ECE 3170 Cryptographic Hardware for Embedded Systems

Fall 2024

Assoc. Prof. Vincent John Mooney III
Georgia Institute of Technology

Homework 7, 85 pts.

Due Friday Oct. 18 prior to 11:55pm

(please turn in homework electronically on Canvas)

As always this semester, you are required to solve this homework alone.

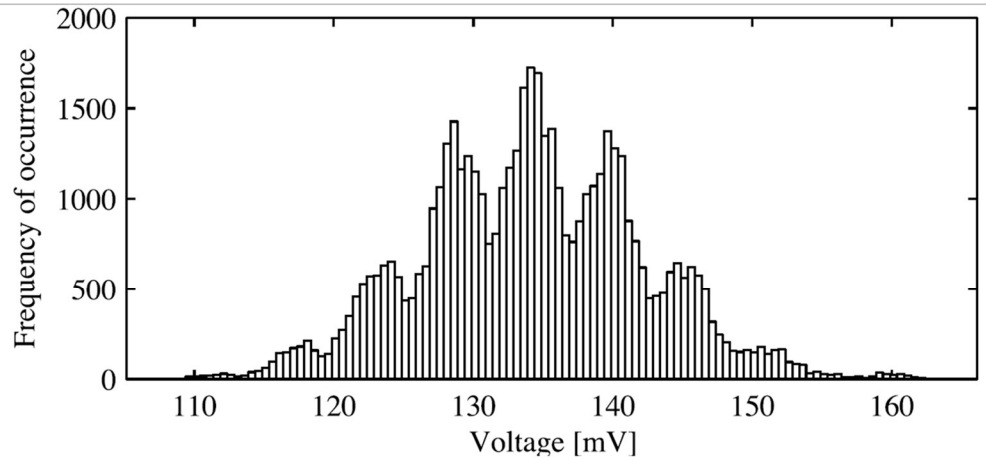

Figure 4.4. Histogram of the power consumption at 362 ns if different data values are transferred from the internal memory to a register.

1) (20 pts.) Explain using your own words what are the nine different Gaussian distributions which compose the energy consumption histogram above. Explain how and why the nine distributions are combined as well as the source of each Gaussian distribution. Include in your explanation the following terms and concepts: $\{D_i\} = D_0 \ldots D_{255}$ = 8-bit data values; HW = Hamming Weight (a.k.a. Hamming Distance from all zeros); μ = mean; and σ = standard deviation. Please define and explain any additional terms you utilize. You may refer to the figure on the next page, Figure 4.5, in your answer to explain the power histogram shown in Figure 4.4 above.

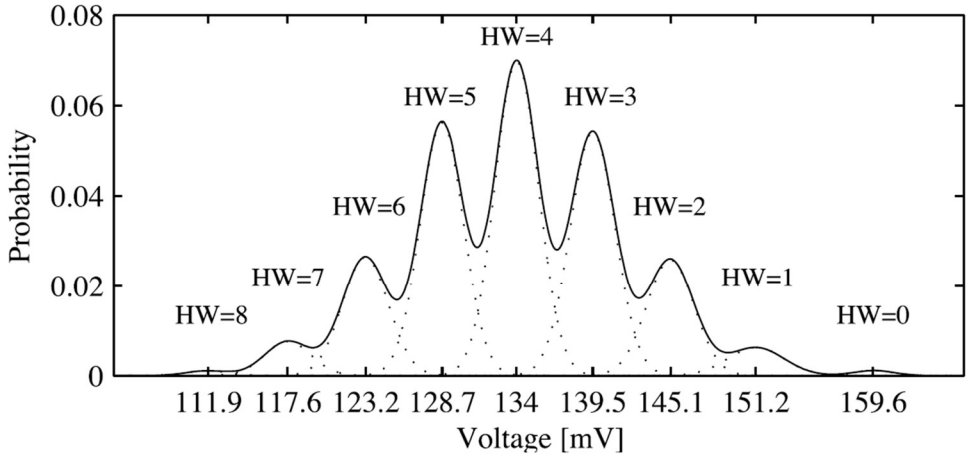

Figure 4.5. The distribution of the power consumption when the microcontroller transfers different data from the internal memory to a register.

At 362 ns, a memory operation is being performed by the microcontroller whose power is being measured. The energy consumption variation is dominated by the Hamming Distance (HD) of the data value D_i from 0xFF (all ones) because the 8-bit data bus used by the memory operation is precharged to all ones. An analysis of the statistical distribution of the data values $D_0 \ldots D_{255}$ (i.e., from 0x00 to 0xFF) shows a binomial distribution of the nine Hamming Weight (HW) values from zero to eight; hence, each collection of data values corresponding to a particular HW form a separate Gaussian weighted by the number of occurrences of that HW (e.g., for HW of 1, the data values are D_1, D_2, D_4, D_8, D_{16}, D_{32}, D_{64} and D_{128}). Note that the smallest HD from all ones (0xFF) is for the HW of 8 (recall that HW is measured from all zeros, i.e., 0x00), hence the Gaussian associated with HW 8 has the least mean (μ) in Figures 4.3 and 4.4. Similarly, the largest mean (μ) is for HW 0, while HW 4 has mean (μ) in the middle.

Finally, since the standard deviation (σ) is due to noise, all of the nine Gaussians have the same standard deviation (σ).

2) (25 pts.) Consider Figures 4.1 and 4.2 from Chapter 4 of *Power Analysis Attacks: Revealing the Secrets of Smart Cards* by Mangard et al. The point of this question is qualitative, not quantitative.

----------GRADED EFFORT/NO EFFORT----------

Please refer to the solutions below to ensure you understand the material

a. (5 pts.) Explain in your own words what the five power traces shown in Figure 4.1 indicate.

The five power traces in Fig. 4.1 indicate that the energy consumption or power trace for the same instructions with the

same data are nearly equal except for slight variations due to noise.

b. (10 pts.) Use your own words to describe the experiment being carried out in Figure 4.2. Make sure that the following items appear in your answer: (i) $P_{el.\ noise}$, (ii) histogram and (iii) 10,000.

Fig. 4.2 shows a histogram of 10,000 power trace measurements at 362 ns. Since exactly the same program (AES) with exactly the same data is executed each time, the information in the histogram is used to calculate $P_{el.noise}$. In other words, the differences between the power traces are not due to the operation being performed nor to the data being processed. Therefore, the differences are accounted for by attributing them to noise sources such as thermal noise.

c. (5 pts.) In your answer to the previous question, 2.b, you used the quantity 10,000. What if this quantity were changed to 100,000: would your answer to 2.b be the same or different? Why or why not? Please note that a correct "yes" or "no" answer without a proper reason will receive zero points.

If the number of power traces used in Fig. 4.2 were 100,000 instead of 10,000, then Fig. 4.2 should look nearly identical except that the y-axis would range from 0 to 8,000 (instead of ranging from 0 to 800). The answer to 2.b would be the same, i.e., the histogram can be used to accurately estimate $P_{el.\ noise}$.

d. (5 pts.) In your answer to question 2.b, you used the quantity 10,000. What if this quantity were changed to 1,000: would your answer to 2.b be the same or different? Why or why not? Please note that a correct "yes" or "no" answer without a proper reason will receive zero points.

If the number of power traces used in Fig. 4.2 were 1,000 instead of 10,000, then it is not clear how Fig. 4.2 would appear. If 1,000 measurements are enough to properly sample the noise, then the version of Fig. 4.2 based on 1,000 power traces would look nearly identical and the answer to question 2.b would be the same. However, if 1,000 traces are not enough, then there may be significant variations and differences due to not enough samples being taken, and thus the answer to 2.b would not be the same.

NOTE: exact statistical techniques and methods to determine when are "enough" samples taken is beyond the scope of this course (at least, as it is being taught this semester, Fall 2024).
However, the fact that there exist sample sizes which are "not enough" versus "enough" to be accurate is within the scope of ECE 3170 CHES.

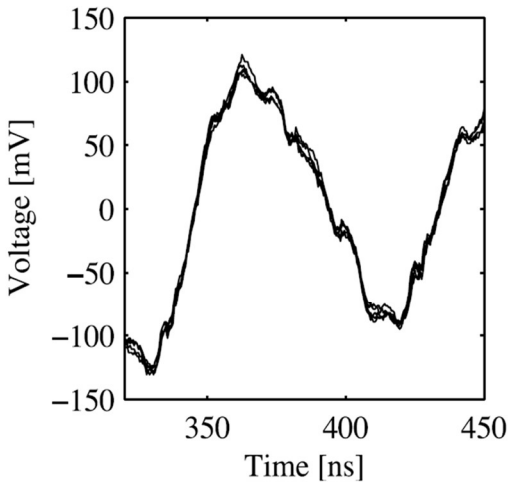

Figure 4.1. Power traces look very similar if the same data is processed.

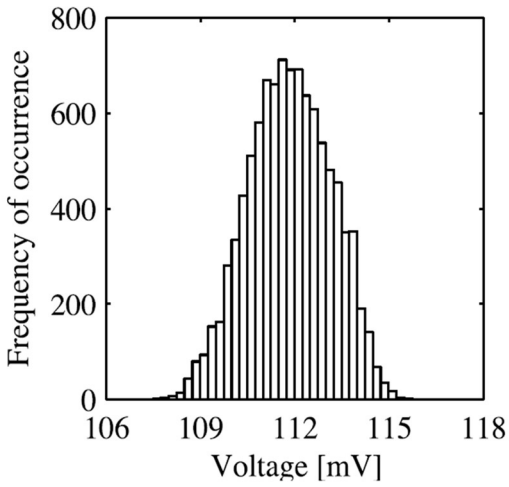

Figure 4.2. Histogram of the power consumption at 362 ns of Figure 4.1.

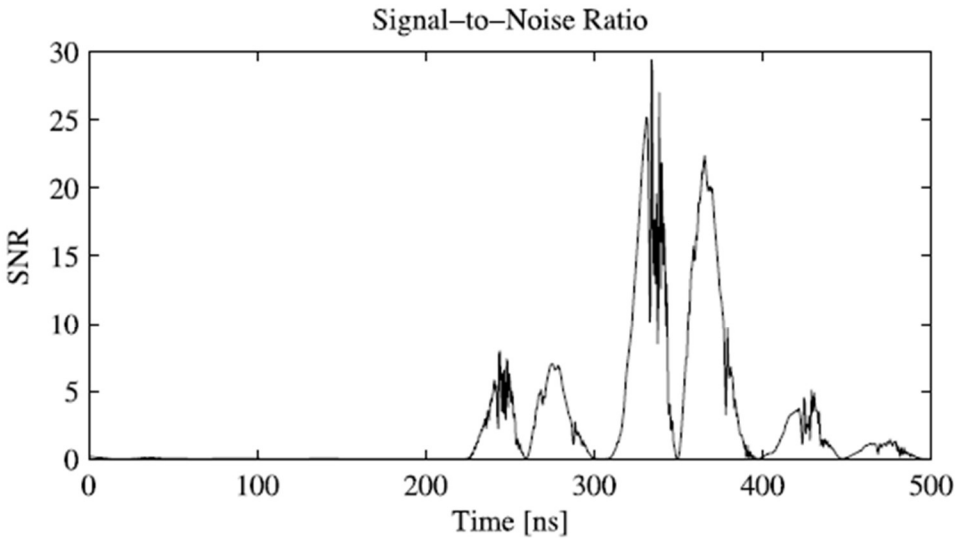

Figure 4.7. The signal levels, the standard deviation of the noise, and the SNR when attacking a uniformly distributed 8-bit data value on our microcontroller.

3) (25 pts.) Consider the third plot of Figure 4.7 from Chapter 4 of *Power Analysis Attacks: Revealing the Secrets of Smart Cards* by Mangard et al. The third plot of Fig. 4.7 shows the resulting Signal-to-Noise Ratio (SNR) for a uniformly distributed 8-bit data value used as a plaintext input to AES encryption with a secret key.

   a. (5 pts.) Write a proper expression for the Signal-to-Noise Ratio (SNR) calculation used to produce the y-axis in the third plot of Figure 4.7. Make sure that the following three variables appear in your answer: P_{exp}, $P_{switching}$ and $P_{el.\,noise}$.

$$SNR = \frac{Var(Signal)}{Var(Noise)} = \frac{Var(P_{\exp})}{Var(P_{switching} + P_{el.noise})}$$

b. (5 pts.) Prior to 200 ns, the SNR in the third plot of Figure 4.7 has a value of approximately zero. Why is this the case and what does this SNR value of zero indicate?

The SNR in Fig. 4.7 prior to 200 ns is zero because the magnitude of P_{exp} is approximately zero. P_{exp} is zero prior to 200 ns because there are no exploitable instructions and/or data being executed; most likely the first 200 ns execute set-up code not specific to the cryptographic algorithm being executed after 200 ns.

c. (15 pts.) In class it was stated that the two largest SNR values in the third plot of Figure 4.7 occur at 330 ns and 362 ns. For the purposes of using power as a side-channel attack, why are the measurements taken at 362 ns preferred over the even larger SNR measurements at 330 ns? You are required to answer this question in three parts as follows:

i. (5 pts.) What is the goal of the attacker, i.e., what information is the attacker trying to gain?

The attacker is trying to figure out the value of the key.

ii. (5 pts.) What assumption is made regarding attacker access to the microcontroller assembly code running AES, and why is this assumption reasonable? To receive full credit for this question, you must give at least one valid explanation for why this assumption is reasonable – but do please note that there are many such equally valid reasons; one good reason is enough for full credit.

The assumption is that the attacker has full access to the assembly code running AES. This is perfectly reasonable since AES is a standard used worldwide.
Furthermore, in a typical system, the AES code is stored in a nonvolatile storage medium such as flash memory or a hard disc drive; such nonvolatile storage can typically be read from independent of the system operation, and hence the specific AES assembly code can be accessed.

iii. (5 pts.) Now, in the context of your answers to 3.c.i and 3.c.ii above, explain why measurements taken at 362 ns are preferred for use by the attacker over the SNR measurements at 330 ns even though the measurements taken at 330 ns have a larger SNR for our quantities of interest (according to the answer to 3.a) as compared to the measurements taken at 362 ns.

As explained in lecture (and in the assigned readings), the

code at 362 ns executes an SBOX operation based, in part, on 8 bits of the key value. Therefore, the information gleaned from the power pin on the microcontroller can potentially be used to obtain information relevant to discovering 8 bits of the key value.

The code at 330 ns is not directly related to the key value and thus is less helpful.

NOTE: a number of students said that at 330 ns the SNR was somehow significantly more difficult to utilize for DPA. This is not true at all! In fact, in lecture Figure 4.8 with an SNR of around 0.15 (instead of between 10 and 30 as shown in Figure 4.7 at 330 ns) was introduced as sufficient for DPA.

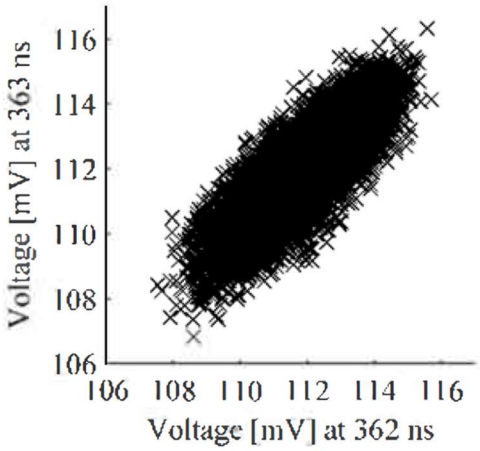

Figure 4.9. Scatter Plot: The power consumption at 362 ns is correlated to the power consumption at 363 ns. $r = 0.82$

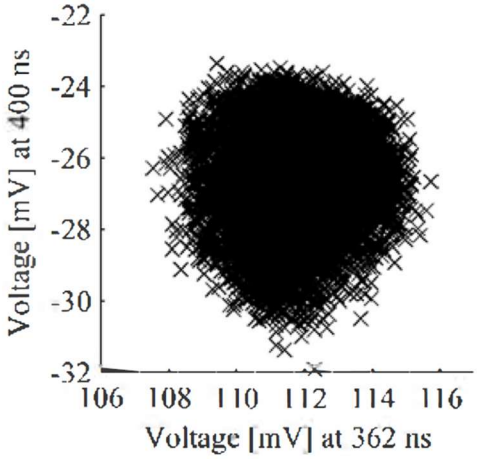

Figure 4.10. Scatter Plot: The power consumption at 362 ns is largely uncorrelated to the power consumption at 400 ns. $r = 0.12$

4) (15 pts.) Consider Figures 4.9 and 4.10 from Chapter 4 of *Power Analysis Attacks: Revealing the Secrets of Smart Cards* by Mangard et al. The point of this question is qualitative, not quantitative.

   a. (5 pts.) Explain the units used in the *x* and *y* axes of Figures 9 and 10. What do the millivolts represent and why is the representation valid?

The millivolts are measured across a resistor connected to the power supply pin to a microcontroller. The millivolts are supposed to represent power. Power equals current times voltage ($P=IV$), so if the current provided by the power supply circuitry is approximately constant, then $P \cong V$.

Furthermore, $V = IR \Rightarrow I = V/R \Rightarrow P = IV = V^2/R$. From $P = V^2/R$ we see that power is proportional to the square of the voltage.

Keep in mind that the units of power are Joules per second, so it is not immediately obvious that $J/s \cong$ Volts.

b. (10 pts.) Explain in your own words the difference between Figures 4.9 and 4.10 in terms of what the plots indicate about the relationship between the measurements at 362 ns as compared to 363 ns versus compared to 400 ns. In your answer, make sure to use the words "covariance" and "correlation" with their appropriate technical meanings clear from the context and usage.

The scatter plot shows two variables which may (or may not) be correlated. Figure 4.9 shows that the energy consumption (power) at 362 ns is in fact correlated to the energy consumption at 363 ns; in particular, as the measurement at 362 ns increases and approaches 116 mV, similarly the energy consumption at 363 ns also increases and the measurement approaches 116 mV as well.

On the other hand, the energy consumption at 400 ns appears to be uncorrelated to the energy consumption (power) at 362 ns.
Figure 4.10 shows, for example, no significant difference in voltage distribution in the y-axis (representing energy consumption or power at 400 ns) when the x-axis (representing power at 362 ns) is high (e.g., 116 mV) versus low (e.g., 108 mV). The calculation of covariance is shown on the bottom right of each figure. A covariance of zero or close to zero indicates no correlation, while a covariance of 1 or close to one indicates a high level of correlation. Clearly, the covariance coefficient of $r = 0.12$ for Figure 4.10 indicates a low level of correlation. Thus, while there does not appear to be any correlation, the covariance calculation does appear to show some type of small correlation.

YOU MAY NOT CONSULT HOMEWORK SOLUTIONS OF THESE EXACT PROBLEMS FROM OTHER COURSES. ALL LABORATORY SUBMISSIONS MUST INCLUDE YOUR NAME, COURSE NUMBER, SECTION, AND THE HOMEWORK SET NUMBER. ALL SHEETS MUST BE STAPLED TOGETHER. ALL WRITING MUST BE EASY TO READ (FOR EXAMPLE, YOU MAY HAVE TO WRITE ONLY ON ONE SIDE OF EACH SHEET OF PAPER THAT YOU SUBMIT AND MAY NOT BE ABLE TO USE RECYCLED PAPER). FAILURE TO FOLLOW INSTRUCTIONS MAY RESULT IN ZERO POINTS. ALL WORK MUST BE YOUR OWN. NO PLAGIARISM IS ALLOWED, AND YOU MUST PROPERLY REFERENCE ALL SOURCES OF YOUR INFORMATION – ALTHOUGH YOU SHOULD NOT LOOK FOR AND MAY NOT CONSULT "SOLUTIONS" AVAILABLE FROM OTHER SOURCES (TO REPEAT, YOU MAY NOT CONSULT HOMEWORK SOLUTIONS OF THESE EXACT PROBLEMS FROM OTHER COURSES!).