Optimized Threshold Implementations: Number of Shares and Area/Latency Trade-off

Ventzi Nikov, NXP Semiconductors

joint work with Dušan Božilov and Miroslav Knežević

12.03.2019

PRINCE cipher

- The fastest low latency cipher [Borghoff et al. 2012]
- PRINCE is a 64-bit block cipher with a 128-bit key
- PRINCE is based on the so-called FX construction, PRINCE$_{core}$ is 12-round block cipher with a 64-bit key
- PRINCE$_{core}$ has a unique *alpha*-reflection property
- Decryption reuses the encryption circuit

$$D_{(k_0||k_0'||k_1)}(\cdot) = E_{(k_0'||k_0||k_1 \oplus \alpha)}(\cdot)$$

- De facto standard for IoT memory encryption

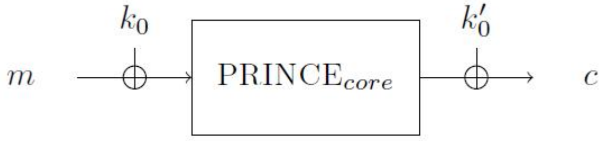

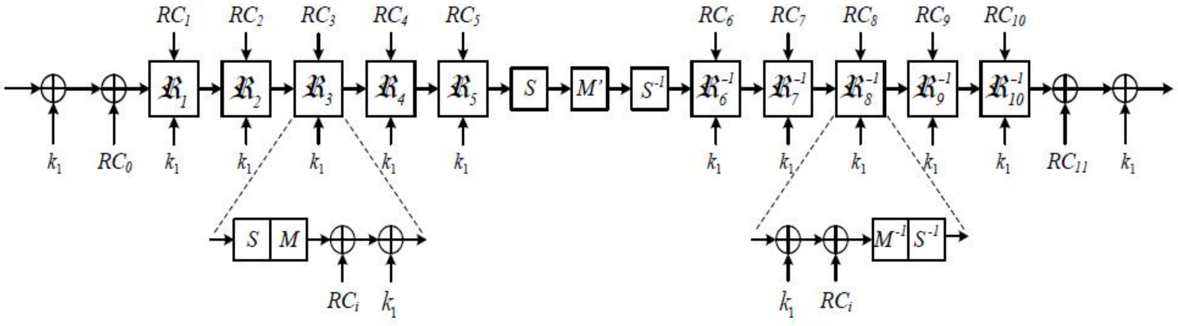

HW implementations

- Compare HW implementations area/power/energy/latency/etc. only when same library and corner case is used!

- [+] A. Moradi, T. Schneider: Side-Channel Analysis Protection and Low-Latency in Action - case study of PRINCE and Midori, ASIACRYPT 2016

- [*] In UMCL18 standard cell library in the typical PVT corner case
  [**] In TSMC90 in the worst PVT corner case i.e. the temperature of +125° C and the supply voltage of 1.0 V

PRINCE - 1st (td+1) TI HDL from [+]	Area [GE]	Clock # [cycles]	Latency [ns]
Min Area [*]	9292	40	160
Min Area [**]	9484 (2%)	40	342 (114%)
Min Latency [*]	11275	40	76
Min Latency [**]	15123 (34%)	40	122 (61%)

Threshold Implementations

- Side-channel power attacks – a problem for IoT devices
- HW side-channel leakage is different than the SW

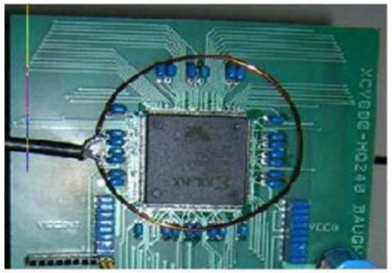

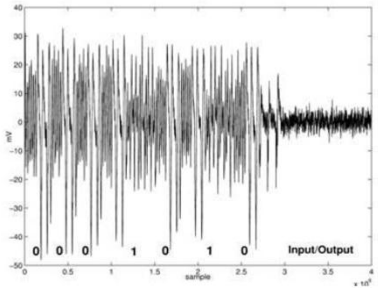

- Provable SW countermeasures [e.g. ISW 2003] can leak when implemented in HW
- TI proposed by Nikova, Rechberger, and Rijmen [2006]
- Provable secure countermeasure against SCA in presence of glitches
- TI main property: non-completeness of the sharing

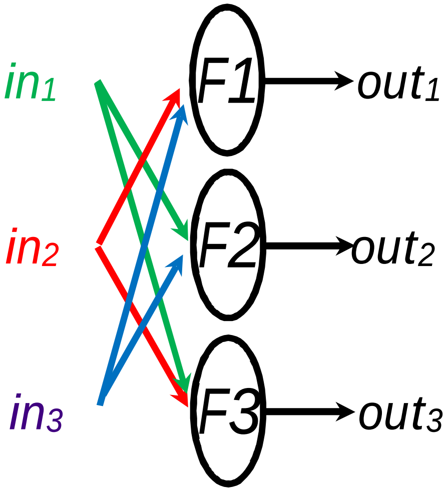

- Many publications followed since then
  - Different ciphers: AES, PRESENT, KECCAK, PRINCE, SHA1, SHA2, etc.
  - Any protection order against SCA
  - Several flavors of TI exist - (td+1) and (d+1)
  - Different optimizations trade offs: (mainly on) area and randomness; (less on) power, latency and energy

State of the art

- SCA resistant 1st order TI for low-latency
- *A. Moradi and T. Schneider: Side-Channel Analysis Protection and Low-Latency in Action - case study of PRINCE and Midori, ASIACRYPT 2016
- ** Our design(s) - how one can achieve a very high level of SCA protection by keeping the latency as low as possible

PRINCE in TSMC90 worst PVT case	Area [GE]	Power [uW]	Energy [pJ]	Rand per cycle [bits]	Clock # [cycles]	f_{max} [MHz]	Latency [ns]	Area [GE]
	@10 MHz						@ f_{max}	
1st (td+1) TI *	9484	66	264	0	40	328	122	15123
1st (d+1) TI **	12220	115	276	112	24	289	83 (47%)	17187
1st (td+1) TI **	31116	576	691	48	12	204	59 (107%)	78281

- Two of our designs achieve better latency

State of the art

- SCA resistant 1st order TI for low-latency
- *A. Moradi and T. Schneider: Side-Channel Analysis Protection and Low-Latency in Action - case study of PRINCE and Midori, ASIACRYPT 2016
- ** Our design(s) - how one can achieve a very high level of SCA protection by keeping the latency as low as possible

PRINCE in TSMC90 worst PVT case	Area [GE]	Power [uW]	Energy [pJ]	Rand per cycle [bits]	Clock # [cycles]	f_{max} [MHz]	Latency [ns]	Area [GE]
	@10 MHz						@ f_{max}	
1st (td+1) TI *	9484	66	264	0	40	328	122	15123
1st (d+1) TI **	12220	115	276	112	24	289	83 (47%)	17187
1st (td+1) TI **	31116	576	691	48	12	204	59 (107%)	78281
Unprotected	3589				1		13 (354%)	27997

- Still compared to an unprotected implementation latency decreases a lot

State of the art

- SCA resistant 1st order TI for low-latency
- *A. Moradi and T. Schneider: Side-Channel Analysis Protection and Low-Latency in Action - case study of PRINCE and Midori, ASIACRYPT 2016
- ** Our design(s) - how one can achieve a very high level of SCA protection by keeping the latency as low as possible

PRINCE in TSMC90 worst PVT case	Area [GE]	Power [uW]	Energy [pJ]	Rand per cycle [bits]	Clock # [cycles]	f_{max} [MHz]	Latency [ns]	Area [GE]
		@10 MHz						@ f_{max}
1st (td+1) TI *	9484	66	264	0	40	328	122	15123 (60%)
1st (d+1) TI **	12220	115	276	112	24	289	83	17187 (41%)
1st (td+1) TI **	31116	576	691	48	12	204	59	78281 (152%)

- Note significant area increase when designs are “pushed” to perform

State of the art

- SCA resistant 1st order TI for low-latency
- *A. Moradi and T. Schneider: Side-Channel Analysis Protection and Low-Latency in Action - case study of PRINCE and Midori, ASIACRYPT 2016
- ** Our design(s) - how one can achieve a very high level of SCA protection by keeping the latency as low as possible

PRINCE in TSMC90 worst PVT case	Area [GE]	Power [uW]	Energy [pJ]	Rand per cycle [bits]	Clock # [cycles]	f_{max} [MHz]	Latency [ns]	Area [GE]
	@10 MHz						@ f_{max}	
1st (td+1) TI *	9484	66	264	0	40	328	122	15123
1st (d+1) TI **	12220	115	276	112	24	289	83	17187
1st (td+1) TI **	31116	576	691	48	12	204	59	78281

- Implementation of Moradi and Schneider is better in area/power/energy/randomness in the unconstrained case

Implementations trade-offs

- Absence of randomness is important for reducing the power – since switching activity diminishes
- Note area of the second design is larger

PRINCE in TSMC90 worst PVT case	Area [GE]	Power [uW]	Energy [pJ]	Rand/cycle [bits]	Clock # [cycles]	In/Out [shares]	Latency [ns]
Unprotected - Round Based	3589	59	71	0	12	1/1	30.5
Unprotected - Min Latency	27997			0	1	1/1	13
1st (d+1) TI - with S-box decomp.	8701	97	698	24	72	2/4	277
1st (td+1) TI - with S-box decomp.	14153	75	270	0	36	3/3	134
1st (d+1) TI - w/o S-box decomp.	12220	115	276	112	24	2/8	83
1st (td+1) TI - w/o S-box decomp.	31116	576	691	48	12	4/4	58.8
2nd (d+1) TI - with S-box decomp.	13421	161	1159	72	72	3/8	288
2nd (td+1) TI - with S-box decomp.	18767	232	1670	40	72	5/10	296
2nd (d+1) TI - w/o S-box decomp.	32444	374	898	432	24	3/27	82.2
2nd (td+1) TI - w/o S-box decomp.	177647	1533	3679	352	24	8/17	85.1

- Adding (or removing) the mask refreshing changes the power up to a factor of 2

Implementations trade-offs

- Power vs Energy – performance is important

PRINCE in TSMC90 worst PVT case	Area [GE]	Power [uW]	Energy [pJ]	Rand/cycle [bits]	Clock # [cycles]	In/Out [shares]	Latency [ns]
Unprotected - Round Based	3589	59	71	0	12	1/1	30.5
Unprotected - Min Latency	27997			0	1	1/1	13
1st (d+1) TI - with S-box decomp.	8701	97	698	24	72	2/4	277
1st (td+1) TI - with S-box decomp.	14153	75	270	0	36	3/3	134
1st (d+1) TI - w/o S-box decomp.	12220	115	276	112	24	2/8	83
1st (td+1) TI - w/o S-box decomp.	31116	576	691	48	12	4/4	58.8
2nd (d+1) TI - with S-box decomp.	13421	161	1159	72	72	3/8	288
2nd (td+1) TI - with S-box decomp.	18767	232	1670	40	72	5/10	296
2nd (d+1) TI - w/o S-box decomp.	32444	374	898	432	24	3/27	82.2
2nd (td+1) TI - w/o S-box decomp.	177647	1533	3679	352	24	8/17	85.1

Implementations trade-offs

- Absence of randomness is also important for reducing the energy, although performance of the first design is worse
- 1st order designs are considerable more energy efficient than 2nd order designs

PRINCE in TSMC90 worst PVT case	Area [GE]	Power [uW]	Energy [pJ]	Rand/cycle [bits]	Clock # [cycles]	In/Out [shares]	Latency [ns]
Unprotected - Round Based	3589	59	71	0	12	1/1	30.5
Unprotected - Min Latency	27997			0	1	1/1	13
1st (d+1) TI - with S-box decomp.	8701	97	698	24	72	2/4	277
1st (td+1) TI - with S-box decomp.	14153	75	270	0	36	3/3	134
1st (d+1) TI - w/o S-box decomp.	12220	115	276	112	24	2/8	83
1st (td+1) TI - w/o S-box decomp.	31116	576	691	48	12	4/4	58.8
2nd (d+1) TI - with S-box decomp.	13421	161	1159	72	72	3/8	288
2nd (td+1) TI - with S-box decomp.	18767	232	1670	40	72	5/10	296
2nd (d+1) TI - w/o S-box decomp.	32444	374	898	432	24	3/27	82.2
2nd (td+1) TI - w/o S-box decomp.	177647	1533	3679	352	24	8/17	85.1

Implementations trade-offs

- As expected: (d+1) designs are smaller in area than (td+1) designs, but use more randomness
- 1st order (td+1) designs are 2 times faster (clock cycles) than the corresponding (d+1) designs

PRINCE in TSMC90 worst PVT case	Area [GE]	Power [uW]	Energy [pJ]	Rand/cycle [bits]	Clock # [cycles]	In/Out [shares]	Latency [ns]
Unprotected - Round Based	3589	59	71	0	12	1/1	30.5
	27997			0	1	1/1	13
	8701	97	698	24	72	2/4	277
	14153	75	270	0	36	3/3	134
	12220	115	276	112	24	2/8	83
	31116	576	691	48	12	4/4	58.8
	13421	161	1159	72	72	3/8	288
	18767	232	1670	40	72	5/10	296
	32444	374	898	432	24	3/27	82.2
	177647	1533	3679	352	24	8/17	85.1

- The higher the order of protection is, the larger the area is and more randomness is required

Conclusions

- Study on how a very high level of SCA protection can be achieved by keeping the latency as low as possible
- Optimized low-latency TI has been shown
- Comparison of different implementation trade-offs
  - Area
  - Power consumption
  - Energy consumption
  - Randomness used
  - Latency

- Optimizing TI only on area or randomness (and therefore only on power) is easier
  Very good results are known
- Optimizing TI on more than one criteria like latency or energy is harder and still an open problem

Questions