on CPU
- Same techniques work for different signal sources, including timing, E&M and RF

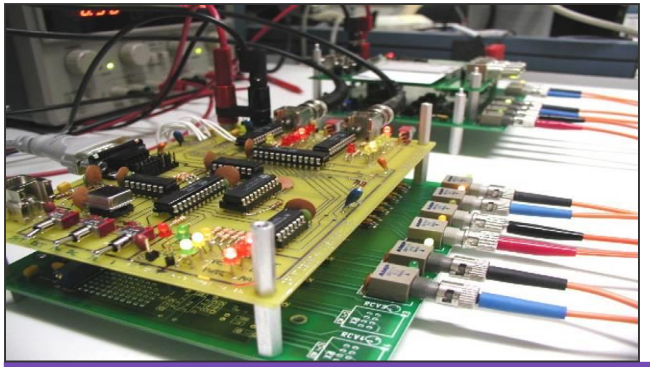

Differential Power Analysis

Paul Kocher, Joshua Jaffe, and Benjamin Jun

Cryptography Research, Inc.
607 Market Street, 5th Floor
San Francisco, CA 94105, USA.
http://www.cryptography.com
E-mail: {paul,josh,ben}@cryptography.com.

Abstract. Cryptosystem designers frequently assume that secrets will be manipulated in closed, reliable computing environments. Unfortunately, actual computers and microchips leak information about the operations they process. This paper examines specific methods for analyzing power consumption measurements to find secret keys from tamper resistant devices. We also discuss approaches for building cryptosystems that can operate securely in existing hardware that leaks information.

Keywords: differential power analysis, DPA, SPA, cryptanalysis, DES

Background

Advances in Cryptology – Crypto 99 Proceedings, LNCS 1666, Springer-Verlag, 1999

Side-Channel Overview – Simple Power Analysis

- Keys can be extracted from a single trace

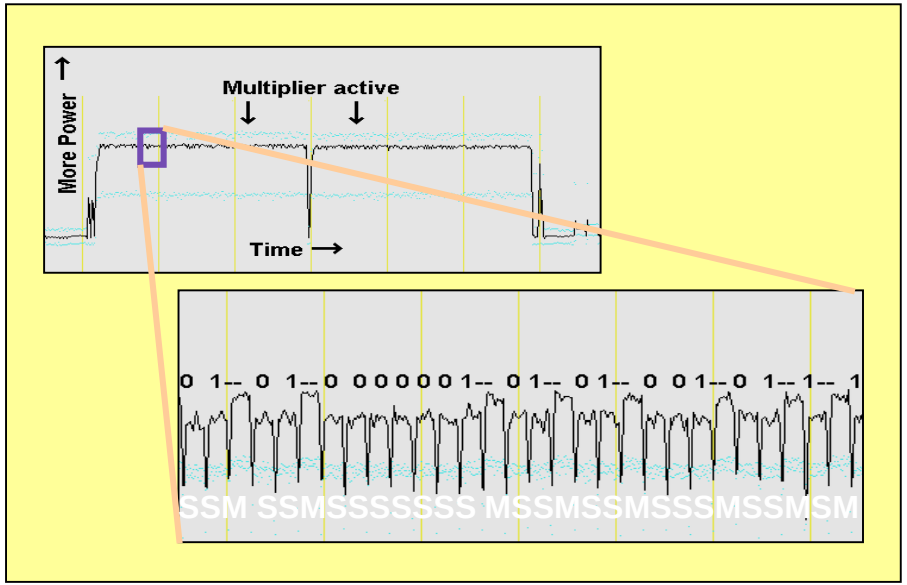

Straightforward RSA Implementation

```
For each bit i of secret d
   perform "Square"
   if (bit i == 1)
      perform "Multiply"
   endif
endfor
```

- Similar analysis also applies to EM

Side-Channel Overview – Differential Power Analysis

- Signal / noise ratio may be very small
  - However, statistical influence remains…

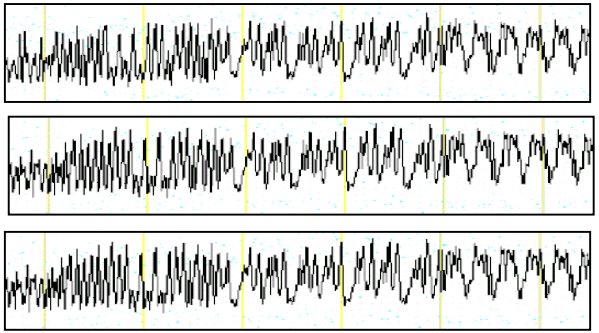

Side-Channel Overview – Differential Power Analysis

- Signal / noise ratio may be very small
  - However, statistical influence remains…

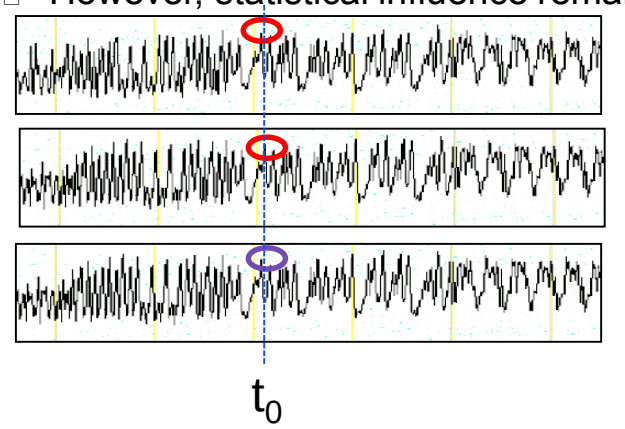

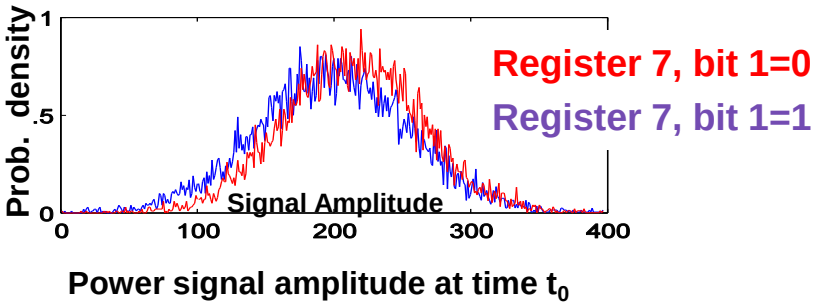

Side-Channel Overview – Differential Power Analysis

- Signal / noise ratio may be very small
  - However, statistical influence remains…

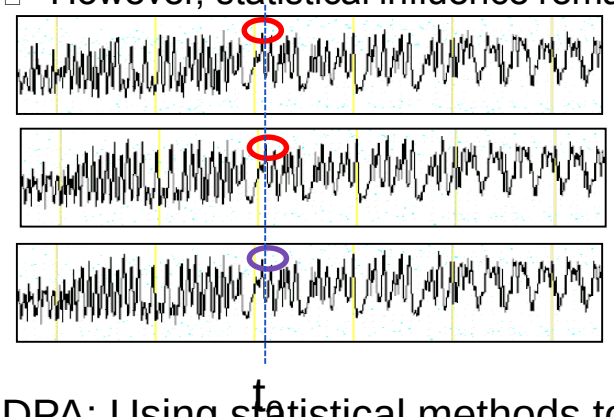

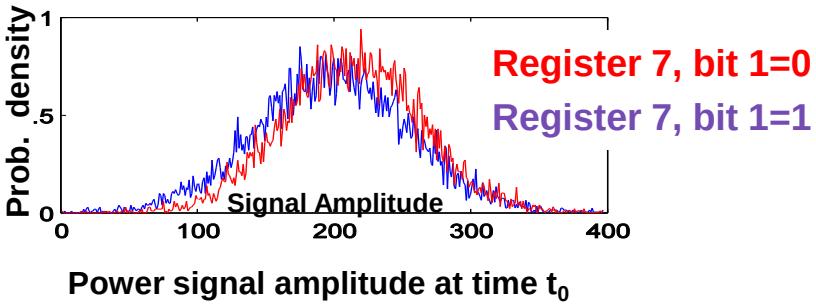

- DPA: Using statistical methods to analyze minute differences in power measurements due to the data being manipulated

Side-Channel Overview – Differential Power Analysis

- Signal / noise ratio may be very small
  - However, statistical influence remains…

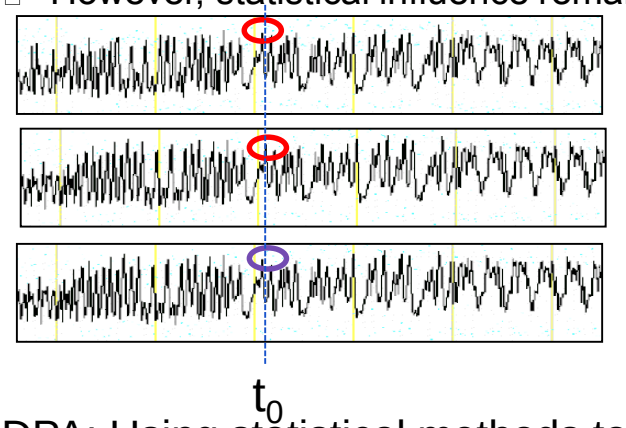

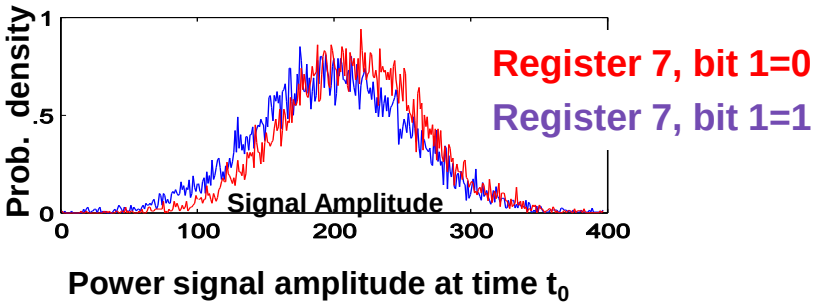

- DPA: Using statistical methods to analyze minute differences in power measurements due to the data being manipulated
- Can be used to extract secret keys and data

Side-Channel Overview – Differential Power Analysis

- Signal / noise ratio may be very small
  - However, statistical influence remains…

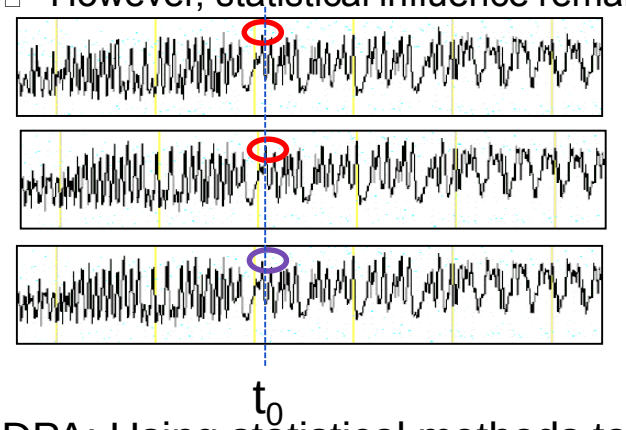

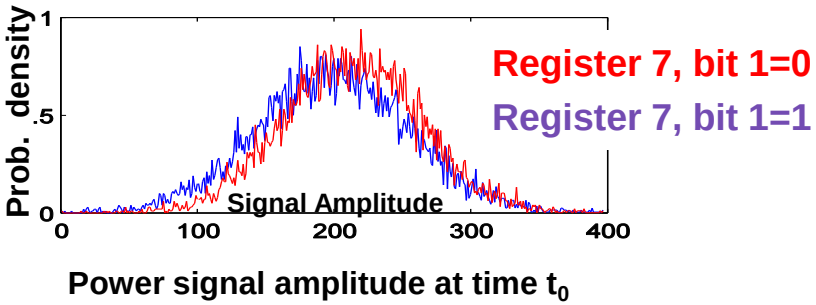

- DPA: Using statistical methods to analyze minute differences in power measurements due to the data being manipulated
- Can be used to extract secret keys and data
- Similar analysis applies to EM and timing measurements

Hypothesis:
The 8 key bits
are **00001111**

DPA - Correct Key Guess

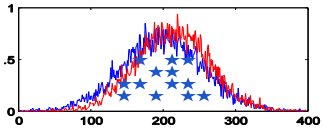

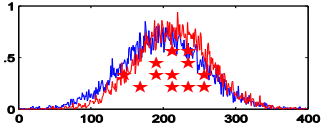

Predicted LSB(I)

Traces with predicted LSB = 1

Traces with predicted LSB = 0

0
1
1
0
1
0
0
1
⋮

Hypothesis:
The 8 key bits
are **00001111**

DPA - Correct Key Guess

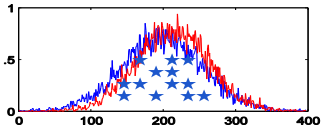

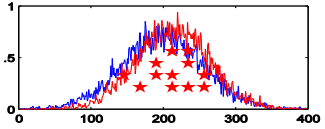

Predicted LSB(I)

Traces with predicted LSB = 1

Traces with predicted LSB = 0

0
1
1
0
1
0
0
1
⋮

Average of traces

DPA - Correct Key Guess

Hypothesis:
The 8 key bits
are **00001111**

Predicted LSB(I)

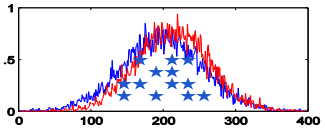

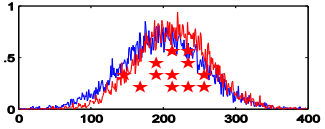

Traces with predicted LSB = 1

Traces with predicted LSB = 0

0
1
1
0
1
0
0
1
⋮

Average of traces

Average of traces

DPA - Correct Key Guess

Hypothesis:
The 8 key bits
are **00001111**

Predicted LSB(I)

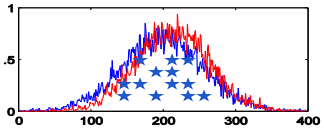

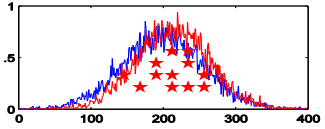

Traces with predicted LSB = 1

Traces with predicted LSB = 0

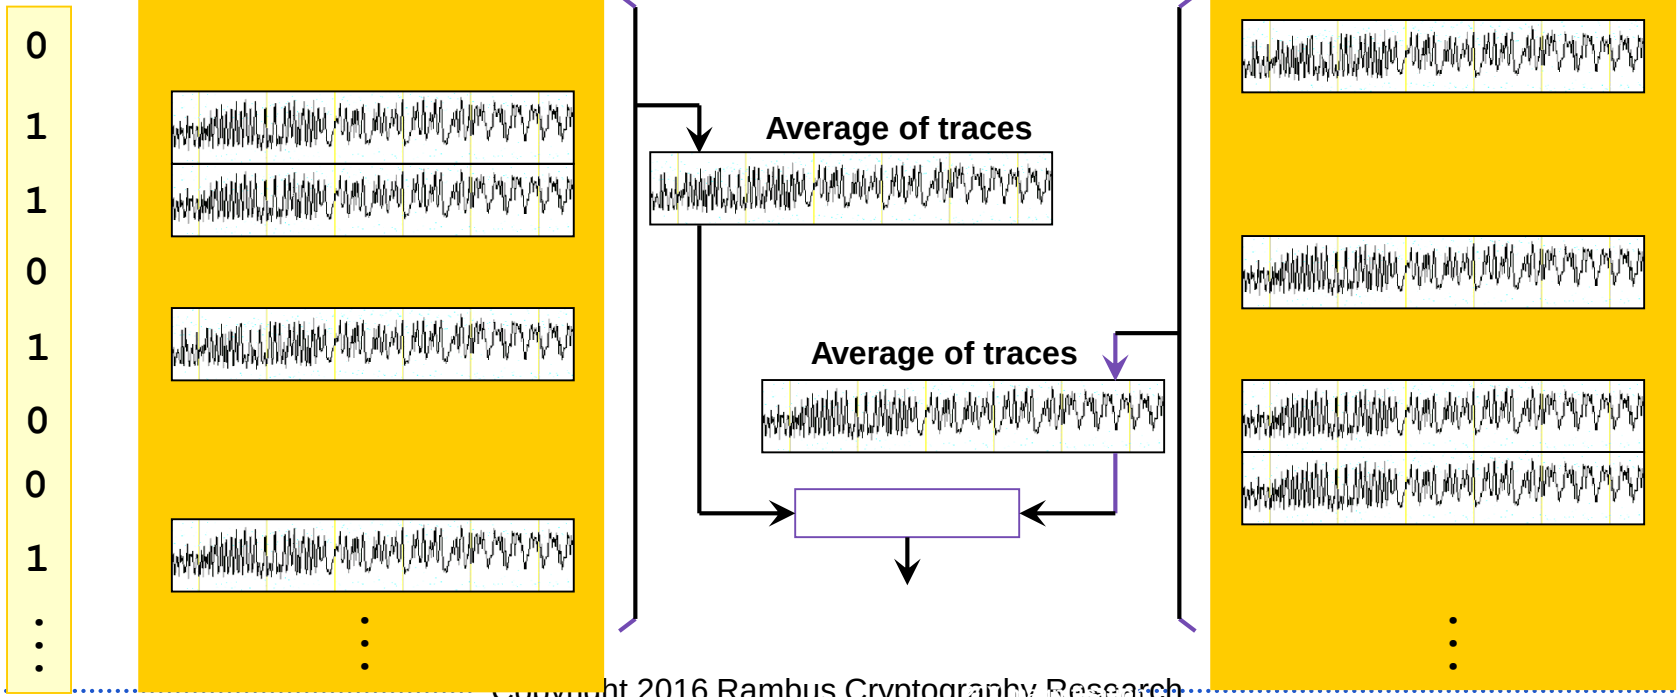

DPA - Correct Key Guess

Hypothesis:
The 8 key bits
are **00001111**

Predicted LSB(I)

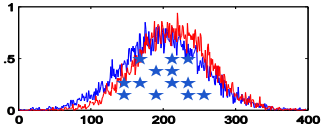

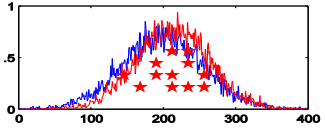

Traces with predicted LSB = 1

Traces with predicted LSB = 0

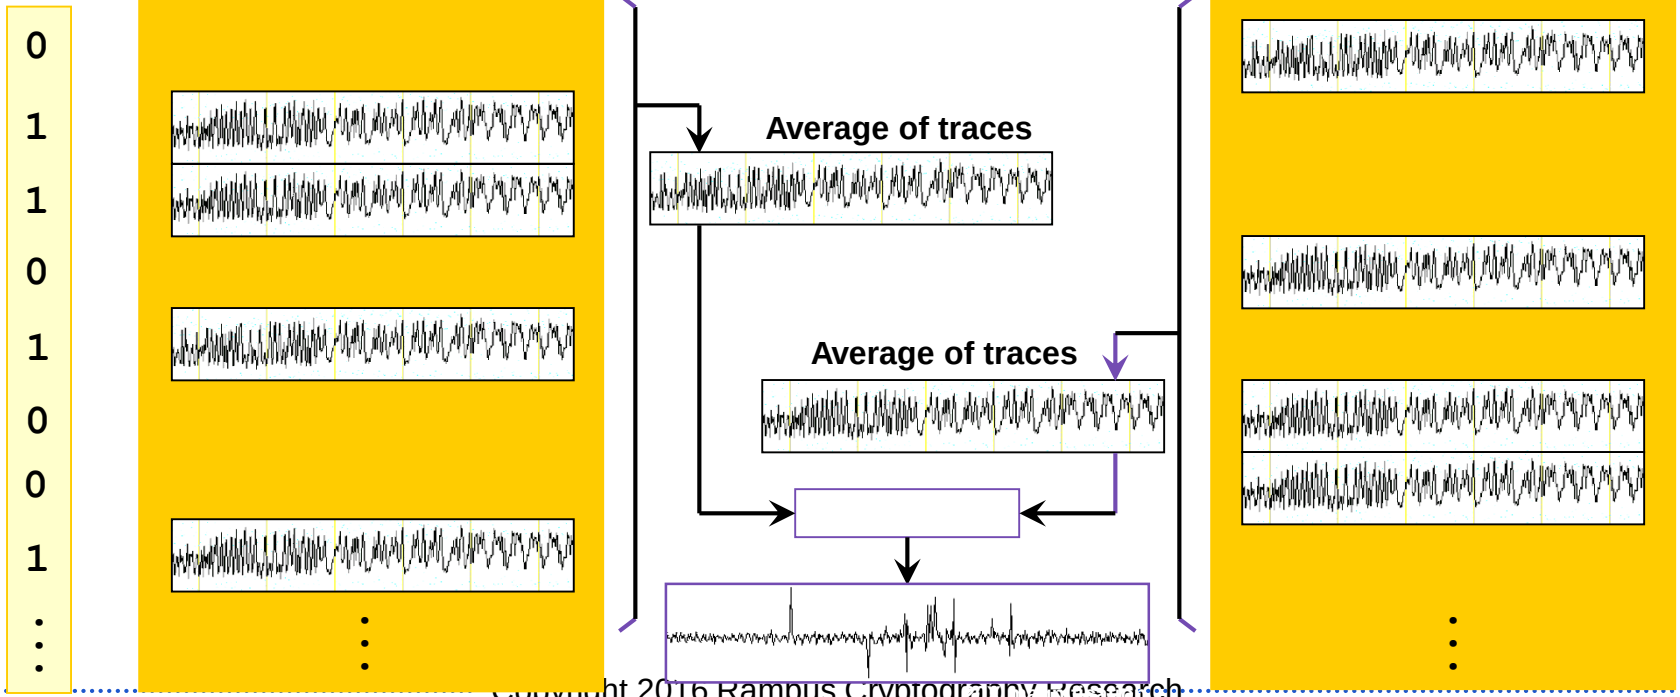

DPA - Incorrect key guess

Hypothesis:
The 8 key bits
are 00010000

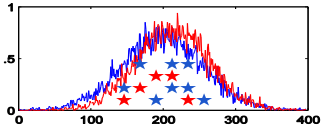

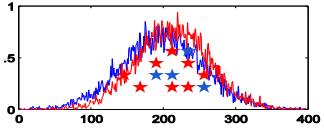

Predicted LSB(I)

Traces with predicted LSB = 1

Traces with predicted LSB = 0

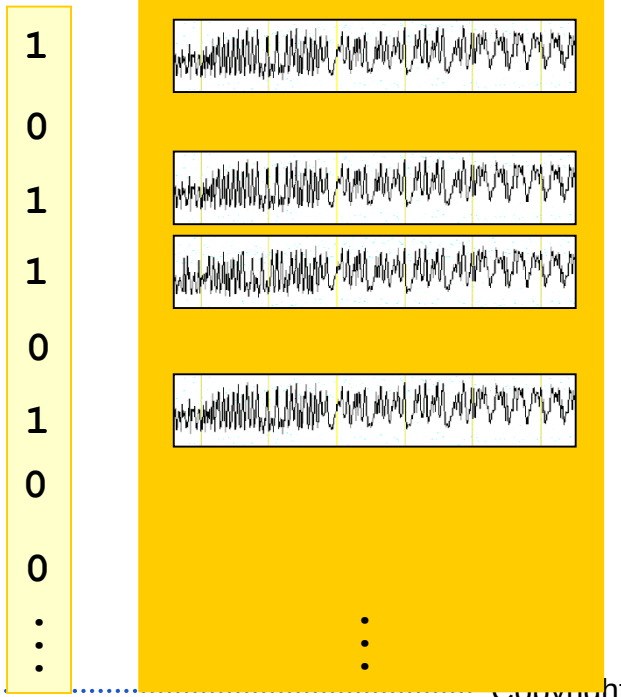

DPA - Correct Key Guess

Hypothesis:
The 8 key bits
are 00010000

Predicted LSB(I)

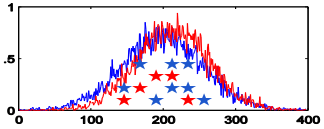

Traces with predicted LSB = 1

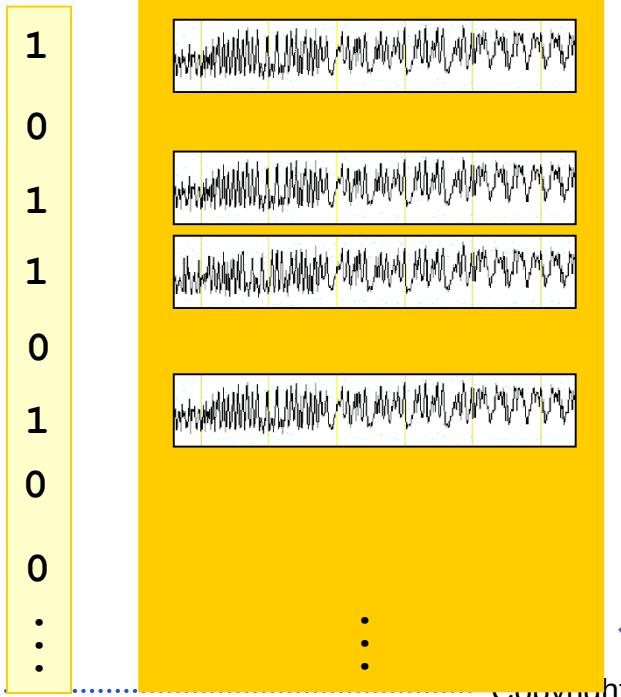

Average of traces

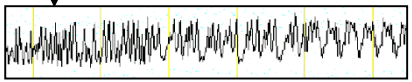

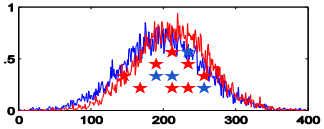

Traces with predicted LSB = 0

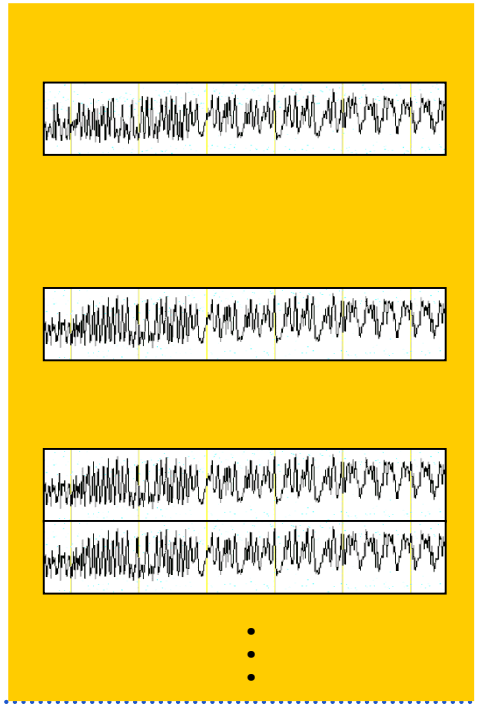

DPA - Incorrect Key Guess

Hypothesis:
The 8 key bits
are 00010000

Predicted LSB(I)

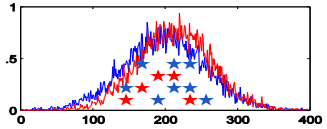

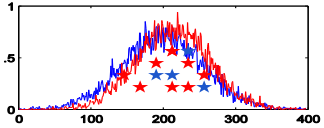

Traces with predicted LSB = 1

Traces with predicted LSB = 0

1
0
1
1
0
1
0
0
⋮

Average of traces

Average of traces

DPA - Incorrect Key Guess

Hypothesis:
The 8 key bits
are 00010000

Predicted LSB(I)

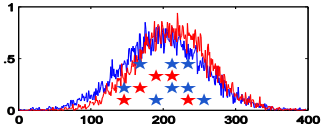

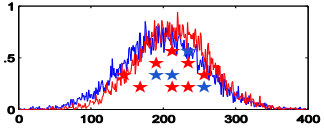

Traces with predicted LSB = 1

Traces with predicted LSB = 0

1
0
1
1
0
1
0
0
⋮

Average of traces

Average of traces

DPA - Incorrect Key Guess

Hypothesis:
The 8 key bits
are 00010000

Predicted LSB(I)

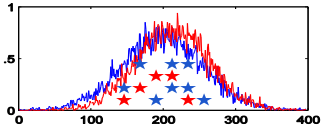

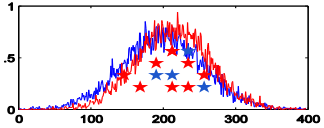

Traces with predicted LSB = 1

Traces with predicted LSB = 0

1
0
1
1
0
1
0
0

Average of traces

Average of traces

Side-Channel Testing Methodology - Roadblocks

- Side-channel analysis has been studied extensively since the late 90's, but….
- Development of formal testing standards and criteria has been lacking
  - Most testing expertise is found in European labs designed to evaluate smart cards
  - US vendors and labs seem to prefer verification-style testing rather than evaluation-style testing favored by Europeans
- Lack of good testing methodology has lead many to ignore the issue
  - Lack of good side channel requirements in standards (e.g. FIPS 140-2)
  - Waivers for DoD contractors in their systems
- Chicken and egg problem
  - Reluctance to levy requirements until there are side-channel capable labs
  - Labs don't want to invest in side-channel capabilities until there are requirements

Side-Channel Testing Methodology – What's needed

- Objective, measurable, reproducible tests for each cryptographic algorithm
- Tests must be efficient and low cost
- Should not require exceptionally skilled testers
  - Could be performed by designers
- Tests should provide good coverage and results be reasonable indicators of resistance achieved
- Failed tests should provide feedback to designer about what went wrong
- Easy to extend to cover new side-channel attacks
- Define requirements in terms of documenting countermeasures and passing tests
- Core idea: Focus on information leakage, not key extraction

Side-Channel Testing Methodology – Key Statistic

- Statistical test: Welch's t-test for significance of "difference of means"

$$t(I) = \frac{X_A(I) - X_B(I)}{\sqrt{\frac{S_A^2(I)}{N_A} + \frac{S_B^2(I)}{N_B}}}$$

- Each test compares two subsets of collected traces
  - Targets sensitive computational intermediates
  - Intermediates will be different if the implementation not properly protected
  - Statistically significant difference between subsets → sensitive information leakage → device fails
- Test performed twice on two independent data sets
  - Failure must occur at the same time-instant in both tests

Side-Channel Testing Methodology – Pass/Fail Criteria

- Device fails if t-statistic exceeds +/- 4.5 for two independent data sets
- For large data set, a single excursion implies confidence of 99.999%
- When test is repeated, probability of excursion occurring randomly at the same point in time is negligible
- Compensates for large trace sizes and large number of tests
- Method called Test Vector Leakage Assessment

AES Example – Test Vectors

- Targeted leakage tests
  - Tester chooses a middle round R to examine
  - Targets 5 specific leakages
    - XOR of round input and output
    - S-box output for round
    - Round output
    - Byte analysis of 1st byte of round output
    - Byte analysis of 2nd byte of round output
- Non-specific leakage test
  - Fixed vs. varying data
  - Examine middle third of operation

AES Example – FPGA Implementations

- Evaluated three different AES implementations on FPGA using the TVLA methodology
  - No DPA countermeasures
  - Incomplete DPA countermeasures
  - Full DPA countermeasures

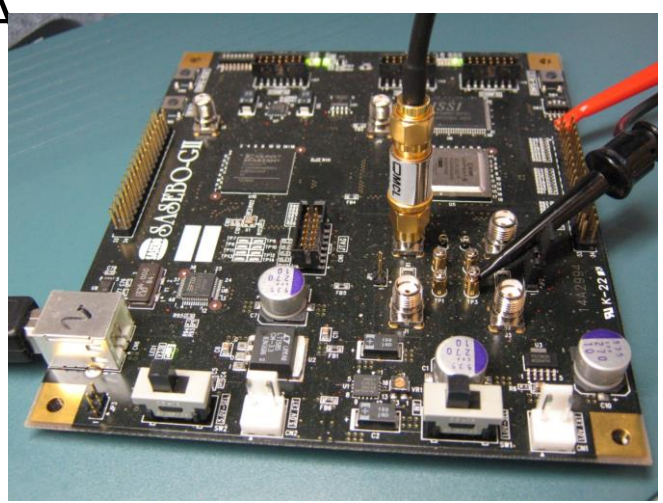

- For each implementation:
  - Definitive result (PASS/FAIL)
  - Less than 24 hours of data-collect + analysis

AES Example - No Countermeasures

- Automated data collection
  - 20 traces/second
  - Early exit condition was reached with 50 minutes of data collection
- Automated application of statistical criteria
  - t-tests detected failure for multiple criteria
  - Failure criteria reached before full data collect

- Result is a **definitive FAIL**
- Total time: 50 minutes collect + 12 minutes analysis = **62 minutes**

AES Example - No Countermeasures

- Example failing test: Round 4 output 1st byte, value 254
- Both subgroups exceed t = ±4.5 at same point in time
- Many other tests, rounds, values fail similarly

t significance thresholds

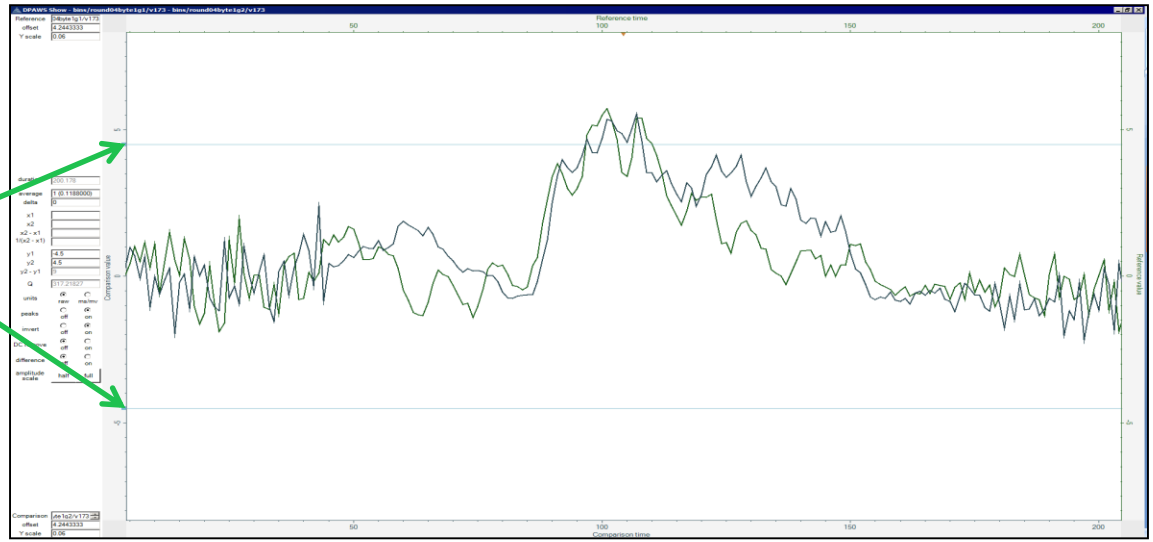

AES Example - Incomplete Countermeasures

- Automated data collection
  - 20 traces/second
  - Bulk ECB encryption allows 10000 ops/2 minutes
  - Overnight data collect using bulk ECB mode: 3 million AES ops
- Countermeasure not fully effective

- Result is a **definitive FAIL**
  - Passed all specific leakage tests
  - Failed non-specific Fixed vs. Random test
- Less than 24 hours data collect + analysis

AES Example - Incomplete Countermeasures

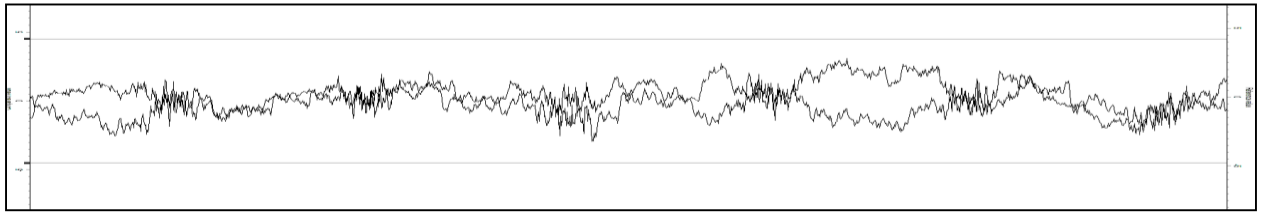

Typical specific test : PASS

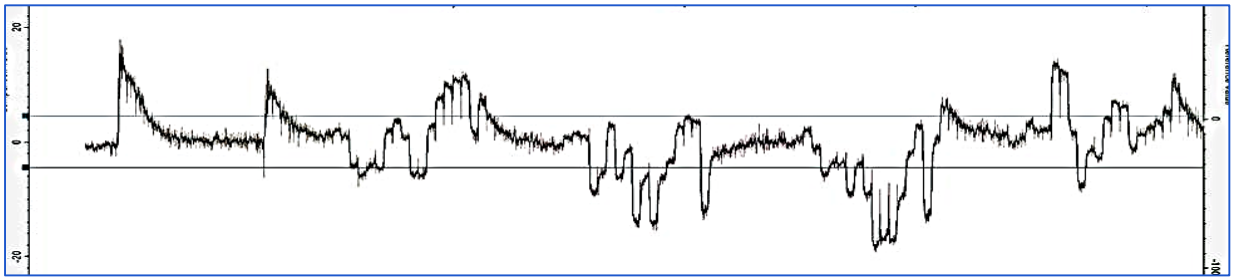

Fixed vs. Random test : FAIL

AES Example - Effective Countermeasures

- Automated data collection
  - 20 traces/second
  - Bulk ECB encryption allows 10000 ops/2 minutes
  - Overnight data collect using bulk ECB mode: 10 million AES ops
- Countermeasure fully effective for over 10 AES ops

- Result is a **definitive PASS**
  - Passed all specific leakage tests
  - Passed non-specific Fixed vs. Random test
- Less than 24 hours data collect + analysis

AES Example - Effective Countermeasures

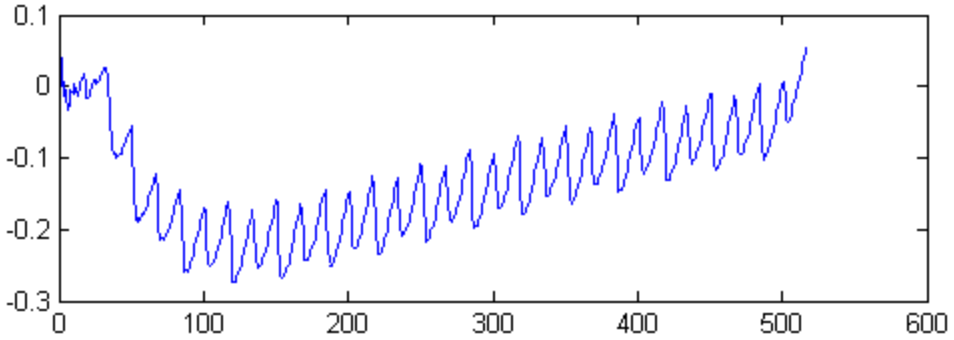

Typical power trace

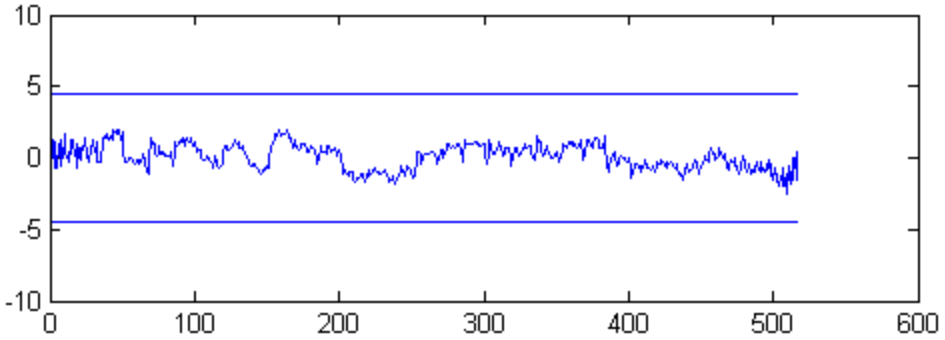

Fixed vs. Random test : PASS

AES Operation Without Countermeasures

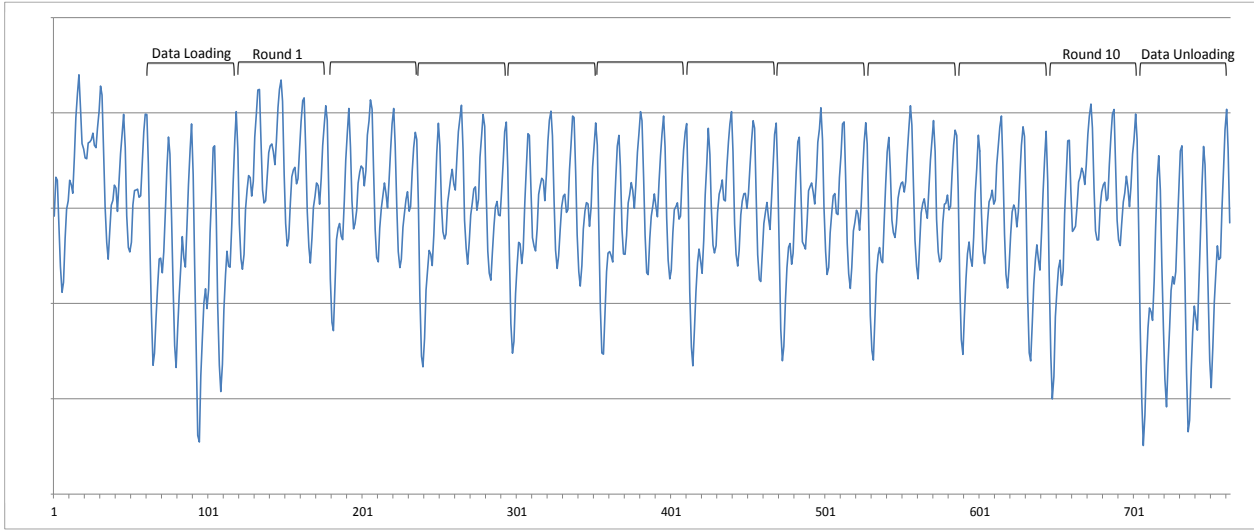

AES without DPA Countermeasures:
Fails TVLA in <100 AES Blocks (Traces)

Courtesy of
Athena Group

AES Operation With Leakage (SNR) Reduction Countermeasures

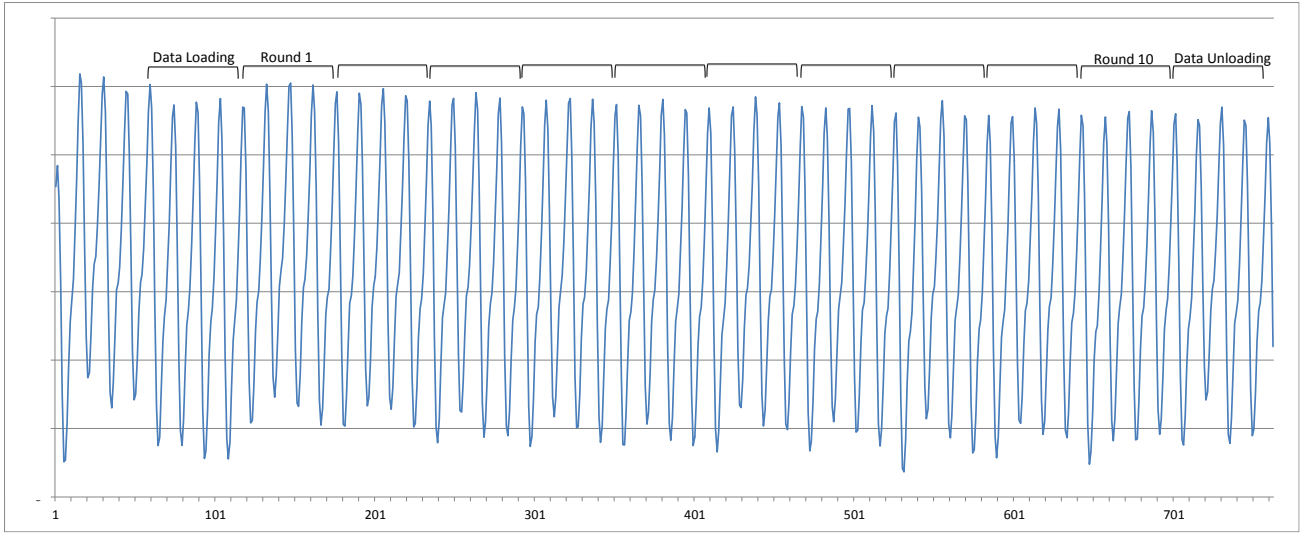

AES with DPA Countermeasures:
Passes TVLA with > 100 Million AES Blocks

Courtesy of
Athena Group

AES:
Countermeasures vs. No Countermeasures

No Countermeasures, 256 AES Blocks vs.
Countermeasures Enabled, 1 Billion AES Blocks vs.
Statistical Leakage Threshold

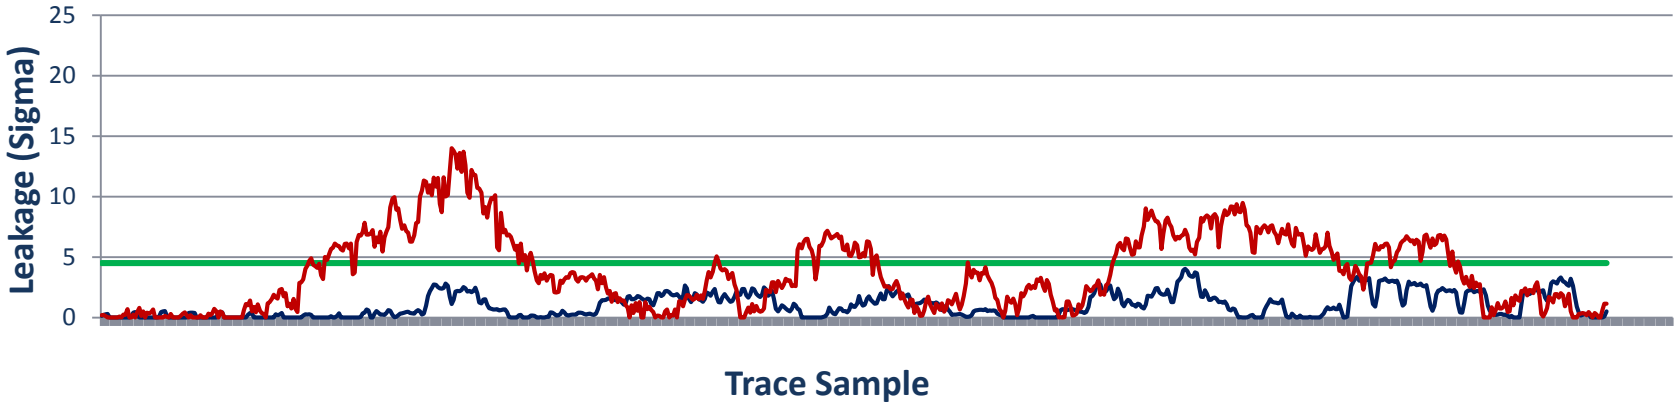

Unprotected cryptography is provably vulnerable.
Athena cryptography is provably resistant.

Courtesy of
Athena Group

Demo

- Live TVLA demo with unprotected and protected AES implementations

Summary

- Side-channel testing need not be difficult or costly
  - High quality, efficient and cost effective testing is possible with modest test operator skill
  - Testing can be performed by designers, enabling quick feedback
- TVLA methodology features
  - Test vectors designed by side-channel experts
  - Focus on measuring information leakage instead key recovery
  - Standardized statistical scoring with objective pass/fail criteria
- Next steps
  - TVLA gaining traction in vendors, labs, academia, government
  - Standardize?