A note on quantum related-key attacks

Martin Rötteler
NEC Laboratories America
4 Independence Way, Suite 200
Princeton, NJ 08540, U.S.A.
mroetteler@nec-labs.com

Rainer Steinwandt
Florida Atlantic University
Department of Mathematical Sciences
Boca Raton, FL 33431
rsteinwa@fau.edu

June 10, 2013

Abstract

In a basic related-key attack against a block cipher, the adversary has access to encryptions under keys that differ from the target key by bit-flips. In this short note we show that for a quantum adversary such attacks are quite powerful: if the secret key is (i) uniquely determined by a small number of plaintext-ciphertext pairs, (ii) the block cipher can be evaluated efficiently, and (iii) a superposition of related keys can be queried, then the key can be extracted efficiently.

1 Introduction

The availability of scalable quantum computers would jeopardize the security of many currently deployed asymmetric cryptographic schemes (Shor, 1997). For symmetric cryptography the expectations for a post-quantum setting tend to be more optimistic, see e.g. (Bernstein, 2009), from which we quote *"quantum computers seem to have very little effect on secret-key cryptography, hash functions, etc. Grover's algorithm forces somewhat larger key sizes for secret-key ciphers, but this effect is essentially uniform across ciphers; today's fastest pre-quantum 256-bit ciphers are also the fastest candidates for post-quantum ciphers at a reasonable security level."*

Related-key attacks are a powerful cryptanalytic tool when exploring block ciphers. In such attacks, the adversary is granted access to encryptions and/or decryptions of messages under secret keys which are related to the target key in a known or chosen way. As argued in (Kelsey et al., 1996), this type of attack is of practical interest, despite the assumptions made. When Winternitz and Hellman described this attack model more than 25 years ago, they focused on key relations given by bit-flips (Winternitz and Hellman, 1987). An illustrative example for an application of this attack model is an attack against 9 rounds of Rijndael with a 256-bit key, invoking 256 related keys with a particular choice of the bit-flips (Ferguson et al., 2001).

Current approaches to formalize related-key attacks allow more general key relations (Bellare and Kohno, 2003; Albrecht et al., 2011), and restricting to bit-flips can be considered to be a rather conservative choice. Below we show that for a quantum adversary such a basic form of related-key attack is quite powerful. We show that the possibility to query a superposition of related keys to a block cipher enables the efficient extraction of the secret key, if some rather mild conditions are met:

1. the block cipher can be implemented efficiently as a quantum circuit, and
2. the secret key is uniquely determined by a small number of available plaintext-ciphertext pairs.

The attack we describe is unlikely to pose a practical threat as querying a superposition of secret keys may not be feasible for a typical implementation. Notwithstanding this, from the structural point of view our observation indicates an interesting limitation for the security guarantees of a block cipher that one can hope to prove in a post-quantum scenario.

2 Preliminaries

A block cipher with *key length* k and *block length* n is a family of 2^k permutations $\{E_K : \{0,1\}^n \longrightarrow \{0,1\}^n\}_{K\in\{0,1\}^k}$ on bitstrings of length n. Popular block ciphers limit the possible choices of the key length k—e. g., for the Advanced Encryption Standard (NIST, 2001) we have $n = 128$ and $k \in \{128, 192, 256\}$. To characterize the efficiency of certain types of attacks, it can nonetheless be convenient to consider families of block ciphers, interpreting the key length k as a scalable security parameter. Measuring the running time of an adversary as a function of k, it is meaningful to speak of an expected polynomial time attack.

2.1 Related-key attacks

The attack model we consider goes back to (Winternitz and Hellman, 1987). After a key $K \in \{0,1\}^k$ has been chosen uniformly at random, the adversary has access to two oracles:

$\mathcal{E}$**:** On input a bitmask $L \in \{0,1\}^k$ and a bitstring $m \in \{0,1\}^n$, this oracle returns the encryption $E_{K\oplus L}(m)$ of m under the key $K \oplus L$.

$\mathcal{E}^{-1}$**:** On input a bitmask $L \in \{0,1\}^k$ and a bitstring $c \in \{0,1\}^n$, this oracle returns the decryption $E^{-1}_{K\oplus L}(c)$ of c under the key $K \oplus L$.

After interacting with these oracles, the adversary has to output a guess K' for K, and it is considered successful if and only if $K = K'$. For our attack we will also assume that the block cipher at hand can be evaluated efficiently, i. e., with a polynomial-size quantum circuit that has the secret key and a plaintext as input. For block ciphers that are actually used this condition is of no concern.

The quantum attack below will not involve $\mathcal{E}^{-1}$, but we will allow the adversary to query the block cipher and also the oracle $\mathcal{E}$ with a superposition of keys. Finally, we require that the adversary has access to a polynomial number of plaintext-ciphertext pairs $(m_1, c_1), \ldots, (m_r, c_r)$ such that there exists exactly one secret key $K \in \{0,1\}^k$ satisfying

$$(c_1, \ldots, c_r) = (E_K(m_1), \ldots, E_K(m_r)).$$

It is easy to come up with a pathological block cipher where the secret key cannot be uniquely determined by any number of plaintext-ciphertext pairs[1], but for typical block ciphers we do not think this to be a concern. In (Menezes et al., 2001, Definition 7.34) the *known plaintext unicity distance* is defined as a measure for the number of (known) plaintext-ciphertext pairs that are needed to determine the secret key of a block cipher uniquely, and with (Menezes et al., 2001, Fact 7.35) it seems plausible to estimate that for an n-bit block cipher with key length k having

$$r > \lceil k/n \rceil \tag{1}$$

plaintext-ciphertext pairs suffices. So for the 128-bit version of AES, where $n = k = 128$, one can think of an r-value as small as 2. Throughout we will assume that r satisfies Inequality (1). Then the main idea to mount a quantum related-key attack is a reduction to a quantum algorithm described in (Simon, 1994).

[1] Encryption and decryption can simply ignore parts of the secret key.

2.2 Simon's problem

Let $f: \{0,1\}^k \longrightarrow \{0,1\}^{k'}$ with $k \le k'$ be a function such that one of the following two conditions holds:

(a) f is injective;

(b) there exists a bitstring $s \in \{0,1\}^k \setminus \{0^k\}$ such that for every two distinct $x, x' \in \{0,1\}^k$ we have

$$f(x) = f(x') \iff x = x' \oplus s.$$

Simon's problem asks to decide for such a function f which of the two conditions holds, and in the case (b) to find s. Allowing the function f to be evaluated at a superposition of inputs, (Simon, 1994) establishes the following result:

Theorem 1 *Let $g(k)$ be an upper bound for the time needed to solve a $k \times k$ linear system of equations over the binary field $\mathbb{F}_2$, and let $t_f(k)$ be an upper bound for the time needed to evaluate the function f on (a superposition of) inputs from $\{0,1\}^k$. Then the above problem can be solved in expected time $\mathrm{O}(k \cdot t_f(k) + g(k))$. In particular, for $t_f = t_f(k)$ being polynomial, the above problem can be solved in expected polynomial time.*

3 Description of the attack

Alluding to the Electronic Code Book mode of operation (Menezes et al., 2001, Section 7.2.2), subsequently we will simply write $E_K(\vec{m})$ for the tuple of ciphertext blocks $(E_K(m_1), \ldots, E_K(m_r)) \in \{0,1\}^{rn}$. For a fixed, unknown secret key $s \in \{0,1\}^k \setminus \{0^k\}$ and messages $\vec{m} \in \{0,1\}^{rn}$ as described in Section 2, we define the function

$$\begin{array}{rccc} f_s: & \{0,1\}^k & \longrightarrow & 2^{\{0,1\}^{2rn}} \\ & x & \longmapsto & \{E_x(\vec{m}), E_{s \oplus x}(\vec{m})\} \end{array}.$$

Remark 1 *For each x in the domain of f_s, the image is comprised of two different ciphertexts, i. e., it does not collapse to a singleton set: because of the choice of the plaintexts $m_1, \ldots, m_r$ the condition $s \neq 0^k$ implies that $E_x(\vec{m}) \neq E_{s \oplus x}(\vec{m})$.*

Meeting the conditions of Simon's problem To argue that f_s meets the conditions of Theorem 1, let us first clarify how to encode the images as elements of $\{0,1\}^{k'}$ for some $k' \ge k$. As we impose the condition (1), with $k' = 2rn$ we clearly have $k' \ge k$ as desired. For instance by interpreting elements in $\{0,1\}^{rn}$ as binary numbers, we can impose a linear order on $\{0,1\}^{rn}$. Then, to store an element $\{c, c'\}$ in the image of f_s, we simply store the ordered pair $(\min(c,c'), \max(c,c'))$ as its unique k'-bit representation.

Now denote by $x \neq x'$ two different k-bit strings satisfying $f_s(x) = f_s(x')$:

- If $E_x(\vec{m}) = E_{x'}(\vec{m})$ then the choice of the plaintexts $m_1, \ldots, m_r$ implies $x = x'$, so this cannot happen.
- If $E_x(\vec{m}) \neq E_{x'}(\vec{m})$, then $E_x(\vec{m}) = E_{s \oplus x'}(\vec{m})$, which by the choice of the plaintexts $m_1, \ldots, m_r$ means that $x = s \oplus x'$.

So we have the implication $f_s(x) = f_s(x') \Longrightarrow x = x' \oplus s$. The converse follows trivially from $s \oplus (x' \oplus s) = x'$.

Next, let us check that the function $f_s(\cdot)$ can be evaluated efficiently.

Evaluating $f_s(\cdot)$ in polynomial time By assumption the underlying block cipher can be evaluated with a polynomial-size quantum circuit, so computing the two values $E_x(\vec{m})$ and $E_{s\oplus x}(\vec{m})$ for a given x can certainly be done in polynomial time. In the actual attack, the value $E_{s\oplus x}(\vec{m})$ is obtained by invoking the encryption oracle $\mathcal{E}$ for each of the plaintexts $m_1,\dots,m_r$, i. e., with a polynomial number of queries to $\mathcal{E}$. This means we can obtain the pair $(E_x(\vec{m}), E_{s\oplus x}(\vec{m}))$ in polynomial time, and we are left to distill our unique k'-bit representation of the *set* comprised by these two elements.

As indicated in the previous paragraph, such representation can be implemented by interpreting the two ciphertexts as integers and then sort them. A quantum circuit to determine this unique representation of a pair of bitvectors consists of swapping $E_x(\vec{m})$ and $E_{s\oplus x}(\vec{m})$ conditioned on the the latter value being smaller than the former. For instance with a reversible circuit to perform addition (Cuccaro et al., 2004; Draper et al., 2006; Takahashi et al., 2010) one can compute the difference of the binary numbers represented by $E_x(\vec{m})$ and $E_{s\oplus x}(\vec{m})$ in polynomial time. The most significant bit of the result then reveals the result of the comparison. The swap operation can be conditioned on this bit, followed by an uncomputation of the garbage introduced by the adder (Bennett, 1973). The overall circuit to compute f_s therefore uses 2 adders and one controlled swap operation of two bitvectors, implying the following result.

Proposition 1 *For every $s \in \{0,1\}^k \setminus \{0^k\}$ the function f_s defined above satisfies the conditions needed to apply Theorem 1, and the bound t_{f_s} can be chosen to be polynomial.*

From Theorem 1 we now obtain the following attack which runs in expected polynomial time:

1. Check if the secret target key s is the all-zero key $s = 0^k$ by computing $E_{0^k}(\vec{m})$ and comparing these ciphertexts with the given ciphertexts $E_s(\vec{m})$.

2. If $s \neq 0^k$ then apply Simon's algorithm—which constitutes the proof of Theorem 1—to recover s.

4 Conclusion

This note shows that in a quantum setting even a basic related-key attack is very powerful: under rather mild assumptions on the attacked block cipher the secret key can be extracted efficiently.

Acknowledgment. RS was supported by the Spanish *Ministerio de Economía y Competitividad* through the project grant MTM-2012-15167.

References

Albrecht, M. R., Farshim, P., Paterson, K. G., and Watson, G. J. (2011). On Cipher-Dependent Related-Key Attacks in the Ideal-Cipher Model. In *Fast Software Encryption – FSE 2011*, volume 6733 of *Lecture Notes in Computer Science*, pages 128–145. International Association for Cryptologic Research, Springer.

Bellare, M. and Kohno, T. (2003). A Theoretical Treatment of Related-Key Attacks: RKA-PRPs, RKA-PRFs, and Applications. In Biham, E., editor, *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 491–506. International Association for Cryptologic Research, Springer.

Bennett, C. H. (1973). Logical reversibility of computation. *IBM J. Res. Develop.*, 17:525–532.

Bernstein, D. J. (2009). *Post-quantum cryptography*, chapter Introduction to post-quantum cryptography, pages 1–14. Springer.

Cuccaro, S. A., Draper, T. G., Kutin, S. A., and Moulton, D. P. (2004). A new quantum ripple-carry addition circuit. arXiv:quant-ph/0410184.

Draper, T. G., Kutin, S. A., Rains, E. M., and Svore, K. M. (2006). A logarithmic-depth quantum carry-lookahead adder. *Quantum Information and Computation*, 6(4&5):351–369.

Ferguson, N., Kelsey, J., Lucks, S., Schneier, B., Stay, M., Wagner, D., and Whiting, D. (2001). Improved Cryptanalysis of Rijndael. In Schneier, B., editor, *Fast Software Encryption – FSE 2000*, volume 1978 of *Lecture Notes in Computer Science*, pages 213–230. Springer.

Kelsey, J., Schneier, B., and Wagner, D. (1996). Key-Schedule Cryptoanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES. In Koblitz, N., editor, *Advances in Cryptology – CRYPTO '96*, volume 1109 of *Lecture Notes in Computer Science*, pages 237–251. Springer.

Menezes, A. J., van Oorschot, P. C., and Vanstone, S. A. (2001). *Handbook of Applied Cryptography*. CRC Press. Sample chapters available at `http://cacr.uwaterloo.ca/hac/`.

NIST (2001). Specification for the ADVANCED ENCRYPTION STANDARD (AES). Federal Information Processing Standards Publication 197.

Shannon, C. (1949). Communication theory of secrecy systems. *Bell Systems Technical Journal*, 28(4):656–715.

Shor, P. W. (1997). Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5):1484–1509.

Simon, D. R. (1994). On the Power of Quantum Computation. In *35th Annual Symposium on Foundations of Computer Science*, pages 116–123. IEEE Computer Society.

Takahashi, Y., Tani, S., and Kunihiro, N. (2010). Quantum addition circuits and unbounded fan-out. *Quantum Information and Computation*, 10(9&10):872–890.

Winternitz, R. and Hellman, M. (1987). Chosen-key attacks on a block cipher. *Cryptologia*, XI(1):16–20.