18.218 Project (Limits of Preprocessing)

Nathan Sheffield and Isabel McGuigan

Apr 2024

1 Introduction

It is known that sub-exponential size AC^0 circuits – i.e. circuit families of constant depth and a subexponential number of gates – can't compute the inner product function. However, known proof techniques (relying on random restrictions) fail if the two inputs to the inner product are allowed arbitrary preprocessing. That is, suppose an all-powerful Alice is given input x and outputs some arbitrary function $f(x)$, and an all-powerful Bob is given input y and outputs some arbitrary function $g(y)$. Is it possible that there exists a family of poly-size, constant depth circuits that computes $\langle x,y\rangle$ on input $f(x),g(y)$? It seems pretty absurd that this could be the case – the difficulty of inner product seems like it comes from combining the two inputs, as opposed to computing a hard function on either half. It doesn't seem especially plausible that there should be an encoding of the inputs that makes it substantially easier. However, proving that AC^0 circuits can't compute inner product under any poly-length input encoding remains an open problem.

We should note that there absolutely *are* functions on two inputs where being able to preprocess the two inputs into arbitrary encodings makes the problem much easier. For instance, the problem of computing the parity of the number of 1s in x and y together is known to require exponential-size AC^0 circuits. However, if arbitrary preprocessing of x and y independently is allowed the problem becomes trivial: let $f(x)=\text{PARITY}(x)$ and $g(y)=\text{PARITY}(y)$, so that the AC^0 circuit on $f(x)$ and $g(y)$ just needs to compute the XOR of its two input bits. The IPPP (inner product with preprocessing) conjecture says that these kind of hacks aren't possible for inner product.

In this project, we'll summarize a recent work of Filmus, Ishai, Kaplan, and Kindler, where they prove a weak version of these sorts of lower bounds [Fil+20]. Specifically, they show that if one input's length is extended by an arbitrary polynomial amount under the preprocessing but the other is extended by only a tiny sublinear stretch (from length n to $n+n^\alpha$ for $0\le\alpha<1$), then no subexponential-size AC^0 circuit can compute inner product even given that encoding. They also show that these hardness results hold for (cryptographic) weak PRFs as well as rounded inner product functions. Finally, they connect the IPPP conjecture to other cryptographic statements by showing that cryptographic assumptions imply either versions of the IPPP conjecture or hardness results for learning AC^0 circuits.

1.1 Communication Complexity

Maybe the most interesting characterization of the goal of this line of work is in communication complexity. A communication protocol for a function on two inputs is a strategy for two all-powerful parties to exchange messages back and forth and eventually compute the function. That is, Alice is given input x, Bob is given input y, and they take turns sending messages until one of them has enough information to report the value of the function on inputs x and y. A protocol is considered "efficient" if Alice and Bob only have to exchange $\mathsf{polylog}$ many bits before they're guaranteed to be able to output the answer – this is the communication complexity analogue of P, and so is denoted P^{cc}.

We can also define a communication version of the polynomial hierarchy, denoted PH^{cc} [BFS86]. Here, two competing provers with knowledge of x and y interact for k rounds, sending $\mathsf{polylog}$ many bits in total, and Alice and Bob must each decide whether to accept or reject the transcript of their interaction. The YES prover should be able to make them both accept whenever their inputs are a yes instance of the problem, and the NO prover should be able to make one of them reject whenever their inputs are a no instance of the problem. The relation to the standard polynomial hierarchy should be clear: the polynomial hierarchy can also be described in terms of verifying transcripts of a constant-round game between two all-powerful provers, except in that case the verification is done by a single time-bounded verifier.

How is PH^{cc} related to the IPPP conjecture? This comes from the following lemma:

Lemma 1. PH^{cc} is precisely the set of functions computable by quasi-polynomial sized AC^0 circuits with arbitrary independent preprocessing of both inputs.

Proof. Suppose there exists a quasi-polynomial size AC^0 circuit whose bottom layer inputs are arbitrary functions of either x or y alone. In the corresponding protocol, we imagine the two provers walking down from the output gate, taking turns deciding which child to step down to, with the YES prover acting on OR gates and the NO player acting on AND gates. This is a constant number of rounds of players saying $\mathsf{polylog}$ length messages, so the transcript has length $\mathsf{polylog}$. If the circuit is satisfied, the YES player has a strategy that ensures this process will end up in a bottom-layer input with value 1, and otherwise the NO player can ensure we end up in an input with value 0. So, all Alice and Bob have to do after seeing this transcript is verify that it's a valid walk down the circuit, then evaluate the function at the bottom-level input it reaches. Since these bottom-level input functions depend on only one of x and y, one of them will be able to compute it and report the answer.

For the other direction, suppose there exists a PH^{cc} protocol. Build a circuit starting from the output gate corresponding to the ***game tree*** of the prover's interaction. That is, the tree has a node for all the 2^{polylog} possible partial transcripts of the interaction, with edges from each node to all nodes reachable in a single player's turn. At the bottom of this tree, the winner of the game is decided by some arbitrary function of x ANDed with some arbitrary function of y. Pre-process x and y to replace each with the list of all those relevant function values. Now, label nodes of the circuit where the YES player moves with OR and nodes where the NO player moves with AND; the output of this circuit will be 1 if and only if the YES player had a winning strategy. □

So, proving that a function $f \notin \mathsf{PH}^{cc}$ is equivalent to proving that it doesn't have quasipolynomial AC^0 circuits with preprocessing. It's a major open question whether inner product lies in PH^{cc}; this paper can be seen as a step towards resolving in the negative, showing impossibility when one of the inputs is processed only very slightly.

1.2 Cryptography

At first, it might be unclear why it's interesting that these results hold for weak PRFs as well as inner product. A ***weak pseudorandom function (PRF)*** is a function of one input and one "key", such that if the key is fixed to one random value and the function is fed a stream of random inputs, no poly-time adversary can distinguish the outputs of the function from a truly random function on those inputs. This sounds like a much, much more difficult sort of thing to compute than the inner product – we don't even know for sure that such functions are computable in $\mathsf{P/poly}$. However, there has been substantial interest in the cryptography world in trying to construct cryptographic primitives in the simplest complexity classes possible, and there are indeed some very simple functions that have been proposed as candidate weak PRFs. For instance, the rounded inner product mod 6 – that is, the function

$$h:\{0,1\}^n\times\{0,1\}^n\to\{0,1\},$$

$$h_k(x)=\begin{cases}0 \text{ if } \langle k,x\rangle \bmod 6 \in\{0,1,2\}\\ 1 \text{ if } \langle k,x\rangle \bmod 6 \in\{3,4,5\}\end{cases} \qquad (*)$$

has been proposed as a weak PRF candidate [Bon+18]. Of course, this is not a ***strong PRF*** – if an adversary is allowed to make arbitrary queries to h_k as opposed to just seeing it on random inputs, they can easily distinguish it from a random function. But it could plausibly be a weak PRF – and for many cryptographic applications weak PRFs suffice. Note that this function is computable by constant-depth, poly-size circuits if the circuits are given mod 6 gates (i.e. the class $\mathsf{AC}^0[6]$), but not computable in AC^0 without mod 6 gates. In fact, there's a known impossibility result for weak PRFs computable by AC^0 circuits, at least if you want better than quasipolynomial security (i.e. you want the function to fool quasipolynomial-size circuits as opposed to just poly-size ones). This impossibility result comes from a ***learning*** algorithm for AC^0: given samples of any AC^0 function on quasipolynomially many inputs, it's possible to compute an approximation that very closely matches the function on future inputs [LMN93] – if the function continues to align with our approximation, we know it probably wasn't truly random.

If you really wanted to do cryptography in AC^0, this impossibility result would make you pretty sad – you might look for some way to circumvent it. One option that's been proposed is that of an ***encoded-input PRF*** [Bon+18]: maybe there exist poly-time computable functions f and g and an AC^0-computable function h such that $h_{f(k)}(g(x))$ is a (strong or weak) PRF. This could be basically as good as a PRF in AC^0 for efficient cryptography reasons: when you're choosing your random key you have to do some arbitrary poly-time computation to generate the AC^0 circuit computing $h_{f(k)}$, but then whenever somebody wants to evaluate your PRF on x, you just ask them to first encode the input with g (which they can do, because g doesn't depend on k and thus can be made public), and then you can do the PRF evaluation with your super efficient AC^0 circuit [1].

The results of Filmus, Ishai, Kaplan, and Kindler show that, if one of f and g extends by only a small sublinear amount, AC^0 can't compute encoded-input PRFs (even weak ones) with exponentially good security (i.e. fooling exponentially large adversaries). If their results could be extended to the case where both f and g have polynomial stretch, it would rule out exponentially-secure encoded-input PRFs in AC^0 altogether.

1.3 Overview of Results

In Section 2, we'll sketch the proof of the main technical tool of the paper, which is the result due to Linial, Mansour, and Nisan, and its later quantititative strengthening by Tal, that functions with small AC^0 circuits are well-approximated by low-degree polynomials (i.e. have small high-degree Fourier mass). This is what will let us thinking about the problem as a Boolean function analysis problem – we won't worry about any of the other details about the circuits, instead just proving the results for functions with small high-degree Fourier mass.

In Section 3, we'll give a simple proof of the IPPP conjecture when one of the inputs isn't extended at all. The idea is that, if we had a small AC_0 circuit for this preprocessed inner product, by hardwiring each possible extended input we can get 2^n many small AC_0 circuits whose corresponding functions are all orthogonal to each other. But then by the LMNT theorem we would know that all of these functions are close to low-degree polynomials, and the space of low-degree polynomials has too small a dimension to have that many orthogonal functions.

In Section 4, we show that this simple proof idea can in fact be applied in much more general settings, by showing that hardwiring the arbitrarily-extended input yields a class of functions in which many are pretty close to orthogonal. This immediately implies the main result for inner product (i.e. inner product doesn't have AC^0 circuits with preprocessing as long as one input's preprocessing only gives a small sublinear stretch). In Section 5.1 and Section 5.2, we show that the hypotheses of this more general statement hold for any weak PRF, or any rounded inner-product function (defined in Section 5.2).

Finally, in Section 6, we mention an interesting argument showing that either

- the rounded mod 6 inner product function we mentioned earlier fails to be a weak PRF,
- the IPPP conjecture holds,
- the IPPP conjecture holds mod 3, or
- AC^0 can't be learned in subexponential time from random samples over a worst-case sampleable distribution.

[1] One caveat of this definition is that security could be totally broken if the PRF is queried on an improperly-encoded input (i.e. something not in the image of g). Boneh et al extend the definition to *protected* encoded-input PRFs, where there also exists an AC^0 circuit to verify whether an input is encoded properly, and give some heuristic evidence that protected encoded-input PRFs are approximately as strong an assumption as encoded-input PRFs.

Any of these hypotheses sound very plausible – but despite some effort none are known individually. So it's interesting that we can prove at least one must hold.

2 Functions in AC^0 are close to low degree

The main ingredient in the analysis is a bound on the high-degree Fourier mass of an AC^0 circuit.

Theorem 1 (Linial, Mansour, Nisan [LMN93], Tal [Tal17])**.** Let f be a Boolean function computable by an AC^0 circuit of depth h and size M. Then for any integer t,

$$||f^{\geq t}||^2 \leq 2 \cdot 2^{-t/O_h(\log M)^{h-1}}.$$

In other words, f is approximated by a function of degree $O_h(\log^{h-1} M)$.

This theorem is a quantitative improvement on the original result of Linial, Mansour and Nisan. We will sketch the proof of the original bounds, and then very roughly outline how the improvement worked. In both cases, the key tool is a variant of the famous ***Håstad switching lemma***.

Lemma 2 (Håstad, see O'Donnell's textbook)**.** Suppose f is computable by either a DNF or CNF of width (i.e. maximum clause size) at most w. Then, for any k, a p-random restriction of f will be computable by a decision tree of depth k except with probability at most $(5pw)^k$.

We will omit the proof of this lemma; a proof by clever counting argument due to Razborov can be found in Arora and Barak's textbook. Instead, we will explain how this can use this to show that the Fourier mass of the function computed by a low-depth circuit is concentrated on low levels.

Theorem 2 (Linial, Mansour, Nisan [LMN93]; see O'Donnell's textbook)**.** Let f be a Boolean function computable by an AC^0 circuit of depth h and size M. Then for any integer t,

$$||f^{\geq t}||^2 \leq M \cdot 2^{-\Omega(t^{1/h})}.$$

Proof. Let $\varepsilon = M \cdot 2^{-\Omega(t^{1/h})}$, so that $t = O(\log(M/\varepsilon))^{h-1} \cdot \log(1/\varepsilon)$. In order to show that f has less than ε mass above level t, we'll need to argue a version of Håstad's switching lemma for larger depth circuits, as opposed to just 2. Suppose we have a depth h, size M, width w (where width refers to the fan-in on the bottom layer) circuit, and we hit it with a $\frac{1}{10w}$-random restriction. Looking at the bottom two layers of the circuit[2], we can observe that any gate on the second-to-bottom layer is computing a depth-2 circuit function of width w. So, fixing a parameter $\ell = \log(M/2\varepsilon)$, By Håstad's switching lemma the probability that a given second-to-bottom gate's output isn't described by a depth-ℓ decision tree is at most $(1/2)^\ell$. As long as this bad event doesn't happen for any gate on the second-to-bottom layer, we could replace all of them with depth-ℓ decision trees.

But now, note that a depth-ℓ decision tree can be represented by both a width-ℓ CNF and a width-ℓ DNF (either take the OR over all accepting paths of the AND of variables on the path, or the AND over all rejecting paths of the OR of negations of variables on the path). So, assuming no bad events happened, we can switch the bottom two layers from ANDs of ORs to ORs of ANDs (or vice-versa). After this switching has been done, we can combine the 2nd and 3rd layers (we're doing either ANDs of ANDs or ORs of ORs, so can merge them) to now have a circuit with depth-$(h-1)$ and width ℓ. We can now do another $\frac{1}{10\ell}$ random restriction to decrease the depth again – if we repeat this process a total of $h-2$ times, assuming no bad event ever occurs, the entire remaining circuit will now be depth 2.

Each gate in the tree has at most one opportunity to cause a bad event, and this event happens with probability at most $(1/2)^\ell$, so by union bound the probability of any bad event occurring is at most $M \cdot 2^{-\ell} = \varepsilon/2$. Now, performing another $\frac{1}{10\ell}$-random restriction will reduce the entire circuit to a depth-$\log(2/\varepsilon)$ decision tree with

[2] We should mention that in this proof we're assuming that the circuit is layered, so that all nots appear at the bottom and layers alternate between all ANDs and all ORs – this is without loss of generality at cost of a constant increase in circuit depth.

failure probability at most $\varepsilon/2$. Note that this series of random restrictions is equivalent to performing one giant $\frac{1}{10w}\left(\frac{1}{10\ell}\right)^{h-2}$-random restriction.

Ok, now let's show our main claim. We can't immediately apply this multi-level switching lemma, because the circuit might have very large width – first step is to reduce the width of the circuit to ℓ. We can do so simply by looking at every gate on the first level with more than ℓ inputs, and cutting all but ℓ of the incoming wires – it can be seen that this has negligible effect on the Fourier mass of the function, because the inputs to such wide gates are each individually very non-influential.

Now, we know that after performing a $\left(\frac{1}{10\ell}\right)^{h-1}$-random restriction, with probability at least $1-\varepsilon$ the whole circuit will reduce to a depth-$\log(2/\varepsilon)$ decision tree. A depth-$\log(2/\varepsilon)$ decision tree computes a function of degree at most $\log(2/\varepsilon)\leq t$, so if it does reduce to such a decision tree the circuit will end up with 0 Fourier mass above level t. Thus, the expected high-degree ($\geq t$) Fourier mass after a $\left(\frac{1}{10\ell}\right)^{h-1}$-random restriction is at most ε. We know that $\mathbb{E}_{f'\sim p\text{-random restriction of } f}[\sum_{|S|\geq k}\widehat{f'}(S)^2]=\sum\Pr[S \text{ has } \geq k \text{ living variables after restriction }]\cdot\widehat{f}(S)^2$. Using a Chernoff bound, we can use this to show that f has Fourier mass at most 3ε on degrees larger than $3\ell/\left(\frac{1}{10\ell}\right)^{h-1}=O(t)$. □

The stronger bounds come from a more powerful version of Håstad's switching lemma, which shows that by performing one random restriction followed by one carefully-chosen restriction we can with high probability reduce a large number of depth-2 circuits to decision trees simultaneously. Plugging this sort of lemma into a similar proof strategy will allow for better probability bounds and thus tighter control on the Fourier mass.

3 A proof of IPPP when only one input length is extended

With the LMNT bound in hand, we can, without too much trouble, obtain the main result on the inner product in the special case where one of the input lengths is not extended at all. (The stronger main result will allow this length to be extended up to $n+n^{\alpha}$; however, we present the proof of this special case because it gives good intuition for how the proof of the main result will work.)

Theorem 3. Let $\mathsf{IP}(x,y)$ be the mod-2 inner product. Suppose there exists an AC^0 circuit C of depth h and size M with arbitrary preprocessing functions A,B such that $C(A(x),B(y))=\langle x,y\rangle$. If B maps $\{0,1\}^n\to\{0,1\}^n$, then C must have exponential size; in particular,

$$M\geq 2^{\Omega_h\left(n^{1/(h-1)}\right)}.$$

Proof. For a fixed $x\in\{0,1\}^n$, define the function $C_x:\{0,1\}^n\to\{0,1\}$ by $C_x(y)=C(A(x),y)$. The proof consists of three steps:

1. Find a collection of orthogonal functions $g_x:\{0,1\}^n\to\{0,1\}$ such that g_x agrees with C_x.
2. Find a specific function g_x with small low-degree Fourier mass.
3. Use the LMNT bound on the high-dimensional Fourier mass of f_x to conclude that M must be exponential in n.

Step 1: Because IP is a ***right one-to-one function*** – i.e. fixing y to any two distinct values will induce two distinct functions of x – the preprocessing function $B:\{0,1\}^n\to\{0,1\}^n$ must be injective, hence bijective. For each $x\in\{0,1\}^n$, let $g_x(y)=\mathsf{IP}(x,B^{-1}(y))$. We have that

$$g_x(y)=\mathsf{IP}(x,B^{-1}(y))=C(A(x),B(B^{-1}(y)))=C(A(x),y)=C_x(y),$$

so g_x agrees with C_x.

For any vector x, let $S_x=\{i\mid x_i=1\}$. Then, g_x is just the character $\chi_{B(S_x)}$, so the g_x's are orthogonal.

Step 2: Let V be the inner product space of functions $\{0,1\}^n\to\mathbb{R}$. The functions g_x form an orthonormal basis for V. Let U be the subspace of functions of degree at most $\frac{n}{4}$; then, U is spanned by the characters χ_S for

$|S| \le \frac{n}{4}$, and has dimension $D = \binom{n}{\le n/4} \le 2^{H(1/4)n} \le 2^{0.9n}$ (where H is the binary entropy function). Let $u_1,...,u_D$ be an orthonormal basis for U. Then, for any function $f \in V$, we have

$$||f^{\le n/4}||^2 = ||\mathrm{proj}_U(f)||^2 = \sum_{k=1}^{D} \langle f,u_k\rangle^2.$$

Because the u_k's are orthonormal, we have $||u_k||^2 = \sum_{x\in\{0,1\}^n} \langle g_x,u_k\rangle^2 = 1$ for each k. Therefore,

$$\mathbb{E}_x[||g_x^{\le n/4}||^2] = \frac{1}{2^n} \sum_{x\in\{0,1\}^n} \sum_{k=1}^{D} \langle g_x,u_k\rangle^2 = \frac{1}{2^n} \sum_{k=1}^{D} \sum_{x\in\{0,1\}^n} \langle g_x,u_k\rangle^2 = \frac{1}{2^n} \sum_{k=1}^{D} 1 = \frac{D}{2^n}.$$

So there is some x such that $||g_x^{\le n/4}||^2 \le \frac{D}{2^n} \le 2^{-0.1n} \le \frac{1}{2}$.

Step 3: For this x, because g_x agrees with C_x, it is computable by the unbounded fan-in circuit of depth at most h and size at most M obtained by plugging $A(x)$ into C. Thus, LMNT tells us that

$$2 \cdot 2^{-t/O_h(\log M)^{h-1}} \ge ||g_x^{\ge n/4}||^2 = 1 - ||g_x^{< n/4}||^2 \ge 1 - \frac{1}{2} = \frac{1}{2}.$$

Rearranging, we get

$$M \ge 2^{\Omega_h\left(n^{1/(h-1)}\right)}$$

as desired.

□

4 A more general version of the theorem

The previous proof relied on the fact that the functions $f_x(y) = \mathsf{IP}(x,y)$ are orthogonal. However, we may relax this condition slightly by requiring only that *a large subset of* the g_x's are *nearly* orthogonal. This relaxation yields the following main theorem, which is a generalization of Theorem 3.

Theorem 4. Let $f : \{-1,1\}^n \times \{-1,1\}^n \to \{-1,1\}$ be a Boolean function, let $0 \le k \le n/2 - 1$, and let $0 \le t \le n+k$ be an integer. Suppose that the following hold:

- f is a right one-to-one function.
- There is a large subset of f_x's which is almost orthogonal. Specifically, there exists a subset $X \subset \{-1,1\}^n$ of size $|X| \ge 13 \cdot 2^{2(k+1)} \cdot \binom{n+k}{\le t}$ such that

$$\mathbb{E}_{x \ne x' \sim X}\left[\langle f_x, f_{x'}\rangle^2\right] \le \frac{2^{2k}}{36|X|^2}.$$

- There exists a depth-h size-M circuit C, and arbitrary functions A,B such that $B : \{-1,1\}^n \to \{-1,1\}^{n+k}$ and $C(A(x),B(y)) = f(x,y)$.

Then M has exponential size; in particular,

$$M \ge 2^{\Omega_h\left(\left[\frac{t}{k}\right]^{1/(h-1)}\right)}.$$

Proof. Again, for a fixed $x \in \{-1,1\}^n$, define the function $C_x : \{-1,1\}^{n+k} \to \{-1,1\}$ by $C_x(y) = C(A(x),y)$. The proof of this theorem follows much the same outline as the proof of Theorem 3. However, at each step, we will have to do slightly more work to accommodate the fact that our functions are *almost* orthogonal rather than perfectly orthogonal. The three steps will be:

1. Find a collection of *almost* orthogonal functions $g_x : \{0,1\}^{n+k} \to \{0,1\}$ such that $g_x(y)$ *almost* agrees with C_x.
2. Find a specific function g_x with small low-degree Fourier mass.
3. Use the LMNT bound on the high-dimensional Fourier mass of g_x to conclude that M must be exponential in n.

Step 1: Because f is a right one-to-one function, the preprocessing function $B:\{-1,1\}^n\to\{-1,1\}^{n+k}$ must be injective. In the proof of Theorem 3, it was critical that the preprocessing function B was a bijection. To adapt the proof to this setting, we will extend B to a bijection $\beta:\{-1,1\}^{n+k}\to\{-1,1\}^{n+k}$. To do this, let $Y=\{y\in\{-1,1\}^{n+k}\mid y_{n+1}=\cdots=y_{n+k}=1\}$; then, the complement $\overline{Y}$ is the set of all $y\in\{-1,1\}^n$ so that at least one of $y_{n+1},...,y_{n+k}=-1$. We define β by

$$\beta(y)=\begin{cases}B(y_{1,...,n}) & y\in Y,\\ \text{arbitary choice} & y\in\overline{Y}\end{cases}.$$

We likewise extend each f_x to a (collection of) functions $\{-1,1\}^{n+k}\to\{-1,1\}$. For any $x\in\{-1,1\}^n$ and $R\subset\{n+1,...,n+k\}$, define

$$f_x^R(y)=\begin{cases}f_x(y_{1,...,n}) & y\in Y,\\ \chi_{S(x)}(y_{1,...,n})\cdot\chi_R(y_{n+1,...,n+k}) & y\in\overline{Y}\end{cases}.$$

We claim that, for every x, there is some R such that f_x^R *almost* agrees with $C_x\circ\beta$:

Claim 1. For every $x\in\{-1,1\}^n$, there exists $R(x)\subset\{n+1,...,n+k\}$ such that $f_x^{R(x)}$ agrees with $C_x\circ\beta$ on at least $\frac{1}{2}+\frac{1}{2^{k+1}}$ fraction of inputs.

Proof. Fix $y\in\{-1,1\}^{n+k}$. We'll consider two cases, based whether or not $y\in Y$.

Case 1. If $y\in Y$, then by definition,

$$f_x^R(y)=f_x(y_{1,...,n})=C(A(x),B(y_{1,...,n}))=C(A(x),\beta(y))=C_x\circ\beta(y).$$

Hence, each f_x^R automatically agrees with $C_x\circ\beta$ on 2^n inputs.

Case 2. Otherwise, if $y\in\overline{Y}$, then there's some $i\in\{n+1,...,n+k\}$ such that $y=-1$. For any R containing i, we have

$$\chi_R(y_{n+1,...,n+k})=-\chi_{R\setminus\{i\}}(y_{n+1,...,n+k}),$$

so $f_x^R(y)=-f_x^{R\setminus\{i\}}(y)$, and thus exactly one of $f_x^R,f_x^{R\setminus\{i\}}$ agrees with $C_x\circ\beta$ on y. Therefore, for a fixed y, we have that $f_x^R(y)=C_x\circ\beta(y)$ for exactly half of all subsets $R\subset\{n+1,...,n+k\}$. Hence,

$$\mathbb{E}_{R\subset\{n+1,...,n+k\}}\left[\Pr_{y\in\overline{Y}} f_x^R(y)=C_x\circ\beta(y)\right]=\mathbb{E}_{y\in\overline{Y}}\left[\Pr_{R\subset\{n+1,...,n+k\}} f_x^R(y)=C_x\circ\beta(y)\right]=\mathbb{E}_{y\in\overline{Y}}\left[\frac{1}{2}\right]=\frac{1}{2}.$$

So there is some $R(x)$ for which $f_x^{R(x)}$ agrees with $C_x\circ\beta$ on at least half of all inputs $y\in\overline{Y}$ – since there are $2^{n+k}-2^n$ such y's, this $f_x^{R(x)}$ agrees with $C_x\circ\beta$ on at least $2^{n+k-1}-2^{n-1}$ inputs.

The total fraction of inputs that $f_x^{R(x)}$ agrees with $C_x\circ\beta$ on is thus at least

$$\frac{2^n+2^{n+k-1}-2^{n-1}}{2^{n+k}}=\frac{1}{2}+\frac{1}{2^{k+1}}$$

as desired.

□

Having found a collection of functions $f_x^{R(x)}$ that *almost* agree with $C_x\circ\beta$, we will now check that they are *almost* orthogonal. To do this, we'll bound $\langle f_x^{R(x)},f_{x'}^{R(x')}\rangle=\mathbb{E}_{y\in\{-1,1\}^{n+k}}\left[f_x^{R(x)}(y)f_{x'}^{R(x)}(y)\right]$ for all $x\neq x'\in X$. Again, we'll look at two cases based on whether or not $y\in Y$.

Case 1. If $y\in Y$, then $f_x^{R(x)}(y)f_{x'}^{R(x')}(y)$ is just $f_x(y_{1,...,n})f_{x'}(y_{1,...,n})$.

Case 2. Otherwise, $f_x^{R(x)}(y)f_{x'}^{R(x')}(y)$ is

$$\chi_{S(x)}(y_{1,...,n})\chi_{S(x')}(y_{1,...,n})\cdot\chi_{R(x)}(y_{n+1,...,n+k})\chi_{R(x')}(y_{n+1,...,n+k}).$$

Because the coordinates $y_1,...,y_n$ and $y_{n+1},...,y_{n+k}$ are disjoint, the expectation of this over all $y\in\overline{Y}$ is

$$\mathbb{E}_{y\in Y}\left[\chi_{S(x)}(y_{1,...,n})\chi_{S(x')}(y_{1,...,n})\right]\cdot\mathbb{E}_{y\in\overline{Y}}\left[\chi_{R(x)}(y_{n+1,...,n+k})\chi_{R(x')}(y_{n+1,...,n+k})\right].$$

But the left factor is just $\langle\chi_{S(x)},\chi_{S(x')}\rangle=0$.

Therefore,

$$\begin{aligned}\langle f_x^{R(x)}, f_{x'}^{R(x')}\rangle &= \frac{1}{2^k}\cdot\mathbb{E}_{y\in Y}\left[f_x^{R(x)}(y)f_{x'}^{R(x')}(y)\right]+\left(1-\frac{1}{2^k}\right)\mathbb{E}_{y\in\overline{Y}}\left[f_x^{R(x)}(y)f_{x'}^{R(x')}\right]\\ &=\frac{1}{2^k}\cdot\mathbb{E}_{y\in Y}[f_x(y_{1,\ldots,n})f_{x'}(y_{1,\ldots,n})]+\left(1-\frac{1}{2^k}\right)\cdot 0\\ &=2^{-k}\langle f_x, f'_x\rangle.\end{aligned}$$

By assumption, $E_{x\neq x'\sim X}\left[\langle f_x, f_{x'}\rangle^2\right]\leq\frac{2^{2k}}{36|X|^2}$, so

$$\mathbb{E}_{x\neq x'\sim X}\left[\langle f_x^{R(x)}, f_{x'}^{R(x')}\rangle^2\right]=2^{-2k}\mathbb{E}_{x\neq x'\sim X}\left[\langle f_x, f'_x\rangle^2\right]\leq\frac{1}{36|X|^2}.$$

So the $f_x^{R(x)}$'s are functions which almost agree with $C_x\circ\beta$ and are close to being orthogonal. The functions $g_x = f_x^{R(x)}\circ\beta^{-1}$, then, are functions which almost agree with C_x and are also close to being orthogonal.

Step 2: The next step is to find a function g_x which has small low-degree Fourier mass. As in step 2 of the proof of Theorem 3, we will do this by finding a function whose projection onto a subspace of small-degree functions is small.

Let V be the inner product space of functions $\{0,1\}^{n+k}\to\mathbb{R}$, and let U be the subspace of functions of degree at most t; then, U is spanned by the characters χ_S for $|S|\leq t$, and has dimension $D=\binom{n}{\leq t}$. We already have a set X such that the expectation of $\langle g_x, g_{x'}\rangle$ over all x,x' is small; the first step is to find a subset $Z\subset X$ such that, for each individual $x\in Z$, the expectation of $\langle g_x, g_{x'}\rangle$ over all x' is small. To do this, note that by Cauchy-Schwarz,

$$\mathbb{E}_{x\neq x'\sim X}[|\langle g_x, g_{x'}\rangle|]\leq\mathbb{E}_{x\neq x'\sim X}\left[\langle g_x, g_{x'}\rangle^2\right]^{1/2}\leq\frac{1}{6|X|}.$$

Let Z be the set of all $x\in X$ such that $\mathbb{E}_{x'\sim X\setminus\{x\}}\left[\langle f_x^{R(x)}, f_{x'}^{R(x')}\rangle^2\right]<\frac{1}{12|X|^2}$ and $\mathbb{E}_{x'\sim X\setminus\{x\}}[|\langle g_x, g_{x'}\rangle|]<\frac{1}{2|X|}$. By Markov's inequality,

$$\Pr_{x\sim X}\left[\mathbb{E}_{x'\sim X\setminus\{x\}}\left[\langle g_x, g_{x'}\rangle^2\right]\geq\frac{1}{12|X|^2}\right]\leq 12|X|^2\cdot\mathbb{E}_{i\neq j\sim X}\left[\langle g_x, g_{x'}\rangle^2\right]\leq\frac{1}{3},$$

and similarly,

$$\Pr_{x\sim X}\left[\mathbb{E}_{x'\sim X\setminus\{x\}}[|\langle g_x, g_{x'}\rangle|]\geq\frac{1}{2|X|}\right]\leq 2|X|\cdot\mathbb{E}_{i\neq j\sim X}[|\langle g_x, g_{x'}\rangle|]\leq\frac{1}{3}.$$

Therefore, Z contains at least one third of the elements in X. Let W be the subspace of V spanned by g_x for all $x\in Z$.

Let $u_1,\ldots,u_D$ be an orthonormal basis for U, and for each $k\in[D]$, let w_k be the projection of u_k onto W. Since $||u_k||=1$ for each k, we'll have $||w_k||\leq 1$. Since w_k lies in W, we can write

$$w_k=\sum_{x\in Z}c_x g_x$$

for some coefficients $c_x \in \mathbb{R}$. We can then compute

$$\begin{aligned}1 &\geq \|w_k\|^2 \\ &= \langle w, w\rangle \\ &= \left\langle \sum_{x\in Z} c_x g_x, \sum_{x'\in Z} c_{x'} g_{x'} \right\rangle \\ &= \sum_{x,x'\in Z} c_x c_{x'} \langle g_x, g_{x'}\rangle \\ &= \sum_{x\in Z} c_x^2 + \sum_{x\neq x'\in Z} c_x c_{x'} \langle g_x, g_{x'}\rangle\end{aligned}$$

For any x,x', we have

$$c_x c_{x'} \langle g_x, g_{x'}\rangle \geq \frac{1}{2}(c_x^2 + c_{x'}^2)\langle g_x, g_{x'}\rangle \geq -\frac{1}{2}(c_x^2 + c_{x'}^2)\langle g_x, g_{x'}\rangle,$$

so we get

$$\begin{aligned}1 &\geq \sum_{x\in Z} c_x^2 - \frac{1}{2}\sum_{x\neq x'\in Z}(c_x^2 + c_{x'}^2)|\langle g_x, g_{x'}\rangle| \\ &= \sum_{x\in Z} c_x^2 \cdot \left(1 - \sum_{x'\in Z\setminus\{x\}} |\langle g_x, g_{x'}\rangle|\right) \\ &= \sum_{x\in Z} c_x^2 \cdot \left(1 - (|X|-1)\mathbb{E}_{x'\in Z\setminus\{x\}}|\langle g_x, g_{x'}\rangle|\right) \\ &\geq \sum_{x\in Z} c_x^2 \cdot \left(1 - (|X|-1)\frac{1}{2|X|}\right) \\ &\geq \frac{1}{2}\sum_{x\in Z} c_x^2,\end{aligned}$$

and thus $\sum_{x\in Z} c_x^2 \leq 2$. We now do a similar calculation for $\langle g_x, g_{x'}\rangle^2$. Because $(a+b)^2 \leq 2a^2 + 2b^2$ for any a,b, we can compute

$$\begin{aligned}\langle w_k, g_x\rangle^2 &= \left\langle \sum_{x\in Z} c_x g_x, g_x \right\rangle^2 \\ &= \left(c_x + \sum_{x'\in Z\setminus\{x\}} c_x \langle g_x, g_{x'}\rangle\right)^2 \\ &\leq 2c_x^2 + 2\left(\sum_{x'\in Z\setminus\{x\}} c_x \langle g_x, g_{x'}\rangle\right)^2.\end{aligned}$$

By Cauchy-Schwarz, this is bounded by

$$\begin{aligned}\langle w_k, g_x\rangle^2 &\leq 2c_x^2 + 2\sum_{x'\in Z\setminus\{x\}} c_{x'}^2 \cdot \sum_{x'\in Z\setminus\{x\}} \langle g_x, g_{x'}\rangle^2 \\ &\leq 2c_x^2 + 2\cdot 2\cdot(|X|-1)\cdot \mathbb{E}_{x'\in Z\setminus\{x\}}\langle g_x, g_{x'}\rangle^2 \\ &\leq 2c_x^2 + 2\cdot 2\cdot(|X|-1)\cdot\frac{1}{12|X|^2} \\ &\leq 2c_x^2 + \frac{1}{3|X|}.\end{aligned}$$

Taking expectation over Z, we get

$$\mathbb{E}_{x\in Z}\langle w_k, g_x\rangle^2 \leq \mathbb{E}_{x\in Z}\left[2c_x^2 + \frac{1}{3|X|}\right] = \frac{2}{|Z|}\sum_{x\in Z} c_x^2 + \frac{1}{3|X|} \leq \frac{4}{|Z|} + \frac{1}{3|X|}.$$

Because $|Z| \geq |X|/3$, this is at least $\frac{12}{|X|} + \frac{1}{3|X|} \leq \frac{13}{|X|}$. This holds for all $k \in D$, so we get

$$\begin{aligned}\mathbb{E}_{x\in Z}\left[||\mathrm{proj}_U(g_x)||^2\right] &= \mathbb{E}_{x\in Z}\left[\sum_{k\in D}\langle g_x, u_k\rangle^2\right] \\ &= \mathbb{E}_{x\in Z}\left[\sum_{k\in D}\langle g_x, w_k\rangle^2\right] \\ &= \sum_{k\in D}\left[\mathbb{E}_{x\in Z}\langle g_x, w_k\rangle^2\right] \\ &\leq D\cdot\frac{13}{|X|}.\end{aligned}$$

So there is some $x \in Z$ such that

$$||g_x^{\leq t}||^2 = ||\mathrm{proj}_U(g_x)||^2 \leq \frac{13D}{|X|} \leq \frac{13\binom{n+t}{\leq t}}{13\cdot 2^{2(k+1)}\cdot\binom{n+k}{\leq t}} = 2^{-2(k+1)},$$

and thus $||g_x^{\leq t}|| \leq 2^{-k-1}$. This is the function with small low-degree Fourier mass that we're looking for.

Step 3: All that remains is to bound the high-degree Fourier mass of this g_x. We know that g_x agrees with the function C_x on at least $\frac{1}{2}+2^{-k-1}$ fraction of inputs. Thus,

$$\langle g_x, C_x\rangle = 2\Pr[g_x = C_x] - 1 \geq 2\cdot\left(\frac{1}{2}+2^{-k-1}\right) - 1 = 2^{-k}.$$

But also, by Cauchy-Schwarz,

$$\langle g_x, C_x\rangle = \langle g_x^{\leq t}, C_x^{\leq t}\rangle + \langle g_x^{>t}, C_x^{>t}\rangle \leq ||g_x^{\leq t}|| + ||C_x^{>t}||.$$

Because C_x is a circuit of depth h and size M, by LMNT,

$$||C_x^{>t}||^2 \leq 2\cdot 2^{-t/O_h(\log M)^{h-1}}.$$

Therefore,

$$2^{-k} \leq ||g_x^{\leq t}|| + ||C_x^{>t}|| \leq 2^{-k-1} + 2\cdot 2^{-t/O_h(\log M)^{h-1}}.$$

Solving for M, we get

$$M \geq 2^{\Omega_h\left(\left[\frac{t}{2k+3}\right]^{1/(h-1)}\right)},$$

as desired.

□

5 Applying the main theorem

In the remainder of the paper, we illustrate the utility of Theorem 4 by presenting a few examples of functions f which satisfy its hypotheses. In order to apply Theorem 4 to a Boolean function f, f must satisfy the technical conditions in the hypothesis:

- f is right one-to-one function.

- (***Almost-orthogonality***) There exist integers $0 \leq k \leq n/2-1$, $0 \leq t \leq n+k$, and a subset $X \subset \{-1,1\}^n$ of size $|X| \geq 13 \cdot 2^{2(k+1)} \cdot \binom{n+k}{\leq t}$ such that
$$\mathbb{E}_{x \neq x' \sim X}\left[\langle f_x, f_{x'} \rangle^2\right] \leq \frac{2^{2k}}{36|X|^2}.$$

Typically, checking the right one-to-one condition is easy (it's in some sense almost without loss of generality, since if f behaves identically on two different y's you might as well be encoding your y's less wastefully). The almost-orthogonality condition is much more stringent. In order to apply Theorem 4, we first translate the technical formulation of this condition into "plain mathematics" by presenting a simplified statement of the condition:

- (***Simplified almost-orthogonality***) There exists a subset $\tilde{X} \subset \{-1,1\}^n$ with $|\tilde{X}| = 2^{\Omega(n)}$ and $\mathbb{E}_{x \neq x' \sim X} \langle f_x, f_{x'} \rangle^2 = 2^{-\Omega(n)}$.

It is not hard to see that simplified almost-orthogonality implies almost-orthogonality.

Proposition 1. Let $f : \{-1,1\}^n \times \{-1,1\}^n \to \{-1,1\}$ be a Boolean function satisfying the simplified almost-orthogonality condition. Then, for any $0 \leq \alpha < 1$, there is some constant a such that f satisfies the almost-orthogonality condition for any $0 \leq k < n^\alpha$, $t = a(n+k)$.

Proof. Suppose that there are constants c_1, c_2 such that, for sufficiently large n, there exists a subset $\tilde{X} \subset \{-1,1\}^n$ with $|\tilde{X}| \geq 2^{c_1 n}$ and $\mathbb{E}_{x \neq x' \sim \tilde{X}} \langle f_x, f_{x'} \rangle^2 \leq 2^{-c_2 n}$. Let s be an integer. By averaging, there is a set $X_s \subset \tilde{X}$ with size s such that $\mathbb{E}_{x \neq x' \sim X_s} \langle f_x, f_{x'} \rangle^2 \leq 2^{-c_2 n}$. So it suffices to choose $s \leq |\tilde{X}|$ to satisfy

$$13 \cdot 2^{2(k+1)} \cdot \binom{n+k}{\leq a(n+k)} \leq 13 \cdot 2^{2(k+1)} \cdot 2^{H(a)(n+k)} \leq s \leq \frac{2^k}{6 \cdot 2^{-c_2 n/2}} \leq \frac{2^k}{6 \cdot \left(\mathbb{E}_{x \neq x' \sim \tilde{X}_s} \langle f_x, f_{x'} \rangle^2\right)^{1/2}};$$

then, X_s will be the desired set. In order for such an s to exist, the left hand side must be strictly smaller than both the right hand side and $|\tilde{X}|$. By choosing a small enough that $H(a) < c_1$ and $H(a) < c_2/2$, we get

$$13 \cdot 2^{2(k+1)} \cdot 2^{H(a)(n+k)} \leq 2^{H(a)n + o(n)} \leq 2^{c_1 n}, 2^{c_2/2n} \leq |\tilde{X}|, \frac{2^k}{6 \cdot 2^{-c_2 n/2}},$$

so both conditions are met and we can choose a satisfactory s. □

Using this formulation of the almost-orthogonality condition, we obtain the following reformulation of Theorem 4.

Theorem 5. Let $f : \{-1,1\}^n \times \{-1,1\}^n \to \{-1,1\}$ be a Boolean function. Let $0 \leq \alpha < 1$, let $0 \leq k \leq n^\alpha$. Suppose that the following hold:
- f is a right one-to-one function.
- There exists a subset $\tilde{X} \subset \{-1,1\}^n$ with $|\tilde{X}| = 2^{\Omega(n)}$ and $\mathbb{E}_{x \neq x' \sim X} \langle f_x, f_{x'} \rangle^2 = 2^{-\Omega(n)}$.
- There exists a depth-h size-M circuit C, and arbitrary functions A, B such that $B : \{-1,1\}^n \to \{-1,1\}^{n+k}$ and $C(A(x), B(y)) = f(x,y)$.

Then M has exponential size; in particular,

$$M \geq 2^{\Omega_h\left(n^{(1-\alpha)/(h-1)}\right)}.$$

We will now see a few examples of functions which satisfy the simplified almost-orthogonality condition, and to which Theorem 5 applies. The most obvious example is the inner product function itself: it satisfies simplified almost-orthogonality by taking $\tilde{X} = \{-1,1\}^n$. So, applied to IP, Theorem 5 gives a strengthening of Theorem 3, where the preprocessing function B is allowed to extend the input length by n^α bits rather than keeping it constant.

The other two examples – weak PRFs and rounded inner product functions – will require more work.

5.1 Applying the main theorem to weak PRFs

Our goal in this section will be to show that exponentially secure PRFs satisfy the near-orthogonality condition of the main theorem. This will imply that subexponential-size AC^0 circuits can't compute them with preprocessing that only extends one input, giving modest evidence against encoded-input PRFs in AC^0. First, a reminder of the definition of a weak PRF:

Definition 1. A function $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ is a weak PRF with security s if, for any size-s circuit C, the probability over random key k, random inputs $x_1,x_2,...$, and the internal randomness of C, that $C(x_1,F(x_1),x_2,F(x_2),...)$ accepts differs by at most $\frac{1}{s}$ from the probability that $C(U)$ accepts, where U is a uniform random bitstring.

Theorem 6. For any constants $c > 1$, $\alpha \in (0,1)$, and any right one-to-one weak PRF[3] F with security $2^{n/c}$, if $A : \{0,1\}^n \to \{0,1\}^{n+n^\alpha}$, $B : \{0,1\}^n \to \{0,1\}^{\mathsf{poly}(n)}$, $C : \{0,1\}^{n+n^\alpha} \times \{0,1\}^{\mathsf{poly}(n)} \to \{0,1\}$ are such that $C(A(k),B(x)) = F_k(x)$ on all inputs, then C cannot be computed by subexponential-size AC^0 circuits.

Proof. Let's move to the $\{-1,1\}$ domain for the purposes of this argument. If we could show that

$$\mathbb{E}_{k\neq k'}[\langle F_k,F_{k'}\rangle^2] \leq 2^{-\Omega(n)},$$

then F would satisfy simplified almost-orthogonality with $X = \{-1,1\}^n$, and by Theorem 5 we'd be done. So, let's assume that

$$\mathbb{E}_{k,k'}[\langle F_k,F_{k'}\rangle^2] > 2^{-n/100c},$$

and derive a contradiction.

If we took a truly random function U, we know

$$\mathbb{E}_{U,k'}[\langle U,F_{k'}\rangle^2] = 2^{-n}.$$

This suggests a approach for distinguishing F_k from a truly random function: estimate the inner product of the function with $F_{k'}$ for a random k', and check whether the inner product magnitude looks suspiciously high. Suppose, for now, that we're able to compute the inner product of two functions with a small circuit. By Hoeffding, we can certainly say

$$\Pr_{U,k'}[\langle U,F_k'\rangle^2 > 2^{-n/50c}] \leq 2^{-n/50c}.$$

So, by assumption that F is a weak PRF of security s, we must have $\Pr_{k,k'}[\langle F_k,F_{k'}\rangle^2 > 2^{-n/50c}] \leq 2^{-n/50c} + 1/s \leq 2^{-n/60}$, because otherwise we could get more than $1/s$ advantage in distinguishing between the PRF and a true random function by choosing a random k' ourself and accepting if the inner product between the mystery function and $F_{k'}$ is larger than $\sqrt{2^{-n/60c}}$ in magnitude. But now, this means that

$$\begin{aligned}\mathbb{E}_{k,k'}[\langle F_k,F_{k'}\rangle^2] &\leq 1 \cdot \Pr_{k,k'}[\langle F_k,F_{k'}\rangle^2 > 2^{-n/50c}] + 2^{-n/50c} \cdot \Pr_{k,k'}[\langle F_k,F_{k'}\rangle^2 \leq 2^{-n/50c}] \\ &\leq 2^{-n/60c} + 2^{-n/50c} \leq 2^{-n/70c},\end{aligned}$$

contradicting our assumption that $\mathbb{E}_{k,k'}[\langle F_k,F_{k'}\rangle^2] > 2^{-n/100c}$. Of course, we haven't quite won yet, because we don't actually have a small circuit to exactly compute the inner product of the mystery function and our randomly chosen $F_{k'}$. But, note that the above analysis still holds even if we only have a circuit that computes something within $2^{-n/10c}$ of the inner product of two functions, with failure probability less than $2^{-n/10c}$. And this we can certainly do: given oracle access to two functions f and g, our circuit will choose $N = 2^{-n/5c}$ random inputs $x_1,...,x_n \in \{-1,1\}^n$, and compute the approximate inner product by taking

$$\frac{1}{N}\sum_{i\in[N]} f(x_i)g(x_i).$$

Again by Hoeffding, this gives with high probability a good approximation to the inner product. Noting that this procedure can be implemented by a circuit much smaller than s, we have shown the desired contradiction. □

[3] The right one-to-one condition can be removed if F is a strong PRF. Also, the theorem can be shown for PRFs of somewhat worse security, with correspondingly weaker impossibility guarantees for AC^0 circuits – we are just presenting this instantiation of the theorem for concreteness.

5.2 Applying the main theorem to rounded inner products

A more concrete example of a class of functions which satisfy the hypotheses of Theorem 4 consists of the "rounded inner product functions".

Definition 2. For an integer $q \geq 2$ and set $R \subset \{0,1,\ldots,q-1\}$, the (q,R)-rounded inner product function $\mathsf{IP}^{[q,R]}: \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ is given by

$$\mathsf{IP}^{[q,R]}(x,y) = \begin{cases} 0 & \sum_{i=1}^n x_i y_i \pmod{q} \in R \\ 1 & \text{otherwise} \end{cases}.$$

So the function $\mathsf{IP}^{[2,\{0\}]}$, for example, is the normal inner product mod 2.

Of note is that some specific rounded inner product functions, like $\mathsf{IP}^{[6,\{0,1,2\}]}$ (presented as (*) in Section 1.2), are conjectured to be weak PRFs [Bon+18]. If this conjecture is true, then these functions would be covered by the result of Theorem 6. However, some rounded inner product functions, like the plain mod 2 inner product, are provably *not* weak PRFs. It turns out that all rounded inner products, regardless of weak PRF status, satisfy the hypotheses of Theorem 4. Thus, we obtain that rounded inner products, too, cannot be computed by subexponential-size AC^0 circuits with preprocessing that only extends one input.

Theorem 7. Let $q \geq 2$ be even, let $0 \leq \alpha < 1$ be a constant, let $R \subset \{0,1,\ldots,q-1\}$ be a subset of size $\frac{q}{2}$, and let $f(x,y) = \mathsf{IP}^{[q,R]}(x,y)$. If $k \leq n^\alpha$ and there exists a depth-h size-M circuit C and arbitrary functions A, B with $B: \{-1,1\}^n \to \{-1,1\}^{n+k}$, then C cannot be computed by subexponential-size AC^0 circuits.

Proof. By Theorem 5, it suffices to find a subset $X \subset \{0,1\}^n$ with $|X| \geq 2^{\Omega(n)}$ and $\mathbb{E}_{x \neq x' \sim X} \langle f_x, f_{x'} \rangle^2 = 2^{-\Omega(n)}$. We will take X to be a set with large minimum Hamming distance. The main step of the proof is to argue that if $x,x' \in \{0,1\}^n$ have large Hamming distance, then the inner product $\langle f_x, f_{x'} \rangle$ is small. To compute this inner product, we will estimate the probability $\Pr_y[f_x(y) = f_{x'}(y)]$. Estimating this probability requires knowing how likely a specific function f_x is to take on certain values, which is accomplished by the lemma below:

Lemma 3. Let $q \geq 2$ be an integer. There exists a constant $0 < c_q < 1$ such that, for any $r \in \{0,1,\ldots,q-1\}$ and large enough n,

$$\Pr_{y \sim \{0,1\}^n}\left[\sum_i y_i = r\right] = \frac{1}{q} \pm O(c_q^n).$$

Sketch of proof. This follows from convergence properties of Markov chains. □

Suppose that x,x' have Hamming distance at least $n/3$. Let $S_x, S_{x'}$ be the subsets of $[n]$ characterized by x and x'. Since $d(x,x') = |S_x \setminus S_{x'}| + |S_{x'} \setminus S_x|$, at least one of these subsets must have size at least $n/6$. Without loss of generality, suppose $|S_x \setminus S_{x'}| \geq n/6$, and let $J = S_x \setminus S_{x'} = \{i \mid x_i = 1, x'_i = 0\}$. We can write the relevant inner products as

$$\sum_{i=1}^n x_i y_i = \sum_{i \in J} x_i y_i + \sum_{i \in \overline{J}} x_i y_i = \sum_{i \in J} 1 \cdot y_i + \sum_{i \in \overline{J}} x_i y_i,$$

and

$$\sum_{i=1}^n x'_i y_i = \sum_{i \in J} x'_i y_i + \sum_{i \in \overline{J}} x'_i y_i = \sum_{i \in J} 0 \cdot y_i + \sum_{i \in \overline{J}} x'_i y_i = \sum_{i \in \overline{J}} x'_i y_i.$$

For any fixed $v \in \{0,1\}^{\overline{J}}$, the restriction

$$(f_{x'})_{\overline{J} \to v}(y) = \begin{cases} 0 & \sum_{i \in \overline{J}} x'_i v_i \in R \\ 1 & \text{otherwise} \end{cases}$$

is constant. Hence, since $|R| = q/2$, there are $q/2$ possible values that $\sum_{i \in J} y_i$ can take on that will make $(f_x)_{\overline{J} \to v}(y)$ agree with $(f_{x'})_{\overline{J} \to v}(y)$. By Lemma 3, this will occur with probability

$$\frac{q}{2} \cdot \left(\frac{1}{q} \pm O(c^{|J|}) \right) = \frac{1}{2} \pm O(c^{n/6}).$$

Therefore, the probability that $f_x = f_{x'}$ is

$$\begin{aligned} \Pr_y[f_x(y) = f_{x'}(y)] &= \mathbb{E}_v \left[\Pr_y[(f_x)_{\overline{J} \to v}(y) = (f_{x'})_{\overline{J} \to v}(y)] \right] \\ &= \mathbb{E}_v \left[\frac{1}{2} \pm O(c^{n/6}) \right] \\ &= \frac{1}{2} \pm O(c^{n/6}), \end{aligned}$$

and so the inner product can be computed as

$$\langle f_x, f_{x'} \rangle = 2 \Pr_y[f_x(y) = f_{x'}(y)] - 1 = \pm O(c^{n/6}).$$

There is thus a constant K such that

$$\langle f_x, f_{x'} \rangle^2 \le K c^{n/3} = 2^{-\Omega(n)}$$

for all x, x' with Hamming distance at least $n/3$.

It follows that any set $X \subset \{0,1\}^n$ with minimum Hamming distance at least $\frac{n}{3}$ will have the f_x's close to orthogonal. By the Gilbert-Varshanov bound, such a set $\tilde{X}$ exists with size at least $|\tilde{X}| \ge 2^{n(1-H(1/3))} = 2^{\Omega(n)}$, so the almost-orthogonality condition is satisfied.

□

6 A relationship between proving IPPP and learning AC^0

As a final note, we will briefly mention an observation about attempts to prove the IPPP conjecture conditional on cryptographic assumptions. Recall the weak PRF candidate of rounded mod 6 inner product, presented as (*) in Section 1.2. Given that the potentially cryptographically-hard function is essentially just an inner product, we might hope that assuming this PRF is secure could be useful in establishing the IPPP conjecture conditionally. Filmus, Ishai, Kaplan, and Kindler were unable to establish such a result, but they did establish an interesting sort of win-win theorem.

Theorem 8. Suppose rounded mod 6 inner product is an exponentially-secure weak PRF. Then, either

1. inner product mod 2 can't be computed by poly-size AC^0 circuits with poly-time preprocessing,
2. inner product mod 3 can't be computed by poly-size AC^0 circuits with poly-time preprocessing, or
3. there exist poly-time samplable distributions on which no subexponential-time algorithm can learn AC^0.

Proof. Suppose that all 3 of these consequent statements fail – we will use this to show a weak PRF distinguisher for rounded mod 6 inner product. First, note that if statements 1 and 2 both fail, this implies an AC^0 circuit for rounded mod 6 inner product with poly-time preprocessing. Simply build a circuit for inner product mod 2, a circuit for inner product mod 3, and then feed those outputs into a constant-size circuit to output rounded inner product mod 6.

For any value of key k, hardwiring the first input of the circuit described above gives an AC^0 circuit for computing $B(x) \mapsto \mathsf{Round}(\langle k, x \rangle \mod 6)$, where B is our poly-time preprocessing of x. So, observing outputs of the rounded mod 6 PRF on random inputs is the same as observing outputs of that AC^0 circuit on inputs sampled from $B(U_n)$. By the failure of our 3rd assumption, we can learn the AC^0 circuit in subexponential time over this distribution, which would be impossible if we were seeing outputs of a true random function – so we have a subexponential-time distinguisher for the supposed PRF. □

This might seem like a silly statement, because all 3 of the consequent statements seem very plausible, and the cryptographic hypothesis seems very strong. But even though this may well be saying "False $\implies$ True $\vee$ True $\vee$ True", despite some effort none of the involved statements have been succesfully shown on their own, so it's of note that we at least know this implication.

References

[BFS86] Laszlo Babai, Peter Frankl, and Janos Simon. “Complexity classes in communication complexity theory”. In: *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*. 1986, pp. 337–347. DOI: 10.1109/SFCS.1986.15.

[Bon+18] Dan Boneh et al. *Exploring Crypto Dark Matter: New Simple PRF Candidates and Their Applications*. Cryptology ePrint Archive, Paper 2018/1218. https://eprint.iacr.org/2018/1218. 2018. URL: https://eprint.iacr.org/2018/1218.

[Fil+20] Yuval Filmus et al. “Limits of preprocessing”. In: *35th Computational Complexity Conference (CCC 2020)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik. 2020.

[LMN93] Nathan Linial, Yishay Mansour, and Noam Nisan. “Constant depth circuits, Fourier transform, and learnability”. In: *J. ACM* 40.3 (July 1993), pp. 607–620. ISSN: 0004-5411. DOI: 10.1145/174130.174138. URL: https://doi.org/10.1145/174130.174138.

[Tal17] Avishay Tal. “Tight Bounds on the Fourier Spectrum of AC0”. In: *32nd Computational Complexity Conference (CCC 2017)*. Ed. by Ryan O’Donnell. Vol. 79. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017, 15:1–15:31. ISBN: 978-3-95977-040-8. DOI: 10.4230/LIPIcs.CCC.2017.15. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.CCC.2017.15.