

AICR · OPEN PROBLEM

No Key Agreement from Garbling One-Way-Function Circuits

Arbitrary polynomial round complexity

Status. Statement: AI-written from Sanjam Garg, Mohammad Hajiabadi, Mohammad Mahmoody and Ameer Mohammed, *Limits on the Power of Garbling Techniques for Public-Key Encryption*, Cryptology ePrint Archive 2018, not yet checked by a human. The two-message case, equivalent to public-key encryption, is proved as a theorem of the paper, and the extension to $2m$ -round key agreement for every constant m is asserted and sketched in the paper's Appendix A but never stated as a theorem; nothing at all is claimed for round complexity that grows with the security parameter, which is exactly what the conjecture asserts.

Category. *Idealized Models & Non-Uniformity.*

Conjecture (informal). *Garbled circuits are the most powerful cryptographic technique we know that needs nothing beyond one-way functions. Garbling a circuit that has one-way-function gates buried inside it is a genuinely non-black-box use of that one-way function, even though the garbling machinery itself is used as a black box, so classical black-box impossibility results say nothing about it. The paper studied here shows that this technique still cannot produce public-key encryption: relative to an idealized oracle providing a random function together with garbling of circuits with gates for that function, secure garbling exists but public-key encryption does not. Public-key encryption is exactly two-message key agreement, and the paper reports that its argument survives for any fixed number of rounds, but breaks once the number of rounds is allowed to grow with the security parameter: the proof removes the garbling-evaluation queries one round at a time, and each removal inflates the protocol's communication by a polynomial factor, so a growing number of rounds compounds into a tower of polynomials. The conjecture is that the separation nonetheless holds with no restriction on round complexity — that garbling circuits with one-way-function gates cannot give any key agreement at all. For plain random oracles the corresponding statement has been known for decades, with no dependence on the number of rounds, so the round restriction here looks like an artifact of the proof rather than a feature of the world.*

1 The setting

Whether public-key encryption can be based on one-way functions is a foundational open question. Impagliazzo and Rudich [IR89] ruled out fully black-box constructions of key agreement — and hence of public-key encryption — from one-way functions, and Barak and Mahmoody [BMG09] sharpened this to an optimal query attack on any key exchange in the random-oracle model. Crucially, neither result restricts the round complexity of the protocol being attacked: for a plain random oracle, the number of rounds was never a parameter of the impossibility.

These results say nothing about *non-black-box* uses of the one-way function, and one of the most powerful classes of non-black-box techniques that can be based on one-way functions alone is Yao's

garbled circuit technique [Yao86, BHR12]. Garbling a circuit that has one-way-function gates inside it uses the circuit description of the one-way function, and is therefore a non-black-box use of it in the taxonomy of [RTVo4], even though the garbling mechanism itself is used only as a black box. That the distinction is not idle was shown by Döttling and Garg [DG17], whose identity-based encryption construction uses precisely this technique to step around an earlier black-box barrier.

The paper studied here closes that route for public-key encryption: there is no fully black-box construction of PKE from a one-way function f together with a garbling scheme for circuits with f -gates. The garbling is *projective* — labels are chosen per input bit, so garbled inputs can be assembled wire by wire — which is what separates the result from the earlier separation of Asharov and Segev [AS15], which covers only non-decomposable garbling, and from the zero-knowledge-based limitations of Brakerski, Katz, Segev and Yerukhimovich [BKSY11]. The proof is a compilation: relative to an idealized garbling oracle, the garbling-evaluation queries are removed from encryption and then from decryption, leaving a construction that is inefficient but still makes polynomially many queries to a plain random oracle, at which point [IR89] applies.

Since public-key encryption is exactly two-message key agreement, the natural strengthening is to key agreement with no restriction on the number of rounds. The paper reports that its argument extends to any constant number of rounds and states exactly what stops it going further: the rounds are compiled out one at a time, and each compilation inflates the protocol's parameters — notably its communication complexity — by a polynomial factor, so a super-constant number of rounds produces a tower of polynomials and the argument dies. Settling the polynomial-round case would put garbling of one-way-function circuits on the same footing as the plain random oracle [IR89, BMG09], and would say that this whole class of non-black-box techniques cannot produce *any* public-key primitive at the key-agreement level, not merely the two-message ones.

What is known. The paper gives the full proof for public-key encryption: an idealized oracle $(f, \text{gc}, \text{gi}, \text{eval}, \text{rev})$ relative to which one-wayness and secure garbling of f -aided circuits both hold, plus a two-step compilation removing garbling-evaluation queries first from

encryption and then from decryption, with correctness and security losses that are inverse-polynomial for suitable parameter choices. Appendix A of the paper sketches the same compilation for one round of an interactive protocol, which yields the constant-round key-agreement case by iteration, and states the polynomial parameter blow-up per iteration as the reason iteration cannot be carried a super-constant number of times. Appendix B additionally shows the same oracles support non-interactive witness-indistinguishable proofs for satisfiability of f -aided circuits, which resolves an open question of Brakerski et al. about protocols with imperfect completeness in the PKE and constant-round key-agreement cases.

2 Notation and parameters

- $\kappa \in \mathbb{N}$ is the security parameter and $\text{negl}(\kappa)$ denotes an arbitrary function that is $\kappa^{-\omega(1)}$; a quantity is *non-negligible* if it is not of this form. *PPT* means probabilistic polynomial time, and an *oracle PPT* is a PPT algorithm that may make oracle calls.
- f denotes a function $\{0,1\}^* \rightarrow \{0,1\}^*$. An oracle algorithm S *breaks the one-wayness of f* if

$$\Pr_{x \leftarrow \{0,1\}^\kappa} [S(1^\kappa, f(x)) \in f^{-1}(f(x))]$$

is non-negligible in κ .

- A *binary-output oracle-aided circuit* C is a circuit built from Boolean gates together with *oracle gates*, whose output is a single bit. Its input size is $\text{inpsize}(C)$, its size $|C|$ is its number of gates and input wires, and C^f denotes C with every oracle gate instantiated by f .
- $\tilde{C}$ denotes a garbled circuit and $\text{label}_{i,b}$ the label for the value $b \in \{0,1\}$ on the i -th input wire; for $x \in \{0,1\}^m$, x_i is its i -th bit and $\{\text{label}_{i,x_i}\}_{i \in [m]}$ is the garbled input for x .
- $X \stackrel{c}{\approx} Y$ means that the two families of random variables indexed by κ are computationally indistinguishable: every PPT distinguisher has advantage $\text{negl}(\kappa)$.

- $\langle A^O(r_a), B^O(r_b) \rangle(1^\kappa)$ denotes an execution of an interactive protocol between oracle algorithms A and B with oracle access to O , private randomness r_a and r_b and no other inputs, and $\text{Trans}\langle A(r_a), B(r_b) \rangle$ denotes the resulting transcript. When A and B are given oracle access to both f and the three algorithms of a garbling scheme L we write $A^{f,L}$ and $B^{f,L}$.

The parameters are:

- κ , the security parameter.
- m : in Definition 2, half the round complexity of the protocol (the number of message pairs); in Definition 1, the input size of the circuit being garbled.
- δ , the key-disagreement probability; the protocol is correct when $\delta(\kappa) \leq 1/p(\kappa)$ for some polynomial p .
- γ , the eavesdropper's probability of outputting the agreed key; the protocol is secure when this is negligible.

Definition 1 (Garbling scheme for oracle-aided circuits relative to f). Fix a function f . A triple of algorithms $L = (\text{Garb}, \text{Eval}, \text{Sim})$ is a *garbling scheme for oracle-aided circuits relative to f* (equivalently, *with f -gates*) if:

- $\text{Garb}(1^\kappa, C)$ takes the security parameter and a binary-output oracle-aided circuit C with $m = \text{inpsize}(C)$, and outputs a garbled circuit $\tilde{C}$ together with a set of labels $\{\text{label}_{i,b}\}_{i \in [m], b \in \{0,1\}}$;
- $\text{Eval}^f(\tilde{C}, \{\text{label}_{i,b_i}\}_{i \in [m]})$ takes a garbled circuit and a garbled input and outputs a value in $\{0,1\}^* \cup \{\perp\}$;
- **Correctness.** For every binary-output oracle-aided circuit C with $m = \text{inpsize}(C)$ and every $x \in \{0,1\}^m$,

$$\Pr[C^f(x) = \text{Eval}^f(\tilde{C}, \{\text{label}_{i,x_i}\}_{i \in [m]})] = 1,$$

the probability being over $(\tilde{C}, \{\text{label}_{i,b}\}_{i \in [m], b \in \{0,1\}}) \leftarrow \text{Garb}(1^\kappa, C)$;

- **Security.** For every polynomial $m = m(\kappa)$, every poly-size binary-output oracle-aided circuit C with $\text{inpsize}(C) = m$ and every $x \in \{0,1\}^m$,

$$(\tilde{C}, \{\text{label}_{i,x_i}\}_{i \in [m]}) \stackrel{c}{\approx} \text{Sim}(1^{|C|}, m, C^f(x)),$$

where $(\tilde{C}, \{\text{label}_{i,b}\}_{i \in [m], b \in \{0,1\}}) \leftarrow \text{Garb}(1^\kappa, C)$.

The labels are per input *bit* (“projective” or “decomposable” garbling): a garbled input for a string is assembled wire by wire out of independently chosen labels. A scheme is *correct* if it satisfies the correctness clause. The pair consisting of a one-way function f and such a scheme L is the primitive GC-OWF.

Definition 2 (Key-agreement protocol). For $m \in \mathbb{N}$, a $2m$ -round *key-agreement protocol* is a pair of PPT interactive algorithms (A, B) such that $\langle A(r_a), B(r_b) \rangle(1^\kappa) = (k_a, k_b)$ for private randomness r_a, r_b ; write $k := k_a$.

- **Correctness.** For a function $\delta(\cdot)$ the protocol is $(1 - \delta)$ -correct if $\Pr[k_a = k_b] \geq 1 - \delta(\kappa)$ for every κ , the probability being over the randomness of A and B . It is *correct* if it is $(1 - \delta)$ -correct with $\delta(\kappa) = 1/p(\kappa)$ for some polynomial p .
- **Security.** For a function $\gamma(\cdot)$ the protocol is γ -secure if for every PPT adversary E and every κ , $\Pr[E(1^\kappa, \text{Tr}) = k] \leq \gamma(\kappa)$, where $\text{Tr} \leftarrow \text{Trans}\langle A(r_a), B(r_b) \rangle$ and the probability is over the randomness of E , A and B . It is *secure* if it is γ -secure for some $\gamma(\kappa) = \text{negl}(\kappa)$. An adversary E *breaks* the protocol if $\Pr[E(1^\kappa, \text{Tr}) = k]$ is non-negligible in κ .

Definition 3 (Fully black-box construction of key agreement from GC-OWF). A *fully black-box construction of a key-agreement protocol from GC-OWF* consists of a pair of oracle PPT interactive algorithms (A, B) together with an oracle PPT security reduction $S = (S_1, S_2)$, such that for every function f and every correct garbling scheme $L = (\text{Garb}, \text{Eval}, \text{Sim})$ for

oracle-aided circuits relative to f (Definition 1) both of the following hold.

- **Correctness.** $\langle A^{f,L}, B^{f,L} \rangle$ is a correct key-agreement protocol in the sense of Definition 2. (The paper's Definition 3.6 states this bar as $(1 - \frac{1}{2^\kappa})$ -correctness and remarks that the choice is without loss of generality by correctness boosting.)
- **Security.** For every adversary E that breaks $\langle A^{f,L}, B^{f,L} \rangle$, either
 - $S_1^{f,L,E}$ breaks the one-wayness of f ; or
 - $S_2^{f,L,E}$ breaks the security of L relative to f : for some binary-output oracle-aided circuit C with $m = \text{inpsize}(C)$ and some $x \in \{0,1\}^m$, the algorithm $S_2^{f,L,E}$ distinguishes $(\tilde{C}, \{\text{label}_{i,x_i}\}_{i \in [m]})$ from $\text{Sim}(1^{|C|}, 1^{|x|}, C^f(x))$ with non-negligible advantage, where $(\tilde{C}, \{\text{label}_{i,b}\}_{i \in [m], b \in \{0,1\}}) \leftarrow \text{Garb}(1^\kappa, C)$.

The round complexity of $\langle A, B \rangle$ is whatever (A, B) prescribe. Since A and B are PPT it is bounded by a polynomial in κ ; it is *not* required to be bounded by a constant.

3 The conjecture

Conjecture 1 (no key agreement from GC-OWF). *There is no fully black-box construction of a key-agreement protocol from GC-OWF in the sense of Definition 3.*

Equivalently and explicitly: for every pair of oracle PPT interactive algorithms (A, B) and every oracle PPT security reduction $S = (S_1, S_2)$, there exist a function $f: \{0,1\}^ \rightarrow \{0,1\}^*$ and a correct garbling scheme $L = (\text{Garb}, \text{Eval}, \text{Sim})$ for oracle-aided circuits relative to f (Definition 1) such that at least one of the following fails:*

1. $\langle A^{f,L}, B^{f,L} \rangle$ is a correct key-agreement protocol (Definition 2);
2. for every adversary E that breaks $\langle A^{f,L}, B^{f,L} \rangle$ (Definition 2),

either $S_1^{f,L,E}$ breaks the one-wayness of f or $S_2^{f,L,E}$ breaks the security of L relative to f (Definition 3).

Remark 1 (round complexity is the whole content). No bound is placed on the round complexity of $\langle A, B \rangle$ beyond the polynomial bound implied by A and B being PPT. The two-message case, which is equivalent to public-key encryption, is a theorem of the paper (Theorem 4.1). The paper further asserts, and sketches in its Appendix A, the extension to $2m$ -round key agreement for every constant m . The content of the conjecture is protocols whose round complexity is $\omega(1)$ as a function of κ .

Remark 2 (why it matters, and where to attack). It is the difference between “garbling one-way-function circuits cannot give you public-key encryption” and “garbling one-way-function circuits cannot give you public-key cryptography at all”. Round complexity is not a restriction in the corresponding random-oracle results [IR89, BMG09], so a gap that exists only for super-constant rounds looks like an artifact of the compilation technique rather than a real boundary; if instead it is real, then there is a genuinely interactive way of exploiting garbling that nobody has found, which would be a much more interesting outcome. Either answer is informative, and the class of techniques covered is broad: the paper notes that all known constructions combining garbling with one-way functions fall inside its model. The obstruction is a single named quantity — the polynomial blow-up in the protocol’s parameters incurred per round of compilation — which can be attacked directly, either by finding a compilation whose cost does not compound, or by proving the separation without round-by-round compilation at all.

Remark 3 (caveats for a reviewer). The paper poses this as a future direction, not as a numbered conjecture, so the *direction* of the answer — that the separation does extend to polynomially many rounds — is the paper’s evident expectation but is never asserted by it; the alternative framing would be a question rather than a conjecture. Second, Conjecture 1 is a transposition: it mirrors the paper’s Definition 3.6 (fully black-box construction of PKE from GC-OWF) with the key-agreement notion of the paper’s Appendix A substi-

tuted for the PKE notion. The paper never writes that combined definition out, so the wording here is faithful but not quoted. Third, the correctness notion of Definition 2 permits inverse-polynomial disagreement, which is unusually weak; anyone attacking the problem should decide whether they want that or negligible error, and the answer might differ between the two. Fourth, the constant-round case is only sketched in the paper’s appendix, so a would-be prover should not assume the constant-round machinery is fully written down anywhere. Finally, no follow-up work settling the polynomial-round case is known to the author of this record, but this has not been verified against the post-2018 literature and should be checked before publication.

4 Bibliography

- [AS15] Gilad Asharov and Gil Segev. *Limits on the power of indistinguishability obfuscation and functional encryption*. Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on, pages 191–209, IEEE, 2015.
- [BHR12] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. *Foundations of garbled circuits*. ACM CCS 12: 19th Conference on Computer and Communications Security, pages 784–796, ACM Press, 2012.
- [BKS11] Zvika Brakerski, Jonathan Katz, Gil Segev, and Arkady Yerukhovich. *Limits on the power of zero-knowledge proofs in cryptographic constructions*. Theory of Cryptography Conference, pages 559–578, Springer, 2011.
- [BMG09] Boaz Barak and Mohammad Mahmoody-Ghidary. *Merkle puzzles are optimal — an $O(n^2)$ -query attack on any key exchange from a random oracle*. Advances in Cryptology — CRYPTO 2009, volume 5677 of Lecture Notes in Computer Science, pages 374–390, Springer, 2009.
- [DG17] Nico Döttling and Sanjam Garg. *Identity-based encryption from the Diffie-Hellman assumption*. Advances in Cryptology — CRYPTO 2017, Part I, volume 10401 of Lecture Notes in Computer Science, pages 537–569, Springer, 2017.
- [IR89] Russell Impagliazzo and Steven Rudich. *Limits on the provable consequences of one-way permutations*. 21st Annual ACM Symposium on Theory of Computing, pages 44–61, Seattle, WA, USA, ACM Press, 1989.
- [RTV04] Omer Reingold, Luca Trevisan, and Salil P. Vadhan. *Notions of reducibility between cryptographic primitives*. TCC 2004: 1st Theory of

Cryptography Conference, volume 2951 of Lecture Notes in Computer Science, pages 1–20, Springer, 2004.

- [Yao86] Andrew Chi-Chih Yao. *How to generate and exchange secrets (extended abstract)*. 27th Annual Symposium on Foundations of Computer Science, pages 162–167, IEEE Computer Society Press, 1986.