On the Non-Existence of Blockwise 2-Local PRGs with Applications to Indistinguishability Obfuscation

Alex Lombardi* MIT

Vinod Vaikuntanathan† MIT

April 6, 2017

Abstract

Lin and Tessaro (Eprint 2017/250) recently proposed indistinguishability obfuscation and functional encryption candidates and proved their security based on a standard assumption on bilinear maps and a non-standard assumption on "Goldreich-like" pseudorandom generators (PRG). In a nutshell, they require the existence of pseudo-random generators $G : \Sigma^n \to \{0,1\}^m$ for some $\mathsf{poly}(n)$-size alphabet Σ where each output bit depends on at most two input alphabet symbols, and which achieve sufficiently large stretch. We show a polynomial-time attack against such generators.

Our attack uses tools from the literature on two-source extractors (Chor and Goldreich, SICOMP 1988) and efficient refutation of 2-CSPs over large alphabets (Allen, O'Donnell and Witmer, FOCS 2015). Finally, we propose new ways to instantiate the Lin-Tessaro construction that do not immediately fall to our attacks. While we cannot say with any confidence that these modifications are secure, they certainly deserve further cryptanalysis.

*E-mail: alexjl@mit.edu. Supported by an Akamai Presidential Fellowship.

†E-mail: vinodv@mit.edu. Research supported in part by NSF Grants CNS-1350619 and CNS-1414119, Alfred P. Sloan Research Fellowship, Microsoft Faculty Fellowship, the NEC Corporation, a Steven and Renee Finn Career Development Chair from MIT. This work was also sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236.

Contents

1 Introduction

There has been much recent progress on indistinguishability obfuscation (IO) culminating with the recent result of Lin and Tessaro [LT17] that constructs an IO candidate from standard assumptions on bilinear maps and non-standard assumptions on "Goldreich-like" pseudorandom generators [Gol00] with locality 2. This is a remarkable development: until recently, we had IO candidates based on constant degree (most recently, degree-5) multilinear maps and "Goldreich-like" PRGs with constant locality (most recently, locality 5) [Lin16b, LV16, LT17, AS16]. There were no secure candidates for the multilinear maps but there are candidates for the locality-5 PRG that resist many classes of attacks [OW14, AL16]. The Lin-Tessaro result dramatically shifted the burden of existence from degree-5 multilinear maps to the existence of pseudorandom generators with (so-called) blockwise locality 2 and polynomial stretch, putting us in a completely different landscape. (For the formal definitions of all these technical terms, see below and Section 2).

In this work, we first show a polynomial-time attack against such pseudorandom generators. As such, this constitutes a break of the Lin-Tessaro IO (as well as functional encryption) constructions that use bilinear maps. As a secondary contribution, we show several potential ways to fix the construction against our attacks. However, these "countermeasures" have so far not been cryptanalyzed with any care and their security should be highly suspect.

Goldreich-Like Generators with Blockwise Local Predicates. We start by describing the object we attack. Let P be a predicate from $\mathbb{Z}_q^2$ to $\{0,1\}$. Let H be a (directed) constraint graph with n vertices and m edges. The pseudorandom generator $G_{H,P} : \mathbb{Z}_q^n \to \{0,1\}^m$ is defined in the following way. Let $e = (i,j)$ be a directed edge in G. Then, the e^{th} bit of the output of the generator is computed as $P(x_i, x_j)$.

Theorem 1.1. *There is a poly(n) time algorithm $\mathcal{D}$ with the following property: for any $m \geq \tilde{\Omega}(q \cdot n)$ and any predicate $P : \mathbb{Z}_q^2 \to \{0,1\}$ in two variables, and almost all $(1 - o(1))$ graphs H with n vertices and m edges, $\mathcal{D}(H, P, z)$ distinguishes (with distinguishing advantage $\Omega(1)$) a random string $z \sim U_m$ from a random output $z \sim G_{H,P}(U_{n,q})$ of Goldreich's pseudorandom generator $G_{H,P} : \mathbb{Z}_q^n \to \{0,1\}^m$ when instantiated with P and H.*

We note here that when $q = \Omega(\frac{n}{\log n})$, Theorem 1.1 is trivially true: for any $m > n^2$, any graph H on n vertices and m edges contains a duplicate pair of edges e and e', so $G_{H,P}$ is not pseudorandom because z_e and $z_{e'}$ will always be equal for any $z = G_{H,P}(\mathbf{x})$. Therefore, we restrict to the case of $q = O(\frac{n}{\log n})$ for the rest of the paper.

This construction can also be thought of as a "block-wise local" pseudo-random generator, a terminology that Lin and Tessaro introduce and use [LT17]. In an (L, w)-block-wise PRG, the input is divided into blocks of size w bits each, and each output bit of the PRG can depend on at most L blocks. It is easy to see that a Goldreich-like PRG as defined above with alphabet size $q = 2^w$ is a $(2, w)$-block-wise PRG; in fact, the candidate block-wise 2-local PRGs described in [LT17] are Goldreich-like PRGs.

The Lin-Tessaro Result and Connection to Goldreich-Like PRGs. Lin and Tessaro [LT17], building on earlier work [BV15, AJ15, Lin16a, LV16, Lin16b, AS16] showed an IO candidate based on the hardness of standard assumptions on bilinear maps and the existence of a Goldreich-like PRG with locality 2 and *sufficiently large* stretch.

Theorem 1.2 ([LT17]). *Under standard assumptions on bilinear maps and the existence of a subexponentially secure (n, m, q)-Goldreich-like PRG with $q = \mathsf{poly}(n)$ and $m = (nq)^{1+\epsilon}$ for some constant $\epsilon > 0$, there is an IO scheme. Assuming the existence of such a generator with quasipolynomial security, there is a compact FE scheme.*

In a nutshell, they pre-compute all possible monomials on the bits of each alphabet symbol $x_i \in \mathbb{Z}_q$ ($i = 1, \ldots, n$) and encrypt it in an FE ciphertext. Computing the PRG output, then, can be written as a degree-2 computation which can be performed using a bilinear map (leveraging on an earlier result of Lin [Lin16b]). Thus, the number of bits being encrypted is $n \cdot q$. To achieve sublinear-compactness (which is necessary to apply the FE-to-IO transformations), they need the output length of the PRG m to be large enough, namely $m = \Omega((nq)^{1+\epsilon})$ for some constant $\epsilon > 0$.

Our main theorem (Theorem 1.1) now implies that a natural class of candidates for such PRGs, proposed and analyzed in [LT17], can be broken in polynomial-time. There are two gaps between our break and a complete break of the [LT17] candidate: First, Theorem 1.1 breaks "Goldreich-like" PRGs where the predicate computing each output bit is the same; this is not necessary for Theorem 1.2. Secondly, our attack works for a $1 - o(1)$ fraction of constraint graphs H whereas Theorem 1.2 only requires the existence of *some good constraint graph*. We discuss these issues further in Section 5.

In the rest of the introduction, we will briefly describe the techniques behind our proof of Theorem 1.1.

Outline of Our Attack. We attack $G_{H,P}$ based on the following interpretation of Goldreich's PRG. Any graph H, predicate P, and string $z \in \{0,1\}^m$ define an instance $\mathcal{I}$ of the constraint satisfaction problem CSP$(P, \neg P)$: for every edge $e = (i, j) \in E(H)$, $\mathcal{I}$ includes the constraint $z_e \oplus P(x_i, x_j) = 1$ (where $z_e \oplus P(x_i, x_j)$ is either $P(x_i, x_j)$ or $\neg P(x_i, x_j)$). The task of breaking $G_{H,P}$ can be thought of as distinguishing instances $\mathcal{I}$ in which the negations of P are chosen uniformly at random from instances $\mathcal{I}$ in which the negations of P are determined by a random planted solution $\mathbf{x} \in \mathbb{Z}_q^n$.

Allen, O'Donnell, and Witmer [AOW15] developed a polynomial time algorithm for a related problem, namely that of *random CSP refutation*: in their setting (specializing to 2CSPs), a random instance $\mathcal{I}$ is generated by choosing a random graph H along with *random negation patterns* $(a_e, b_e) \in \mathbb{Z}_q^2$ for each edge $e = (i, j) \in E(H)$, and including constraints $P(x_i + a_e, x_j + b_e) = 1$ in $\mathcal{I}$. Their algorithm can certify, for example, that $\mathrm{Opt}(\mathcal{I}) < 1$ provided given at least $\tilde{\Omega}(n \cdot \mathrm{poly}(q))$ constraints. Intuitively, we would like to break $G_{H,P}$ by attempting to refute the satisfiability of the CSP instance associated to (G, H, z).

However, there are two main obstacles to this approach: first, the algorithm in [AOW15] requires $n \cdot \mathrm{poly}(q)$ clauses where $\mathrm{poly}(q)$ is some high (constant) degree polynomial in q; we need to reduce this to $n \cdot q$ for our purposes. Moreover, the random CSP setting in [AOW15] includes q^2 different negation patterns (chosen at random for each clause); it turns out that the naive reduction from our setting to theirs incurs a loss of q^2 (in addition to the large polynomial in q already present) which we additionally need to avoid.

These problems would both be resolved if we could convert our CSP on an alphabet of size $q = \omega(1)$ to a related CSP on a constant sized alphabet. We achieve this by showing that any predicate $P : \mathbb{Z}_q^2 \to \{0,1\}$ is $(\frac{1}{2} + \Omega(\frac{1}{\sqrt{q}}))$-correlated to another predicate $P' : \mathbb{Z}_q^2 \to \{0,1\}$ which "depends on only a constant number of bits of each input". This builds off of a lower bound due

to Chor and Goldreich [CG88] on 2-source extractors. Then, we show that to break $G_{H,P}$ it suffices to certify that a random CSP with *constant-sized predicate* Q is not $\frac{1}{2} + \Omega(\frac{1}{\sqrt{q}})$-satisfiable, which can be done using the algorithm of [AOW15] with only $\tilde{\Omega}(n \cdot q)$ clauses.

2 Preliminaries

Notation. We let U_n denote the uniform distribution on n bits, i.e., on the set $\{0,1\}^n$. Additionally, we let $U_{n,q}$ denote the uniform distribution on the set $\mathbb{Z}_q^n$. Let $\mathsf{negl}(n) : \mathbb{N} \to \mathbb{R}$ denote any function that is smaller than any inverse polynomial in n. That is, we require that for every polynomial p, there is an $n_p \in \mathbb{N}$ such that for all $n > n_p$, $\mathsf{negl}(n) < 1/p(n)$.

2.1 Pseudorandom Generators

We say that a function $G : \{0,1\}^n \to \{0,1\}^m$ is a *pseudorandom generator* (PRG) if it has the following properties: (1) G is computable in (uniform) time poly(n), and (2) any probabilistic polynomial time adversary $A : \{0,1\}^m \to \{0,1\}$ has the property that

$$\left| \mathop{\mathbf{E}}_{x \leftarrow U_n} [A(G(x))] - \mathop{\mathbf{E}}_{y \leftarrow U_m} [A(y)] \right| = \mathsf{negl}(n)$$

We say that a PRG $G : \{0,1\}^n \to \{0,1\}^m$ has *stretch* $m - n = m(n) - n$. In this paper, we focus on the *polynomial stretch* regime, namely where $m = O(n^c)$ for some constant $c > 1$.

If G is computable in NC^0, then we define the *locality* of G to be the maximum number of input bits on which any output bit of G depends.

2.2 Constraint Satisfaction Problems (CSPs)

Let $P : \mathbb{Z}_q^k \to \{0,1\}$ be a q-ary predicate. We denote by CSP(P) the set of instances of the constraint satisfaction problem (CSP) with predicate P and *arbitrary negation patterns*, as defined below. An instance $\mathcal{I} = \{C_1, C_2, ..., C_m\}$ of CSP(P) consists of m equations "$C_i(\mathbf{x}) = 1$" in n variables $\mathbf{x} = (x_i)_{i=1}^n, x_i \in \mathbb{Z}_q$. Each equation (also called a constraint), is of the form $C(\mathbf{x}) := P(x_{i_1} + a_1, x_{i_2} + a_2, ..., x_{i_k} + a_k)$ for some collection of k indices $(i_1, ..., i_k)$ and k "negation patterns" $(a_1, ..., a_k) \in \mathbb{Z}_q^k$. Equivalently, an instance $\mathcal{I}$ is a labelled k-uniform hypergraph H, where a clause $C = ((i_1, i_2, ..., i_k), a_1, a_2, ..., a_k))$ corresponds to a hyperedge $(i_1, i_2, ..., i_k)$ with label $(a_1, a_2, ..., a_k)$.

Given an instance $\mathcal{I} \in$ CSP(P), an important quantity of interest is the *maximum fraction of constraints* simultaneously satisfiable by some input $\mathbf{x} \in \mathbb{Z}_q^n$, which is denoted Opt$(\mathcal{I})$. For a particular $\mathbf{x} \in \mathbb{Z}_q^n$, we let $\mathrm{Val}_{\mathcal{I}}(\mathbf{x}) = \mathop{\mathbf{E}}_{i \sim [m]} [C_i(\mathbf{x})]$, so that $\mathrm{Opt}(\mathcal{I}) = \max_{\mathbf{x} \in \mathbb{Z}_q^n} (\mathrm{Val}_{\mathcal{I}}(\mathbf{x}))$.

Following [AOW15], we consider a model for random CSP(P) instances in which a random instance $\mathcal{I} \in$ CSP(P) is generated by including each possible constraint $C(\mathbf{x}) = P(x_{i_1} + a_1, ..., x_{i_k} + a_k) = 1$ with probability p. In this model, a random instance $\mathcal{I}$ typically has $m = \overline{m}(1 \pm O(\frac{1}{\sqrt{m}}))$ constraints, where $\overline{m} = p \cdot q^k n^k$. We call a sample instance $\mathcal{I}$ from this model "a random instance of density p".

The above definitions can be naturally extended to constraint satisfaction problems with $t > 1$ predicates $P_1, ..., P_t$. In particular, we denote by CSP$(P_1, ..., P_t)$ the set of instances of such a CSP, and an instance $\mathcal{I} = (\mathcal{I}_1, ..., \mathcal{I}_t)$ to consist of an instance $\mathcal{I}_j$ of CSP(P_j) for each j.

2.3 Refutation of random CSPs

One of the key tools used in our algorithm that breaks the PRGs of Lin and Tessaro is a *refutation algorithm* for random CSPs developed by Allen, O'Donnell, and Witmer in [AOW15]. We state the relevant definitions and results here.

Definition 1. Let $P : \mathbb{Z}_q^k \to \{0,1\}$ be a predicate, and let $\overline{P} = \underset{(x_1,\ldots,x_k)\sim U_{k,q}}{P(x_1,\ldots,x_k)}$. An algorithm $\mathcal{A}$ is said to $(\overline{P}+\epsilon)$-refute random density p instances of CSP(P) if $\mathcal{A}$ has the following properties: (1) $\mathcal{A}$ returns "fail" for any instance $\mathcal{I} \in$ CSP(P) with Opt$(\mathcal{I}) \geq \overline{P}+\epsilon$, and (2) with probability $1-o(1)$, $\mathcal{A}$ returns "Opt$(\mathcal{I}) < \overline{P}+\epsilon$" for a random instance $\mathcal{I} \in$ CSP(P) of density p.

In [AOW15], the problem of refuting random CSPs is solved by certifying that a random CSP satisfies a property, called ϵ-quasirandomness, which is strictly stronger than "Opt$(\mathcal{I}) < \overline{P}+\epsilon$."

Definition 2. An instance $\mathcal{I} \in$ CSP(P) is *ϵ-quasirandom* if for all $\mathbf{x} \in \mathbb{Z}_q^n$, the probability distribution

$$\pi_{\mathbf{x},\mathcal{I}}(y) = \Pr_{((i_1,i_2,\ldots,i_k),\mathbf{a})\sim E(H)}[x_{i_l} + a_l = y_l \text{ for all } 1 \leq l \leq k]$$

is ϵ-close to the uniform distribution $U_{m,q}$ in statistical distance, where H denotes the labelled hypergraph associated to $\mathcal{I}$.

Note that quasirandomness of an instance $\mathcal{I}$ is solely a property of the hypergraph H associated to $\mathcal{I}$ and has nothing to do with the predicate P.

Definition 3. Let $P : \mathbb{Z}_q^k \to \{0,1\}$ be a predicate. An algorithm $\mathcal{A}$ is said to certify ϵ-quasirandomness of random density p instances of CSP(P) if $\mathcal{A}$ has the following properties: (1) $\mathcal{A}$ returns "fail" for any instance $\mathcal{I} \in$ CSP(P) which is not ϵ-quasirandom, and (2) with probability $1-o(1)$, $\mathcal{A}$ returns "$\mathcal{I}$ is ϵ-quasirandom" for a random instance $\mathcal{I} \in$ CSP(P) of density p.

Theorem 2.1 ([AOW15], Theorem B.6)**.** *Let $P : \mathbb{Z}_q^k \to \{0,1\}$. Then, there is a poly$(n)$-time algorithm $\mathcal{A}_\epsilon$ which certifies ϵ-quasirandomness of random density p instances of CSP(P), as long as $\overline{m} = pq^k n^k \geq \tilde{\Omega}(\frac{n^{k/2}q^{O(k)}}{\epsilon^2})$. Moreover, this algorithm is oblivious to the predicate P.*

2.4 Goldreich's Candidate PRG

Goldreich's candidate pseudorandom generator, first introduced in [Gol00] (then as a candidate one-way function), can be instantiated with any k-ary predicate $P : \{0,1\}^k \to \{0,1\}$ and any k-uniform (directed) hypergraph H on n vertices and m hyperedges. (To the best of our knowledge, the generalization to predicates P that take symbols from a larger alphabet was first considered by Lin and Tessaro under the name of "block-wise local" PRGs). Given H and P, we identify each vertex in H with an index in $[n]$ and each hyperedge with an index $i \in [m]$. For each $i \in [m]$, let $\Gamma_H(i) \in [n]^k$ be the sequence of k vertices in the ith hyperedge. Then, Goldreich's PRG is the function from $\{0,1\}^n$ to $\{0,1\}^m$ defined by

$$G_{H,P}(x) = \left(P(x|_{\Gamma_H(i)})\right)_{i\in[m]}.$$

That is, the ith bit of $G_{H,P}(x)$ is the output of P when given the $\Gamma_H(i)$-restriction of x as input.

Goldreich's generator is often instantiated with a *uniformly random* choice of hypergraph H; in this setting, we say that "Goldreich's generator instantiated with P is a PRG" for some predicate P if for a random k-uniform hypergraph H, $G_{H,P}$ is a PRG with high probability (say, probability $1 - o(1)$). Often (see [AL16], [OW14], [ABR12]) instead of proving hardness results for random hypergraphs it suffices to use hypergraphs with "good expansion" for varying definitions of expansion. For a more in-depth survey and discussion of Goldreich's PRG, see [App16].

In this paper, we use a slightly different model of "random hypergraph", which we later call $\mathcal{G}_{\text{dir}}(n,p)$ in the case of $k = 2$, in which instead of fixing the number of hyperedges m and choosing a uniformly random hypergraph with m edges, we fix an average number of edges $\overline{m}$ and sample H by including each hyperedge independently with probability $p = \frac{\overline{m}}{n^k}$. This results in a hypergraph H whose number of hyperedges m is highly concentrated around $\overline{m}$ but is not constant. Our main result, Theorem 1.1, applies for both of these random models; we formally prove our result in the $\mathcal{G}_{\text{dir}}(n,p)$ model, but it easily extends to the standard Goldreich random model. Essentially, if we are handed a random graph with exactly $2\overline{m}$ edges, we can simulate the distribution $\mathcal{G}_{\text{dir}}(n,p)$ for $p = \frac{\overline{m}}{n^2}$ with negligible error; see [AOW15] Appendix D for an analogous transfer theorem in the setting of random CSP refutation.

Note that given a hypergraph H and predicate P used to instantiate Goldreich's PRG, the pair (H, P) along with any string $y \in \{0,1\}^m$ define an instance $\mathcal{I} \in \text{CSP}(P, \neg P)$ with clauses $G_i(\mathbf{x}) = y_i$ (where $G_i(\mathbf{x})$ denotes the ith bit of the output of $G(x)$) giving either the constraint $P(x|_{\Gamma_H(i)}) = 1$ (if $y_i = 1$) or $\neg P(x|_{\Gamma_H(i)}) = 1$ (if $y_i = 0$). The task of breaking Goldreich's PRG can be thought of distinguishing a "random instance" (in the sense that y is chosen uniformly at random in $\{0,1\}^m$ from a "random planted instance" (where y is chosen from $G(U_m)$). However, even when the graph H is chosen uniformly at random, these instances do not directly correspond to "random instances" in the sense of Section 2.2. For each clause, instead of having random negation patterns applied to $\mathbf{x}$, we randomly negate the entire predicate P with probability $\frac{1}{2}$. This distinction is especially relevant in the generalization to large alphabet size, as then only a $\frac{1}{q^k}$ fraction of clauses in a random CSP (as defined in Section 2.2) have the (non-)negation pattern appearing in Goldreich's PRG.

3 A Structural Result on Predicates $P : \mathbb{Z}_q^2 \to \{0,1\}$

Our first result says that any predicate $P : \mathbb{Z}_q^2 \to \{0,1\}$ is non-trivially correlated to a *decomposed predicate* P' where $P'(x,y) = Q(f(x), g(y))$ for some functions f, g, and Q where f maps from $\mathbb{Z}_q$ to $\{0,1\}^4$, and Q maps from $\{0,1\}^4 \times \{0,1\}^4$ to $\{0,1\}$. Our result is a consequence of a lower bound on the possible error of two-source extractors, due to Chor and Goldreich [CG88].

Theorem 3.1. *Suppose that $P : \mathbb{Z}_q^2 \to \{0,1\}$ is a balanced predicate and $q = 16q'$ is divisible by 16. Then, there is a pair of q'-to-one maps $f_1, f_2 : \mathbb{Z}_q \to \mathbb{Z}_{16}$, a balanced predicate $Q : \mathbb{Z}_{16} \times \mathbb{Z}_{16} \to \{0,1\}$, and some constant c such that $P(x,y)$ is at least $\frac{1}{2} + \frac{c}{\sqrt{q}}$-correlated to $P'(x,y) := Q(f_1(x), f_2(y))$.*

Proof. We start with the result of Chor and Goldreich [CG88], which says that given any $P : \mathbb{Z}_q^2 \to \{0,1\}$, there exists a set $S \subset \mathbb{Z}_q$ of size q' and a set $\tilde{T} \subset \mathbb{Z}_q$ of size $\frac{q}{2} = 8q'$ such that

$$\mathop{\mathbf{E}}_{x,y \leftarrow S \times \tilde{T}}[P(x,y)] \geq \frac{1}{2}\left(1 + \frac{1}{\sqrt{8q'}}\right).$$

This implies the existence of a set $T \subset \tilde{T}$ of size q' such that

$$\mathop{\mathbf{E}}_{x,y \leftarrow S \times T}[P(x,y)] \geq \frac{1}{2}\left(1 + \frac{1}{\sqrt{8q'}}\right).$$

The intuition behind the rest of the solution is as follows: we have identified a $q' \times q'$ rectangle $S \times T$ on which P is biased towards 1 on average. The decomposed predicate P' we construct will map each input x, y to a rectangle (one of which is $S \times T$) and then output 0 or 1 depending on the expected value of P inside that rectangle.

More formally, let $\mathcal{S} = \{S_1, S_2, ..., S_{16}\}$ and $\mathcal{T} = \{T_1, T_2, ..., T_{16}\}$ be partitions of $\mathbb{Z}_q$ into 16 equally sized subsets with $S_1 = S$ and $T_1 = T$. Let $f_1 : \mathbb{Z}_q \to \mathbb{Z}_{16}$ send x to the unique i such that $x \in S_i$, and let $f_2 : \mathbb{Z}_q \to \mathbb{Z}_{16}$ send y to the unique j such that $y \in S_j$. By construction, every $i \in \mathbb{Z}_{16}$ has a preimage of size exactly q' under both f_1 and f_2.

Finally, for each i, j, consider the quantity

$$E_{ij} := \mathop{\mathbf{E}}_{x,y \leftarrow S_i \times T_j}[P(x,y)],$$

and let $Q : \mathbb{Z}_{16} \times \mathbb{Z}_{16} \to \{0,1\}$ be defined by $Q(i,j) = 1$ if and only if E_{ij} is one of the 128 largest elements of the multiset $\{E_{ij}, (i,j) \in \mathbb{Z}_{16} \times \mathbb{Z}_{16}\}$. We claim that $P(x,y)$ is at least $\frac{1}{2} + \Omega(\frac{1}{\sqrt{q}})$-correlated to $Q(f_1(x), f_2(y))$. To see this, we note that

$$\Pr[P(x,y) = Q(f_1(x), f_2(y))] = \frac{1}{2} + \frac{1}{256}\sum_{i=1}^{16}\sum_{j=1}^{16}(-1)^{Q(i,j)+1}\text{disc}_{S_i,T_j}(P)$$

where $\text{disc}_{S_i,T_j}(p) := E_{ij} - \frac{1}{2}$. Now, we claim that

$$\sum_{i=1}^{16}\sum_{j=1}^{16}(-1)^{Q(i,j)+1}\text{disc}_{S_i,T_j}(P) \geq \frac{1}{2}\sum_{i=1}^{16}\sum_{j=1}^{16}\left|\text{disc}_{S_i,T_j}(P)\right|, \quad (*)$$

in which case we are done, since

$$\sum_{i=1}^{16}\sum_{j=1}^{16}\left|\text{disc}_{S_i,T_j}(P)\right| \geq \text{disc}_{S,T}(P) = \Omega(\frac{1}{\sqrt{q}}).$$

To prove the claim, note that if every term in $(*)$ is non-negative, the claim is true; otherwise, suppose without loss of generality that $E_{ij} > \frac{1}{2}$ for some (i,j) such that $Q(i,j) = 0$ (implying that $Q(i,j) = 0$ for all i, j such that $E_{ij} < \frac{1}{2}$). Then, we note that by definition of Q,

$$\sum_{(i,j):Q(i,j)=0, E_{ij}>\frac{1}{2}}\left(E_{ij} - \frac{1}{2}\right) \leq \sum_{(i,j):Q(i,j)=1, E_{ij}>\frac{1}{2}}\left(E_{ij} - \frac{1}{2}\right),$$

which implies that

$$\sum_{i,j}(-1)^{Q(i,j)+1}\text{disc}_{S_i,T_j}(P) \geq \sum_{i,j:E_{ij}<\frac{1}{2}} -\text{disc}_{S_i,T_j}(P) = \frac{1}{2}\sum_{i,j}\left|\text{disc}_{S_i,T_j}(P)\right|,$$

as desired. □

4 Reduction to CSPs in the AOW15 setting

We now prove Theorem 1.1 by combining our structure theorem for predicates (Theorem 3.1) with quasirandomness certification (Theorem 2.1) over *constant-sized* alphabets. Recall the statement of Theorem 1.1.

Theorem 4.1. *There is a poly(n) time algorithm $\mathcal{D}$ with the following property: for any $m \geq \tilde{\Omega}(q \cdot n)$ and any predicate $P : \mathbb{Z}_q^2 \to \{0,1\}$ in two variables, and almost all $(1 - o(1))$ graphs H with n vertices and m edges, $\mathcal{D}(H, P, z)$ distinguishes (with distinguishing advantage $\Omega(1)$) a random string $z \sim U_m$ from a random output $z \sim G_{H,P}(U_{n,q})$ of Goldreich's pseudorandom generator $G_{H,P} : \mathbb{Z}_q^n \to \{0,1\}^m$ when instantiated with P and H.*

We already mentioned in Section 2.4 that it is sufficient to prove the theorem when H is drawn from the distribution $\mathcal{G}_{\text{dir}}(n,p)$, i.e. by independently including each directed edge $(i,j) \in [n] \times [n]$ with probability $p = \frac{\overline{m}}{n^2}$, so we will focus on this modified theorem.

Fix any predicate $P : \mathbb{Z}_q^2 \to \{0,1\}$. Let $G_{H,P} : \mathbb{Z}_q^n \to \{0,1\}^m$ be a randomly drawn instance of Goldreich's PRG, where the graph H is sampled according to the distribution $\mathcal{G}_{\text{dir}}(n,p)$. We now describe an algorithm $\mathcal{D}$ which, given P, H, and a string $z \in \{0,1\}^m$ (where m is the number of edges in H), outputs a bit $b \in \{0,1\}$ such that $\mathop{\mathbf{E}}_{z \sim U_m}[\mathcal{D}(P,H,z)]$ is noticeably different from $\mathop{\mathbf{E}}_{z \sim G_{H,P}(U_n)}[\mathcal{D}(P,H,z)]$. The algorithm $\mathcal{D}$ does the following:

1. Set $H^{(0,0)} = H$, $z^{(0,0)} = z$.
2. Draw an additional $16^4 - 1$ graphs $H^{(a,b)}$ (for $(a,b) \neq ((0,0),(0,0)) \in \mathbb{Z}_{16}^4 \simeq \mathbb{Z}_{16}^2 \times \mathbb{Z}_{16}^2$) independently from the distribution $\mathcal{G}_{\text{dir}}(n,p)$, and sample $16^4 - 1$ strings $z^{(a,b)} \in \{0,1\}^{m_{ab}}$ uniformly at random (where m_{ab} denotes the number of edges in $H^{(a,b)}$).
3. Define CSP "instances" $\mathcal{I}_1$ and $\mathcal{I}_2$ (with the predicate unspecified) based on the graphs $H^{(a,b)}$ and strings $z^{(a,b)}$, in which $\mathcal{I}_1$ contains the clauses corresponding to $z_i^{(a,b)} = 1$ (included with negation pattern (a,b)) and $\mathcal{I}_2$ contains the clauses corresponding to $z_i^{(a,b)} = 0$ (included with negation pattern (a,b)).
4. Call $\mathcal{A}_\epsilon(\mathcal{I}_1)$ and $\mathcal{A}_\epsilon(\mathcal{I}_2)$, where $\mathcal{A}$ is the algorithm from Theorem 2.1, $\epsilon = \frac{c}{2^{20}\sqrt{q}}$, and c is the constant in Theorem 3.1. Recall that $\mathcal{A}$ is predicate-oblivious, so this can be done.
5. Return 1 if and only if both calls to $\mathcal{A}$ return "this instance is ϵ-quasirandom".

Let Q be a predicate satisfying the conclusion of Theorem 3.1 applied to P, and let $\tilde{Q} : (\mathbb{Z}_{16} \times \mathbb{Z}_{16})^2 \to \{0,1\}$ be the 2-variable predicate $\tilde{Q}((x,y),(x',y')) = Q(x,y')$. We will be thinking of $\mathcal{I}_1$ above as an instance of CSP$(\tilde{Q})$ and $\mathcal{I}_2$ as an instance of CSP$(\neg\tilde{Q})$, so that the graphs $H^{(a,b)}$ and strings $z^{(a,b)}$ define an instance $(\mathcal{I}_1, \mathcal{I}_2)$ of CSP$(\tilde{Q}, \neg\tilde{Q})$.

We now need to analyze the behavior of $\mathcal{D}(P,H,z)$ when (1) $z \sim U_m$ is truly random, and when (2) $z \sim G_{H,P}(U_{n,q})$ is pseudorandom.

Lemma 4.2. *Suppose that $\overline{m} \geq \tilde{\Omega}(qn)$. Then, for a $1 - o(1)$ fraction of graphs H as drawn from $\mathcal{G}_{dir}(n,p)$, the following holds: with probability $1 - o(1)$ over $z \sim U_m$ and the randomness of $\mathcal{D}$, both $\mathcal{I}_1$ and $\mathcal{I}_2$ as constructed in $\mathcal{D}(P,H,G_{H,P}(z))$ will be certified to be $\frac{c}{2^{20}\sqrt{q}}$-quasirandom.*

Proof. We prove the lemma separately for $\mathcal{I}_1$ and $\mathcal{I}_2$, which suffices by a union bound. To see that $\mathcal{I}_1$ will almost surely be certified quasirandom, consider the distribution

$$\{z^{(a,b)} \sim U_m, H^{(a,b)} \sim \mathcal{G}_{\text{dir}}(n, \frac{\overline{m}}{n^2}) : \mathcal{I}_1\}.$$

This is exactly a random instance in CSP($\tilde{Q}$) where each clause-negation pattern pair is chosen with probability $\frac{\overline{m}}{2n^2}$. Therefore, [AOW15] tells us that with probability $1 - o(1)$ over the choice of graphs $(H^{(a,b)})$ and bit strings $(z^{(a,b)})$ (which together determine the graph defining $\mathcal{I}_1$), as long as $\overline{m} \geq \tilde{\Omega}(qn)$, $\mathcal{I}_1$ will successfully be certified to be $\frac{c}{2^{20}\sqrt{q}}$-quasirandom. Suppose that the certification is successful for a $1 - \epsilon(n)$ fraction of collections of graphs and bit strings. Then, by a union bound, we conclude that for at least a $1 - \sqrt{\epsilon(n)}$ fraction of graphs $H = H^{(0,0)}$, the quasirandomness certification succeeds with probability at least $1 - \sqrt{\epsilon(n)}$ over the randomness of $\mathcal{D}$ and input $z \sim U_m$.

The same argument as above applies to $\mathcal{I}_2$ as well (as the quasirandomness certification algorithm is oblivious to the predicate), so we have obtained the desired result. □

Lemma 4.3. *Suppose that $\overline{m} \geq \tilde{\Omega}(qn)$. Then, for a $1 - o(1)$ fraction of graphs H as drawn from $\mathcal{G}_{dir}(n, p)$, the following holds: with constant probability over $\mathbf{x} = (x_i)_{i=1}^n \sim U_{n,q}$ and the randomness of $\mathcal{D}$, at least one of $\mathcal{I}_1$ and $\mathcal{I}_2$ as constructed by $\mathcal{D}(P, H, G_{H,P}(\mathbf{x}))$ is* not *$\frac{c}{2^{20}\sqrt{q}}$-quasirandom.*

Proof. We will show that with constant probability over $\mathbf{x} \sim U_{n,q}$ and the randomness of $\mathcal{D}$, the particular input $\tilde{\mathbf{x}} = ((f_1(x_i), f_2(x_i))_{i=1}^n \in \mathbb{Z}_{256^2}^n$ will satisfy at least a $\frac{1}{2} + \frac{c}{2^{20}\sqrt{q}}$ fraction of the combined instance $(\mathcal{I}_1, \mathcal{I}_2) \in \text{CSP}(\tilde{Q}, \neg\tilde{Q})$, which suffices to prove the lemma. We first focus on the clauses defined by $H^{(0,0)} = H$. Suppose that $\mathbf{E}[P(x, y) \oplus Q(x, y)] = \alpha$. Then, in expectation over $\mathbf{x}$, the fraction of clauses from H in $(\mathcal{I}_1, \mathcal{I}_2)$ satisfied by $\tilde{\mathbf{x}}$ is given by

$$E := \mathop{\mathbf{E}}_{\mathbf{x} \sim U_{n,q}} \left[\Pr_{(i,j) \sim E(H)} [Q(x_i, x_j) = P(x_i, x_j)] \right] \geq \alpha - \frac{n}{m} \geq \frac{1}{2} + \frac{c}{\sqrt{q}} - \frac{n}{m}.$$

where the $\frac{n}{m}$ term comes from the fraction of edges in H which are self loops (we cannot say that $P(x_i, x_i)$ is necessarily correlated to $Q(x_i, x_i)$). Now, we compute the variance (over $\mathbf{x}$) of this fraction of clauses to be

$$\text{Var} = \mathop{\mathbf{E}}_{\mathbf{x} \sim U_{n,q}} \left[\left(\Pr_{(i,j) \sim E(H)} [Q(x_i, x_j) = P(x_i, x_j)] \right)^2 \right] - E^2$$

$$= \mathop{\mathbf{E}}_{\mathbf{x} \sim U_{n,q}} \left[\frac{1}{m^2} \sum_{(i,j) \in E(H), (k,l) \in E(H)} \chi\left(Q(x_i, x_j) = P(x_i, x_j)\right) \chi\left(Q(x_k, x_l) = P(x_k, x_l)\right) \right] - E^2$$

$$= \frac{1}{m^2} \sum_{(i,j) \in E(H), (k,l) \in E(H)} \Pr_{\mathbf{x} \sim U_{n,q}} [Q(x_i, x_j) = P(x_i, x_j) \text{ and } Q(x_k, x_l) = P(x_k, x_l)] - E^2$$

Note that if the edges $(i, j), (k, l) \in E(H)$ have no vertices in common, the events "$Q(x_i, x_j) = P(x_i, x_j)$" and "$Q(x_k, x_l) = P(x_k, x_l)$" are independent. This means that our variance is bounded by

$$\mathrm{Var} \leq \frac{1}{m^2} \sum_{(i,j)\in E(H),(k,l)\in E(H)} \Pr_{\mathbf{x}\sim U_{n,q}}[Q(x_i,x_j)=P(x_i,x_j)] \Pr_{\mathbf{x}\sim U_{n,q}}[Q(x_k,x_l)=P(x_k,x_l)] + \frac{m_{\mathsf{bad}}}{m^2} - E^2 = \frac{m_{\mathsf{bad}}}{m^2},$$

where

$$m_{\mathsf{bad}} = |\{((i,j),(k,l)) \in E(H)\times E(H) : (i,j) \text{ and } (k,l) \text{ have a vertex in common}\}| \leq \sum_{i\in[n]} \deg_H(i)^2.$$

Now, we note that since $\deg_H(i) \leq 2n$ for all i, we obtain the bound

$$m_{\mathsf{bad}} \leq \sum_{i\in[n]} \deg_H(i)^2 \leq 2n \cdot \sum_{i\in[n]} \deg_H(i) = 4mn.$$

Moreover, with probability $1 - \mathrm{negl}(n)$ over the choice of H we have that $m \geq \frac{\overline{m}}{2}$, in which case the variance is bounded by

$$\mathrm{Var} \leq \frac{4mn}{m^2} = \frac{4n}{m} \leq \frac{8n}{\overline{m}} \leq \frac{8}{q\log(n)},$$

as we are assuming that $\overline{m} \geq \tilde{\Omega}(nq)$. By Chebyshev's inequality, this means that with constant probability over $\mathbf{x} \sim U_{n,q}$, we have that

$$\Pr_{(i,j)\sim E(H)}[Q(x_i,x_j)=P(x_i,x_j)] \geq \alpha - \frac{n}{m} - \frac{4}{\sqrt{q\log(n)}} \geq \frac{1}{2} + \frac{c}{\sqrt{q}} - \frac{2n}{\overline{m}} - \frac{4}{\sqrt{q\log(n)}} \geq \frac{1}{2} + \frac{c}{2\sqrt{q}}.$$

This in turn implies that with constant probability over $\mathbf{x}$, the string $\tilde{\mathbf{x}}$ will satisfy at least $\frac{1}{2} + \frac{c}{2\sqrt{q}}$ of the clauses in $(\mathcal{I}_1, \mathcal{I}_2)$ corresponding to H.

Finally, we consider the other graphs $H^{(a,b)}$ and their corresponding clauses. Since we chose the negations $z^{(a,b)}$ uniformly at random, we see that for *any* $\mathbf{x}$ and all collections of graphs $(H^{(a,b)})$ with sufficiently many (at least $\frac{\overline{m}}{2}$) edges (all but a negligible fraction of collections of graphs), with constant probability over the choice of $(z^{(a,b)}) \in \{0,1\}^{\sum_{a,b} m_{ab}}$ we have that the fraction of non-H clauses satisfied by $\tilde{\mathbf{x}}$ is at least $\frac{1}{2} - \frac{1}{100\sqrt{\overline{m}}}$.

Finally, for all collections of graphs $(H^{(a,b)})$ with sufficiently few (at most $2\overline{m}$) edges (all but a negligible fraction of collections of graphs), we conclude that with constant probability over $\mathbf{x}$ and the randomness of $\mathcal{D}$,

$$\mathrm{Opt}(\mathcal{I}_1,\mathcal{I}_2) \geq \mathrm{Val}_{\mathcal{I}_1,\mathcal{I}_2}(\tilde{\mathbf{x}}) \geq \frac{1}{2} + \frac{1}{2^{17}} \cdot \frac{c}{2\sqrt{q}} - \frac{1}{100\sqrt{\overline{m}}} \geq \frac{1}{2} + \frac{c}{2^{20}\sqrt{q}},$$

since we are assuming that $\overline{m} \geq \tilde{\Omega}(qn)$. This completes the proof of Lemma 4.3. □

To conclude, we summarize how Lemma 4.2 and Lemma 4.3 together imply Theorem 1.1. Lemma 4.2 tells us that with probability $1 - o(1)$ over the choice of H, $\mathcal{A}_\epsilon(\mathcal{I}_1)$ and $\mathcal{A}_\epsilon(\mathcal{I}_2)$ will both output "this instance is quasirandom" with probability $1 - o(1)$, so $\mathcal{D}(P,H,z)$ will output 1 with probability $1 - o(1)$. Lemma 4.3 tells us that with probability $1 - o(1)$ over the choice of H,

either $\mathcal{I}_1$ or $\mathcal{I}_2$ will *not* be ϵ-quasirandom with constant probability, which by Theorem 2.1 means that $\mathcal{A}_\epsilon(\mathcal{I}_1)$ or $\mathcal{A}_\epsilon(\mathcal{I}_2)$ will *not* output "this instance is quasirandom", in which case $\mathcal{D}(P, H, z)$ will return 0. Thus, by a union bound, we conclude that with probability $1 - o(1)$ over H, $\mathcal{D}$ achieves a constant distinguishing advantage between the "truly random z" case and the "pseudorandom z" case, as desired.

5 Potential Strategies for Repairing the IO Candidate

We outline three possible strategies for fixing the blockwise 2-local PRG and thus the IO candidate.

Non-Uniform Constraint Graphs. We showed that for any block-wise 2-local predicate P and a *uniformly random* constraint graph, the corresponding Goldreich-like PRG is broken, as long as its stretch meets the minimum requirement from the Lin-Tessaro construction.

On the other hand, for the construction to work, it is sufficient that there is *some* constraint graph and some 2-local predicate P for which the Goldreich-like PRG is secure. This possibility remains open.

We note that even in the case of the standard Goldreich PRG (with Boolean predicates), we know of few attacks known to work on *all* graphs. Mossel, Shpilka and Trevisan in [MST06] succeed in breaking Goldreich's PRG for Boolean predicates of locality at most 4 and all hypergraphs with sufficiently many hyperedges, but later attacks [AL16, OW14, ABR12, BQ12] focus on uniformly random hypergraphs.

Different Predicates for Each Output Bit. Our algorithm breaks the PRG as long as each output bit is computed by applying the *same predicate* on some subset of input bits. A potential fix is to allow for different predicates, one for each output bit.

A potential way to come up with such a family of predicates is to fix a good predicate P and consider the family $\mathcal{P} = \{P_{c_1,c_2} : c_1, c_2 \in \mathbb{Z}_q \times \mathbb{Z}_q\}$ where $P_{c_1,c_2}(x, y) = P(x + c_1, y + c_2)$ with addition done mod q.

Local PRGs with Lazy Evaluation? A third and final strategy is to consider a generalization of blockwise local PRGs into what we call "preprocessed locality-2 PRGs". Consider a PRG from n bits to m bits which permits computation in two stages: in the first stage, there is an arbitrary algorithm that takes the n bits of input and produces s bits of output, with $s = m^{1-\epsilon}$ for some constant $\epsilon > 0$, and in the second stage, there is a degree-2 function that takes the s bits to the actual m bits of PRG output.

Clearly, a blockwise 2-local PRG with appropriate alphabet size and stretch would have given us such a preprocessed locality-2 PRG. It seems to us that preprocessed locality-2 PRGs are the minimal primitive necessary for the Lin-Tessaro framework to bear fruit. The existence of such PRGs remains open.

None of the three potential countermeasures we suggest have been subject to any rigorous cryptanalysis. Doing so and coming up with polynomial-time attacks would definitively lay to rest the question of building IO from bilinear maps.

Acknowledgements. We thank Gil Cohen, Dana Moshkovitz and Prasad Raghavendra for their quick responses to our oracle calls about two-source extractors and CSPs.

References

[ABR12] Benny Applebaum, Andrej Bogdanov, and Alon Rosen. A dichotomy for local small-bias generators. In *Theory of Cryptography Conference*, pages 600–617. Springer, 2012.

[AJ15] Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, volume 9215 of *Lecture Notes in Computer Science*, pages 308–326. Springer, 2015.

[AL16] Benny Applebaum and Shachar Lovett. Algebraic attacks against random local functions and their countermeasures. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1087–1100. ACM, 2016.

[AOW15] Sarah R. Allen, Ryan O'Donnell, and David Witmer. How to refute a random csp. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 689–708. IEEE, 2015.

[App16] Benny Applebaum. Cryptographic hardness of random local functions. *Computational complexity*, 25(3):667–722, 2016.

[AS16] Prabhanjan Ananth and Amit Sahai. Projective arithmetic functional encryption and indistinguishability obfuscation from degree-5 multilinear maps. *IACR Cryptology ePrint Archive*, 2016:1097, 2016.

[BQ12] Andrej Bogdanov and Youming Qiao. On the security of Goldreich's one-way function. *Computational complexity*, 21(1):83–127, 2012.

[BV15] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 171–190. IEEE Computer Society, 2015.

[CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.

[Gol00] Oded Goldreich. Candidate one-way functions based on expander graphs. *IACR Cryptology ePrint Archive*, 2000:63, 2000.

[Lin16a] Huijia Lin. Indistinguishability obfuscation from constant-degree graded encoding schemes. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part I*, pages 28–57, 2016.

[Lin16b] Huijia Lin. Indistinguishability obfuscation from DDH on 5-linear maps and locality-5 PRGs. *Preprint: http://eprint.iacr.org/2016/1096.pdf*, 2016.

[LT17] Huijia Lin and Stefano Tessaro. Indistinguishability obfuscation from bilinear maps and block-wise local prgs. *IACR Cryptology ePrint Archive*, 2017:250, 2017.

[LV16] Huijia Lin and Vinod Vaikuntanathan. Indistinguishability obfuscation from DDH-like assumptions on constant-degree graded encodings. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 11–20. IEEE, 2016.

[MST06] Elchanan Mossel, Amir Shpilka, and Luca Trevisan. On ε-biased generators in NC^0. *Random Structures & Algorithms*, 29(1):56–81, 2006.

[OW14] Ryan O'Donnell and David Witmer. Goldreich's PRG: evidence for near-optimal polynomial stretch. In *Computational Complexity (CCC), 2014 IEEE 29th Conference on*, pages 1–12. IEEE, 2014.