Breaking One-Round Key-Agreement Protocols in the Random Oracle Model

Miroslava Sotáková*

January 30, 2008

Abstract

In this work we deal with one-round key-agreement protocols, called Merkle's Puzzles, in the random oracle model, where the players Alice and Bob are allowed to query a random permutation oracle n times. We prove that Eve can always break the protocol by querying the oracle $O(n^2)$ times. The long-time unproven optimality of the quadratic bound in the fully general, multi-round scenario has been proven recently by Barak and Mahmoody-Ghidary. The results in this paper have been found independently of their work.

1 Introduction

In this work we prove the tight upper-bound on the number of queries needed to break a key-agreement protocol in the random oracle model. The key-agreement protocol called Merkle's puzzles, developed by Merkle in 1974 a published in 1978 [3] is one of the earliest example of public-key encryption.

Following the protocol, two parties can agree on a secret-key by exchanging messages, assuming that they share no secrets beforehand. Informally, Alice creates a message for Bob in the following way - she constructs a large number of puzzles of moderate difficulty, each of them being possible to solve with Bob's computational resources. All of them are in the form of an encrypted message with an unknown key that is short enough to allow the brute force attack. After receiving the message from Alice, Bob chooses one puzzle uniformly at random and solves it. The solution contains an identifier and a key. Bob encrypts the identifier with the key, and announces it back to Alice. The solution of the puzzle solved by Bob becomes Alice's and Bob's secret-key. Since the puzzle's identifier is sent to Alice as a message encrypted with a key that is unknown to Eve, the eavesdropper's best strategy to attack the key-agreement protocol is to solve as many puzzles as possible. To achieve constant probability of success, she has to solve a constant fraction of them, which might require much more computational power than Alice and Bob have.

In a similar way we construct a key-agreement protocol in the random oracle scenario, where the computational difficulty of key-agreement is expressed by the number of oracle queries that

*Basic Research in Computer Science, University of Aarhus, mirka@daimi.au.dk

Alice and Bob make to agree on a secret-key. Instead of creating many puzzles, Alice queries the oracle in many positions that are unknown to both Bob and Eve, and sends the images of the queried elements to Bob. Bob queries the oracle in sufficiently many positions to get a collision with Alice's set of queries with high probability. He recognizes the collision from Alice's message and reports it back to Alice by its identifier – the oracle image. The pre-image becomes Alice's and Bob's secret-key. The communication gives Eve almost no information about the key, since the oracle is random. With the same number of queries as Bob, she would find a collision with Alice's set of queries with high probability, but not necessarily the one found by Bob. Hence, finding the right element might require significantly more oracle queries than Alice and Bob needed to agree on the secret-key.

Until recently, the best upper-bounds on Eve's number of queries needed to break such protocols have been proven by Impagliazzo and Rudich [2]. They have shown that in any key-agreement protocol based on a random permutation-oracle, where Alice and Bob agree on the secret-key in n rounds in such a way that they query only one query per round, Eve needs $O(n^3)$ oracle queries to output a secret-key guess that matches with Bob's secret-key with the same probability as Alice's key does. They call this form of protocols the normal form, and proved that every protocol can be put into the normal form with at most quadratic blow-up in the number of oracle queries used. Here the question is studied in the larger context to show that any proof that secure key-agreement relative to some random permutation oracle is possible implies P $\neq$ NP. In other words, no proof that the existence of one-way function implies the existence of secure key-agreement relativizes.

The bound from [2] has been improved recently by Barak and Mahmoody-Ghidary [1] to the optimal $O(n^2)$.

In our paper we deal with one-round key-agreement protocols where Alice and Bob query the oracle a and b times, respectively, is a subset of protocols whose normal form consists of $a+b$ rounds. We prove the tight – $O((a+b)^2)$ upper-bound on the number of queries Eve needs to break the protocol.

2 One-Round Key-Agreement Protocols

In this section, we model one round key-agreement protocols between Alice and Bob. We assume that Alice, Bob, and an eavesdropper Eve have access to an oracle computing a random permutation f on $\{1,\ldots,n\}$. We define a one-round key-agreement protocol as follows:

Protocol 1
Given $n \in \mathbb{N}$ and an oracle computing a random permutation f on $\{1,\ldots,n\}$,

1. Alice queries the oracle f in positions $A_1 \in \{1,\ldots,n\}^{\leq a}$, computes a message $\mathcal{C}_A$ and sends it to Bob.

2. Bob, given $\mathcal{C}_A$, queries the oracle f in positions $B \in \{1,\ldots,n\}^{\leq b}$, computes message $\mathcal{C}_B$ and sends it to Alice. Bob generates the secret-key $k_B \in \{0,1\}^{\ell}$, $k_B = g_B(B, f(B), \mathcal{C}, R_B)$, where $\mathcal{C} = (\mathcal{C}_A, \mathcal{C}_B)$, and R_B denotes some randomness.

3. Alice, given $\mathcal{C}$, queries the oracle in positions $A_2 \subseteq \{1,\ldots,n\}$ such that for $A := A_1 \cup A_2$, $|A| \leq a$, and generates the secret-key $k_A \in \{0,1\}^\ell$, $k_A = g_A(A, f(A), \mathcal{C}, R_A)$, where R_A denotes some randomness.

We denote by (a,b,ε)–key-agreement any one-round key-agreement protocol defined as above and satisfying the following condition: $\Pr[k_A \neq k_B] \leq \varepsilon$ where $\varepsilon < 1$ is a constant.

In fact, a, and b are functions of n, but for simplicity we keep this notations instead of $a(n)$ and $b(n)$, if the latter one is not explicitly needed. Since key-agreement protocols take place between players Alice and Bob sharing no initial secret, the key generation mechanism must involve common queries to the oracle f. We will say that Eve breaks the protocol, if she outputs the string that agrees with Bob's secret-key with the same probability as Alice does.

Let E denote the set of oracle queries of an attacker Eve.

Lemma 2.1 *In order to break an (a,b,ε)–key-agreement protocol it is sufficient for Eve to query all intersection queries of Alice and Bob used for the generation of Alice's secret-key.*

Proof Eve querying all elements in $A_1 \cap B$ can construct a permutation f' matching with f on E, and a set A_1' of queries to the oracle computing f' such that $\mathcal{C}_A = \mathcal{C}_{A_1'}$ and f' is consistent with $\mathcal{C}_B$. Therefore, after querying B, Bob has exactly the same view about A_1 as he has about A_1'. Eve constructs the set A_2' according to A_1' and $\mathcal{C}$ and then "queries" the f'-oracle on the positions in A_2'. Finally, she generates her secret-key $k_E = g_A(A', f(A'), \mathcal{C}, R_{A'})$ for $A' := A_1' \cup A_2'$. From Bob's point of view, both k_E and k_A are generated from the same set $K \subseteq A \cap B$, i.e. $\Pr[k_B = k_E] = \Pr[k_B = k_A]$. ∎

3 Proof of the Quadratic Upper-Bound

We will consider the following attack of an (a,b,ε)–key-agreement protocol:

1. Eve repeats Bob's querying strategy γa times for some constant γ (γab oracle queries) in order to query all queries in $A_1 \cap B$ with constant probability
2. Eve extracts the position of the A_2-queries from $\mathcal{C}_B$ and queries the oracle on these positions (a oracle queries)

Next we prove that with the proposed strategy Eve breaks the protocol with constant probability.

Lemma 3.1 *By repeating Bob's strategy independently $5a$ times, Eve finds all elements in $A_1 \cap B$ with constant probability.*

Proof Let $\mathbf{A}$, $\mathbf{B}$ denote the random variables associated with Alice querying the elements in A_1 and Bob querying the elements in B, respectively. Let $\mathbf{E}$ denote the random variable associated with the set of Eve's queries. Assume that for $x, y \in \{1,\ldots,n\}, x \leq y$, $P_{\chi_{\mathbf{B}}=1|\mathcal{C}_A}(x) \leq P_{\chi_{\mathbf{B}}=1|\mathcal{C}_A}(y)$.

Define $A_1^0 := A_1$, $B^0 := B$, $\mathbf{A}^0 = \mathbf{A}$, $\mathbf{B}^0 := \mathbf{B}$, $s_0 := \langle |A_1 \cap B| \rangle_{\mathcal{C}_A}$, $n_0 := n$. In the i-th step (starting with $i = 0$), let us consider n_{i+1}, A_1^{i+1}, B^{i+1}, $\mathbf{A}^{i+1}$, $\mathbf{B}^{i+1}$, s_{i+1} with the following properties:

$$\forall x \in \{n_{i+1}+1, \ldots, n_i\} : \; P_{\chi_{\mathbf{B}}=1|\mathcal{C}_A}(x) \geq \frac{s_i}{2a},$$

$A_1^{i+1} := A_1 \backslash \{n_{i+1}+1, \ldots, n\}$, $B^{i+1} := B \backslash \{n_{i+1}+1, \ldots, n\}$, $\mathbf{A}^{i+1}$, $\mathbf{B}^{i+1}$ denote the corresponding random variables, $s_{i+1} := \langle |A_1^{i+1} \cap B^{i+1}| \rangle_{\mathcal{C}_A}$.

Furthermore, consider u such that

$$\Pr[A_1 \cap B \subseteq \{n_u + 1, \ldots, n\} | \mathcal{C}_A] \geq \frac{1}{2}.$$

First we prove that

1. there exists $u \in \mathbb{N}$ with the desired property
2. $n_{i+1} < n_i$ for $i \in \{0, \ldots, u-1\}$
3. $s_i - s_{i+1} \geq 1$ for $i \in \{0, \ldots, u-1\}$
4. $s_u \geq 1$

Proof of 1.,2.,3.,4.

$$s_i = \langle |A_1^i \cap B^i| \rangle_{\mathcal{C}_A} = \sum_{|A| \leq a} P_{\mathbf{A}^i|\mathcal{C}_A}(A) \sum_{|B| \leq b} P_{\mathbf{B}^i|\mathcal{C}_A} |A \cap B|,$$

thus there exists at least one $\bar{A} \subseteq \{1, \ldots, n\}^{\leq a}$ such that $\sum_{|B| \leq b} P_{\mathbf{B}^i|\mathcal{C}_A}(B) |\bar{A} \cap B| \geq s_i$.

Let us choose one such $\bar{A}$. Then

$$s_i \leq \sum_{|B| \leq b} P_{\mathbf{B}^i|\mathcal{C}_A}(B) |\bar{A} \cap B| = \sum_{|B| \leq b} \sum_{x \in |\bar{A} \cap B|} P_{\mathbf{B}^i|\mathcal{C}_A}(B) = \sum_{x \in \bar{A}} \sum_{B: \; x \in \bar{A} \cap B} P_{\mathbf{B}^i|\mathcal{C}_A}(B) = \sum_{x \in \bar{A}} P_{\chi_{\mathbf{B}^i}=1|\mathcal{C}_A}(x).$$

Since $|\bar{A}| \leq a$, there is an $x \in \{1, \ldots, n_i\}$ such that $P_{\chi_{\mathbf{B}^i}=1|\mathcal{C}_A}(x) \geq \frac{s_i}{a}$. If we remove $x \in \{1, \ldots, n_i\}$ such that $P_{\chi_{\mathbf{B}^i}=1|\mathcal{C}_A}(x) \geq \frac{s_i}{2a}$, then $s_{i+1} \leq \frac{s_i}{2}$.

Since in every step we remove at least one $x \in \{1, \ldots, n\}$, the procedure terminates after finitely many steps, thus u is well-defined and is at most n. Clearly, for $s_i < 1$ we have $\Pr[A_1^{i+1} \cap B^{i+1} = \emptyset | \mathcal{C}_A] > \frac{1}{2}$, thus with probability at least $1/2$ we have $A_1 \cap B \subseteq \{n_i+1, \ldots, n\}$, therefore $s_u \geq 1$ and for $i \in \{0, \ldots, u-1\}$:

$$s_i - s_{i+1} \geq \frac{s_i}{2} \geq \frac{s_{u-1}}{2} \geq 1.$$

Now let us finish the proof of the statement by showing that by repeating Bob's strategy $5a$ times independently Eve does not query all elements in $A_1 \cap B$ with probability at most $7/8$.

For $x \in \{n_{i+1}+1, \ldots, n_i\}$ Eve does not query x with probability

$$P_{\chi_{\mathbf{E}}=0|\mathcal{C}_A}(x) \leq \left(1 - \frac{s_i}{2a}\right)^a \approx e^{-s_i/2}.$$

That means that in the case when

$$|A_1^i \cap B^i \cap \{n_{i+1}+1, \ldots, n_i\}| \leq \frac{e^{s_i/2}}{2s_i^2}$$

the probability that Eve does not query at least one element in $\{n_{i+1}, \ldots, n_i\} \cap A_1^i \cap B^i$ is

$$\Pr\left[\prod_{x \in \{n_{i+1},\ldots,n_i\} \cap A_1^i \cap B^i} \chi_E(x) = 0 | \mathcal{C}_A\right] \leq \frac{e^{s_i/2}}{2s_i^2} \cdot e^{-s_i/2} = \frac{1}{2s_i^2}.$$

Since the expected number of elements in $A_1^i \cap B^i \cap \{n_{i+1}+1, \ldots, n_i\}$ is s_i, Markov's inequality tells us that this happens with probability at most $\frac{2s_i^3}{e^{s_i/2}}$. This implies that there is at least one $0 \leq i < u$ such that $|A_1^i \cap B^i \cap \{n_{i+1}+1, \ldots, n_i\}| > \frac{e^{s_i/2}}{2s_i^2}$ with probability at most $\sum_{i=0}^{u} \frac{2s_i^3}{e^{s_i/2}}$. The function $\frac{2x^3}{e^{x/2}}$ is decreasing for $x \geq 6$, thus

$$\sum_{i=0}^{u'-1:s_{u'} \geq 6} \frac{2s_i^3}{e^{s_i/2}} \leq \sum_{i=0}^{u'-1:s_{u'} \geq 6} (s_i - s_{i+1}) \frac{2s_i^3}{e^{s_i/2}} \leq \int_{x=s_{u'}}^{\infty} \frac{2x^3}{e^{x/2}} \mathrm{d}x.$$

Then for $s_{u'} \geq 28$ we have

$$\sum_{i=0}^{u'-1:s_{u'} \geq 28} \frac{2s_i^3}{e^{s_i/2}} < \frac{1}{8}.$$

Furthermore, for $s_i < 28$ (there are at most 5 of them, since $s_{i+1} \leq s_i/2$ and $s_u \geq 1$), the probability that $A_1^i \cap B^i \cap \{n_{i+1}+1, \ldots, n_i\}$ contains more than $40s_i$ elements is at most $1/40$, by Markov's inequality.

Thus the probability that there exists an $i, 0 \leq i < u$ such that

$$|A_1 \cap B \cap \{n_{i+1}, \ldots, n_i\}| > \max\{40s_i, \frac{e^{s_i/2}}{2s_i^2}\}$$

is at most $\frac{1}{8} + \frac{5}{40} = \frac{1}{4}$. This is the probability that $A_1 \cap B$ has a "bad structure" for finding all its elements by Eve.

It is sufficient for Eve to repeat Bob's algorithm $(\log 80 + 3\log s_i)a/s_i \leq 5a$ times to get all elements in $A_1 \cap B \cap \{n_{i+1}, \ldots, n_i\}$, $i \geq u'$, assuming that there are no more than $40s_i$ of them, with probability at least $1 - \frac{1}{2s_i^2}$.

That means that with $5a$ independent iterations of Bob's strategy Eve does not query at least one element of $A_1 \cap B \cap \{n_u+1, \ldots, n\} = A_1 \cap B$ for well-structured $A_1 \cap B$ with probability

$$\begin{aligned} \Pr & \left[\prod_{x \in \{n_u+1,\ldots,n\} \cap A_1 \cap B} \chi_E(x) = 0 | \mathcal{C}_A \right] \leq \frac{1}{2} \cdot \sum_{i=0}^{u} \frac{1}{s_i^2} \\ \leq & \sum_{i=0}^{u-1} (s_i - s_{i+1}) \cdot \frac{1}{2s_i^2} + \frac{1}{2s_u^2} \leq \frac{1}{2} \cdot \int_{x=s_u}^{\infty} \frac{\mathrm{d}x}{x^2} = \frac{1}{2s_u} \leq \frac{1}{2}. \end{aligned}$$

Since $A_1 \cap B \subsetneq \{n_u + 1, \ldots, n\}$ with probability at most $\frac{1}{2}$, and $A_1 \cap B$ is malstructured with probability at most $\frac{1}{4}$, $A_1 \cap B \subseteq \{n_u + 1, \ldots, n\}$ and is well-structured with probability at least $\frac{1}{4}$ must hold. In this case Eve queries all intersection elements with probability at least $\frac{1}{2}$, thus Eve finds all intersection queries of A_1 and B with probability at least $\frac{1}{8}$. ∎

Theorem 3.2 *Eve can break an (a, b, ε)–key-agreement protocol with $O((a + b)^2)$ queries with constant probability.*

Proof As was claimed in the proof of Lemma 2.1, Eve querying all queries in $A_1 \cap B$ needs at most $|A_2| \leq a$ queries more to generate the key that matches with Bob's secret-key with the same probability as Alice's key does. Lemma 3.1 shows that Eve can always query all elements in $A_1 \cap B$ with probability $1/8$ with at most $5ab$ queries. Therefore, Eve can break the protocol with constant probability with $5ab + a \in O((a + b)^2)$ oracle queries. ∎

4 Optimality of the Bound

Consider the following protocol:

Protocol 2

1. Alice chooses a set $A \subseteq \{1, \ldots, n\}$, $|A| = a = \lceil \sqrt{n} \rceil$ uniformly at random, queries the oracle for the elements of A, and sends $\mathcal{C}_A = \{f(x) : x \in A\}$ to Bob.

2. Bob chooses a set $B \subseteq \{1, \ldots, n\}$, $|B| = b = \lceil \sqrt{n} \rceil$ uniformly at random, queries the elements of B, chooses a collision element $k \in \{f(y) : y \in B\} \cap \mathcal{C}_A$ at random, and sends $\mathcal{C}_B = \{f(k)\}$ to Alice. He outputs his secret-key k.

3. Alice recognizes k according to $\mathcal{C}_B$ and A, and outputs her secret-key k.

Attack: Bob finds at least one collision with Alice's set of queries due to the birthday argument, therefore, the given protocol is an example of $(\sqrt{n}, \sqrt{n}, \varepsilon)$-key-agreement protocol for a constant $\varepsilon < 1$. Given $\mathcal{C}$, the secret-key is uniformly distributed on $\{1, \ldots, n\}$ and furthermore, since the oracle is random, Eve knowing the oracle image for $o(n)$ elements still has $(1 - o(1)) \log n$ entropy about $f(x)$ for $x \notin E$. Hence, Eve has to query the oracle in $\Theta(n)$ positions to get the right secret-key with constant probability, i.e the best Eve's attack to break the protocol with constant probability has to involve $O(n) = O((a + b)^2)$ oracle queries.

5 Conclusion

We provided an analysis of the most commonly considered attack of these type of key-agreement protocols where the attacker iterates the players' strategies with gradually updated information in the case of one-round protocols. We were hoping to generalize the result to apply in the multi-round scenario, which has been done very recently by Barak and Mahmoody-Ghidary.

References

[1] Barak, B., and Mahmoody-Ghidary, M. Merkle puzzles are optimal, 2008.

[2] Impagliazzo, R., and Rudich, S. Limits on the provable consequences of one-way permutations. In *STOC '89: Proceedings of the twenty-first annual ACM symposium on Theory of computing* (New York, NY, USA, 1989), ACM Press, pp. 44–61.

[3] Merkle, R. C. Secure communications over insecure channels. *Communications of the Association for Computing Machinery 21*, 4 (Apr. 1978), 294–299.