3.5 Secure AES Engine with A Local Switched-Capacitor Current Equalizer

Carlos Tokunaga, David Blaauw

University of Michigan, Ann Arbor, MI

Hardware implementations of the popular AES encryption algorithm [1,2] provide attackers with important side-channel information (delay, power consumption or EM radiation) that can be used to disclose the secret key of the encryption device. Differential power analysis (DPA) [3-5] is one of the most common side-channel attacks because of its simplicity and effectiveness (Fig. 3.5.1). It performs a statistical analysis of supply-current measurements and either the plaintext or ciphertext to disclose the secret key. These two elements can be easily recorded externally without probing internal signals on the chip. Either the plaintext or ciphertext is used to build a model of the current consumption (e.g., during 0 to 1 transition) using knowledge of the AES algorithm and a key guess. By calculating the correlation between the model and the measured current for each possible key guess the key is discovered. In the AES algorithm, the key consists of 16 blocks of 8b, each of which can be attacked independently since AES is a block cipher. For the 128b secret key, the DPA search space is only 16×2^8, as opposed to 2^{128} for a brute-force attack.

Key disclosure has been demonstrated in both ASIC and processor AES implementations [6,7]. It has been demonstrated in [1,8] that masking the current signature by adding additional random noise or randomly adding dummy instructions only moderately increases the number of measurements to disclosure (MTD). A charge-pump circuit that provides current to a DES engine was presented in [9] but its MTD is not reported. In [10], an ASIC implementation that increases the security of the secret keys is demonstrated. This implementation uses dual-rail logic and completely balanced interconnect to equalize the current in rising and falling transitions. The method raises the MTD of the first block of the secret key to 20k runs. However, it incurs a 3× area, 4× power and 4× performance overhead, in addition to significantly changing the design flow and requiring custom synthesis and modified routing algorithms.

This work implements a switched-capacitor block, show in Fig. 3.5.2, that isolates the switching activity by equalizing the current drawn from the encryption core to secure an AES engine. An array of capacitors provides the supply current for the sensitive blocks of the encryption engine while non-sensitive blocks run directly from the supply. Each capacitor is charged from the supply and is then isolated while it provides charge to the encryption core. The key idea is that the capacitor is then discharged to a known voltage before it is recharged in order to equalize the amount of charge provided by the external power supply.

To demonstrate this idea, the current equalizer is implemented in a 128b AES encryption engine in 0.13μm CMOS. The test-chip also contains an unprotected version for comparison. The switched-capacitor block incurs a 7.2% area, 33% power and 2× performance overhead. In addition, the approach allows the designer to use any logic family, follow traditional design flows and protect any encryption engine. Using DPA, the key of the unprotected core is disclosed with a minimum of 4k ciphertexts. To date, the secured core has been subjected to 10M encryptions, 2500× more than the unprotected core, and none of its secret key blocks have yet been disclosed.

The current equalizer block is composed of an array of capacitor modules, show in Fig. 3.5.2. Each capacitor module has 3 different switching states: (S1) replenish charge from the supply, (S2) provide charge for encryption, and (S3) continue discharging to a pre-programmed value. Since there are 3 distinct switching states, a minimum of 3 independent capacitor modules is needed for uninterrupted operation of the encryption core. Different switching configurations operate the different modules with the only restriction that at least one of the modules is in S2 at any given time to provide current for the protected logic.

The switching cycle starts with S1, where the capacitor is charged to a full potential by connecting it directly to the supply and disconnecting it from the core and shunt path. In S2, the capacitor is then disconnected from the supply and connected to the encryption core. At the end of S2, each capacitor contains a variable charge that depends on the encryption activity. If at this point the capacitor is connected to the supply, the amount of current drawn would carry encryption information to the external pads. Instead, in S3, the capacitor is disconnected from the core and is shunted until it reaches a known value. This state enforces that the amount of charge replenished by the supply in S1 is the same each time. During the shunt state, the current loop is local to the capacitor and shunt and charge and is not externally observable. Although the GND net is not isolated from the internal core, the operations in S2 and S3 are locally isolated and any observable current does not carry significant side-channel information.

Figure 3.5.3 shows the current-equalizer block, which is implemented with 3 capacitor modules, sized and measured to support up to 100MHz of encryption frequency and more than 500MHz of capacitance switching frequency. Each capacitor module contains a 100pF capacitor; the supply, core, and shunting switches; and a comparator that monitors the voltage of the capacitor and triggers when the predefined reference level is reached. The comparator voltage and current references are provided externally for experimentation, but can be implemented on-chip. These modules incur a 25% area overhead compared to the unprotected core and 7.2% when considering the area of necessary modules for encryption such as input/output buffers.

The bar graph of Fig. 3.5.4 shows the correlation coefficients of the guessed keys (grey for incorrect, black for correct) for the DPA attack on the unprotected block at 100k encryptions. The right plot shows that the MTD occurs at 4k runs when the correlation coefficient of the correct key becomes higher than the maximum correlation coefficient of all incorrect keys. The measured current transient for an encryption cycle is shown to illustrate the information that attackers have available to them by probing the supply externally.

The same attack applied to the unprotected core as well as a modified attack (targeting possible side-channel leakage from the MixColumns block) is attempted on the protected core but has not disclosed any block of the key at this point. A total of 10M ciphertexts have been tried and Fig. 3.5.5 shows that the correlation of the correct key is substantially lower than that of the unprotected core. When running the encryption core at 1.2V and 100MHz with the current equalizing block operating at 200MHz, the power overhead is 33% compared to the unprotected core. Operating at 1.2V the performance of the unprotected core is 2× that of the protected core. However, if higher performance is required, area overhead could be traded off at design time by increasing the size of the capacitor module's capacitor and switches, or increasing the number of capacitor modules used. Alternatively, a high switching-capacitor frequency can be used to trade off power for performance. Figure 3.5.6 summarizes the features and performance of the the protected and unprotected AES engines and Fig. 3.5.7 shows a die micrograph of the fabricated chip.

Acknowledgements:
We thank Jerry Kao, Sanjay Pant and Visvesh Sathe for their help. Funding is provided in part by ARM Ltd.

References:

[1] *Advanced Encryption Standard (AES)*, FIPS 197, National Institute of Standards and Technology, Nov. 2001.

[2] J. Daemen, V. Rijmen, *The Design of Rijndael (AES - The Advanced Encryption Standard)*, Springer, 2002.

[3] P. Kocher, J. Jaffe, B. Jun, "Differential Power Analysis", *Proc. Advances in Cryptology (CRYPTO 1999)*, pp. 388-397, 1999.

[4] T. Messerges, "Using Second-Order Power Analysis to Attack DPA resistant Software", *CHES 2000, LNCS 1965*, pp 238-251, 2000.

[5] R. Bevan, E. Knudsen, "Ways to Enhance Differential Power Analysis", *ICISC 2002, LCNS 2587*, pp. 327-342, 2003.

[6] T. Messerges, E. Dabbish, R. Sloan, "Examining Smart-Card Security under the Threat of Power Analysis Attacks", *IEEE Trans. Computers*, vol. 51, no. 5, May 2002.

[7] S. Mangard, T. Popp, B. Gammel, "Side-Channel Leakage of Masked CMOS Gates", *CT-RSA 2005, LNCS 3376*, pp. 351-365, 2005.

[8] C. Clavier, J. Coron, N. Daboous, "Differential Power Analysis in the Presence of Hardware Countermeasures", *CHES 2000, LNCS 1965*, pp. 252-263, 2000.

[9] P. Corsonello, S. Perri, M. Margala, "An Integrated Countermeasure against Differential Power Analysis for Secure Smart-Cards", ISCAS, 2006.

[10] D. Hwang, K. Tiri, A. Hodjat, B. Lai, S. Yang, P. Schaumont, I. Verbauwhede, "AES-Based Security Coprocessor IC in 0.18μm CMOS with Resistance to Differential Power Analysis Side-Channel Attacks", *IEEE J. Solid-State Circuits*, vol. 41, no. 4, 2006.

[11] M. Nagata, T. Okumoto, K. Taki, "A Built-in Technique for Probing Power Supply and Ground Noise Distribution within Large-Scale Digital Integrated Circuits", *IEEE J. Solid-State Circuits*, vol. 40, no. 4, Apr. 2005

978-1-4244-3457-2/09/$25.00 ©2009 IEEE

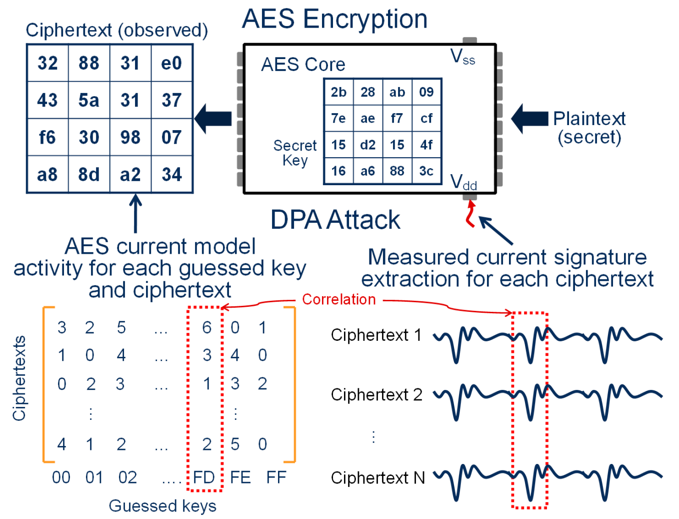

Figure 3.5.1: AES encryption basics and differential power analysis (DPA) side-channel attack.

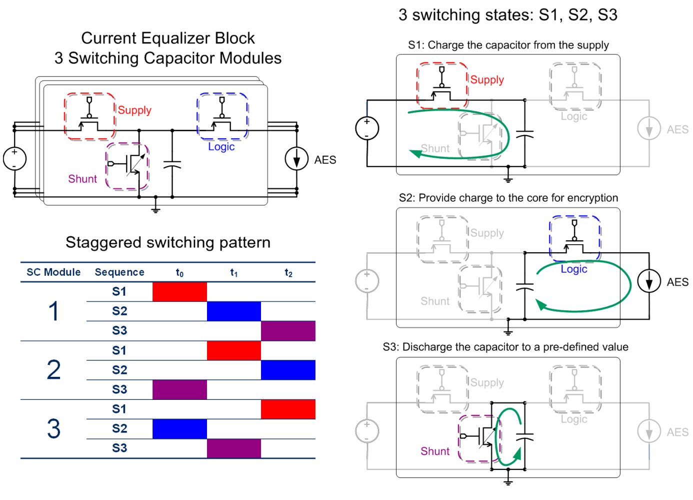
Current Equalizer Block
3 Switching Capacitor Modules
Supply
Logic
Shunt
AES
3 switching states: S1, S2, S3
S1: Charge the capacitor from the supply
S2: Provide charge to the core for encryption
S3: Discharge the capacitor to a pre-defined value
Staggered switching pattern
SC Module
Sequence
t_0
t_1
t_2
1
2
3
S1
S2
S3

Figure 3.5.2: Switching-capacitor current-equalizer block with 3 capacitor modules implemented with logic, supply and shunt switches and a comparator to disable the shunting mechanism.

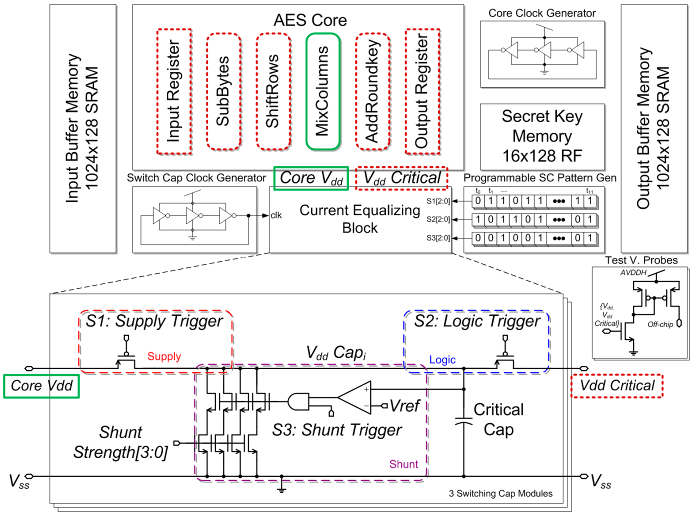

Figure 3.5.3: Circuit diagram and implementation of the current equalizer block and secure AES engine.

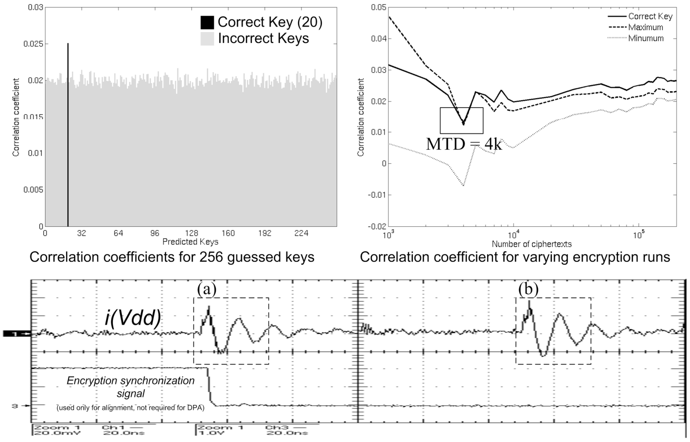

Measured transient response of the current supplied to the core V_{dd} $i(V_{dd})$. The 2 boxes (a,b) highlight 2 different core clock cycles and it is clear that the current information in the transient response is different.

Figure 3.5.4: DPA attack on the unprotected AES engine. Correlation coefficients for 256 possible key values for one block (top left). Correlation coefficient as a function of encryption runs (top right). Transient response of the V_{dd} current during encryption (bottom). Key disclosure occurs at 4k ciphertexts.

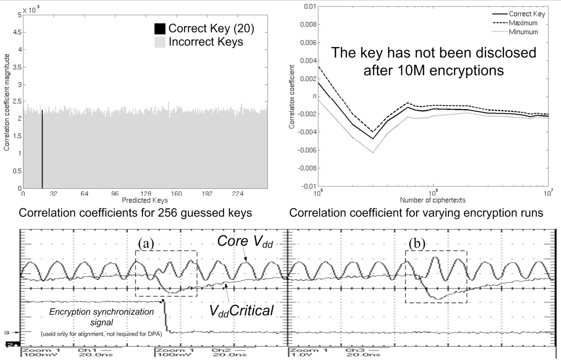

Measured transient response of the *Core* V_{dd} and *Critical* V_{dd} inside the protected AES core. The current equalizer block guarantees that the maximum/minimum peaks of the *Core* V_{dd} are uniform.

Figure 3.5.5: DPA attack on the protected AES engine. The two top graphs show that this block's key is not disclosed after 10M encryptions. The bottom graph shows the transient response during encryption (captured with V-I probes [11]). The Core V_{dd} signal has 100MHz frequency transients due to the switching capacitors. Boxes (a) and (b) show where current for the unprotected non-critical portions of the chip superimpose noise on the constant fluctuations from the equalizer block. VddCritical is the supply voltage at the protected portions of the core and show different behavior from cycle to cycle, as expected due to changing encryption activity.

Parameter	Unprotected	Protected
Area [mm^2]		
AES Core	0.35	0.364
I/O buffers + Clk gen	0.93	0.93
Current equalizer block	-	0.079
Total	1.28	1.37 (+7.2%)
Operating Range (V)	0.6 – 1.2	0.78 – 1.2
Power (1.2V, 100MHz) [mW]	33.32	44.34 (+33%)
Maximum Throughput (1.2V) [Gb/s]	220MHz 2.56	110MHz 1.28 (-50%)
Measurements to Disclosure of 1st block	6k	(not yet disclosed) 10M
Maximum number of blocks disclosed (out of 16)	16	0

Figure 3.5.6: Performance and comparison table between unprotected and protected AES engines.

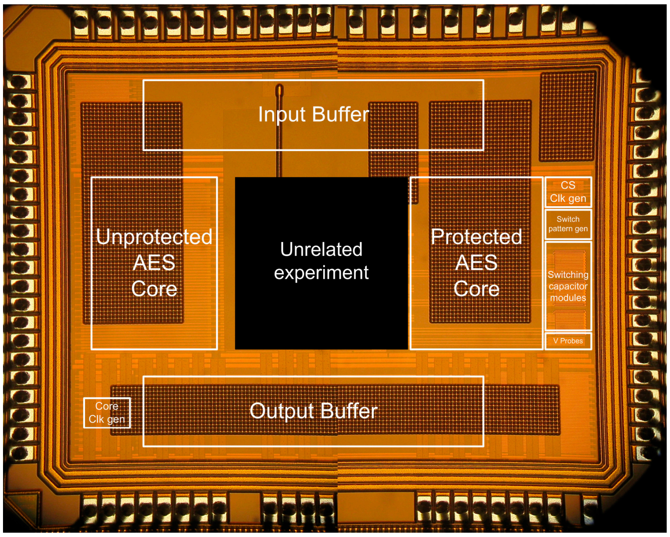

Figure 3.5.7: Die micrograph.