Efficient resilient functions*

Peter Ivanov Raghu Meka Emanuele Viola

November 11, 2022

Abstract

An n-bit boolean function is resilient to coalitions of size q if no fixed set of q bits is likely to influence the value of the function when the other $n-q$ bits are chosen uniformly at random, even though the function is nearly balanced. We construct explicit functions resilient to coalitions of size $q = n/(\log n)^{O(\log\log n)} = n^{1-o(1)}$ computable by linear-size circuits and linear-time algorithms. We also obtain a tight size-depth tradeoff for computing such resilient functions.

Constructions such as ours were not available even non-explicitly. It was known that functions resilient to coalitions of size $q = n^{0.63\cdots}$ can be computed by linear-size circuits [BL85], and functions resilient to coalitions of size $q = \Theta(n/\log^2 n)$ can be computed by quadratic-size circuits [AL93].

One component of our proofs is a new composition theorem for resilient functions.

A function $f : \{0,1\}^n \to \{0,1\}$ is *resilient to coalitions of size* q if, informally, no adversary controlling q input bits can noticeably influence the output of the function, when the other $n-q$ bits are chosen uniformly at random. A large number of works, many of which are discussed below, has been devoted to constructing and analyzing resilient functions. Indeed, the study of resilient functions is fundamental in the analysis of boolean functions [O'D14] and has found many applications, ranging from the original ones about *collective coin-flipping protocols* [BL85, AL93, RZ98] to the construction of *randomness extractors* [KZ07, GVW15, CZ16, Mek17, CL18, HIV22], to *correlation bounds for polynomials* [CHH+20].

Given that resilient functions are such powerful objects, the main question we address is: what are the minimal resources needed to compute resilient functions? This question was explicitly raised in [HIV22].

Before we discuss our answers to this question, we give some background on previous constructions. A standard example of a resilient function is the majority function, which is resilient to coalitions of size $\Theta(\sqrt{n})$. Ben-Or and Linial [BL85] proved that one can improve the resilience by recursively composing majority on three bits. This yields a function which is resilient to coalitions of size $\Theta(n^{\log_3 2}) = \Theta(n^{0.63\cdots})$ and computable by a linear-size circuit. Jumping ahead, our techniques are flexible enough to recover this result (and show something much stronger).

*This paper subsumes an unpublished work by Meka. PI and EV are partially supported by NSF grant CCF-2114116.

There is a beautiful result by Ajtai and Linial [AL93] which remains essentially state-of-the-art to this day. They prove *the existence* of functions computable by quadratic-size circuits which are resilient to coalitions of size $\Theta(n/\log^2 n)$. This is nearly the optimal resilience one can achieve, as any boolean function on n bits can be controlled by coalitions of size $\Theta(n/\log n)$ [KKL88].

The fact that resilient functions with parameters as in the Ajtai-Linial result exist at all may be quite shocking. Furthermore, given that the quadratic-size bound for non-explicit circuits which are resilient to coalitions of size $n^{1-o(1)}$ has not been improved for 30 years, one might naively conjecture this is optimal.

One might also consider the less demanding conjecture that super-linear circuit size is required, as was conjectured for a list of combinatorial objects including super-concentrators [Val77] and universal hash functions [IKOS08] – conjectures which were disproved with significant impact on our intuition of what can be computed super efficiently. For more discussion on surprising bounds for seemingly difficult combinatorial objects see [Vio18].

Our contribution is adding resilient functions to the list of objects which can be computed super efficiently. In other words, we show the existence of functions with almost the same resiliency as the Ajtai-Linial circuits that can be computed by *linear-time algorithms* and *linear-size circuits*. Furthermore, our techniques also yield *explicit* functions with similar parameters. We highlight the fact that before our work, it was not clear how to achieve this even non-explicitly.

Next we define resiliency. As in the literature, we will in fact give *tradeoffs* between the coalition size and how much it can influence the function. We put forth a somewhat different definition, which is closely related to the standard definition of resiliency (see Definition 9 and Lemma 10) but connects better to our proof strategy.

Definition 1. Let $f : \{0,1\}^n \to \{0,1\}$. For $G \subseteq \{1,2,\ldots,n\}$ we let $f_G(x,y)$ denote the output of f when the bits indexed by G are set by $x \in \{0,1\}^{|G|}$ and the others by $y \in \{0,1\}^{n-|G|}$. We say that f_G is *fixed* by x to $b \in \{0,1\}$ if $f_G(x,y) = b$ for every y.

Then f is (ρ,β)-*fix resilient* if for any set $G \subseteq \{1,2,\ldots,n\}$ and any $b \in \{0,1\}$, with probability $\geq 1/2 - (n-|G|)\rho - \beta$ over a uniform $x \in \{0,1\}^{|G|}$ we have that f_G is fixed by x to b.

One should think of β as the bias of f. Note if $\beta = 0$ this implies that f is balanced (by taking $G = \{1,2\ldots,n\}$). When this is the case, we will write that f is ρ-fix resilient. One should think of ρ^{-1} as the maximal coalition size that f is resilient to.

To summarize the discussion in the introduction, all previous constructions either require quadratic size (and time), or else have fix resiliency $\geq n^{\Omega(1)}/n$. On the other hand, we prove the following:

Theorem 2. *For all sufficiently large n and b, there are $\left((\log n)^{O(\log\log_b n)}/n, 1/n^{\omega(1)}\right)$-fix resilient functions computable by circuits of size $n \cdot \mathrm{poly}(b)$. In particular:*

1. *There are explicit $\left((\log n)^{O(\log\log n)}/n, 1/n^{\omega(1)}\right)$-fix resilient functions computable by linear-size circuits (set $b := O(1)$).*
2. *There are explicit $\left((\log n)^{O(\log(1/\epsilon))}/n, 1/n^{\omega(1)}\right)$-fix resilient functions computable by $n^{1+\epsilon}$ size circuits for any constant $\epsilon > 0$ (set $b := n^{\epsilon}$).*

We give a number of related constructions. In particular in Theorem 7 we give non-explicit constructions as above but with bias zero (the bias in the theorem above is quasi-polynomially small, and can be traded with the other parameters as will be apparent in the proof).

The depth of the circuits in Theorem 2 is of interest. To discuss this it is convenient to work in the unbounded fan-in model (whereas the size bounds in Theorem 2 are with respect to the bounded fan-in model). In this model, the circuits in Item (1) have depth $O(\log \log n)$. Those in Item (2) have depth $O(\log(1/\epsilon))$. The latter matches a result by Chaudhuri and Radhakrishnan [CR96] which says any circuit of depth $\log(1/\epsilon)/2$ and size $n^{1+\epsilon}$ can be made constant by fixing $O(n^{1-\epsilon^2}) = n^{1-\Omega(1)}$ bits. In other words, there is a coalition of size $n^{1-\Omega(1)}$ that controls the circuit.

Define "resiliency loss" as ρn for a ρ-fix resilient function. Recall that by [KKL88] the smallest possible resiliency loss of any function is $\Omega(\log n)$, while Theorem 3 yields a function with $O(\log^2 n)$ resiliency loss. Since both of these quantities are $2^{\Theta(\log \log n)}$, the "resiliency loss" in Item (1) in Theorem 7 is quasi-polynomial in the optimal. At the other end, Item (2) shows that for any ϵ we can compute functions with resiliency loss polynomial in the optimal by circuits of size $n^{1+\epsilon}$. A natural open problem is to exhibit linear-size circuits computing a function with resiliency loss polynomial in the optimal.

The complexity of extractors Resilient functions have been used to construct extractors in [KZ07, GVW15, CZ16] and subsequent works. We briefly discuss the relevance of our results to this line of works. Resilient functions computable by circuits with *small depth* are important for the approach in [CZ16], and we note that the depth of the circuits in Item (1) is small enough to be used in their framework. This could lead to affine and two-source extractors computable by linear-size circuits. Also, combined with the work [HIV22] our results give non-explicit affine extractors computable by constant-depth circuits composed with a layer of parity gates, a model which lies at the frontier of our understanding of circuit lower bounds. Specifically, we obtain such extractors computable in size $n^{1+\epsilon}$ and depth $O(\log 1/\epsilon)$. This matches a lower bound for computing such extractors in [CGJ$^+$18].

1 Zero-bias constructions

For ease of presentation, we begin in this section with some simpler constructions based on *balanced* functions (i.e., with zero bias). This keeps the parameters to a minimum, while conveying the main ideas. In addition, we will be able to construct explicit functions which are *monotone* and *balanced*, something which is not given by Theorem 2. Moreover, the results in this section will recover and generalize classic results in this area (see below). The drawback is that for some interesting range of parameters the constructions are not explicit. In the proceeding Section 2, by building on these ideas and applying a result by Meka [Mek17], we achieve explicit constructions.

The basic building block of the constructions in this section is the classical Ajtai-Linial construction [AL93], stated next.

Theorem 3. *For sufficiently large n, there are $O(\frac{\log^2 n}{n})$-fix resilient circuits of size $O(n^2)$.*

This result appears to be folklore but as stated does not seem to be in the literature. The works we are aware of [BL85, AL93, RZ98, CZ16, Mek17, Wel20] either don't prove a full tradeoff, or don't achieve bias zero, or have worse resilience. However because the proof is somewhat technical and not needed for our main result, we do not include it in this paper.

1.1 A composition lemma

A new tool we introduce is a generic *composition lemma* for resiliency. The proof in hindsight is not involved given our definition of fix resiliency. One can speculate that the lack of a "correct" definition served as a barrier to proving the result below, despite recursive constructions based on specific functions already existing in the literature [BL85].

Lemma 4. *Let $f' : \{0,1\}^{n'} \to \{0,1\}$ be ρ'-fix resilient, and let $f'' : \{0,1\}^{n''} \to \{0,1\}$ be ρ''-fix resilient. Then $f := f' \circ f'' : \{0,1\}^{n' \cdot n''} \to \{0,1\}$ is $2\rho'\rho''$-fix resilient.*

Proof. Let $n := n' \cdot n''$ and fix a set $G \subseteq \{1, 2, \ldots, n\}$ of size $n - q$. This induces sets G_i of sizes $n - q_i$ for the n' copies of f''.

For $1 \leq i \leq n'$ we let $A_{i,0}, A_{i,1} \subseteq \{0,1\}^{|G_i|}$ be sets of maximal equal density such that for every $x_i \in A_{i,b}$ we have that f''_{G_i} is fixed to b by x_i. And let $B_i := \{0,1\}^{|G_i|} - (A_{i,0} \cup A_{i,1})$.

Note that if $x_i \in \{0,1\}^{|G_i|}$ is uniformly sampled from $A_{i,0} \cup A_{i,1}$ then the output of f''_{G_i} is a uniform bit independent from the rest of the inputs. Furthermore, B_i has density $\leq 2\rho'' q_i$ since f'' is ρ''-fix resilient. We next describe a sampling process which is equivalent to sampling a uniform $x \in \{0,1\}^{|G|}$.

1. For every $1 \leq i \leq n'$ decide independently whether $x_i \in \{0,1\}^{|G_i|}$ is sampled from $A_{i,0} \cup A_{i,1}$ (in which case x_i is *good*) or from B_i with probability equal to the corresponding densities.

2. Sample x_i uniformly from the set picked in the first step.

After the first step, let $G' \subseteq \{1, \ldots, n'\}$ index the set of of good x_i. The second step can be viewed as inputting f' with $|G'|$ uniform independent bits. The locations of these bits depend on the first step, but are fixed before the second step. Then by the resiliency of f', the probability after the second step that x fixes f'_G to 1 is $\geq 1/2 - (n' - |G'|)\rho'$.

So the probability that an x sampled as above fixes f_G to 1 is

$$\geq \sum_{j=0}^{n'} \mathbb{P}[|G'| = j](1/2 - (n' - |G'|)\rho') = 1/2 - \mathbb{E}[n' - |G'|]\rho'.$$

Note that $n' - |G'| = \sum_{i=1}^{n'} 1_{x_i \in B_i}$, where 1_E is the indicator of the event E. By the aforementioned density of B_i, $\mathbb{P}[1_{x_i \in B_i}] \leq 2\rho'' q_i$ which implies that $\mathbb{E}[n' - |G'|] \leq 2\rho'' q$.

Combining everything together, the probability that a uniform $x \in \{0,1\}^{|G|}$ fixes f_G to 1 is $\geq 1/2 - 2\rho'\rho''q$. The case of fixing f_G to 0 is done identically. □

Recursive majority. Lemma 4 allows us to recover the classic result of Ben-Or and Linial about the influence of recursive-majority (Theorem 4(a) in [BL90]). Let $f : \{0,1\}^3 \to \{0,1\}$ be the majority function on 3 bits. Note that f is 1/4-fix resilient. Indeed, when $|G| = 2$ the the function is 1 with probability $\geq 1/4 = 1/2 - 1 \cdot (1/4)$, and when $|G| = 1$ then the function is 1 with trivial probability $\geq 1/2 - 2 \cdot (1/4) = 0$. If we make a balanced ternary tree with 3^t leaves where each node corresponds to f we obtain a function on $3^{t+1} =: n$ bits which by Lemma 4 is fix resilient with resilience $2^t \cdot (1/4)^t = 2^{-t} = 2^{-(\log_3 n - 1)} = 2 \cdot n^{-\log_3 2} = 2 \cdot n^{-0.6309\ldots}$.

The proof in [BL90] is tailored to the specific function and does not easily extend to other functions. On the other hand, our approach allows us to compose arbitrary resilient functions.

1.2 A regular construction

By Theorem 3 and repeatedly applying Lemma 4 we arrive at the next construction.

Theorem 5. *For all sufficiently large n and b, there are $\left(n^{O(\log\log b)/\log b}/n\right)$-fix resilient functions computable by circuits of size $O(nb)$ which can be constructed in time $O(nb) + 2^{O(b^3)}$.*

For some intuition on the parameters, consider setting b to some large enough constant. Then we obtain a function which is $1/n^{1-\delta}$-fix resilient and computable in linear time for some constant $\delta > 0$. Alternatively, we can set $b = \log^{1/3} n$ and obtain a function computable in time $\leq n + O(nb) < n \log n$ with resiliency $n^{O(\log\log\log n)/\log\log n}/n = 1/n^{1-o(1)}$.

Before we start the proof we need the following.

Fact 6. *The circuit from Theorem 3 can be brute-forced in time $2^{O(n^3)}$.*

Proof of Theorem 5. Theorem 3 gives a $O(\log^2 b)/b$-fix resilient function on b bits computable by a circuit of size $O(b^2)$, which itself can be computed in time $2^{O(b^3)}$ by Fact 6.

We recursively compose the previous function t times and obtain a function on $n := b^t$ inputs. The size is asymptotically dominated by the size of the input layer. This layer consists of n/b functions each taking size $O(b^2)$ for a cost of $O(nb)$.

By Lemma 4 the composed function is $O((\log^2 b)/b)^t$-fix resilient. As $t = \log n / \log b$, this equals $O(\log b)^{2\log n/\log b}/b^t = n^{O(\log\log b)/\log b}/n$. □

1.3 An irregular construction

We are able to significantly improve the parameters by modifying the fan-ins of the functions in each layer. For intuition on how this is done, consider applying Theorem 5 with $b = \sqrt{n}$; this results in the composed circuit $C := f_{out}(f_1, f_2, \ldots, f_{\sqrt{n}})$, where each f_i, f_{out} has fan-in $\sqrt{n}$. Note the size of C is dominated by the size of the inner layer, which will be $O(b^2) \cdot \sqrt{n} = O(n^{3/2})$.

Now consider the composition $C' := f_{out}(f_1, f_2, \ldots, f_{n^{2/3}})$ where each f_i has fan-in $n^{1/3}$ and f_{out} has fan-in $n^{2/3}$. Note the size of both layers are balanced and furthermore, they are both $O(n^{4/3})$. Additionally, Lemma 4 implies that C' achieves the same resilience as C.

Generalizing this idea allows us to greatly improve on the resilience, size, and depth of the resulting composition.

Theorem 7. *For all sufficiently large n and b, there are $\left((\log n)^{O(\log\log_b n)}/n\right)$-fix resilient functions computable by circuits of size $O(nb)$.*

Proof. Consider a tree of functions of depth t where each function at level i has fan-in b_i, where $i = 0$ is the level closest to the input with $b_0 = b$. We set $b_{i+1} := b_i^\gamma$ for some constant γ we fix later. Note $b_i = b^{\gamma^i}$ and $n = \prod_{i=0}^{t-1} b_i$.

At level i we use the ρ_i-fix resilient circuits from Theorem 3 on b_i bits, where $\rho_i = O(\log^2 b_i)/b_i$. Additionally, there is some universal constant c such that the circuits have size $\leq cb_i^2$ at every level i.

First we bound the size of the tree. Let s_i denote the size of the layer at level i. We have

$$\frac{s_{i+1}}{s_i} \leq \frac{c(b_{i+1})^2}{b_{i+1} \cdot c(b_i)^2} = \frac{b^{\gamma^{i+1}}}{b^{2\gamma^i}} = b^{\gamma^i(\gamma-2)} \leq 1/2$$

as long as $1 < \gamma < 2$ and b is sufficiently large. Hence, the size $\sum_{i=0}^{t-1} s_i$ of the tree is $O(s_0)$. Note that $s_0 = (n/b) \cdot cb^2 = O(nb)$.

Next we compute the resilience. At level i we use functions which are ρ_i-fix resilient. By Lemma 4 the final resilience is

$$\prod_{i=0}^{t-1} 2\rho_i = \prod_{i=0}^{t-1} O(\log^2 b_i)/b_i = n^{-1} \prod_{i=0}^{t-1} O(\log^2 b_i) \leq n^{-1}(\log n)^{O(t)}.$$

The last inequality follows as $b_i \leq n$ for every i.

To conclude we compute the depth t of the tree in terms of n. We have

$$n = \prod_{i=0}^{t-1} b_i = b^{\sum_{i=0}^{t-1}\gamma^i} = b^{(\gamma^t-1)/(\gamma-1)}.$$

This implies $t = O(\log\log_b n)$. □

2 Proof of Theorem 2

In this section we construct explicit and efficient resilient functions building on the ideas developed in the previous section. The main technical challenge is dealing with the fact that all existing constructions of explicit tradeoff resilient functions have some amount of bias. Hence we will need to generalize the previous results which only dealt with balanced functions.

First we give an analogous version of Lemma 4 that works for unbalanced functions. The proof is similar to before.

Lemma 8. *Let $f' : \{0,1\}^{n'} \to \{0,1\}$ be (ρ', β')-fix resilient, and let $f'' : \{0,1\}^{n''} \to \{0,1\}$ be (ρ'', β'')-fix resilient. Then $f := f' \circ f'' : \{0,1\}^{n' \cdot n''} \to \{0,1\}$ is $(2\rho'\rho'', 2n'\rho'\beta'' + \beta')$-fix resilient.*

For intuition on the parameters, consider some f on m bits which is $(\frac{d}{m}, \frac{d}{m})$-fix resilient. The composition $f \circ f$ will be $(\frac{3d^2}{m^2}, \frac{3d^2}{m})$-fix resilient.

Proof of Lemma 8. Let $n := n' \cdot n''$ and fix a set $G \subseteq \{1, 2, \ldots, n\}$ of size $n - q$. This induces sets G_i of sizes $n - q_i$ for the n' copies of f''.

Next, for $1 \leq i \leq n'$ we let $A_{i,0}, A_{i,1} \subseteq \{0,1\}^{|G_i|}$ be sets of maximal equal density such that for every $x_i \in A_{i,b}$ we have that f''_{G_i} is fixed to b by x_i. And we let $B_i := \{0,1\}^{|G_i|} - (A_{i,0} \cup A_{i,1})$.

Note that if $x_i \in \{0,1\}^{|G_i|}$ is uniformly sampled from $A_{i,0} \cup A_{i,1}$ then the output of f''_{G_i} is a uniform bit independent from the rest of the inputs. Furthermore, B_i has density $\leq 2(\rho'' q_i + \beta'')$ since f'' is (ρ'', β'')-fix resilient. We next describe a sampling process which is equivalent to sampling a uniform $x \in \{0,1\}^{|G|}$.

1. For every $1 \leq i \leq n'$ decide independently whether $x_i \in \{0,1\}^{|G_i|}$ is sampled from $A_{i,0} \cup A_{i,1}$ (in which case x_i is *good*) or from B_i with probability equal to the corresponding densities.

2. Sample x_i uniformly from the set picked in the first step.

After the first step, let $G' \subseteq \{1, \ldots, n'\}$ index the set of good x_i. The second step can be viewed as inputting f' with $|G'|$ uniform independent bits. The locations of these bits depend on the first step, but are fixed before the second step. Then by the resilience of f', the probability after the second step that x fixes f'_G to 1 is $\geq 1/2 - (n' - |G'|)\rho' - \beta'$.

So the probability that an x sampled as above fixes f_G to 1 is

$$\geq \sum_{j=0}^{n'} \mathbb{P}[|G'| = j](1/2 - (n' - |G'|)\rho' - \beta') = 1/2 - \mathbb{E}[n' - |G'|]\rho' - \beta'.$$

Note that $n' - |G'| = \sum_{i=1}^{n'} 1_{x_i \in B_i}$. By the aforementioned density of B_i, $\mathbb{P}[1_{x_i \in B_i}] \leq 2(\rho'' q_i + \beta'')$ which implies $\mathbb{E}[n' - |G'|] \leq 2(\rho'' q + n'\beta'')$.

Combining everything together, the probability that a uniform $x \in \{0,1\}^{|G|}$ fixes f_G to 1 is $\geq 1/2 - 2\rho'\rho''q - 2n'\rho'\beta'' - \beta'$. The case of fixing f_G to 0 is done similarly. □

We will need another way to compose resilient functions. We seek to drive the bias down while not damaging the resiliency too much. We accomplish this by simply taking the XOR of r independent copies. We will use this in conjunction with a result in [Mek17] that gives explicit tradeoff resilient functions albeit with poor bias.

However, the result in [Mek17] is not stated for fix resilience. So first we state a different definition of tradeoff resilience (referred to as 'strong resilience' in the literature) and then we relate it to fix resilience.

Definition 9. We say that f is (ρ, β)-*change resilient* if $|\mathbb{P}[f = 1] - 1/2| \leq \beta$ and if for any set $G \subseteq \{1, 2, \ldots, n\}$, with probability $\geq 1 - (n - |G|)\rho$ over $x \in \{0,1\}^{|G|}$ we have that f_G is fixed by x to either 0 or 1.

Lemma 10. *If f is (ρ, β)-change resilient then it is (ρ, β)-fix resilient. Moreover, if f is (ρ, β)-fix resilient then f is $(2(\rho + \beta), \beta)$-change resilient.*

Proof. Suppose f is (ρ, β)-change resilient, and fix some set $G \subset [n]$ of size $n - q$. Then the probability over $x \in \{0,1\}^{|G|}$ that f_G is fixed by x to 1 is at least

$$\mathbb{P}[f = 1] - \mathbb{P}_{x \sim \{0,1\}^{|G|}}[f_G \text{ not fixed by } x] \geq 1/2 - \beta - \rho q.$$

The same proof shows that f_G is fixed to 0 with the same probability.

Now suppose that f is (ρ, β)-fix resilient. First note this immediately implies $|\mathbb{P}[f = 1] - 1/2| \leq \beta$. Now fix any G of size $n-q$ where $q \geq 1$. Over a uniform $x \in \{0,1\}^{|G|}$, f_G is fixed to either 0, 1 with probability $\geq (1/2 - q\rho - \beta) + (1/2 - q\rho - \beta) = 1 - 2q\rho - 2\beta \geq 1 - 2q(\rho + \beta)$. □

Lemma 11 ([Mek17], Theorem 1.2)**.** *There are explicit $(O(\log^2 m)/m, 1/20)$-change resilient functions on m bits computable by circuits of size $poly(m)$.*

Applying the strategy above we obtain the following.

Lemma 12. *For every m, r there are explicit $(O(\log^2 m)/m, 2^{-r})$-fix resilient functions on $O(mr)$ bits computable by circuits of size $poly(m, r)$.*

Proof. We compose the functions in Lemma 11 on input length m with the XOR on cr bits, with c a constant set so that the resulting bias is 2^{-r}.

Now fix a coalition $Q \subseteq \{1, \ldots, cmr\}$ of size q. Q induces coalitions of sizes $q_1, q_2, \ldots, q_{cr}$ for each of the cr subfunctions on m bits. By the definition of change resilience, the probability the i-th coalition can change the output of the i-th function is $\leq q_i \cdot O(\log^2 m)/m$.

So the probability that Q can change the final output is by a union bound at most $q \cdot O(\log^2 m)/m$. We conclude by applying Lemma 10 which implies fix resilience. □

Finally we can prove our main result.

Proof of Theorem 2. It suffices to prove the theorem for $b \geq \log^3 n$. The theorem for smaller b then follows by padding. That is, to obtain the theorem for arbitrary n and $b < \log^3 n$ use the construction for $n' := n/\log^a n$ and $b' := \log^3 n'$, padded to n input bits. The circuit size is $n'\text{poly}(b') < n$ for all sufficiently large a. This change does not affect the resiliency parameter because

$$(\log n')^{O(\log \log_{b'} n')}/n' \leq (\log n)^{O(\log \log_b n)}/n$$

up to the constant in the big-Oh. The bias parameter is similarly unaffected.

From this point we assume that $b \geq \log^3 n$. We construct a tree of functions of depth t where each function at level i has fan-in b_i, where $i = 0$ is the level closest to the input with $b_0 = b$. We set $b_{i+1} := b_i^\gamma$ for some constant γ we fix later. Note $b_i = b^{\gamma^i}$ and $n = \prod_{i=0}^{t-1} b_i$.

On input length b_i we apply the explicit $(r\rho_i, 2^{-r})$-fix resilient function on b_i bits given by Lemma 12, for $r := \log^2 n$ and $\rho_i := O(\log^2 b_i)/b_i$. Here we use that $b \geq \log^3 n$.

We can fix some universal constant c so that every circuit above on fan-in b_i has size $\leq b_i^c$. Using this we first analyze the total size of the composed circuit. Let s_i be the circuit size of the layer at distance i from the input. We have

$$\frac{s_{i+1}}{s_i} \leq \frac{b_{i+1}^c}{b_{i+1} \cdot {b_i}^c} = \frac{(b^{\gamma^{i+1}})^{c-1}}{(b^{\gamma^i})^c} = b^{\gamma^i((c-1)\gamma - c)} \leq 1/2$$

for a universal γ and all sufficiently large b. Hence, the total size $\sum_{i=0}^{t-1} s_i$ of the circuit is $O(s_0)$. We have $s_0 \leq (n/b) \cdot b^c = nb^{c-1}$.

Now we compute the resulting resilience. By Lemma 8, the final resiliency is

$$\prod_{i=0}^{t-1} 2r\rho_i = n^{-1} \prod_{i=0}^{t-1} O(1) r \log^2 b_i \leq n^{-1} \log^{O(t)} n.$$

The last inequality follows since $b_i \leq n$ for every i and $r = \log^2 n$.

Next we deal with the bias parameter. Let β_i denote the bias at level i of the tree. Recall $\beta_0 = 2^{-r}$. We bound β_i for $i > 0$ by Lemma 8 which says

$$\beta_i \leq 2 \cdot \beta_{i-1} \cdot b_i \cdot (r\rho_i) + 2^{-r} \leq O(\beta_{i-1} \cdot \log^4 n).$$

The last inequality follows as β_i is non-decreasing as i increases, and $\beta_i \geq \beta_0 = 2^{-r}$. Hence the final bias β_{t-1} is

$$\leq \beta_0 \cdot \prod_{i=0}^{t-1} O(\log^4 n) \leq 2^{-\log^2 n} \cdot (\log n)^{O(t)}.$$

Next we bound the depth t of the tree. By construction,

$$n = \prod_{i=0}^{t-1} b_i = b^{(\gamma^t - 1)/(\gamma - 1)}.$$

This implies $t = O(\log \log_b n)$.

Plugging this into the bound above for the resiliency parameter yields the desired result. For the bias parameter we note that $2^{-\log^2 n} \cdot (\log n)^{O(\log \log_b n)} \leq 1/n^{\omega(1)}$. □

References

[AL93] Miklos Ajtai and Nathan Linial. The influence of large coalitions. *Combinatorica*, 13:129–145, 1993.

[BL85] Michael Ben-Or and Nathan Linial. Collective coin flipping, robust voting schemes and minima of Banzhaf values. In *26th Symposium on Foundations of Computer Science*, pages 408–416, Portland, Oregon, 21–23 October 1985. IEEE.

[BL90] Michael Ben-Or and Nathan Linial. Collective coin-flipping. In Silvio Micali, editor, *Randomness and Computation*, pages 91–115. Academic Press, New York, 1990.

[CGJ+18] Mahdi Cheraghchi, Elena Grigorescu, Brendan Juba, Karl Wimmer, and Ning Xie. $AC^0 \circ \mathrm{mod}_2$ lower bounds for the boolean inner product. *J. Comput. Syst. Sci.*, 97:45–59, 2018.

[CHH+20] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, Shachar Lovett, and David Zuckerman. XOR lemmas for resilient functions against polynomials. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *ACM Symp. on the Theory of Computing (STOC)*, pages 234–246. ACM, 2020.

[CL18] Kuan Cheng and Xin Li. Randomness extraction in AC0 and with small locality. In Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*, volume 116 of *LIPIcs*, pages 37:1–37:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[CR96] Shiva Chaudhuri and Jaikumar Radhakrishnan. Deterministic restrictions in circuit complexity. In *28th ACM Symp. on the Theory of Computing (STOC)*, pages 30–36, 1996.

[CZ16] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. In *ACM Symp. on the Theory of Computing (STOC)*, pages 670–683, 2016.

[GVW15] Oded Goldreich, Emanuele Viola, and Avi Wigderson. On randomness extraction in AC0. In *IEEE Conf. on Computational Complexity (CCC)*, 2015.

[HIV22] Xuangui Huang, Peter Ivanov, and Emanuele Viola. Affine extractors and ac0-parity. In *Workshop on Randomization and Computation (RANDOM)*, 2022.

[IKOS08] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with constant computational overhead. In *40th ACM Symp. on the Theory of Computing (STOC)*, pages 433–442, 2008.

[KKL88] Jeff Kahn, Gil Kalai, and Nathan Linial. The influence of variables on Boolean functions. In *29th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 68–80, 1988.

[KZ07] Jesse Kamp and David Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. *SIAM J. Comput.*, 36(5):1231–1247, 2007.

[Mek17] Raghu Meka. Explicit resilient functions matching ajtai-linial. In *ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 1132–1148, 2017.

[O'D14] Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.

[RZ98] A. Russell and D. Zuckerman. Perfect information leader election in log*n+o(1) rounds. In *Proceedings 39th Annual Symposium on Foundations of Computer Science*, pages 576–583, 1998.

[Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In *6th Symposium on Mathematical Foundations of Computer Science*, volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1977.

[Vio18] Emanuele Viola, 2018. https://emanueleviola.wordpress.com/2018/02/16/i-believe-pnp/.

[Wel20] Jake Wellens. *Assorted results in boolean function complexity, uniform sampling and clique partitions of graphs*. PhD thesis, Massachusetts Institute of Technology, 2020.