

Query and Communication Lower Bounds for Key-Exchange Protocols: a Survey

Noah Golowich

May 2, 2018

Abstract

In this survey we review recent work on the optimality of Merkle’s puzzles in the random oracle model. Specifically, we cover Barak and Mahmoody’s (CRYPTO, 2009) proof that there is no ℓ -query protocol secure against $\omega(\ell^2)$ adversaries, as well as the recent proof of Haitner et al. (2018) showing the optimality of the communication complexity of Merkle’s puzzles in two specific cases.

1 Introduction

In 1974, Merkle [7] presented a key exchange protocol in the random oracle model by which two parties, A and B can negotiate a shared secret key over an insecure channel. In particular, if each of A and B make at most ℓ queries to some random oracle H , then any attacker that makes $o(\ell^2)$ queries to H cannot recover the key with constant probability. The main drawback of Merkle’s protocol is that it only provides polynomial (i.e. quadratic) security, and in the following years, many authors [3, 9, 8] have introduced algebraic assumptions under which one can obtain authenticated key exchange with exponential security. However, these assumptions remain unproven, and to date, it is still unknown, for instance, whether one can achieve public-key encryption using only the assumption of pseudo-random generators. Moreover, it is known [10] that quantum computers can efficiently break discrete logarithm-based [3] and factoring-based [9] crypto systems, as well as their elliptic-curve analogues, meaning that if quantum computers become sufficiently powerful, much will rely upon the (as yet unproven) assumptions underlying lattice-based systems.

Even though security against a quadratic adversary may seem somewhat weak, as the computational power of computers increases, such a polynomial gap only becomes more useful [2]. Moreover, as noted in [5], results proving that any protocol using a random function or random permutation (which is a one-way permutation) as a black-box achieve at most polynomial security imply that any proof that a protocol has super-polynomial security must be based on the hardness of a particular problem in NP . In particular, if $P = NP$, the adversary E in such negative results is generally polynomial-time (as is the case in the proofs presented in [5, 2]). Thus, the proof of super-polynomial security of a protocol using a random function or permutation oracle would amount to a proof that $P \neq NP$.

For these reasons, it is valuable to consider what level of security one can guarantee in the absence of such algebraic assumptions, and assuming only access to a random oracle that is shared between A, B. Impagliazzo and Rudich [5] showed that if A, B each make at most ℓ queries to a random oracle H (modeled as a random function), then an eavesdropper E can always break

the protocol with $O(\ell^6 \log \ell)$ queries to H ; they also showed that if H is a random permutation, then there is a successful E that makes $O(\ell^{12} \log^2 \ell)$ queries. It remained an open question whether Merkle’s puzzles were tight until Barak and Mahmoody [2] resolved this question in the affirmative, showing the existence of an E with $O(\ell^2)$ queries in the case that H is a random function. By the same reduction as in [5], their result implies an E for a random permutation H that makes $O(\ell^4)$ queries.

An alternative question raised recently by Haitner et al. [4] is on the *communication complexity* of secure key-exchange protocols in the random oracle model. Merkle’s puzzles uses $O(\ell)$ communication to agree on a constant-length key, and it is natural to conjecture that this is in fact optimal for protocols that are secure against $o(\ell^2)$ adversaries. Haitner et al. [4] show that this is indeed the case for protocols which satisfy either of the below properties:

- **Uniform-queries:** Both parties A, B choose their queries uniformly and independently from a fixed set \mathcal{S} before the protocol begins.
- **2-round, non-adaptive:** Both parties A, B choose their queries independently (but not necessarily uniformly) from a fixed set \mathcal{S} before the protocol begins, and the protocol has at most 2 rounds of communication.

In Section 2, we introduce Merkle’s puzzles and give an overview of Barak and Mahmoody’s [2] proof that they achieve optimal security in the random oracle model. In Section 3 we give an overview of the proof of optimality [4] of Merkle’s puzzles with respect to communication complexity for the two cases above. Finally, in Section 4, we discuss further observations and directions for future work on this topic.

2 Merkle’s puzzles

In this section we give an overview of some relevant prior work on key exchange protocols in the random oracle model.

2.1 Communication model, protocols

We consider the setting outlined in [2, 4]. In particular, we suppose that a protocol between Alice and Bob is given by a pair of interactive probabilistic polynomial time Turing machines $\Pi = (A, B)$. We suppose that A, B are given access to a random oracle $H : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$, for some $\ell \in \mathbb{N}$. The protocol begins with A and B tossing coins r_A and r_B , respectively. The protocol then consists of a finite number of rounds, beginning with round 1: in round j , if j is odd, then Alice makes some number of queries to H and then sends a message m_j to Bob, and if j is even, then Bob makes some number of queries to H and then sends a message m_j to Alice. The **view** of party A consists of the tuple $v_A = (i_A, r_A, m)$, where i_A is party A ’s input and m is the concatenation of all messages of the protocol. Similarly, the view of party B consists of $v_B = (i_B, r_B, m)$, where i_B is B ’s input. At the end of the protocol A obtains an output out_A and B obtains an output out_B . Note that out_A, out_B , as well as the transcript $trans$ are deterministic functions of the **joint view** $v := (i_A, i_B, r_A, r_B)$. We write these functions as $out_A(v), out_B(v), trans(v)$, respectively.

A protocol Π is said to be ℓ -**oracle-aided** if there is an oracle (function) $H : \mathcal{S} \rightarrow \mathcal{T}$, for some sets \mathcal{S}, \mathcal{T} , such that parties A, B have oracle access to H and each makes at most ℓ queries to f .

We will often consider oracle-aided protocols with respect to a family \mathcal{H} of functions $H : \mathcal{S} \rightarrow \mathcal{T}$, where $H \leftarrow_R \mathcal{H}$ uniformly at random.

We will mostly consider protocols without inputs; in this case, the views are given by $v_A = (r_A, m)$, $v_B = (r_B, m)$, and $v = (r_A, r_B, m)$. A protocol Π is said to use **public randomness** if parties A, B have access to a common infinite string r_P of random bits. This definition is useful when we discuss the communication complexity of set disjointness in Section 3.

2.2 Security of oracle-aided protocols

Next we define the security of an oracle-aided protocol, which generalizes the definition given in [4]:

Definition 2.1. *Suppose $\Pi = (A, B)$ outputs a keys in $\{0, 1\}^n$. Suppose $\ell \in \mathbb{N}$, and that $\alpha = \alpha(n), \beta = \beta(n) \in [0, 1]$. Then Π is a (q, α, γ) -**secure** key agreement protocol relative to a family of oracles \mathcal{H} , if we have:*

- $(1 - \alpha)$ -**accuracy**, i.e. for every $H \in \mathcal{F}$,

$$\mathbb{P}_{v \leftarrow_R \Pi^H} [\text{out}_A(v) = \text{out}_B(v)] \geq 1 - \alpha.$$

- γ -**secrecy**, i.e. for any algorithm E that receives as input the transcript of communication and that makes at most q queries to the oracle H ,

$$\mathbb{P}_{H \leftarrow_R \mathcal{H}, v \leftarrow_R \Pi^H} [E^H(\text{trans}(v)) = \text{out}_A(v) \wedge \text{out}_A(v) \neq \perp_A] \leq \gamma.$$

Here \perp_A, \perp_B are distinct symbols denoting failure of the protocol for A, B , respectively.

Note that the adversary E in the above definition need not be computationally efficient.

2.3 Description of Merkle's puzzles

Consider the following oracle-aided protocol, originally introduced by Merkle [7]: ℓ is a security parameter, and $H : \{0, 1\}^m \rightarrow \{0, 1\}^m$ is the random oracle, where $m \gg \log \ell$. We may consider $[\ell^2] \subset \{0, 1\}^m$ via any fixed embedding, e.g. the base-2 representation of integers. Let $c > 1$ be a large constant.

1. A chooses $x_1, \dots, x_{c\ell} \in [\ell^2]$ uniformly at random, and sets $a_i := H(x_i)$, $1 \leq i \leq c\ell$, i.e. it makes $c\ell$ queries to H . A then sends B all of the answers $(a_1, \dots, a_{c\ell})$.
2. B chooses $y_1, \dots, y_{c\ell} \in [\ell^2]$ uniformly at random and then sets $b_i := H(y_i)$, $1 \leq i \leq c\ell$, i.e. it makes $c\ell$ queries to H . B then sends A all of the answers $(b_1, \dots, b_{c\ell})$.
3. A then picks a pair (j, j') , lexicographically, as small as possible such that $a_j = b_{j'}$; if such a (j, j') does not exist, then A halts the protocol with output \perp_A . If such a (j, j') does exist, then A sets $\text{out}_A = x_j$.
4. B similarly picks (j, j') lexicographically as small as possible such that $a_j = b_{j'}$ (halting with \perp_B if (j, j') does not exist), and then sets $\text{out}_B = y_{j'}$.

Note that as long as $2^m > c'\ell^4$ for a large constant c' , then $H(1), \dots, H(\ell^2)$ are distinct with probability $1 - \frac{1}{c'}$,¹ so with all but probability $1/c'$, $a_j = b_j$ implies that $x_j = y_j$. Moreover, the probability that the (j, j') such that $a_j = b_{j'}$ exists, as above is at least $1 - \frac{16}{c^2}$ for $c \leq \ell^2$, meaning that the probability of success (namely that $out_A = out_B$) is at least $1 - \frac{16}{c^2} - \frac{1}{c'}$, which can be made close to 1 by taking c, c' sufficiently large.

Conditioned on success of the above protocol and the transcript, the secret key $out_A = out_B$ is a uniform string in $[\ell^2]$, so for sufficiently large c, c' , for $0 < \delta < 1/2$, the above protocol is a $(\delta\ell^2, 0.01, \Theta(\delta))$ -secure key agreement protocol. In particular, if the adversary E is limited to $o(\ell^2)$ queries, then it cannot succeed in determining the shared key with any constant probability.

2.4 Optimality of Merkle's puzzles

Merkle [7] left open the question of whether there is a protocol $\Pi = (A, B)$, where A, B are ℓ -oracle aided, but such that the protocol is secure against $\omega(\ell^2)$ -aided adversaries, i.e. whether an adversary that is allowed to make $\omega(\ell^2)$ queries can guess out_A with constant probability. This question was open for 35 years until Barak et al. [2] resolved it, showing that Merkle's puzzles are essentially optimal:

Theorem 2.2 (Barak et al. [2]). *Suppose that $\Pi = (A, B)$ is a protocol where A, B each make at most ℓ queries to the random oracle H (which is a random function), and such that $\mathbb{P}_{r_A, r_B, H}[out_A = out_B] \geq \rho$. Then for every $0 < \delta < 1$, there is an E making $(\frac{16\ell}{\delta})^2$ queries to H such that $\mathbb{P}_{H, v \leftarrow R^{\Pi^H}}[E^H(\text{trans}(v)) = out_A(v) \wedge out_A(v) \neq \perp_A] \geq \rho - \delta$.*

We now give an overview of the proof of Theorem 2.2, focusing on the parts that are relevant to the lower bounds presented in Section 3. The proof presented here assumes that the protocol Π is in **normal form**, meaning that there are a total of 2ℓ rounds, in which Alice and Bob alternatively make their queries to the oracle and then send a message to the other recipient; we refer the reader to [2] for the reduction from the general case to the normal form case. Note, however, that any normal form protocol necessarily involves $\Omega(\ell)$ communication since there are 2ℓ rounds, during each of which at least 1 bit must be sent. (In fact, the reduction in [2] makes use of an intermediate reduction to “seminormal protocols”, which may incur up to $\Theta(\ell^2)$ additional communication.) Therefore, one cannot directly make use of the same reduction when proving communication complexity lower bounds.

¹To see this, let, for $1 \leq i, j \leq \ell^2$, W_{ij} denote the indicator of the event that $H(i) = H(j)$. Then $\mathbb{E}[W_{ij}] = 1/2^m \leq 1/(c'\ell^4)$, meaning that if $W = \sum_{i \neq j} W_{ij}$, then $\mathbb{E}[W] \leq 1/c'$. Since W is integer-valued, $\mathbb{P}[W > 0] = \mathbb{P}[W \geq 1] \leq 1/c'$, by Markov's inequality.

²To see this, let $Z_{jj'}$ denote the indicator that $x_j = y_{j'}$, and $Z = \sum_{j, j'} Z_{jj'}$. Then Z denotes the number of pairs (j, j') such that $x_j = y_{j'}$, which is certainly a lower bound on the number of pairs (j, j') such that $a_j = b_{j'}$. Then $\mathbb{E}[Z] \geq c^2\ell^2 \cdot \frac{1}{\ell^2} = c^2$, and

$$\text{Var}[Z] = \mathbb{E} \left[\left(\sum_{j, j'} Z_{jj'} \right)^2 \right] - c^2 \leq c^4\ell^4 \cdot \frac{1}{\ell^4} + 2c^3\ell^3 \cdot \frac{1}{\ell^4} + c^2 - c^4 \leq 3c^2,$$

where the first term accounts for those pairs of tuples $(j, j'), (j'', j''')$ such that j, j' are both distinct from each of (j'', j''') , the second term accounts for those pairs of tuples $(j, j'), (j, j'')$ or $(j', j), (j'', j)$ where $j' \neq j''$, and the last term accounts for those pairs of tuples $(j, j'), (j, j')$. The last inequality holds for $\ell \geq c$. We then use Chebyshev's inequality to get that $\mathbb{P}[Z \geq 1] \geq \mathbb{P}[Z \geq c^2/2] \geq 1 - 16/c^2$.

Proof. Suppose that $H : \{0, 1\}^n \rightarrow \{0, 1\}^n$. Let $EXEC$ denote the distribution over full executions of the protocol, over the randomness in r_A, r_B, H . We will consider **partial executions** of the protocol, namely the execution up to round (and query) i , for $1 \leq i \leq 2\ell$. Let the random variable $\mathcal{X}_{\leq i} \subset \{0, 1\}^n$ denote the queries by **A** in the execution up to round i , $\mathcal{Y}_{\leq i} \subset \{0, 1\}^n$ denote the queries by **B** in the execution up to round i , and $\mathcal{Z}_{\leq i} \subset \{0, 1\}^n$ denote the queries of **E** up to round i (we may consider **E** as making queries to H as it observes more of the transcript each round). Let M_i denote the message transferred at round i , and $M_{\leq i}$ denote the transcript up to round i .

Following [2], for $1 \leq i \leq \ell$, consider some fixed $M_{\leq i}, \mathcal{Z}_{\leq i}, H(\mathcal{Z}_{\leq i})$; we drop the subscripts for convenience. Then let $EXEC(M, \mathcal{Z}, H(\mathcal{Z}))$ denote the distribution of (partial, up to the i th step) views

$$(\mathcal{A}, \mathcal{B}) := ((r_A, (\mathcal{X}, H(\mathcal{X})), M), (r_B, (\mathcal{Y}, H(\mathcal{Y})), M)) \quad (1)$$

conditioned on the partial transcript M and **E**'s queries $\mathcal{Z}, H(\mathcal{Z})$. Let $GOOD(M, \mathcal{Z}, H(\mathcal{Z}))$ be the event over pairs $(\mathcal{A}, \mathcal{B})$ in the support of $EXEC(M, \mathcal{Z}, H(\mathcal{Z}))$ that $\mathcal{X} \cap \mathcal{Y} \subset \mathcal{Z}$. Finally let $GEXEC(M, \mathcal{Z}, H(\mathcal{Z}))$ be the distribution of $EXEC(M, \mathcal{Z}, H(\mathcal{Z}))$ conditioned on $GOOD(M, \mathcal{Z}, H(\mathcal{Z}))$.

The attacker **E** now works as follows:

1. **E** is given a parameter $\epsilon < 1/10$, which is specified below.
2. At each round i , **E** does the following:
 - (a) While there is $x \in \{0, 1\}^n$ such that $\mathbb{P}_{GEXEC(M, \mathcal{Z}, H(\mathcal{Z}))}[x \in \mathcal{X} \cap \mathcal{Y}] \geq \epsilon/\ell$, **E** sets $\mathcal{Z} \leftarrow \mathcal{Z} \cup \{x\}$, and queries $H(x)$.
 - (b) If $|\mathcal{Z}| \geq \ell^2/\epsilon^2$, then abort (this will turn out to be a low probability event).
3. At the end (i.e. after round 2ℓ), **E** samples $(\hat{\mathcal{A}}_{\leq 2\ell}, \hat{\mathcal{B}}_{\leq 2\ell}) \leftarrow_R GEXEC(M_{\leq 2\ell}, \mathcal{Z}_{\leq 2\ell}, H(\mathcal{Z}_{\leq 2\ell}))$, and outputs $out_A(\hat{\mathcal{A}}_{2\ell})$, as is determined (deterministically) from **A**'s sampled view $\hat{\mathcal{A}}_{2\ell}$.

For odd $i \in [2\ell]$, let $FAIL_i$ denote the event that $(\mathcal{Z}_{\leq i-1} \subset \mathcal{X}_{\leq i-1} \cap \mathcal{Y}_{\leq i-1}) \wedge (\mathcal{X}_i \in \mathcal{Y}_{\leq i-1}) \wedge (\mathcal{X}_i \notin \mathcal{Z}_{\leq i})$, namely that the query made by **A** in the i th round was previously made by **B** and has not been made by \mathcal{Z} . Let $FAIL_i$ denote the analogous event for even i . Let $FAIL = \bigvee_i FAIL_i$, which certainly contains the event that $\mathcal{Z}_{\leq 2\ell} \not\subset \mathcal{X}_{\leq 2\ell} \cap \mathcal{Y}_{\leq 2\ell}$.

The following lemma is key in the proof:

Lemma 2.3. $\mathbb{P}_{EXEC}[FAIL] \leq 3\epsilon$.

Proof. By the union bound, it suffices to show that for each i , $\mathbb{P}_{EXEC}[FAIL_i | \mathcal{X}_{\leq i-1} \cap \mathcal{Y}_{\leq i-1} \subseteq \mathcal{X}_{\leq i-1}] \leq \frac{3\epsilon}{2\ell}$. For simplicity we will write $\mathcal{X}_{\leq i-1} = \mathcal{X}_{< i}$, and similarly for $\mathcal{Y}_{< i}, \mathcal{Z}_{< i}$. Note that for each $(\mathcal{A}_{< i}, \mathcal{B}_{< i}) \in \text{Supp}(GEXEC(M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i})))$, then $GOOD(M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i}))$ holds, so

$$\begin{aligned} \mathbb{P}_{GEXEC(M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i}))}[(\mathcal{A}_{< i}, \mathcal{B}_{< i})] &= \frac{\mathbb{P}_{EXEC(M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i}))}[(\mathcal{A}_{< i}, \mathcal{B}_{< i})]}{\mathbb{P}_{EXEC(M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i}))}[GOOD(M_{< i}, \mathcal{Z}_{< i})]} \\ &= \frac{\mathbb{P}_{EXEC}[(\mathcal{A}_{< i}, \mathcal{B}_{< i}, M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i}))]}{\mathbb{P}_{EXEC}[(M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i}))] \cdot \mathbb{P}_{EXEC(M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i}))}[GOOD(M_{< i}, \mathcal{Z}_{< i})]} \\ &= \frac{2^{-|r_A|} \cdot 2^{-|r_B|} \cdot 2^{-\ell \cdot |\mathcal{X}_{< i} \cup \mathcal{Y}_{< i} \cup \mathcal{Z}_{< i}|}}{c(M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i}))} \\ &= \frac{\alpha(\mathcal{A}_{< i}) \cdot \beta(\mathcal{A}_{< i}) \cdot 2^{-\ell |\mathcal{Z}_{< i}|}}{c(M_{< i}, \mathcal{Z}_{< i}, H(\mathcal{Z}_{< i}))}, \end{aligned}$$

where c is some function, and $\alpha(\mathcal{A}_{<i}) = 2^{-|r_A|} \cdot 2^{-\ell \cdot |\mathcal{X}_{<i} \setminus \mathcal{Z}_{<i}|}$, $\beta(\mathcal{B}_{<i}) = 2^{-|r_B|} \cdot 2^{-\ell \cdot |\mathcal{Y}_{<i} \setminus \mathcal{Z}_{<i}|}$ depend only on $\mathcal{A}_{<i}, \mathcal{B}_{<i}$ (for fixed $M_{<i}, \mathcal{Z}_{<i}, H(\mathcal{Z}_{<i})$), and the last equality holds since the event $GOOD(M_{<i}, \mathcal{Z}_{<i}, H(\mathcal{Z}_{<i}))$ holds. In particular, there are distributions $\mathcal{D}_{\mathcal{A}, <i}, \mathcal{D}_{\mathcal{B}, <i}$ over $\mathcal{A}_{<i}, \mathcal{B}_{<i}$ (namely, they are proportional to $\alpha(\cdot), \beta(\cdot)$) such that $GEXEC(M_{<i}, \mathcal{Z}_{<i}, H(\mathcal{Z}_{<i})) = \mathcal{D}_{\mathcal{A}, <i} \otimes \mathcal{D}_{\mathcal{B}, <i} | GOOD(M_{<i}, \mathcal{Z}_{<i}, H(\mathcal{Z}_{<i}))$.

Next, for fixed $M_{<i}, \mathcal{Z}_{<i}, H(\mathcal{Z}_{<i})$, we may define the following bipartite graph $G = (V_A, V_B, E)$. For each view $\mathcal{A}_{<i}$ there are a collection of nodes $u \in V_A$, for which we will write $\mathcal{A}_u = \mathcal{A}_{<i}$, and the number of these corresponds to the distribution $\mathcal{D}_{\mathcal{A}, <i}$ from above. Therefore, the uniform distribution $u \leftarrow_R V_A$ corresponds to drawing $\mathcal{A}_{<i} \leftarrow_R \mathcal{D}_{\mathcal{A}, <i}$. Similarly, the nodes $v \in V_B$ are defined with respect to $\mathcal{D}_{\mathcal{B}, <i}$. For $u \in V_A$, define $\mathcal{Q}_u = \mathcal{X}(\mathcal{A}_u) \setminus \mathcal{Z}_{<i}$, where $\mathcal{X}(\mathcal{A}_u)$ is the set $\mathcal{X}_{<i}$ corresponding to \mathcal{A}_u (see (1)). Similarly, for $v \in V_B$, define $\mathcal{Q}_v = \mathcal{Y}(\mathcal{B}_v) \setminus \mathcal{Z}_{<i}$. We define E by: $(u, v) \in E$ iff $\mathcal{Q}_u \cap \mathcal{Q}_v = \emptyset$, i.e. if the corresponding views of \mathcal{A}, \mathcal{B} do not have any intersection queries not in $\mathcal{Z}_{<i}$. Then by definition of $\mathcal{D}_{\mathcal{A}, <i}, \mathcal{D}_{\mathcal{B}, <i}$, the distribution of $GEXEC(M_{<i}, \mathcal{Z}_{<i}, H(\mathcal{Z}_{<i}))$ is equal to the distribution obtained by choosing $(\mathcal{A}_u, \mathcal{B}_v)$ for $v \leftarrow_R E$.

We next make the following claim.

Claim 2.4. *For every $u \in V_A$, $d_G(u) \geq |V_B|(1 - 2\epsilon)$, and for every $v \in V_B$, $d_G(v) \geq |V_A|(1 - 2\epsilon)$.*

To prove the above claim, note that if for some $u \in V_A$, $\sum_{v \in V_B, (u,v) \notin E} d_G(v) > \epsilon|E|$, then $\mathbb{P}_{(u,v) \leftarrow_R E}[\mathcal{Q}_v \cap \mathcal{Q}_u \neq \emptyset] > \epsilon$, so since $|\mathcal{Q}_u| \leq |\mathcal{X}(\mathcal{A}_u)| \leq \ell$, there is some $x \in \mathcal{Q}_u$ that has conditional probability at least ϵ/ℓ of being in $\mathcal{X}_{<i} \cap \mathcal{Y}_{<i}$, which is impossible, since then $x \in \mathcal{Z}_{<i}$ by the definition of E . Similarly, for each $v \in V_B$, $\sum_{u \in V_A, (u,v) \notin E} d_G(u) > \epsilon|E|$. The claim now follows by a simple counting argument (we refer to reader to [2] for details).

We now complete the proof of Lemma 2.3 by showing that $\mathbb{P}_{EXEC}[FAIL_i | \mathcal{X}_{\leq i-1} \cap \mathcal{Y}_{\leq i-1} \subseteq \mathcal{X}_{\leq i-1}] \leq \frac{3\epsilon}{2\ell}$. We assume that i is odd for simplicity. In fact, we prove even more, namely that for each $v \in V_A$, if we let $\mathcal{A}_{<i} = \mathcal{A}_v$, then

$$\mathbb{P}_{EXEC}[(\mathcal{X}_i \in \mathcal{Y}_{\leq i-1}) \wedge (\mathcal{X}_i \notin \mathcal{Z}_{\leq i}) | \mathcal{X}_{\leq i-1} \cap \mathcal{Y}_{\leq i-1} \subseteq \mathcal{X}_{\leq i-1}, \mathcal{A}_{<i}] \leq \frac{3\epsilon}{2\ell}.$$

Choosing a random edge of G conditioned on one of its vertices being u is the same as choosing a random neighbor of u , so we may note that, if $\mathcal{S} = \{v \in V_B : \mathcal{X}_i \in \mathcal{Y}(\mathcal{B}_v)\}$,

$$\mathbb{P}_{v \leftarrow_R N_G(u)}[\mathcal{X}_i \in \mathcal{Y}(\mathcal{B}_v)] \leq \frac{|\mathcal{S}|}{d_G(u)} \leq \frac{|\mathcal{S}|}{(1 - 2\epsilon)|V_A|} \leq \frac{|\mathcal{S}||V_B|}{(1 - 2\epsilon)|E|} \leq \frac{\sum_{v \in \mathcal{S}} d_G(v)}{(1 - 2\epsilon)^2|E|} \leq \frac{\epsilon}{(1 - 2\epsilon)^2\ell} \leq \frac{3\epsilon}{2\ell},$$

where we have used Claim 2.4 in the second and fourth inequalities, and the fact that if $\frac{1}{|E|} \sum_{v \in \mathcal{S}} d_G(v) > \epsilon/\ell$, then E would have queried \mathcal{X}_i , which also belongs to $\mathcal{Y}(\mathcal{B}_v)$ for each $v \in \mathcal{S}$, in the fifth inequality. \square

To complete the proof of Theorem 2.2, the following lemma is also needed:

Lemma 2.5. *The probability, over $EXEC$, that $|\mathcal{Z}_{\leq i}| \geq \ell^2/\epsilon^2$ at any step i (and therefore E has to abort), is at most 10ϵ .*

We do not prove the above lemma here.³

³Note that for the 2-round protocols considered in the following section, an analogous lemma is needed, though its proof is much simpler since it can simply be shown that at each round no more than $O(\ell^2)$ queries are made. The same is not true of this proof, as there can be $\Omega(\ell)$ rounds.

Finally, consider the output $out_{\mathbf{E}} = out_{\mathbf{A}}(\hat{\mathcal{A}}_{\leq 2\ell})$, as was defined in the definition of the adversary \mathbf{E} . For convenience we will denote views of \mathbf{E} by $\mathcal{E} = (M, \mathcal{Z}, H(\mathcal{Z}))$. Write $GOOD(\mathcal{A}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})$ to be the event that $\mathcal{X}(\mathcal{A}_{\leq 2\ell}) \cap \mathcal{Y}(\mathcal{B}_{\leq 2\ell}) \subseteq \mathcal{Z}(\mathcal{E}_{\leq 2\ell})$.

By Lemma 2.3 and Lemma 2.5, $\mathbb{P}_{EXEC}[GOOD(\mathcal{A}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})] \geq 1 - 13\epsilon$. Next, note that $\hat{\mathcal{A}}_{\leq 2\ell}$ is independent of the pair $(\mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})$, and each of the ℓ elements of $\mathcal{Y}(\mathcal{B}_{\leq 2\ell})$ has at most ϵ/ℓ chance of being in $\hat{\mathcal{A}}_{\leq 2\ell}$, so $\mathbb{P}[GOOD(\hat{\mathcal{A}}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})] \geq 1 - \epsilon$. Thus, the proof of the theorem follows from the following inequalities:

$$\begin{aligned}
& \mathbb{P}_{EXEC(\mathcal{E}_{\leq 2\ell})}[out_{\mathbf{A}}(\mathcal{A}_{\leq 2\ell}) = out_{\mathbf{B}}(\mathcal{B}_{\leq 2\ell}) | GOOD(\mathcal{A}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})] \\
& - \mathbb{P}_{EXEC(\mathcal{E}_{\leq 2\ell})}[out_{\mathbf{A}}(\hat{\mathcal{A}}_{\leq 2\ell}) = out_{\mathbf{B}}(\mathcal{B}_{\leq 2\ell}) | GOOD(\hat{\mathcal{A}}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})] \\
\leq & \Delta((\mathcal{A}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}) | GOOD(\mathcal{A}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}), (\hat{\mathcal{A}}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}) | GOOD(\hat{\mathcal{A}}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})) \\
= & \mathbb{E}_{\mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}}[\Delta(\mathcal{A}_{\leq 2\ell} | \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}, GOOD(\mathcal{A}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}), \hat{\mathcal{A}}_{\leq 2\ell} | \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}, GOOD(\hat{\mathcal{A}}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}))] \\
\leq & 2\epsilon,
\end{aligned}$$

where in the last two lines we are abusing notation slightly, and actually mean the statistical distance between the distributions, conditioned on the $GOOD(\cdot)$ events and on $\mathcal{E}_{\leq 2\ell}$. The second-to-last inequality is a general fact, so the only non-obvious step is the last inequality. To show this, we consider the graph $G = (V_{\mathbf{A}}, V_{\mathbf{B}}, E)$, for step $i = 2\ell$, as defined above. Then the distribution of $\mathcal{A}_{\leq 2\ell} | \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}, GOOD(\mathcal{A}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})$ is given by choosing $v \in V_{\mathbf{B}}$ with $\mathcal{B}_v = \mathcal{B}_{\leq 2\ell}$ and then choosing $u \leftarrow_R N_G(v)$ and outputting \mathcal{A}_u . On the other hand, the distribution $\hat{\mathcal{A}}_{\leq 2\ell} | \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell}, GOOD(\hat{\mathcal{A}}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})$ is given by choosing $v \in V_{\mathbf{B}}$ with $\mathcal{B}_v = \mathcal{B}_{\leq 2\ell}$, and then choosing $u \in N_G(v)$ with probability in proportion to $d_V(u)$ (note that we choose $u \in N_G(v)$, not $u \in V_{\mathbf{A}}$, since we condition on $GOOD(\hat{\mathcal{A}}_{\leq 2\ell}, \mathcal{B}_{\leq 2\ell}, \mathcal{E}_{\leq 2\ell})$). But by Claim 2.4, each $u \in V_{\mathbf{A}}$ satisfies $(1 - 2\epsilon)|V_{\mathbf{B}}| \leq d_G(u) \leq |V_{\mathbf{B}}|$, from which it follows that the statistical distance (i.e. sum of absolute value of difference of probabilities for each $u \in N_G(v)$) is at most 2ϵ . \square

3 Communication complexity of key agreement protocols

Haitner et al. [4] prove lower bounds on the communication complexity of oracle-aided key agreement protocols, and as such, it suffices to prove these bounds for the simplest case where the key length of the agreed-upon key is 1, i.e. $n = 1$ in Definition 2.1.

An oracle-aided protocol is **non-adaptive** if the queries that \mathbf{A}, \mathbf{B} make to H depend only on their respective random coins $r_{\mathbf{A}}, r_{\mathbf{B}}$ (i.e. they fix these queries before the start of the protocol). A protocol is **uniform query** if there are sets $S_{\mathbf{A}}, S_{\mathbf{B}}$, such that \mathbf{A} 's queries are chosen uniformly from $S_{\mathbf{A}}$, and \mathbf{B} 's queries are chosen uniformly from $S_{\mathbf{B}}$. Note that uniform query protocols are a special case of non-adaptive protocols.

3.1 Uniform query protocols

The main result of [4] for uniform query protocols is the following. For any finite set \mathcal{S} , let $\mathcal{H}_{\mathcal{S}} := \{H : \mathcal{S} \rightarrow \{0, 1\}^*\}$. Note that for any fixed protocol $\Pi = (\mathbf{A}, \mathbf{B})$ with bounded time, both \mathbf{A}, \mathbf{B} will only use a bounded number of bits from $H(x)$, for $H \in \mathcal{H}, x \in \mathcal{S}$, meaning that we may always assume without loss of generality that $\mathcal{H}_{\mathcal{S}} = \{H : \mathcal{S} \rightarrow \{0, 1\}^m\}$, for sufficiently large m (given Π), so it makes sense to draw a uniformly random elements $H \leftarrow_R \mathcal{H}_{\mathcal{S}}$.

The following theorem is well-known [1]:

Theorem 3.1 (Set disjointness is hard, Bar-Yossef et al., 2004 [1]). *There is a distribution \mathcal{D}_{disj} over pairs of sets $(\mathcal{X}, \mathcal{Y})$, with $\mathcal{X}, \mathcal{Y} \subset [\ell]$, such that any protocol $\Pi = (\mathbf{A}, \mathbf{B})$ with public randomness $r_{\mathbf{P}}$ and that satisfies*

$$\mathbb{P}_{(\mathcal{X}, \mathcal{Y}) \leftarrow_R \mathcal{D}_{disj}, r_{\mathbf{P}} \leftarrow_R \{0,1\}^*} [\text{out}(\mathbf{A}(\mathcal{X}, r_{\mathbf{P}}), \mathbf{B}(\mathcal{Y}, r_{\mathbf{P}})) = \mathbb{I}[\mathcal{X} \cap \mathcal{Y} \neq \emptyset]] \geq 1 - \epsilon,$$

also has $CC(\Pi) \geq \Omega(\ell)$.

The most-commonly used distribution \mathcal{D}_{disj} for set-disjointness is the following [6]: suppose that we can write $\ell = 4k - 1$ for an integer k . Then choose a random partition $\mathcal{T} = (\mathcal{T}_1, \mathcal{T}_2, \{x\})$ of $[\ell]$ into 3 disjoint sets such that $|\mathcal{T}_1| = |\mathcal{T}_2| = 2k - 1$. Choose at random sets $\mathcal{X} \subset \mathcal{T}_1 \cup \{x\}, \mathcal{Y} \subset \mathcal{T}_2 \cup \{x\}$ such that $|\mathcal{X}| = |\mathcal{Y}| = k$. Then output $(\mathcal{X}, \mathcal{Y})$. Note that for each $(\mathcal{X}, \mathcal{Y}) \in \text{Supp}(\mathcal{D}_{disj})$, $|\mathcal{X} \cap \mathcal{Y}| \in \{0, 1\}$. Note, however, that the specific structure of distribution \mathcal{D}_{disj} for which set-disjointness is hard will not be used in the proof below.

Theorem 3.2 (Haitner et al. [4]). *If $\Pi = (\mathbf{A}, \mathbf{B})$ is a uniform query protocol relative to $\mathcal{H}_{\mathcal{S}}$ where \mathbf{A}, \mathbf{B} make at most ℓ queries each, then if Π is (q, α, γ) -secure, we have that $CC(\Pi) \geq \Omega\left(\frac{(1-\alpha-\gamma)^4 q^2}{\ell^3}\right)$.*

In particular, if Π is secure against an oracle making $c\ell^2$ queries, for some c , and the probability of success for any given oracle H is greater than the probability of any adversary \mathbf{E} discovering $\text{out}_{\mathbf{A}}$ by a constant, the communication complexity of Π must be $\Omega(\ell)$.

We next give the proof of Theorem 3.2, leaving out some minor details for the sake of space.

Proof. The crux of the proof is to introduce the following two protocols which use public and private randomness, respectively, to emulate the execution of Π , with the additional constraint that \mathbf{A} and \mathbf{B} are given as input ℓ -element sets $\mathcal{X} \subseteq \mathcal{S}, \mathcal{Y} \subseteq \mathcal{S}$ describing which queries they make in the execution of the protocol.

- Protocol $\Lambda_{pub} = (\mathbf{A}_{pub}, \mathbf{B}_{pub})$ is defined as follows: \mathbf{A}_{pub} and \mathbf{B}_{pub} use their public randomness $r_{\mathbf{P}}$ to generate a (common) description of a function $H : \mathcal{S} \rightarrow \{0, 1\}^*$. Then $\mathbf{A}_{pub}, \mathbf{B}_{pub}$ emulate the execution of Π , taking roles \mathbf{A}, \mathbf{B} respectively, and making queries from \mathcal{X}, \mathcal{Y} , respectively, using the public randomness H to emulate the oracle's responses, $H(\mathcal{X}), H(\mathcal{Y})$.
- Protocol $\Lambda_{pri} = (\mathbf{A}_{pri}, \mathbf{B}_{pri})$ is defined as follows: \mathbf{A}_{pri} and \mathbf{B}_{pri} use their private randomness $r_{\mathbf{A}}, r_{\mathbf{B}}$ to generate independent descriptions of functions $H_{\mathbf{A}}, H_{\mathbf{B}} : \mathcal{S} \rightarrow \{0, 1\}^*$. Then $\mathbf{A}_{pri}, \mathbf{B}_{pri}$ emulate the execution of Π , taking roles \mathbf{A}, \mathbf{B} respectively, and making queries from \mathcal{S}, \mathcal{Y} respectively, using the private randomness $H_{\mathbf{A}}, H_{\mathbf{B}}$ to emulate the oracles' responses, $H_{\mathbf{A}}(\mathcal{X}), H_{\mathbf{B}}(\mathcal{Y})$.

Let \mathcal{D}_U be the distribution of pairs $(\mathcal{X}, \mathcal{Y})$, where \mathcal{X}, \mathcal{Y} are independent and uniform ℓ -element subset of \mathcal{S} ; note that \mathcal{D}_U is the distribution of query sets made by parties \mathbf{A}, \mathbf{B} in execution of Π . Let $\Lambda_{pri}(\mathcal{D}_U), \Lambda_{pub}(\mathcal{D}_U)$ denote the distribution of the joint view v where $(\mathcal{X}, \mathcal{Y}) \leftarrow_R \mathcal{D}_U$. Note that since Λ_{pri} is a no-oracle protocol, the views $v_{\mathbf{A}}, v_{\mathbf{B}}$ (and therefore the outputs $\text{out}_{\mathbf{A}}, \text{out}_{\mathbf{B}}$) are conditionally independent given the transcript $m = \text{trans}(v)$ (where v denotes the joint view), since they are deterministic functions of the messages m and the random bits $(r_{\mathbf{A}}, \mathcal{X}), (r_{\mathbf{B}}, \mathcal{Y})$, respectively (which are independent). Therefore, given the transcript an adversary \mathbf{E} can sample a pair $(r_{\mathbf{B}}, \mathcal{Y})$ consistent with the transcript, and \mathbf{E} can then compute $\mathbf{E}(\text{trans}(v)) := \text{out}_{\mathbf{B}}(r_{\mathbf{B}}, \mathcal{Y}, m)$

as a deterministic function of the sampled r_B, \mathcal{Y} as well as m . Moreover, this output will be equal to the output of $A_{pri}, out_{A_{pri}}(v)$ with the same probability that $out_{B_{pri}}(v) = out_{A_{pri}}(v)$, where $v \leftarrow_R \Lambda_{pri}(\mathcal{D}_U)$. It is also clear that $\Lambda_{pub}(\mathcal{D}_U)$ is exactly equal to the distribution of joint views v of the protocol $\Pi = (A, B)$, so since Π satisfies $(1 - \alpha)$ -accuracy and γ -secrecy, it must be the case that either

$$\mathbb{P}_{v \leftarrow_R \Lambda_{pub}(\mathcal{D}_U)} [out_{B_{pub}}(v) = out_{A_{pub}}(v)] - \mathbb{P}_{v \leftarrow_R \Lambda_{pri}(\mathcal{D}_U)} [out_{B_{pri}}(v) = out_{A_{pri}}(v)] \geq \frac{1 - \alpha - \gamma}{2}, \quad (2)$$

or

$$\mathbb{P}_{v \leftarrow_R \Lambda_{pri}(\mathcal{D}_U)} [E(\text{trans}(v)) = out_{A_{pri}}(v)] - \mathbb{P}_{v \leftarrow_R \Lambda_{pub}(\mathcal{D}_U)} [E(\text{trans}(v)) = out_{A_{pub}}(v)] \geq \frac{1 - \alpha - \gamma}{2}. \quad (3)$$

In particular, the second term in (2) and the first term in (3) are equal, and the remaining two terms must differ by at least $1 - \alpha - \gamma$ by $(1 - \alpha)$ -accuracy and γ -secrecy. Moreover, if (3) holds, then we may consider the protocols $\Lambda'_{pri} := (A'_{pri}, B'_{pri}), \Lambda'_{pub} := (A'_{pub}, B'_{pub})$, where $\Lambda'_{pri}, \Lambda'_{pub}$ are exactly identical to $\Lambda_{pri}, \Lambda_{pub}$, respectively, except for the fact that B'_{pri} and B'_{pub} output $\neg E(\text{trans}(v)) \in \{0, 1\}$ as the shared key, where v is the joint view from either the private or public randomness protocol, respectively. Then it follows directly from (3) that

$$\mathbb{P}_{v \leftarrow_R \Lambda_{pub}(\mathcal{D}_U)} [out_{B'_{pub}}(v) = out_{A'_{pub}}(v)] - \mathbb{P}_{v \leftarrow_R \Lambda_{pri}(\mathcal{D}_U)} [out_{B'_{pri}}(v) = out_{A'_{pri}}(v)] \geq \frac{1 - \alpha - \gamma}{2},$$

which is exactly (2) for protocols $\Lambda'_{pri}, \Lambda'_{pub}$. Therefore we may assume without loss that (2) holds.

Next, recall that a joint view v in the support of Λ_{pub} has the form $v = (\mathcal{X}, r_A, \mathcal{Y}, r_B, r_P)$, of which the transcript and outputs are deterministic functions. Write $\mathcal{X}(v) = \mathcal{X}, \mathcal{Y}(v) = \mathcal{Y}$ for the inputs corresponding to joint view v , and let, for $0 \leq i \leq \ell$,

$$Acc_{pub}(i) := \mathbb{P}_{v \leftarrow_R \Lambda_{pub}(\mathcal{D}_U)} [out_{B_{pub}}(v) = out_{A_{pub}}(v) | |\mathcal{X}(v) \cap \mathcal{Y}(v)| = i]$$

be the probability of agreement conditioned on an intersection size of i . We then have the following lemma, which follows from (2).

Lemma 3.3. *There is an integer $0 < u < \frac{4\ell^2}{|\mathcal{S}| \cdot (1 - \alpha - \gamma)}$ such that*

$$\max \{ Acc_{pri}(u) - Acc_{pri}(u - 1), Acc_{pub}(u) - Acc_{pub}(u - 1) \} \geq \frac{(1 - \alpha - \gamma)^2 |\mathcal{S}|}{32\ell^2}.$$

We refer the reader to [4] for the details of the proof of Lemma 3.3, but note here that it follows from Markov's inequality, basic rules of probability, and the fact that $Acc_{pub}(0) = Acc_{pri}(0)$ since both the public and private coin protocols have the same distribution when there are no intersections in the queries of A, B . Assuming the proof of the above lemma, we continue with the proof. We may assume without loss that the first term $Acc_{pri}(u) - Acc_{pri}(u - 1)$ in the above is the maximum; the proof for the other case (i.e. public randomness) is nearly identical. We now claim that the following protocol $\Pi_{disj} = (A_{disj}, B_{disj})$ solves set-disjointness with error at most ϵ , for some $\epsilon > 0$

1. A_{disj}, B_{disj} receive as inputs $\mathcal{X}, \mathcal{Y} \subset [\ell]$, such that $|\mathcal{X}| = |\mathcal{Y}| = \ell/4$ and $|\mathcal{X} \cap \mathcal{Y}| \in \{0, 1\}$.
2. Set $k = \frac{2^{13} \ell^4 \log(1/\epsilon)}{|\mathcal{S}|^2 (1 - \alpha - \gamma)^4}$. For k iterations $1 \leq j \leq k$, A_{disj}, B_{disj} perform the following interaction:

- (a) A_{disj}, B_{disj} simulate the interaction of A_{pri}, B_{pri} respectively on inputs $\sigma_j(\mathcal{X}'), \sigma_j(\mathcal{Y}')$, where $\sigma_1, \dots, \sigma_k$ are uniformly independent permutations of $[\ell]$, and

$$\mathcal{X}' = \mathcal{X} \cup \{\ell + 1, \dots, \ell + u - 1\} \cup \{2\ell + 1, \dots, 3\ell - \ell/4 - u + 1\},$$

and where \mathcal{Y}' is defined similarly with respect to \mathcal{Y} .

- (b) B_{disj} sends the output of this protocol $out_{B_{disj}}$ to A_{disj} , which increments a counter C if their outputs are equal (i.e. $out_{B_{disj}} = out_{A_{disj}}$).

3. A_{disj} produces the output of the overall protocol, which is $\mathbb{1} \left[\frac{C}{k} > \frac{Acc_{pri}(u) + Acc_{pri}(u-1)}{2} \right]$.

Assume next that $|\mathcal{S}| > \frac{3\ell}{1-\alpha-\gamma}$ (we will see below that assuming this is without loss), meaning that by Lemma 3.3, $u \leq 3\ell/4$, so that $|\mathcal{X}'| = |\mathcal{Y}'| = \ell$. Since the σ_j are independent and uniform permutations in S_ℓ , the probability that $out_{A_{disj}} = out_{B_{disj}}$ on each inner iteration of the above protocol is $Acc_{pri}(|\mathcal{X}' \cap \mathcal{Y}'|)$. By the Chernoff bound, it follows that

$$\mathbb{P} \left[\left| \frac{C}{k} - Acc_{pri}(|\mathcal{X}' \cap \mathcal{Y}'|) \right| > \frac{(1-\alpha-\gamma)^2 |\mathcal{S}|}{2^7 \ell^2} \right] < \epsilon, \quad (4)$$

by our choice of $k \geq \lceil \log(1/\epsilon)/2 \cdot \frac{(2^7)\ell^2}{(1-\alpha-\gamma)^4 |\mathcal{S}|^2} \rceil$ above.

Now, suppose that \mathcal{X}, \mathcal{Y} are in the support of \mathcal{D}_{disj} ; this means that $|\mathcal{X} \cap \mathcal{Y}| \in \{0, 1\}$. Also note that $|\mathcal{X}' \cap \mathcal{Y}'| = |\mathcal{X} \cap \mathcal{Y}| + u - 1$. Then Π_{disj} obtains the correct answer on inputs \mathcal{X}, \mathcal{Y} for set disjointness if and only if

$$\left| Acc_{pri}(u - 1 + |\mathcal{X} \cap \mathcal{Y}|) - \frac{C}{k} \right| < \frac{Acc_{pri}(u) + Acc_{pri}(u-1)}{4}.$$

By (4) and Lemma 3.3, this has probability at least $1 - \epsilon$ of happening, for each pair $(\mathcal{X}, \mathcal{Y})$. Therefore, the success probability of Π_{disj} is at least $1 - \epsilon$ on the distribution \mathcal{D}_{disj} . Moreover, its communication complexity is $O(k \cdot CC(\Lambda_{pri})) = O(k \cdot CC(\Pi))$. But by the choice of \mathcal{D}_{disj} we must have that this quantity is $\Omega(\ell)$, meaning that

$$CC(\Pi) \geq \Omega(\ell/k) = \Omega \left(\frac{|\mathcal{S}|^2 (1-\alpha-\gamma)^4}{\ell^3 \log(1/\epsilon)} \right).$$

Since an E that can query all of $H(\mathcal{S})$ can guess the key with probability $1 - \alpha$ by simulating A (conditioned on the transcript and the values of $H(\mathcal{S})$), we must have that $q < |\mathcal{S}|$ as long as $\gamma < 1 - \alpha$, meaning that in a (q, α, γ) -secure protocol Π , $CC(\Pi) \geq \Omega \left(\frac{(1-\alpha-\gamma)^4 q^2}{\ell^3} \right)$. \square

Note that this argument does not work for general non-adaptive protocols since in general, the probability of a pair of sets $(\mathcal{X}, \mathcal{Y})$ being the pair that is queried by A, B , does not only depend on $|\mathcal{X} \cap \mathcal{Y}|$, i.e. it may depend on particular elements in \mathcal{X} or \mathcal{Y} . Therefore, given input sets \mathcal{X}, \mathcal{Y} for set-disjointness, it is not simple to create independent samples $\mathcal{X}', \mathcal{Y}'$ of sets that can be queries of A, B , where $|\mathcal{X}' \cap \mathcal{Y}'|$ is related to $|\mathcal{X} \cap \mathcal{Y}|$ in some easily controllable way.

3.2 Non-adaptive protocols

For non-adaptive protocols in general, [4] proves a lower bound on the communication complexity of 2-round protocols, that is, protocols in which A sends B a message, then B sends A a message, and then A, B produce out_A, out_B , respectively. They consider oracle families \mathcal{H}_n , consisting of the set of all functions $H : \{0, 1\}^n \rightarrow \{0, 1\}^n$:

Theorem 3.4 (Haitner et al. [4]). *For any $n \in \mathbb{N}$, if Π is a 2-message non-adaptive ℓ -query protocol that is (q, α, γ) -secure, then $CC(\Pi) \geq \Omega\left(\frac{(1-\alpha-\gamma)^2 q}{\ell}\right)$.*

Proof. The proof proceeds via the construction of an E which is similar to the E used in [2], which queries all elements of $\mathcal{S} = \{0, 1\}^n$ which have high conditional probability of being in either \mathcal{X}, \mathcal{Y} , the sets of queries made by A, B, respectively. The difference is that the probability cutoff above which E queries elements is proportional to $1/CC(\Pi)$ as opposed to $1/\ell$, since the goal is to prove communication complexity lower bounds, rather than lower bounds on the number of queries A, B need to make to guarantee a given level of security.

Let $\mathcal{D}_A, \mathcal{D}_B$ be the distribution of queries $\mathcal{X}, \mathcal{Y} \subset \{0, 1\}^n$ made by A, B respectively.

1. Given query access to an oracle h , a parameter $\delta := 4\ell/q$, and a transcript $m = (m_1, m_2)$, E behaves as follows. Here we use H to denote a random variable $H \leftarrow_R \mathcal{H}_n$, and h to denote a particular draw of this random variable; the same notation is used for the transcript (M_1, M_2) in relation to the particular messages (m_1, m_2) .
2. E queries H on all elements in $\mathcal{E}_0 \cup \mathcal{E}_1$, where

$$\begin{aligned} \mathcal{E}_0 &= \{x \in \{0, 1\}^n : \mathbb{P}[x \in \mathcal{X} \cup \mathcal{Y}] \geq \delta\} \\ \mathcal{E}_1 &= \{x \in \{0, 1\}^n : \mathbb{P}[x \in \mathcal{X} \cup \mathcal{Y} | M_1 = m_1, H|_{\mathcal{E}_0} = h|_{\mathcal{E}_0}] \geq \delta\}, \end{aligned}$$

where the above probabilities are taken with respect to $(\mathcal{X}, \mathcal{Y}) \leftarrow_R \mathcal{D}_A \otimes \mathcal{D}_B$, where \otimes denotes a product distribution. Where necessary we will write $\mathcal{E}_0(\Pi), \mathcal{E}_1(\Pi)$ to emphasize dependence on the protocol Π .

3. E then outputs a random sample $k \leftarrow_R out_B|_{M=m, H|_{\mathcal{E}_0 \cup \mathcal{E}_1} = h|_{\mathcal{E}_0 \cup \mathcal{E}_1}}$ by simulating B conditioned on her queries to h and the transcript.

By γ -secrecy of the original protocol Π , there must exist some function $\tilde{h} : \mathcal{E}_0 \rightarrow \{0, 1\}^n$ such that if $\mathcal{H}_n^{\tilde{h}}$ denotes the set of all $h \in \mathcal{H}$ such that $h|_{\mathcal{E}_0} = \tilde{h}$, the protocol $\Pi^{\mathcal{H}_n^{\tilde{h}}}$, where the oracle is chosen as $H \leftarrow_R \mathcal{H}_n^{\tilde{h}}$, is $(q - |\mathcal{E}_0|, \alpha, \gamma)$ -secure. Indeed, if no such \tilde{h} existed, then we would have the following attacker \mathbf{E}_0 for the original protocol Π : \mathbf{E}_0 first queries \mathcal{E}_0 , receiving responses consistent with some function $\hat{h} : \mathcal{E}_0 \rightarrow \{0, 1\}^n$. \mathbf{E}_0 then simulates an attacker for the protocol $\Pi^{\mathcal{H}_n^{\hat{h}}}$, which we have assumed not to be $(q - |\mathcal{E}_0|, \alpha, \gamma)$ -secure. Note that in the protocol $\Pi^{\mathcal{H}_n^{\hat{h}}}$, A, B do not need to query $H|_{\mathcal{E}_0}$, meaning that $\mathcal{E}_0(\Pi') = \emptyset$.

Next, in any ℓ -query (q, α, γ) -secure protocol Π , consider the protocol $\Pi' = (A', B')$ where A', B' are the same as A, B, respectively, except for the following differences:

- B' makes one additional uniform query $\mathcal{Y}_{\ell+1}$ that is independent of all other queries.
- B' appends the bit $b := out^B \oplus (\mathcal{Y}_{\ell+1})_1$ (i.e. the first bit of the last query $\mathcal{Y}_{\ell+1}$) to the message M_2 .

- B' sets $out^{B'} = (\mathcal{Y}_{\ell+1})_1$.
- A' sets $out^{A'} = b \oplus out^A$.

Clearly the protocol Π' has the same success probability for each $H \in \mathcal{H}_n$ as does Π . It is also γ -secure since an adversary E which satisfies $\mathbb{P}_{v \leftarrow_R (\Pi')^H, H \leftarrow_R \mathcal{H}_n} [E(\text{trans}(v)) = out_{A'}(v)]$ can be converted into an adversary E for Π by drawing $b \leftarrow_R \{0, 1\}$ uniformly at random, simulating E using this bit as the additional message bit in the protocol Π , and outputting $b \oplus out_E$, where out_E denotes the output of E .

By the two paragraphs above, we can convert an ℓ -query, c -communication, (q, α, γ) -secure protocol Π into an $(\ell + 1)$ -query, $(c + 1)$ -communication, $(q - |\mathcal{E}_0(\Pi)|, \alpha, \gamma)$ -secure protocol Π' , which satisfies $\mathcal{E}_0(\Pi') = \emptyset$ and the first bit of the last query of B' is equal to $out_{B'}$.

Next, from the point of view of E , the distribution of the sets $(\mathcal{X}, H(\mathcal{X}))$ and \mathcal{Y} are “almost independent”. Formally, we have the following. Let $v_E = (M, H(\mathcal{E}_1))$ be the random variable representing the view of E in execution of the protocol Π , and let Π_E^H be the distribution of v_E , where $H \leftarrow_R \mathcal{H}_n$. (Recall that we may assume that $\mathcal{E}_0 = \emptyset$.) For distributions $\mathcal{D}_1, \mathcal{D}_2$, let $\Delta(\mathcal{D}_1, \mathcal{D}_2)$ be the statistical distance between $\mathcal{D}_1, \mathcal{D}_2$. Moreover, given random variables Z, W , we let $\mathcal{D}(Z)$ be the distribution taken by Z , and for a specific value w in the sample space of W , let $\mathcal{D}(Z|_w)$ be the distribution of Z conditioned on the event $W = w$. Then:

Lemma 3.5. *We have*

$$\mathbb{E}_{H \leftarrow_R \mathcal{H}_n, v_E \leftarrow_R \Pi_E^H} [\Delta(\mathcal{D}(\mathcal{X}, H(\mathcal{X}), \mathcal{Y}|_{v_E}), \mathcal{D}(\mathcal{X}, H(\mathcal{X})|_{v_E}) \otimes \mathcal{D}(\mathcal{Y}|_{v_E}))] \leq 25\sqrt{\delta \cdot (CC(\Pi) + 5)}.$$

As the proof of Lemma 3.5 is quite tedious and is unlikely to extend (without significant modifications) to non-adaptive protocols with more than 2 rounds, we omit the proof here; the reader is referred to [4] (and the appendix therein) for details.

Assuming Lemma 3.5, we complete the proof of the theorem. We first claim that for the value $\delta = 4\ell/q$ chosen above, E does not make more than q queries (for the original protocol Π , as opposed to the modified protocol discussed in the above paragraphs). We will then show that if the communication of the protocol Π is too low, then E can recover out_A with probability higher than γ . To see our first claim, note that $|\mathcal{E}_0| \leq |\mathcal{X} \cup \mathcal{Y}|/\delta \leq 2\ell/(4\ell/q) = q/2$ since both $|\mathcal{X}| = |\mathcal{Y}| = \ell$. The same argument applies to $|\mathcal{E}_1|$ for all values of m_1 and $h|_{\mathcal{E}_0}$, meaning that $|\mathcal{E}_0 \cup \mathcal{E}_1| \leq q$ with probability 1, as desired.

Since we have assumed that out_B is the first bit of $\mathcal{Y}_{\ell+1}$, it is a deterministic function of \mathcal{Y} . Similarly, conditioned on $v_E = (M_1, M_2, H(\mathcal{E}_1))$, out_A is a deterministic function of $\mathcal{X}, H(\mathcal{X})$, and r_A . Then using the data processing inequality for statistical distance as well as the fact that conditioned on v_E , E produces out_E by sampling $out_B|_{v_E}$ independently,

$$\begin{aligned} 1 - \alpha - \gamma &\leq \mathbb{E}_{H \leftarrow_R \mathcal{H}_n, v_E \leftarrow_R \Pi_E^H} [|\mathbb{P}[out_A = out_B|_{v_E}] - \mathbb{P}[out_A = out_E|_{v_E}]|] \\ &\leq \mathbb{E}_{H \leftarrow_R \mathcal{H}_n, v_E \leftarrow_R \Pi_E^H} [\Delta(\mathcal{D}(out_A, out_B|_{v_E}), \mathcal{D}(out_A|_{v_E}) \otimes \mathcal{D}(out_B|_{v_E}))] \\ &\leq \mathbb{E}_{H \leftarrow_R \mathcal{H}_n, v_E \leftarrow_R \Pi_E^H} [\Delta(\mathcal{D}(\mathcal{X}, H(\mathcal{X}), out_B|_{v_E}), \mathcal{D}(\mathcal{X}, H(\mathcal{X})|_{v_E}) \otimes \mathcal{D}(out_B|_{v_E}))] \\ &\leq \mathbb{E}_{H \leftarrow_R \mathcal{H}_n, v_E \leftarrow_R \Pi_E^H} [\Delta(\mathcal{D}(\mathcal{X}, H(\mathcal{X}), \mathcal{Y}|_{v_E}), \mathcal{D}(\mathcal{X}, H(\mathcal{X})|_{v_E}) \otimes \mathcal{D}(\mathcal{Y}|_{v_E}))] \\ &\leq 25\sqrt{\delta(CC(\Pi) + 5)}, \end{aligned}$$

where we have used α -accuracy and γ -security in the first inequality, and Lemma 3.5 in the last inequality. Recalling that $\delta = 4\ell/q$, this immediately implies the desired lower bound on $CC(\Pi)$. \square

4 Observations, further questions

In light of Merkle's puzzles, both Theorems 3.2 and 3.4 prove bounds on the communication complexity of key-agreement protocols that are tight when $q = \delta \cdot \ell^2$ for sufficiently small δ . In particular, note that Merkle's puzzles are a 2-round uniform protocol, so the hypotheses of both theorems apply. However, it is natural to ask whether they are tight for lower levels of security, i.e. $\ell < q \ll \ell^2$. In fact, Theorem 3.2 is definitely not tight for 2-round uniform protocols for q in this range, in light of Theorem 3.4, as

$$\frac{(1 - \alpha - \gamma)^2 q}{\ell} \gg \frac{(1 - \alpha - \gamma)^4 q^2}{\ell^3},$$

if we hold α, γ constant and let $\ell, q \rightarrow \infty$ with $\ell < q \ll \ell^2$.

We do have, however, the following 2-round, uniform-queries protocol generalizing Merkle's puzzles, which achieves $\Theta(q/\ell)$ communication, showing that Theorem 3.4 is tight (at least for 2-round protocols). We use the same notation as in Section 2.3.

1. Let t, λ be parameters (describing the communication/security tradeoff, as we describe further below).
2. A chooses $x_1, \dots, x_{c\lambda\sqrt{t}} \in [\lambda^2]$ uniformly at random, and sets $a_i = H(x_i)$, $1 \leq i \leq c\lambda$. A then sends B the set $\mathcal{T}_A := \{a_i : 0 \leq a_i < 2^m/t\}$, where strings $a \in \{0, 1\}^m$ are interpreted as base-2 integers in $\{0, \dots, 2^m - 1\}$. If the size of this set is greater than $c''c\lambda/\sqrt{t}$, for some large constant c'' to be specified below, then A instead aborts the protocol with \perp_A .
3. B choose $y_1, \dots, y_{c\lambda\sqrt{t}} \in [\lambda^2]$ uniformly at random, and sets $b_i = H(y_i)$, $1 \leq i \leq c\lambda$. B then sends A the set $\mathcal{T}_B := \{b_i : 0 \leq b_i < 2^m/t\}$. If the size of this set is greater than $c''c\lambda/\sqrt{t}$, then B instead aborts the protocol with \perp_B .
4. A then picks j as small as possible so that $a_j \in \mathcal{T}_A$, and there is $b \in \mathcal{T}_B$ such that $b = a_j$. If such a j does not exist, then A halts without output \perp_A . If such a j does exist, then A sends $out_A = x_j$.
5. B similarly picks j as small as possible such that $b_j \in \mathcal{T}_B$ and there is $a \in \mathcal{T}_A$ such that $b_j = a$ (halting with \perp_B if j does not exist). B then sets $out_B = y_j$.

The following is now straightforward:

Proposition 4.1. *For sufficiently large c, c' , and $0 < \delta < 1/2$, the above protocol is a $c\lambda\sqrt{t}$ -query, $(\delta\lambda^2, 0.01, \Theta(\delta))$ -secure key agreement protocol with communication bounded by $O(\lambda/\sqrt{t})$.*

For fixed δ , if we let $q = \delta\lambda^2$ and $\ell = \lambda\sqrt{t}$, then the communication is $\Theta(q/\ell)$, which holds (in particular) for $\ell < q \ll \ell^2$, as can be obtained by varying t . This shows tightness of the bound of Theorem 3.4.

Proof of Proposition 4.1. First, it is evident by construction of the protocol that the communication is at most $2c''c\lambda/t$. As is the case for Merkle's puzzles, if $2^m > c'\lambda^4$, then $H(1), \dots, H(\lambda^2)$ are distinct with probability at least $1 - \frac{1}{c'}$, so with all but probability $1/c'$, the existence of j, j' such that $a_j = b_{j'}$ implies that $x_j = y_{j'}$.

Next we upper bound the probability that either A cannot find j in step 3 or B cannot find j in step 4 of the above protocol. The proof is very similar to that for Merkle's puzzles: let $Z_{jj'}$

denote the indicator that $x_j = y_{j'}$ and that $H(x_j) < 2^m/t$. Assuming that t is a power of 2 for simplicity, note that $\mathbb{E}[Z_{jj'}] = \frac{1}{\lambda^2} \cdot \frac{1}{t}$. Now let $Z = \sum_{j,j'} Z_{jj'}$ denote the number of pairs (j, j') such that $x_j = y_{j'}$ and $H(x_j) = H(y_{j'}) < 2^m/t$, which is certainly a lower bound on the number of pairs (j, j') such that $a_j = b_{j'} < 2^m/t$. Then

$$\mathbb{E}[Z] = \sum_{j,j'} \mathbb{E}[Z_{jj'}] = (c\lambda\sqrt{t})^2 \cdot \frac{1}{t\lambda^2} = c^2.$$

Similarly, by the same calculation as in Section 2.3, replacing λ with $\lambda\sqrt{t}$ gives that $\text{Var}[Z] \leq 3c^2$ as long as $\lambda\sqrt{t} \geq c$. Then Chebyshev's inequality gives us that $\mathbb{P}[Z \geq 1] \geq 1 - 16/c^2$.

Next we upper bound the probability that either A or B aborts because either $|\mathcal{T}_A| > c''c\lambda/\sqrt{t}$ or $|\mathcal{T}_B| > c''c\lambda/\sqrt{t}$, respectively. For A, since the probability that any given $x_i \in \mathcal{T}_A$ is $1/t$, $\mathbb{E}[|\mathcal{T}_A|] = c\lambda/\sqrt{t}$, so this probability is

$$\mathbb{P}[|\mathcal{T}_A| > c''c\lambda/\sqrt{t}] < \exp(-(c'' - 1)/3),$$

as long as $c\lambda/\sqrt{t} > 1$, by the Chernoff bound. The same probability holds for \mathcal{T}_B .

Assuming that H is injective on $[\lambda^2]$ and that A, B do not abort in any of the ways described and bounded above, the protocol is successful; so, the probability of success is at least

$$1 - \frac{1}{c'} - \frac{16}{c^2} - 2\exp(-(c'' - 1)/3),$$

which can be made arbitrarily close to 1 (i.e. greater than 0.99) by choosing c, c', c'' to be sufficiently large.

Finally, conditioned on success of the protocol and the transcript, the key $out_A = out_B$ is uniformly random in $[\lambda^2]$, so an attacker with $\delta\lambda^2$ queries can recover the key with probability $\Theta(\delta)$. Note that there are a few additional details needed to make this rigorous: namely, we must show that even if the attacker E makes a query x such that $H(x)$ is among the messages sent in the transcript, then E will learn little. This can be accomplished by noting that by a symmetry argument, for any $\mathcal{S} \subset [\lambda^2]$, such that if j, j' is the lexicographically smallest pair with $a_j \in M_A, b_{j'} \in M_B, a_j = b_{j'} \notin H(\mathcal{S})$, the distribution of out_A conditioned on $M_A, M_B, (\mathcal{S}, H(\mathcal{S}))$ is uniform on $[\lambda^2] \setminus \mathcal{S}$. That is, for such sets \mathcal{S} ,

$$H(out_A | M_A, M_B, (\mathcal{S}, H(\mathcal{S}))) = \log(\lambda^2 - |\mathcal{S}|).$$

The proof now follows by induction on the number of queries made by the attacker. Suppose the attacker E has already made ν queries, $\nu \geq 0$, and that the queries are denoted by some set $\mathcal{S}_\nu \subset [\lambda^2]$, $|\mathcal{S}_\nu| \leq \nu$. Let $x_{\nu+1}$ denote the next query of E. Then

$$\mathbb{P}[H(x_{\nu+1}) = a_j | M_A, M_B, (\mathcal{S}_\nu, H(\mathcal{S}_\nu))] \leq \frac{1}{\lambda^2 - |\mathcal{S}_\nu|} \leq 1/(\lambda^2 - \nu),$$

where $a_j = H(out_A)$ is as defined above (as a function of M_A, M_B). If $x_{\nu+1} \neq a_j$, then we set $\mathcal{S}_{\nu+1} = \mathcal{S}_\nu \cup \{x_{\nu+1}\}$ and repeat the same argument. Using the union bound, the probability that $H(x_{\nu+1}) \neq a_j$ at each query ν of the attacker, where the attacker makes a total of $q = \delta \cdot \lambda^2$ queries is then bounded above by

$$\delta\lambda^2 \cdot \frac{2}{\lambda^2} = \delta/2,$$

where we assumed that $\delta < 1/2$. Since we can assume without loss of generality that the attacker's last query is equal to its output (by increasing q by 1), this completes the proof of security of the protocol. \square

The following seems to be true:

Conjecture 4.2. *If $\Pi = (A, B)$ is an ℓ -query aided protocol relative to \mathcal{H}_n that is (q, α, γ) secure, then $CC(\Pi) \geq \Omega\left(\frac{\text{poly}(1-\alpha-\gamma)q}{\ell}\right)$.*

Obviously, the above is not even known yet for non-adaptive protocols of more than 2 rounds, and trying to solve it for simpler cases (e.g. non-adaptive protocols of 3 rounds, or even all uniform-query protocols) is a natural next step.

References

- [1] Ziv Bar-Yossef, T.S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences*, 68(4):702–732, June 2004.
- [2] Boaz Barak and Mohammad Mahmoody-Ghidary. Merkle Puzzles Are Optimal: An $O(n^2)$ -Query Attack on Any Key Exchange from a Random Oracle. In *Advances in Cryptology - CRYPTO 2009*, volume 5677, pages 374–390. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- [3] W. Diffie and M. Hellman. New Directions in Cryptography. *IEEE Trans. Inf. Theor.*, 22(6):644–654, 1976.
- [4] Iftach Haitner, Noam Mazon, Rotem Oshman, Amir Yehudayo, and Omer Reingold. On the Communication Complexity of Key-Agreement Protocols. page 37, 2018.
- [5] R. Impagliazzo and S. Rudich. Limits on the Provable Consequences of One-way Permutations. In *Proceedings of the Twenty-first Annual ACM Symposium on Theory of Computing, STOC '89*, pages 44–61, New York, NY, USA, 1989. ACM.
- [6] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, New York, NY, USA, 1997.
- [7] Ralph C. Merkle. Secure communications over insecure channels. *Communications of the ACM*, 21(4):294–299, April 1978.
- [8] Oded Regev. On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing, STOC '05*, pages 84–93, New York, NY, USA, 2005. ACM.
- [9] R L Rivest, A Shamir, and L Adleman. A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.
- [10] Peter W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997. arXiv: quant-ph/9508027.