Substitution-Permutation Networks, Pseudorandom Functions, and Natural Proofs

ERIC MILES and EMANUELE VIOLA, Northeastern University

This article takes a new step towards closing the gap between pseudorandom functions (PRF) and their popular, bounded-input-length counterparts. This gap is both quantitative, because these counterparts are more efficient than PRF in various ways, and methodological, because these counterparts usually fit in the substitution-permutation network paradigm (SPN), which has not been used to construct PRF.

We give several candidate PRF $\mathcal{F}_i$ that are inspired by the SPN paradigm. Most of our candidates are more efficient than previous ones. Our main candidates are as follows.

- $\mathcal{F}_1 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is an SPN whose S-box is a random function on b bits given as part of the seed. We prove that $\mathcal{F}_1$ resists attacks that run in time $\leq 2^{\epsilon b}$.
- $\mathcal{F}_2 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is an SPN where the S-box is (patched) field inversion, a common choice in practical constructions. We show that $\mathcal{F}_2$ is computable with boolean circuits of size $n \cdot \log^{O(1)} n$ and that it has exponential security $2^{\Omega(n)}$ against linear and differential cryptanalysis.
- $\mathcal{F}_3 : \{0, 1\}^n \rightarrow \{0, 1\}$ is a nonstandard variant on the SPN paradigm, where “states” grow in length. We show that $\mathcal{F}_3$ is computable with TC^0 circuits of size $n^{1+\epsilon}$, for any $\epsilon > 0$, and that it is almost 3-wise independent.
- $\mathcal{F}_4 : \{0, 1\}^n \rightarrow \{0, 1\}$ uses an extreme setting of the SPN parameters (one round, one S-box, no diffusion matrix). The S-box is again (patched) field inversion. We show that $\mathcal{F}_4$ is computable by circuits of size $n \cdot \log^{O(1)} n$ and that it fools all parity tests on $\leq 2^{0.9n}$ outputs.

Assuming the security of our candidates, our work narrows the gap between the Natural Proofs barrier and existing lower bounds in three models: circuits, TC^0 circuits, and Turing machines.

CCS Concepts: • **Theory of computation** → **Cryptographic primitives**

Additional Key Words and Phrases: Pseudorandom function, natural proofs, Advanced Encryption Standard, substitution-permutation network

ACM Reference Format:

Eric Miles and Emanuele Viola. 2015. Substitution-permutation networks, pseudorandom functions, and natural proofs. *J. ACM* 62, 6, Article 46 (December 2015), 29 pages.

DOI: <http://dx.doi.org/10.1145/2792978>

1. INTRODUCTION

This article takes a new step towards closing the gap between pseudorandom functions ([Goldreich et al. 1986], cf. [Goldreich 2001, §3.6]) and their popular, bounded-input-length counterparts. Recall that a pseudorandom function family (hereafter, PRF) is a family of functions on n -bit inputs such that (1) a uniformly random member of the family can be selected and evaluated in time $\text{poly}(n)$, and (2) no $\text{poly}(n)$ -time adversary with black-box access can distinguish a uniformly random member of the family from a truly random function, except with some small advantage. Notably, PRF are defined

E. Miles and E. Viola were supported by NSF grants CCF-0845003 and CCF-1319206.

Authors' addresses: E. Miles (corresponding author) (current address), Boelter Hall, Room 3714, Department of Computer Science, UCLA, Los Angeles, CA, 90095; email: enmiles@cs.ucla.edu; E. Viola, 440 Huntington Ave., West Village H, Room 338, Boston, MA 02115.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2015 ACM 0004-5411/2015/12-ART46 \$15.00

DOI: <http://dx.doi.org/10.1145/2792978>

on arbitrarily large input lengths n , and their security grows with n ; we will mostly be interested in PRF with exponential security 2^n .

The aforementioned gap is both quantitative and methodological. It is quantitative because essentially all candidate PRF with security 2^n based on complexity-theoretic assumptions (e.g., [Goldreich et al. 1986; Håstad et al. 1999; Naor and Reingold 2004; Naor et al. 2002; Haitner et al. 2010; Vadhan and Zheng 2012; Banerjee et al. 2012]) have seed length at least quadratic in the input length n , which also implies a quadratic lower bound on the circuit size of such PRF. (Subsequent to the initial publication of this work [Miles and Viola 2012], Banerjee and Peikert [2014] gave a PRF construction with quasilinear seed length, but which still has quadratic circuit size.) In contrast, bounded-input-length constructions often have seed length which equals the input length. This is, for example, the case with the 128-bit version of the widely-used block cipher AES by Daemen and Rijmen [2002].

The gap is methodological because the popular counterparts to PRF, namely, bounded-input-length hash functions and block ciphers, often use the *substitution-permutation network* (SPN) structure. An SPN is computed over a number of rounds, where each round “confuses” the input by dividing it into bundles and applying a substitution function (S-box) to each bundle, and then “diffuses” the bundles by applying a matrix with certain “branching” properties (cf. [Shannon 1949]). For example, the SPN structure is used in two of the finalists for the recently concluded SHA-3 cryptographic hash function competition [Gauravaram et al. 2011; Wu 2011], and also in the AES block cipher. No piece of this structure appears to have been used to construct PRF. In fact, until this work, no asymptotic analysis of the SPN structure has been given. This is in stark contrast with the seminal work of Luby and Rackoff [1988], which gave such an analysis for the so-called *Feistel network* structure (which in particular was the basis for the block cipher DES, the predecessor to AES). Moreover, the SPN structure is tailored to resist two general attacks on block ciphers which appear to be ignored in the PRF literature, namely, linear and differential cryptanalysis.

In this article, we give several candidate PRF that are inspired by the SPN structure. Each of the many hash functions and block ciphers based on the SPN structure (e.g., those mentioned previously) suggests different choices for the parameters, S-boxes, and diffusion matrices. As a first step, we choose to follow the design considerations behind the AES block cipher, particularly its S-box. We do this for two reasons. First, it is a well-documented, widely-used block cipher that has been around for over a decade. Second, the algebraic structure of its S-box lends itself to an asymptotic generalization; we exploit this fact in some of our results. We hope that future work will systematically address other available bounded-input-length constructions.

Some of our candidates have better parameters than previous candidates, where by “parameters” we refer to the seed length and the resources required to compute each function in various computational models. Candidates 1, 2, and 5 output n bits, while Candidates 3 and 4 output 1 bit. While there is a generic transformation from 1-bit PRF to n -bit PRF, we note that this transformation incurs a multiplicative increase of $\Omega(n)$ in the computational resources required.

- (1) We first consider an SPN with a random S-box (specified as part of the seed). We prove unconditionally that this resists attacks that run in time less than the seed length. For example, we can set the seed length to n^c and withstand attacks running in time $n^{c'}$ for sufficiently large c and $c' = \Theta(c)$. (Note that being a PRF means that the seed length is n^c and that the function withstands all attacks running in time $n^{c'}$ for any c' .)

This result is analagous to that of Luby and Rackoff, who analyzed the Feistel network structure when a certain component is instantiated with a random

function, and indeed, we prove the same level of security (exponential in the input size of the random function). The techniques used are similar to those in the work by Naor and Reingold [1999] that followed Luby and Rackoff's. To our knowledge, this is the first construction of a (provably secure, inefficient) PRF using the SPN structure.

- (2) Using the AES S-box and a strengthened version of the AES diffusion matrix, we give a candidate with seed length $O(n \log n)$ that is computable with Boolean circuits of size $n \cdot \log^{O(1)} n$. We prove that this candidate has exponential security $2^{\Omega(n)}$ against linear and differential cryptanalysis by extending a result of Kang et al. [2001].
- (3) Again using the AES S-box and a different diffusion matrix, we give a candidate computable with size $n^{1+\epsilon}$, for any $\epsilon > 0$, in the restricted circuit class TC^0 of unbounded fan-in majority circuits of constant-depth. The diffusion matrix used here blows up the state to size $O(n)$, and we output a single bit by taking the inner product of this state with a random string. We prove that this candidate is almost 3-wise independent.
- (4) We give another single-bit output candidate, computable with Boolean circuits of size $n \cdot \log^{O(1)} n$, that uses an extreme setting of the SPN parameters (one round, one S-box, no diffusion matrix). This can be viewed as a slightly modified version of the Even-Mansour cipher [Even and Mansour 1997] that uses the AES S-box in place of a random permutation. We prove that this candidate fools all parity tests that look at $\leq 2^{0.9n}$ outputs.
- (5) Our final candidate is a straightforward generalization of AES and may be folklore. We show that it is computable by size- $O(n^2)$, depth- $O(n)$ Boolean circuits, and we further show that, for each fixed seed k , it is computable in time $O(n^2)$ by a single-tape Turing machine with $O(n^2)$ states (that in particular encode k). We do not have any proof of security, but the (heuristic) arguments underlying AES's security also apply to this candidate.

For context, we mention that Hoory et al. [2005] and Brodsky and Hoory [2008], building on work by Gowers [1996], study the random composition of a family of permutations. The SPN structure can be seen as falling into this framework, by taking each round as an individual permutation chosen randomly by the key. However, the permutations constructed in these works do not have the form of an SPN round, and furthermore, the circuit complexity of the composed permutations is not of interest to them (their constructions have size and depth $\Omega(n^3)$).

Natural Proofs. The landscape of circuit lower bounds remains bleak, despite exciting recent results [Williams 2011]. Researchers however have been successful in explaining this lack of progress by pointing out several “barriers,” that is, establishing that certain proof techniques will not give new lower bounds [Baker et al. 1975; Razborov and Rudich 1997; Aaronson and Wigderson 2009].

Of particular interest to us is the Natural Proofs work by Razborov and Rudich [1997]. They make the following two observations. First, most lower-bound proofs that a certain function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ cannot be computed by circuits C (e.g., $C =$ circuits of size n^2) entail an algorithm that runs in time polynomial in $N := 2^n$ and can distinguish truth-tables of n -bit functions $g \in C$ from truth-tables of random functions (i.e., random strings of length N). (For example, the algorithm corresponding to the restriction-based proof that Parity is not in AC^0 , given $f : \{0, 1\}^n \rightarrow \{0, 1\}$, checks if there is one of the $2^{O(n)} = N^{O(1)}$ restrictions of the n variables that makes f constant.) Informally, any proof that entails such an algorithm is called “natural.”

The second observation is that, under standard hardness assumptions, no algorithm such as the preceding one exists when C is a sufficiently rich class. This follows from

the existence of PRF with security $2^{s^{\Omega(1)}}$, where s is the seed length (e.g., [Goldreich et al. 1986; Håstad et al. 1999; Naor and Reingold 2004; Haitner et al. 2010; Vadhan and Zheng 2012; Banerjee et al. 2012]) and by setting $s := n^c$ for a sufficiently large c .

The combination of the two observations is that no natural proof exists against circuits of size n^c , for some constant $c \geq 2$.

Moreover, the PRF constructions by Naor and Reingold [2004] and Banerjee et al. [2012] are implementable in TC^0 , pushing the preceding second observation “closer” to the frontier of known circuit lower bounds. For completeness, we also mention that these constructions achieve seed length $s = O(n^2)$ and can be shown to have hardness $2^{\Omega(n)}$ under certain conjectures related to elliptic curves (for [Naor and Reingold 2004]) and lattices (for [Banerjee et al. 2012]).

The Gap between Lower Bounds and PRF. However, the natural proofs barrier still has a significant gap with known lower bounds, due to the lack of sufficiently strong PRF. For example, there is no explanation as to why one cannot prove superlinear-size circuit lower bounds. For this one, would need a PRF $f_k : \{0, 1\}^n \rightarrow \{0, 1\}$ that is computable by linear-size circuits (hence in particular with $|k| = O(n)$) and with exponential hardness 2^n (so that, given n , if one had a distinguisher running in time $2^{O(n)}$, one could pick a PRF on inputs of length bn for a large enough constant b , to obtain a contradiction).

A work by Allender and Koucký [2010] brings to the forefront another setting where the natural proofs barrier does not apply: proving lower bounds on TC^0 circuits of size $n^{1+\epsilon}$ and depth d , for any $\epsilon > 0$ and large enough $d = d(\epsilon)$. (As previously mentioned, the PRFs in [Naor and Reingold 2004; Banerjee et al. 2012] require larger size.) This setting is especially interesting because Allender and Koucký [2010] show that, if one can prove such a lower bound for functions satisfying a certain self-reducibility property, this would imply a “full-fledged” lower bound for the same function against all TC^0 circuits of polynomial size. Moreover even if the first lower bound were natural, the latter would not be, thus circumventing the natural proofs barrier imposed by the TC^0 PRFs.

Another long-standing problem is that at exhibiting a candidate PRF in ACC^0 . Subsequent to the initial publication of this work [Miles and Viola 2012], Akavia et al. [2014] gave a candidate “weak PRF” computable in $\text{AC}^0[\oplus]$. (Weak PRF security is defined for a uniform set of queries, as opposed to standard PRF security, which is defined for adaptive, adversarially-chosen queries.)

Of course, circuit models such as the preceding ones are only some of the models in which the gap between candidate PRF and lower bounds is disturbing. Other such models include various types of Turing machines and small-space branching programs. For example, there is no explanation as to why the lower bounds for single-tape Turing machines stop at quadratic time (cf., [Kushilevitz and Nisan 1997, §12.2]).

Assuming the (exponential) security of some of our candidates, our work narrows this gap in three ways. First, Candidate 2 is computable by quasilinear-size Boolean circuits. Second, Candidate 3 is computable by TC^0 circuits of size $n^{1+\epsilon}$ and depth $d = d(\epsilon)$ for any $\epsilon > 0$. Third, for each fixed seed k , Candidate 5 is computable in time $O(n^2)$ by a single-tape Turing machine with $O(n^2)$ states. (Note that in the fixed-seed setting, the Turing machine’s states encode the key k , and that this setting is sufficient to apply the preceding natural proofs argument.)

Block Cipher Modes. For context, we note that common methods of extending fixed-input-length block ciphers to domains of arbitrary size, the so-called “modes of operation,” are not sufficient to construct PRF starting from secure block ciphers. This is

Fig. 1. One round of an SPN.

because once the input length of these extensions becomes larger than a certain constant (related to the input length of the underlying block cipher), the function can be distinguished from uniform in polynomial time. For concreteness, we focus here on the widely-used CBC-MAC mode (cf., [Bellare et al. 2000]), though similar attacks hold for all modes of which we are aware.

For a block cipher $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$, the function $\text{CBC-MAC}_f : \{0, 1\}^{n\ell} \rightarrow \{0, 1\}^\ell$ is computed as follows on input $x = (x_1, \dots, x_n) \in \{0, 1\}^{n\ell}$: first set $y_0 := 0$, then compute $y_i := f(x_i \oplus y_{i-1})$ for $i = 1, \dots, n$, and finally output y_n . The point now is that regardless of f 's security as a PRF, and even if independent keys are used for each of the n evaluations of f , CBC-MAC_f has hardness $\leq 2^{O(\ell)} \cdot n^{O(1)}$ as a PRF, which is $\text{poly}(n\ell)$ for $n = 2^{\Omega(\ell)}$.

The attack works as follows. First find two inputs $x \neq x'$ that induce the same value of $y_{n-1} \in \{0, 1\}^\ell$ in this computation. This takes time $2^{O(\ell)} \cdot n^{O(1)}$ by the pigeonhole principle. Then, for any w , we have $\text{CBC-MAC}_f(x_1, \dots, x_{n-1}, w) = \text{CBC-MAC}_f(x'_1, \dots, x'_{n-1}, w)$, while for a truly random function this holds only with probability $2^{-\ell}$.

Organization. In Section 2, we review the SPN structure. In Section 3, we construct our PRF candidates. We conclude and mention some future directions in Section 4. Appendix A contains some details regarding linear and differential cryptanalysis that are omitted from Section 2, and Appendix B presents an attack on low-degree PRF.

2. SUBSTITUTION-PERMUTATION NETWORKS

In this section, we review the necessary background on the SPN structure (refer to Figure 1). The notation introduced here will be used throughout.

An SPN $C_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is indexed by a key $k = (k_0, \dots, k_r) \in (\{0, 1\}^n)^{r+1}$ and is specified by the following three parameters and two functions.

- $r \in \mathbb{N}$, the number of rounds.
- $b \in \mathbb{N}$, the S -box input size.

- $m \in \mathbb{N}$, the number of S -box invocations per round.
- $S : \text{GF}(2^b) \rightarrow \text{GF}(2^b)$, the S -box.
- $M : (\text{GF}(2^b))^m \rightarrow (\text{GF}(2^b))^m$, the linear transformation.

The input/output size of C_k is given by $n := mb$. By way of illustration, AES uses parameters $b = 8$, $n = 128$, and $r \in \{10, 12, 14\}$. Our candidates use a variety of settings, ranging from $b \approx r \approx \log n$, to $b = n^{\Omega(1)}$ and $r = O(1)$, to $b = O(1)$ and $r = n$. (See Section 3.1 for details.) Throughout, we assume a fixed canonical mapping between $\{0, 1\}^b$ and $\text{GF}(2^b)$, and we use M to refer to both the linear transformation and its matrix representation in $(\text{GF}(2^b))^{m \times m}$.

C_k is computed over r rounds. The i th round ($1 \leq i \leq r$) is computed over three steps: (1) m parallel applications of S ; (2) application of M to the entire state; (3) XOR of the entire state with the round key k_i . Note that each round is identical except for step (3).¹

On input x , $C_k(x)$ gives $x \oplus k_0$ as input to the first round; the output of round i becomes the input to round $i + 1$ (for $1 \leq i < r$), and $C_k(x)$'s output is the output of the r th round.

Security against Linear and Differential Cryptanalysis. We now briefly review how the security of an SPN is evaluated against two general attacks on block ciphers: linear and differential cryptanalysis. Resistance to these attacks is typically seen as the main security feature of SPNs. Full details are deferred to Appendix A. Note that we consider here the basic versions of these attacks, and we leave to future work understanding the resistance of our candidates to more sophisticated attacks (such as those considered by Knudsen [1994]).

For both linear and differential cryptanalysis, a crucial property in the security proof is that the linear transformation M has maximal *branch number*. The branch number bounds the minimum total number of nonzero elements in any input/output pair $(x, M(x))$, where x is any nonzero vector. Note that the branch number can be at most $m + 1$, because x and $M(x)$ both have length m , and x may have only one nonzero element.

Definition 2.1. Let $M : \mathbb{F}^m \rightarrow \mathbb{F}^m$ be a linear transformation acting on vectors over a field $\mathbb{F}$. The *branch number* of M is

$$\text{Br}(M) = \min_{\alpha \in \mathbb{F}^m \setminus 0^m} (w(\alpha) + w(M(\alpha))),$$

where $w(\cdot)$ denotes the number of nonzero elements.

Linear cryptanalysis Matsui [1994] exploits the existence of linear correlations to attack a block cipher C_k . For a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and input/output parities $\Gamma_{\text{in}}, \Gamma_{\text{out}} \in \{0, 1\}^n$ with $\Gamma_{\text{out}} \neq 0^n$, define the *correlation* of f with respect to Γ_{in} and Γ_{out} as

$$\text{Cor}_{\Gamma_{\text{in}}, \Gamma_{\text{out}}}(f) := 2 \cdot \Pr_x[\langle \Gamma_{\text{in}}, x \rangle = \langle \Gamma_{\text{out}}, f(x) \rangle] - 1 \in [-1, 1].$$

For a block cipher C_k , the parameter of interest for linear cryptanalysis is

$$p_{\text{LC}}(C_k) := \max_{\Gamma_{\text{in}}, \Gamma_{\text{out}} \neq 0^n} (\mathbb{E}_k[\text{Cor}_{\Gamma_{\text{in}}, \Gamma_{\text{out}}}(C_k)^2]).$$

Specifically, the attack in Matsui [1994] requires an expected number of input/output pairs proportional to $1/p_{\text{LC}}(C_k)$.

Differential cryptanalysis [Biham and Shamir 1991] attacks a block cipher C_k by exploiting the relationship between the XOR of two inputs to C_k and the XOR of the

¹SPNs are sometimes defined more generally, e.g., by allowing the S -box to vary across rounds or by allowing a more complex interaction with k than XOR.

corresponding outputs. For a function $f_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$ parameterized by a key k , and input/output differences $\Delta_{\text{in}}, \Delta_{\text{out}} \in \{0, 1\}^n$ with $\Delta_{\text{out}} \neq 0^n$, define the *difference propagation probability* (DPP) of f_k with respect to Δ_{in} and Δ_{out} as

$$\text{DPP}_{\Delta_{\text{in}}, \Delta_{\text{out}}}(f_k) := \Pr_{x, k}[f_k(x) \oplus f_k(x \oplus \Delta_{\text{in}}) = \Delta_{\text{out}}].$$

(If f is not parameterized by a key, k is ignored in this definition). For a block cipher C_k , the parameter of interest for differential cryptanalysis is

$$p_{\text{DC}}(C_k) := \max_{\Delta_{\text{in}}, \Delta_{\text{out}} \neq 0^n} (\text{DPP}_{\Delta_{\text{in}}, \Delta_{\text{out}}}(C_k)).$$

Specifically, the attack in Biham and Shamir [1991] requires an expected number of input/output pairs proportional to $1/p_{\text{DC}}(C_k)$.

The following theorem [Kang et al. 2001] gives a bound on p_{LC} and p_{DC} for 2-round SPNs with maximal branch number.

THEOREM 2.2 ([KANG ET AL. 2001], THMS. 5 & 6). *Let $C_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an SPN with $r = 2$ rounds and S-box S . Let $q := \max_{\Gamma_{\text{in}}, \Gamma_{\text{out}} \neq 0^n} (\text{Cor}_{\Gamma_{\text{in}}, \Gamma_{\text{out}}}(S)^2)$ and $p := \max_{\Delta_{\text{in}}, \Delta_{\text{out}} \neq 0^n} (\text{DPP}_{\Delta_{\text{in}}, \Delta_{\text{out}}}(S))$. If $\text{Br}(M) = m + 1$, then $p_{\text{LC}}(C_k) \leq q^m$ and $p_{\text{DC}}(C_k) \leq p^m$.*

For typical S-boxes, such as the one used in AES, one can have $q = p = 2^{-b+2}$, and so the theorem guarantees security exponential in $n = mb$. (For completeness, we note that one cannot directly apply the preceding theorem to AES because it is a more complicated SPN.)

We extend this result to $r > 2$ rounds in the following theorem.

THEOREM 2.3. *Let $C_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an SPN with $r = 2\ell \geq 2$ rounds and S-box S . Let $q := \max_{\Gamma_{\text{in}}, \Gamma_{\text{out}} \neq 0^n} (\text{Cor}_{\Gamma_{\text{in}}, \Gamma_{\text{out}}}(S)^2)$ and $p := \max_{\Delta_{\text{in}}, \Delta_{\text{out}} \neq 0^n} (\text{DPP}_{\Delta_{\text{in}}, \Delta_{\text{out}}}(S))$. If $\text{Br}(M) = m + 1$, then $p_{\text{LC}}(C_k) \leq q^{\ell m} \cdot 2^{(\ell-1)n}$ and $p_{\text{DC}}(C_k) \leq p^{\ell m} \cdot 2^{(\ell-1)n}$.*

Intuitively, the S-box provides security q (resp. p) against linear (resp. differential) cryptanalysis, and this security multiplies across “active” S-boxes (instances of S that are evaluated with a nonzero input). The branch number $\text{Br}(M)$ guarantees that there exist $\geq m + 1$ such active S-boxes in any pair of consecutive rounds, hence the term $q^{\ell m} = q^{(r/2)m}$. We note that the factor $2^{(\ell-1)n}$ seems to be an artifact of our extension of Kang et al. [2001], and it is open to getting a tighter bound on p_{LC} and p_{DC} for $r > 2$ rounds ([Kang et al. 2001] only consider $r = 2$). Such an extension has been considered before, (e.g., [Keliher et al. 2001; Cho et al. 2004]), but their results only apply in the fixed-parameter setting because they require extensive computer calculation. We are not aware of any other “closed form” bound for $r > 2$.

Security against Degree-Exploiting Attacks. While resistance to linear and differential cryptanalysis is the main security feature of the SPN structure (and indeed, “the most important criterion in the design” of AES [Daemen and Rijmen 2002, p. 81]), considerations are usually also taken to prevent attacks that would exploit algebraic structure in the cipher. In our Candidates 2–5, we adopt essentially the same S-box that is used in AES. (In the following section, we discuss the differences between AES and our candidates.) This S-box is defined by $S(x) := x^{2^b-2}$ and was chosen to allow the computation to have high degree when considered as a multivariate polynomial over $\text{GF}(2)$. Specifically, the use of $x \mapsto x^{2^b-2}$ results in each of S ’s output bits having (near-maximum) degree $b - 1$. Using instead $x \mapsto x^3$ would not diminish resistance to linear and differential cryptanalysis, but it would result in degree of (only) 2 [Pieprzyk 1991; Nyberg 1993; Kopparty 2011].

We need the degree of each output bit of our candidates (as a multivariate $\text{GF}(2)$ -polynomial) to be $\geq \epsilon n$, for some constant ϵ , to resist attacks that exploit the degree of this polynomial. For completeness, we present such an attack, showing that a PRF that has degree $o(n)$ cannot have hardness 2^n . The proof of the following theorem is deferred to Appendix B.

THEOREM 2.4. *Let $F = \{f_k : \{0, 1\}^n \rightarrow \{0, 1\}\}_k$ be any set of functions such that, for each key k , the polynomial representation of f_k over $\text{GF}(2)$ has degree $o(n)$. Then there is an adversary that runs in time $2^{O(n)}$ and distinguishes a random $f_k \in F$ from a random function with advantage $1 - 2^{-2^{\Omega(n)}}$.*

We note that more “fine-grained” versions of this theorem can be proved, for example, showing that degree- $O(1)$ PRF can be distinguished in time $n^{O(1)}$. As we are primarily concerned with exponential hardness, we omit the details.

The only nonlinear operation in the entire cipher is the b -bit S-box. Candidate 1 uses a random S-box, and thus each of its output bits has degree say $0.99b$ with high probability. Candidates 2–5 use the inversion S-box in which each output bit has degree $b - 1$. Note that any SPN has degree $\leq (\text{degree of S-box})^{\# \text{rounds}}$, and hence we ensure that

$$b^r \geq n$$

in each of our candidates. (The distinction between $(b - 1)^r \geq \epsilon n$ and $b^r \geq n$ is unimportant, as in our candidates, we can always increase r by a constant factor, except in Candidate 4 where we have $b = n$ and $r = 1$.) We do not know if $b^r \geq n$ is sufficient to guarantee degree $\Omega(n)$, and it is an interesting research direction to understand what restrictions (if any) on the SPN parameters ensure that the function has high degree.

Finally, although a block cipher’s security is often measured against *key-recovery* attacks, we share many researchers’ viewpoint that *distinguishing* attacks are the correct model. We also note that there is often an intimate connection between the two types, as many key recovery techniques—including linear and differential cryptanalysis—construct a distinguishing algorithm which is then used to select the correct round keys from a set of potential keys.

3. OUR CANDIDATES

In this section, we construct our new PRF candidates. Candidates 1, 2, and 5 output n bits, while Candidates 3 and 4 output 1 bit. We use $\mathcal{F}_i$ to refer to the function computing Candidate i .

Differences between Our Candidates and AES. The most obvious difference between our candidates and AES is that AES uses a fixed set of parameters: $b = 8$, $m = 16$, $r \in \{10, 12, 14\}$. Besides this, there are a few other differences which we note here.

- (1) In AES, the S-box is computed by first mapping $x \mapsto x^{2^b-2}$ in $\text{GF}(2^b)$ and then by applying a fixed $\text{GF}(2)^b$ -affine transformation. In Candidates 2–5, the S-box omits the affine transformation and simply computes $x \mapsto x^{2^b-2}$. Adding an affine transformation would not affect the (asymptotic) circuit size of our candidates, and to our knowledge, there are no known attacks against the AES variant that uses this “reduced” S-box. In Candidate 1, the S-box is instead a uniformly random b -bit function (chosen as part of the seed).
- (2) In Candidates 1–3, the linear transformation consists of multiplication with a maximal-branch-number matrix $M \in \text{GF}(2^b)^{m \times m}$ (cf. Definition 2.1). In AES, the transformation first permutes the bundles in a certain way (see Section 3.6 for details), and then applies $m/4$ parallel copies of a 4×4 maximal-branch-number

matrix; this same transformation is also used by Candidate 5. Candidate 4 consists of just one round and has no linear transformation.

- (3) In each of our candidates, the $(r + 1)$ n -bit round keys are chosen independently and uniformly at random. AES, and most other popular constructions, employ a so-called “key schedule” that generates the round keys from a key of size $\ll n(r + 1)$.
- (4) AES, being a block cipher, computes a *permutation* function, which is necessary for unambiguous decryption. We do not impose this constraint on our PRF candidates, although Candidates 2 and 5 do compute permutation functions.

We note that the use of maximal-branch-number linear transformation and independent round keys in our constructions (items 2 and 3) should only improve security as compared to AES.

3.1. Overview

Candidate 1. Our first candidate $\mathcal{F}_1$ is an r -round SPN with an S-box that is chosen uniformly at random (i.e., specified as part of $\mathcal{F}_1$'s key) from the set of all functions mapping $\text{GF}(2^b)$ to itself. (Analyzing this candidate when S is a random *permutation* is a natural research direction which we do not address here.) The only restriction we make on $\mathcal{F}_1$'s linear transformation M is that it is invertible and has all entries $\neq 0$; we observe that this holds for any M with maximal branch number. We show that any adversary A has small advantage in distinguishing $\mathcal{F}_1$ from a random function F .

THEOREM 3.1. *If A makes at most q total queries to its oracle, then*

$$\left| \Pr_F[A^F = 1] - \Pr_{\mathcal{F}_1}[A^{\mathcal{F}_1} = 1] \right| < O(r^2 m^3 q^3) \cdot 2^{-b}.$$

The bound achieved here is similar to that of Luby and Rackoff [1988] in the sense that it is exponentially small in the size of the random function, with a polynomial loss in the number of queries. (The fact that the security bound degrades with the number of rounds, contrary to what one might expect, seems to be an artifact of the proof, and in fact the proof only requires $r \geq 2$.) The proof of this theorem is very similar to that of Naor and Reingold [1999, Thm. 3.2] and proceeds by bounding the collision probability between any two inputs to S in the final round. However, we face an additional hurdle, namely, that the inputs to the random function S in the final round depend on outputs of S in previous rounds.

By setting $b = \omega(\log n)$, we get an inefficient PRF (with security $n^{\omega(1)}$). We also note that by setting $b = c \log n$ for some sufficiently large constant c , $\mathcal{F}_1$ is computable in time $n^{O(c)}$ and has security $n^{c'}$ for some $c' = \Omega(c)$.

Finally, note that Theorem 3.1 implies corresponding bounds on $p_{\text{LC}}(\mathcal{F}_1)$ and $p_{\text{DC}}(\mathcal{F}_1)$.

Candidate 2. In this candidate, we set $b = \Theta(\log n)$, and we use the AES S-box on b bits (recall that it maps $x \mapsto x^{2^b-2}$). We use a linear transformation M with maximal branch number, and M is constructed from an error-correcting code in a similar manner to the linear transformation in AES. (AES's linear transformation does not have maximal branch number, however, a choice that was made to reduce computation time.) We set the number of rounds $r = \Theta(\log n)$ (observe that $b' \geq n$).

We prove that Candidate 2 is computable by Boolean circuits of quasilinear-size $\tilde{O}(n) := n \cdot \log^{O(1)} n$. To show this, note that since r is logarithmic, it is enough to show how to compute each round with these resources. Moreover, since b is logarithmic, computing the S-boxes comes at little cost.

Our main technical contribution with this candidate is to show how to efficiently compute the linear transformation M ; specifically, we show that it can be computed

with size $\tilde{O}(n)$, for a total circuit size of $r \cdot (b^{O(1)} + \tilde{O}(n)) = \tilde{O}(n)$. A common method for constructing maximal-branch-number linear transformations is to use the generator matrix G of an $m \rightarrow 2m$ maximum distance separable (MDS) code; specifically, if $G^T = [I | A]$, then $M := A$ has maximal branch number. Our method for computing M efficiently has two parts. First, we use a result by Roth and Seroussi [1985] that if G generates a Reed-Solomon code (which is well-known to be MDS), then M forms a $t \times t$ Cauchy matrix (a type of matrix specified by $O(t)$ elements). We then use a result by Gerasoulis [1988] to compute the product of a vector (consisting of bundles of the state) and a Cauchy matrix in quasilinear time; this requires a simple adaptation of the algorithm in Gerasoulis [1988] to fields of characteristic 2.

By combining Theorem 2.3 with a theorem of Nyberg [1993], we show that this candidate has exponential security against linear and differential cryptanalysis.

THEOREM 3.2. (1) $p_{LC}(\mathcal{F}_2) \leq 2^{-\Omega(n)}$. (2) $p_{DC}(\mathcal{F}_2) \leq 2^{-\Omega(n)}$.

We do not know how to get a candidate computable by circuits of size $O(n)$.

Candidate 3. In the previous candidate, the components S and M remain essentially unchanged from AES. In Candidate 3, we also keep S the same (aside from the increase in input/output size), but we modify the linear transformation M .

Our observation is that the rationale for using a linear transformation with maximal branch number is just that it allows one to lower bound the number $\mathcal{A}$ of so-called “active” S-boxes, which can be defined as follows. Let C be an SPN which uses the identity permutation for S and which has $k_i := 0$ for $0 \leq i \leq r$. Let $w_b : (\{0, 1\}^b)^m \rightarrow \mathbb{N}$ be the function that counts the number of nonzero b -bit bundles in its input. Then,

$$\mathcal{A} := \min_{0^n \neq x \in \{0, 1\}^n} \sum_{i=1}^r w_b(\text{state of } C(x) \text{ at the beginning of round } i).$$

This number $\mathcal{A}$ is crucial in evaluating the security of SPNs against linear and differential cryptanalysis (cf. [Kang et al. 2001; Daemen and Rijmen 2002]). With a simple modification to M , we get that a constant fraction of the S-boxes in each round are active. Specifically, we use the full generator matrix of an error correcting code with minimum distance $\Omega(n)$, which comes at the expense of expanding the state from n bits to $O(n)$ bits at each round. To counteract the fact that such codes may have some output positions fixed to constant values (leading to a simple distinguishing attack), the computation of Candidate 3 concludes by taking the inner product of the state with a uniform $O(n)$ -bit vector that is given as part of the seed. Candidate 3 therefore outputs a single bit.

We take $b = n^\epsilon$ and $r = O(1/\epsilon)$ for arbitrarily small $\epsilon > 0$, and so each round is computable in size

$$\frac{n}{b} \cdot \text{poly}(b) = n^{1+O(\epsilon)},$$

and the whole circuit also in size $n^{1+O(\epsilon)}$.

We further show that Candidate 3 is computable even by TC^0 circuits of size $n^{1+\epsilon}$ for any $\epsilon > 0$, with depth depending on ϵ (cf. paragraph “The gap between lower bounds and PRF” in Section 1). The main technical difficulty in implementing this candidate with the required resources is that the S-box requires computing *inversion* in a field of size 2^b (recall $b = n^{\Omega(1)}$). To implement this in TC^0 , we note (cf. [Healy and Viola 2006])

that inverting the field element $\alpha(x)$ can be accomplished as

$$\alpha(x)^{2^b-2} = \alpha(x)^{\sum_{i=1}^{b-1} 2^i} = \prod_{i=1}^{b-1} \alpha(x)^{2^i} = \prod_{i=1}^{b-1} \alpha(x^{2^i}),$$

where the last equality follows from the fact that we are working in characteristic 2. By hard-wiring the $\leq b$ powers $x, x^2, \dots, x^{2^{b-1}}$ of x in the circuit and using the fact that the iterated product of $\text{poly}(n)$ field elements is computable by $\text{poly}(n)$ -size TC^0 circuits (see e.g. [Hesse et al. 2002, Corollary 6.5] and cf. [Healy and Viola 2006]), we obtain a TC^0 circuit.

Because Candidate 3 deviates somewhat from the SPN structure, we cannot use Theorem 2.2, and indeed it is not clear how to define differential cryptanalysis for functions which output only one bit. However, we are able to leverage a technique from differential cryptanalysis to prove that Candidate 3 is almost 3-wise independent. We are unable to determine if this candidate is 4-wise independent.

Definition 3.3. A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ parameterized by a key k is (d, ϵ) -wise independent if for any distinct $x_1, \dots, x_d \in \{0, 1\}^n$, the distribution $(f(x_1), \dots, f(x_d))$ induced by a uniform choice of k is ϵ -close to U_d in statistical distance.

THEOREM 3.4. $\mathcal{F}_3$ is $(3, 2^{-\Omega(n)})$ -wise independent.

Finally, we mention that implicit in an assumption that Candidate 3 is indeed hard is the assumption that field inversion cannot be computed by unbounded fan-in constant depth circuits with parity gates $\text{AC}^0[\oplus]$. For otherwise, it can be shown that the whole candidate would be in that class, in contradiction with an algorithm in Razborov and Rudich [1997, §3.2.1] which distinguishes truth tables of $\text{AC}^0[\oplus]$ functions from random ones in quasipolynomial time. (M can be seen to be a linear operation over $\text{GF}(2)$, hence it can be computed easily with parity gates.) The question of whether field inversion is in $\text{AC}^0[\oplus]$ was raised by Healy and Viola [2006]. Their work, and later Kopparty's [2011], does show that several functions related to field inversion are not in $\text{AC}^0[\oplus]$.

Candidate 4. In this candidate, we use the extreme setting of parameters $b = n$ and $r = 1$. In other words, Candidate 4 consists of one round, and this round contains only a single S-box (and in particular no linear transformation). This construction can be seen as a concrete instantiation of the Even-Mansour block cipher [Even and Mansour 1997], using the AES S-box in place of the random permutation oracle. While this setting does indeed preserve resistance to linear and differential cryptanalysis, we exhibit a simple attack, inspired by Jakobsen and Knudsen [2001], in which we exploit the algebraic structure to recover the key with just four queries.

We then put forth a related candidate $\mathcal{F}'_4$ where we only output the Goldreich-Levin bit [Goldreich and Levin 1989]: $\mathcal{F}'_4(x) := \langle (x + k_0)^{2^b-2}, k_1 \rangle$. We prove that this candidate is a d -wise small-bias generator with error $d/2^n$ (cf. [Naor and Naor 1993; Alon et al. 1992]), that is, that it fools all parity tests that look at $\leq 2^{0.9n}$ outputs.

THEOREM 3.5. For any choice of $d \leq 2^n$, $\mathcal{F}'_4$ is a d -wise small-bias generator with error $d/2^n$. That is, for any distinct $a_1, \dots, a_d \in \{0, 1\}^n$,

$$\left| \Pr_{k_0, k_1} \left[\sum_{i=1}^d \mathcal{F}'_4(a_i) = 0 \right] - \frac{1}{2} \right| < \frac{d}{2^n}.$$

Using a result of Braverman [2009] (cf. [Bazzi 2009; Razborov 2009]), we obtain that this candidate also fools small-depth AC^0 circuits of any size $w = 2^{n^{o(1)}}$ (that look at only w fixed output bits of the candidate).

This candidate is computable by circuits of quasilinear size $O(n \log^2 n \log \log n)$ using the inversion algorithm in Gao et al. [2000]. Using the same ideas for Candidate 3, this candidate is also computable by poly-size TC^0 circuits.

Candidate 5. Our final candidate is a straightforward generalization of AES and may be folklore. We set $b = 8$ as in AES and we again use AES's S-box. We also use the same linear transformation as in AES (which is slightly different from that of Candidate 2, cf. §3.6), except for the necessary increase in the input/output size. We set the number of rounds $r = n$, and thus the size of the seed is $|k| = n(n + 1)$.

Candidate 5 is computable by size- $O(n^2)$, depth- $O(n)$ Boolean circuits. For each fixed seed k , Candidate 5 is also computable in time $O(n^2)$ by a single-tape Turing machine with $O(n^2)$ states.

We do not know how to get a candidate computable in time $O(n)$ on a 2-tape Turing machine.

3.2. Candidate 1

For our first candidate, we analyze the pseudorandomness of the SPN structure when the S-box is a uniformly random function. The results of this section are of a similar flavor and use similar techniques as those of Luby and Rackoff [1988] and the following work by Naor and Reingold [1999]. One notable difference is that we study SPNs as pseudorandom *functions*, and in particular, we do not allow inverse queries to the SPN. (Indeed, if the S-box is not a permutation, then the SPN may not be either, in which case inverse queries are not well-defined.) Adapting this proof to handle bidirectional queries is a natural research direction which is not addressed here.

Our analysis in this section holds for SPNs in which the matrix M defining the linear transformation is invertible and has all entries $\neq 0$. We observe that this includes all matrices with maximal branch number.

CLAIM 3.6. *Let $M \in (\text{GF}(2^b))^{m \times m}$ be any matrix with maximal branch number $m + 1$. Then, all entries of M are nonzero and M is invertible.*

PROOF. If $M_{i,j} = 0$ for some $i, j \leq m$, let $x \in (\text{GF}(2^b))^m$ be the vector such that $x_j = 1$ and $x_{j'} = 0$ for $j' \neq j$. Then $(Mx)_i = 0$, and so $\text{Br}(M) \leq w(x) + w(Mx) \leq m$.

If M is not invertible, then $\exists x \neq 0^m : Mx = 0$, but this implies $\text{Br}(M) \leq w(x) \leq m$. $\square$

For the remainder of this section, fix any invertible $M \in (\text{GF}(2^b))^{m \times m}$ such that all entries are nonzero. For any function $S : \text{GF}(2^b) \rightarrow \text{GF}(2^b)$ and any set of round keys $(k_0, \dots, k_{r-1}) \in (\{0, 1\}^n)^r$, let $\mathcal{F}_1 = \mathcal{F}_1(S, k_0, \dots, k_{r-1})$ be the r -round SPN on $n := mb$ bits defined by these components, where the final round consists only of S-boxes (i.e., the final round omits the linear transformation and the key addition).

Let $A : \mathbb{N} \times \mathbb{N} \rightarrow \{0, 1\}$ denote an adversary with oracle access to a function mapping $(\text{GF}(2^b))^m$ to itself; A 's input is simply $(1^m, 1^b)$, which we omit from now on. We show that A has small advantage in distinguishing between the case when its oracle is a uniformly random function F and when its oracle is $\mathcal{F}_1$ for a uniform choice of $(S, k_0, \dots, k_{r-1})$.

THEOREM 3.1. *If A makes at most q total queries to its oracle, then*

$$\left| \Pr_F[A^F = 1] - \Pr_{\mathcal{F}_1}[A^{\mathcal{F}_1} = 1] \right| < O(r^2 m^3 q^3) \cdot 2^{-b}.$$

3.2.1. Proof Overview. The proof proceeds in two stages. In the first stage, we consider any set of distinct queries $x_1, \dots, x_q$, and we show that there is a low-probability event

BAD over the choice of $(S, k_0, \dots, k_{r-1})$ such that, conditioned on $\neg \text{BAD}$, $\{\mathcal{F}_1(x_i)\}_{i \leq q}$ is uniformly distributed. Essentially, BAD is the event that any two SPN queries induce the same input to some S-box in the final round.

In the second stage, we consider the distribution over transcripts of A 's interaction with its oracle; we use the results of the first stage in a probability argument to show that the transcripts are distributed nearly identically in either setting, and thus that A 's distinguishing advantage is small. This framework has been used in a number of other works [Naor and Reingold 1999; Ramzan and Reyzin 2000; Gentry and Ramzan 2004].

The first stage actually shows that $\mathcal{F}_1$ is almost q -wise independent, or alternatively, that it is pseudorandom against adversaries that make $\leq q$ nonadaptive queries. The technique used in the second stage is a rather generic way of extending the proof to adaptive queries; however, we note that it crucially relies on the existence of the event BAD, and indeed it is not the case that any almost q -wise independent function is pseudorandom against adversaries making q adaptive queries.² A different method (that does not give a useful bound in our setting) for obtaining adaptive security from nonadaptive security is given by Hoory et al. [2005, Prop. 3].

We will analyze the first $(r - 2)$ rounds of $\mathcal{F}_1$ in a different way from the final two rounds, and to this end, we define the following two functions. Let $\rho = \rho(S, k_0, \dots, k_{r-3})$ compute everything in $\mathcal{F}_1$ before the XOR with k_{r-2} , and let $\rho' = \rho'(S, k_{r-2}, k_{r-1})$ compute the remainder of $\mathcal{F}_1$. So, $\mathcal{F}_1(x) = \rho'(\rho(x))$. As handling ρ' will be the more involved part of the analysis, we note that it can be written as

$$\rho'(x) := S^*(M \cdot S^*(x + k_{r-2}) + k_{r-1}).$$

where for any $x = x^{(1)} \dots x^{(m)}$, we define $S^*(x) := (S(x^{(1)}), \dots, S(x^{(m)}))$.

3.2.2. Stage 1. Fix distinct $x_1, \dots, x_q \in (\text{GF}(2^b))^m$. We view $(S, k_0, \dots, k_{r-1})$ as being chosen as follows.

- (1) Uniformly choose $k_0, \dots, k_{r-3}$.
- (2) Run the computation of $\rho(x_i)$ for all $i \leq q$, and each time the S-box is evaluated on a previously-unseen input, choose the output uniformly at random.
- (3) Uniformly choose k_{r-2} .
- (4) Uniformly choose the output of S on each block of each $(\rho(x_i) + k_{r-2})$ whose output is not already determined.
- (5) Uniformly choose k_{r-1} .
- (6) Uniformly choose the output of S on all remaining inputs.

It is clear that, for any $x_1, \dots, x_q$, this distribution is uniform. Our analysis will use the state of the SPN's computation immediately before the final round of S -boxes, and we denote this state for the i th query by

$$z_i := M \cdot S^*(\rho(x_i) + k_{r-2}) + k_{r-1}.$$

We now define the event BAD. Informally, BAD holds if, after step (5), any of the S -inputs that need to be evaluated (i.e., the blocks of the z_i) collide either with each other or with one of the inputs selected in steps (2) and (4). To reduce notation, we use the following definition.

Definition 3.7. Let $x, y \in (\text{GF}(2^b))^m$, and denote $x = x^{(1)} \dots x^{(m)}$ and $y = y^{(1)} \dots y^{(m)}$. Then, we say that x and y *collide* if $\exists \ell, \ell' : x^{(\ell)} = y^{(\ell')}$. Further, for any $H \subseteq \text{GF}(2^b)$, we

²This can be seen, for example, by considering the distribution over functions $f : [N] \rightarrow [N]$ in which each output is selected uniformly and independently with the restriction that $f(f(0)) = 0$. This is almost pairwise-independent, but trivially distinguishable with two adaptive queries.

say that x and H *collide* if $\exists \ell \leq m, t \in H : x^{(\ell)} = t$. Finally, we say that x *self-collides* if $\exists$ distinct $\ell, \ell' : x^{(\ell)} = x^{(\ell')}$.

Let $H \subseteq \text{GF}(2^b)$ be the set of at most $qm(r-2)$ S -inputs whose output is determined after step (2) in the preceding process for choosing $(S, k_0, \dots, k_{r-1})$. Then, $\text{BAD} = \text{BAD}(x_1, \dots, x_q)$ is the set of all $(S, k_0, \dots, k_{r-1})$ such that at least one of the following holds.

- (1) $\exists h, h' \in H : S(h) = S(h')$.
- (2) $\exists i < q : z_i$ and H collide.
- (3) $\exists i, i' \leq q : z_i$ and $(\rho(x_i) + k_{r-2})$ collide.
- (4) $\exists i \leq q : z_i$ self-collides.
- (5) $\exists$ distinct $i, i' \leq q : z_i$ and $z_{i'}$ collide.

It is crucial for us that determining whether BAD holds can be checked after step (5) in choosing $(S, k_0, \dots, k_{r-1})$.

We now prove two lemmas showing that BAD occurs with low probability and that the query answers are uniformly distributed when conditioned on $\neg \text{BAD}$. In the remainder of this section, we will simply use BAD to mean $(S, k_0, \dots, k_{r-1}) \in \text{BAD}$.

LEMMA 3.8. $\Pr_{S, k_0, \dots, k_{r-1}} [\text{BAD}] < O(r^2 m^3 q^3) \cdot 2^{-b}$.

PROOF. We start by bounding the probability of items (a)–(d) individually.

First, we have $\Pr_{S, k_0, \dots, k_{r-3}} [(a)] < (qm(r-2))^2 \cdot 2^{-b}$ by a union bound over all pairs of S -box instances in the first $r-2$ rounds.

We analyze (b)–(e) starting after step (2), so for these, let $k_0, \dots, k_{r-3}, H$, and $S(H)$ be fixed arbitrarily.

Fix any k_{r-2} and the outputs of S on the blocks of $(\rho(x_i) + k_{r-2})$ for all i , which fixes $\tilde{z}_i := M \cdot S^*(\rho(x_i) + k_{r-2})$; note that $z_i = \tilde{z}_i + k_{r-1}$. Then, $\Pr_{k_{r-1}} [(b)] \leq qm \cdot qm(r-2) \cdot 2^{-b}$ by a union bound over each block of each $\tilde{z}_i$ and each element of H .

By the same argument as for (b), $\Pr_{k_{r-1}} [(c)] \leq (qm)^2 \cdot 2^{-b}$, where the union bound is now over each block of each z_i and each block of each $(\rho(x_i) + k_{r-2})$.

Let the $\tilde{z}_i$ be defined as before, then for each i , $\Pr_{k_{r-1}} [z_i = \tilde{z}_i + k_{r-1} \text{ self collides}] < m^2 \cdot 2^{-b}$ by a union bound over pairs of blocks. So, $\Pr_{k_{r-1}} [(d)] < qm^2 \cdot 2^{-b}$.

We will now bound $\Pr[(e) | \neg(a)]$. Note that $\neg(a)$ implies that each component of ρ is injective on H , and thus that each $\rho(x_i)$ is distinct (this is where we use M 's invertibility).

Fix any distinct $i, i' \leq q$ and any $\ell, \ell' \leq m$. We will show that $\Pr[z_i^{(\ell)} = z_{i'}^{(\ell')} | \neg(a)] < O(rmq) \cdot 2^{-b}$, and then a union bound over i, i', ℓ, ℓ' gives $\Pr[(e) | \neg(a)] < O(rm^3 q^3) \cdot 2^{-b}$. (We remark that the nontrivial case is when $\ell = \ell'$, i.e., when comparing the same final-round S -box for distinct $x_i, x_{i'}$, because in this case, the same block of k_{r-1} affects both S -inputs. If $\ell \neq \ell'$, then one can proceed similarly to (a)–(d), but the following works for either case.)

From the definition of z_i , we have $z_i^{(\ell)} = z_{i'}^{(\ell')}$ if and only if

$$h_{r-1}^{(\ell)} + \sum_{s=1}^m M_{\ell,s} \cdot S(\rho(x_i)^{(s)} + k_{r-2}^{(s)}) = h_{r-1}^{(\ell')} + \sum_{s=1}^m M_{\ell',s} \cdot S(\rho(x_{i'})^{(s)} + k_{r-2}^{(s)}). \quad (1)$$

Let t be such that $\rho(x_i)^{(t)} \neq \rho(x_{i'})^{(t)}$, which must exist because $\rho(x_i) \neq \rho(x_{i'})$. Arbitrarily fix $k_{r-2}^{(s)}$ for all $s \neq t$, the outputs of S on the input set $I := \{(\rho(x_i) + k_{r-2})^{(s)}, (\rho(x_{i'}) + k_{r-2})^{(s)}\}_{s \neq t}$, and $k_{r-1}^{(\ell)}$ and $k_{r-1}^{(\ell')}$. This fixes an $\alpha \in \text{GF}(2^b)$ such that Eq. (1) holds

if and only if

$$M_{\ell,t} \cdot S(\rho(x_i)^{(t)} + k_{r-2}^{(t)}) + M_{\ell',t} \cdot S(\rho(x_{i'})^{(t)} + k_{r-2}^{(t)}) = \alpha. \quad (2)$$

If neither $\rho(x_i)^{(t)} + k_{r-2}^{(t)}$ nor $\rho(x_{i'})^{(t)} + k_{r-2}^{(t)}$ are in $(I \cup H)$, then (2) holds with probability 2^{-b} over the choice of S on these two inputs (this is where we use the fact that all entries of M are nonzero). Further, one or both of these two inputs fall inside $(I \cup H)$ with probability $\leq 2 \cdot (qm(r-2) + 2(m-1)) \cdot 2^{-b}$ over the choice of $k_{r-2}^{(t)}$, by a union bound over the elements of $(I \cup H)$. Thus, $\Pr[z_i^{(\ell)} = z_{i'}^{(\ell')} \mid \neg(a)] = O(rmq) \cdot 2^{-b}$ as promised.

Finally, we have $\Pr[(a) \vee \dots \vee (e)] \leq \Pr[(a)] + \dots + \Pr[(d)] + \Pr[(e) \mid \neg(a)] < O(r^2 m^3 q^3) \cdot 2^{-b}$. $\square$

In the following lemma, we only directly use that $\neg\text{BAD}$ implies $\neg((b) \vee (c) \vee (d) \vee (e))$. However, note that the event (a) was used in the bound on $\Pr[(e)]$ in the preceding lemma.

LEMMA 3.9. *For any distinct $x_1, \dots, x_q$ and any $y_1, \dots, y_q$:*

$$\Pr_{S, k_0, \dots, k_{r-1}} [\forall i \leq q : \mathcal{F}_1(x_i) = y_i \mid \neg\text{BAD}] = 2^{-qmb}.$$

PROOF. After running steps 1 through 5 in the process of choosing $(S, k_0, \dots, k_{r-1})$, if we condition on $\neg\text{BAD}$, then the qm elements of the set $\{z_i^{(\ell)}\}_{i,\ell}$ are distinct and are not used as inputs to S in steps 2 or 4. Thus, each element has a 2^{-b} probability (independent from the other elements) of being mapped by S to the corresponding output (i.e., a block of a y_i), and the lemma follows. $\square$

3.2.3. Stage 2. We now show that even adversaries that make adaptive queries have small distinguishing advantage, i.e., we prove Theorem 3.1. Without loss of generality, we make the standard assumption that the adversary A is deterministic, computationally unbounded, and never queries an oracle twice with the same input.

To prove Theorem 3.1, we extend the results of the previous section by considering the distribution over *transcripts* of A 's interaction with its oracles. A transcript is the sequence $\sigma = (y_1, \dots, y_q)$ of query answers that A received from its oracle; note that the corresponding queries are uniquely determined by σ because A is deterministic. We use T_F to denote the transcript of A^F , and we use $A(\sigma)$ to denote A 's output after seeing transcript σ . (So note for instance that $\Pr_F[A^F = 1]$ and $\Pr_F[A(T_F) = 1]$ are semantically equivalent.)

We now prove Theorem 3.1 with a probability argument similar to that of Naor and Reingold [1999, Thm. 3.2].

PROOF (THEOREM 3.1). Let $\Gamma \subseteq (\{0, 1\}^n)^q$ be the set of transcripts such that $A(\sigma) = 1 \Leftrightarrow \sigma \in \Gamma$. Then,

$$\begin{aligned} & \left| \Pr_F[A^F = 1] - \Pr_{\mathcal{F}_1}[A^{\mathcal{F}_1} = 1] \right| \\ &= \left| \sum_{\sigma \in \Gamma} \left(\Pr_F[T_F = \sigma] - \Pr_{\mathcal{F}_1}[T_{\mathcal{F}_1} = \sigma] \right) \right| \\ &\leq \left| \sum_{\sigma \in \Gamma} \Pr_{\mathcal{F}_1}[\text{BAD}] \cdot \left(\Pr_F[T_F = \sigma] - \Pr_{\mathcal{F}_1}[T_{\mathcal{F}_1} = \sigma \mid \text{BAD}] \right) \right| \end{aligned} \quad (3)$$

$$+ \left| \sum_{\sigma \in \Gamma} \Pr_{\mathcal{F}_1}[\neg\text{BAD}] \cdot \left(\Pr_F[T_F = \sigma] - \Pr_{\mathcal{F}_1}[T_{\mathcal{F}_1} = \sigma \mid \neg\text{BAD}] \right) \right|. \quad (4)$$

Lemma 3.9 implies that (4) = 0, because $\Pr_F[T_F = \sigma \mid \neg \text{BAD}] = 2^{-qmb}$ for any transcript σ . We rewrite (3) as

$$\left| \sum_{\sigma \in \Gamma} \left(\Pr_{\mathcal{F}_1}[\text{BAD}] \cdot \Pr_F[T_F = \sigma] \right) - \sum_{\sigma \in \Gamma} \left(\Pr_{\mathcal{F}_1}[\text{BAD}] \cdot \Pr_{\mathcal{F}_1}[T_{\mathcal{F}_1} = \sigma \mid \text{BAD}] \right) \right|.$$

Each of the two summations is bounded by $\alpha := \max_{\sigma \in \Gamma} (\Pr_{\mathcal{F}_1}[\text{BAD}])$, since each is a convex combination of numbers that are bounded by α . Thus, the absolute value of their difference is bounded by α as well, and $\alpha < O(r^2 m^3 q^3) \cdot 2^{-b}$ by Lemma 3.8. $\square$

3.3. Candidate 2

Our next candidate PRF $\mathcal{F}_2 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is parameterized by a key k of length $O(n \log n)$ and is computable by circuits of size $\leq n \log^{O(1)} n$. We will show that it has security $2^{-\Omega(n)}$ against linear and differential cryptanalysis via Theorem 2.3. The SPN defining $\mathcal{F}_2$ closely follows AES.

Definition of $\mathcal{F}_2$. $\mathcal{F}_2$ is an SPN as defined in Section 2; our parameter choices are as follows. For any $b \in \mathbb{N}$, let $m := 2^{b-1}$, $r := \lceil b/10 \rceil$ and $n := mb$.

For the S-box, we use essentially the same function used in AES. Namely, $S : \text{GF}(2^b) \rightarrow \text{GF}(2^b)$ is defined by $S(x) := x^{2^{b-2}}$. Note that $x \mapsto x^{2^{b-2}}$ is simply inversion in $\text{GF}(2^b)$ with $0^{-1} := 0$. The bounds on p_{LC} and p_{DC} from Theorems 2.2 and 2.3 are stated in terms of bounds on the correlation and the DPP, respectively, of the S-box. The results of Nyberg [1993] and the references therein establish these bounds, stated in the following theorem.

THEOREM 3.10. *Let $S : \text{GF}(2^b) \rightarrow \text{GF}(2^b)$ be defined by $S(x) := x^{2^{b-2}}$. Then,*

$$(1) \max_{\Gamma_{\text{in}}, \Gamma_{\text{out}} \neq 0} (\text{Cor}_{\Gamma_{\text{in}}, \Gamma_{\text{out}}}(S)^2) \leq 2^{b-2}.^3 \quad (2) \max_{\Delta_{\text{in}}, \Delta_{\text{out}} \neq 0} (\text{DPP}_{\Delta_{\text{in}}, \Delta_{\text{out}}}(S)) \leq 2^{b-2}.$$

For the linear transformation $M : \text{GF}(2^b)^m \rightarrow \text{GF}(2^b)^m$, the crucial property is that it has maximal branch number $\text{Br}(M) = m + 1$. Let G be the $2m \times m$ generator matrix of a Reed-Solomon code over $\text{GF}(2^b)$. (Note that $2^b \geq 2m$ is sufficient to guarantee the existence of such a code.) Take G to be in reduced echelon form, that is, take $G^T = [I \mid M]$, where I is the $m \times m$ identity matrix. Then, because G generates a maximum-distance-separable (MDS) code, it can be verified [Daemen 1995, Section 7.2] that the operation defined by left multiplication with M has branch number $m + 1$. This use of MDS codes to create maximal-branch-number transformations is widespread and dates at least to Daemen [1995].

Security of $\mathcal{F}_2$. The security of $\mathcal{F}_2$ is given by the following theorem (restated).

THEOREM 3.2. (1) $p_{\text{LC}}(\mathcal{F}_2) \leq 2^{-\Omega(n)}$. (2) $p_{\text{DC}}(\mathcal{F}_2) \leq 2^{-\Omega(n)}$.

Given the choices of b , r , and m , Theorem 3.2 follows immediately from Theorem 3.10 and Theorem 2.3, restated next. (We defer the proof of Theorem 2.3 to Appendix A, as it requires a more extensive technical analysis of the SPN structure.)

THEOREM 2.3. *Let $C_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an SPN with $r = 2\ell \geq 2$ rounds and S-box S . Let $q := \max_{\Gamma_{\text{in}}, \Gamma_{\text{out}} \neq 0^n} (\text{Cor}_{\Gamma_{\text{in}}, \Gamma_{\text{out}}}(S)^2)$ and $p := \max_{\Delta_{\text{in}}, \Delta_{\text{out}} \neq 0^n} (\text{DPP}_{\Delta_{\text{in}}, \Delta_{\text{out}}}(S))$. If $\text{Br}(M) = m + 1$, then $p_{\text{LC}}(C_k) \leq q^{\ell m} \cdot 2^{(\ell-1)n}$ and $p_{\text{DC}}(C_k) \leq p^{\ell m} \cdot 2^{(\ell-1)n}$.*

³Nyberg [1993] actually bounds a related quantity known as the *nonlinearity* of S , but it translates directly to the stated result.

By varying the constant 10 in $r := \lceil b/10 \rceil$, $p_{LC}(\mathcal{F}_2)$ and $p_{DC}(\mathcal{F}_2)$ can be bounded by $2^{-(1-\epsilon)n}$ for any fixed $\epsilon > 0$.

We also note that a bound of $p_{LC}(C_k) \leq 2^{-\Omega(n)}$ incorporates the same bound on each Fourier coefficient of each output bit of C_k . In turn, this implies that each output bit depends on $\Omega(n)$ input bits. (Otherwise, it can be verified that it would have too large a Fourier coefficient, by Parseval's identity.)

Efficiency of $\mathcal{F}_2$. We now explain how to compute $\mathcal{F}_2$ in quasilinear size. The “tricky” component is multiplication by M . Roth and Seroussi [1985, Thm. 1] show that when the Reed-Solomon matrix G is put into reduced echelon form, that is, when $G^T = [I \mid M]$, then M is a generalized Cauchy matrix, defined as follows.

Definition 3.12. Let $\mathbb{F}$ be a field. A matrix $C \in \mathbb{F}^{m \times m}$ is a *Cauchy matrix* if there exist $2m$ distinct values $\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_m \in \mathbb{F}$ such that $C_{i,j} = (\alpha_i + \beta_j)^{-1}$. Furthermore, a matrix $M \in \mathbb{F}^{m \times m}$ is a *generalized Cauchy matrix* if it can be written as $M = BCD$, where C is a Cauchy matrix and $B, D \in \mathbb{F}^{m \times m}$ are diagonal, nonsingular matrices.

Gerasoulis [1988] shows that multiplication of a vector by an $m \times m$ Cauchy matrix can be done with $\tilde{O}(m)$ operations when the underlying field is $\mathbb{C}$. (Multiplication with B and D in the preceding definition can be done with $O(m)$ operations, so we will focus on multiplication by C .) This algorithm can also be made to work over $\text{GF}(2^b)$, as we now show. We stress that we are using the same algorithm from Gerasoulis [1988]; the purpose here is to show that it works over $\text{GF}(2^b)$.

THEOREM 3.13. *Let $C \in \text{GF}(2^b)^{m \times m}$ be a Cauchy matrix defined by the (distinct) elements $\{\alpha_j, \beta_j\}_{j \in [m]}$. Then, given any vector $z \in \text{GF}(2^b)^m$, the product $C \cdot z$ can be computed with $O(m \cdot \log^2 m \cdot \log \log m)$ operations over $\text{GF}(2^b)$.*

PROOF. Define the following polynomial.

$$f(x) := \sum_{j=1}^m z_j (x + \beta_j)^{2^b-2}.$$

Then we have $C \cdot z = (f(\alpha_1), \dots, f(\alpha_m))$, and so it suffices to evaluate f at the points $\{\alpha_i\}_i$. Now define the following three polynomials.

$$\begin{aligned} g(x) &:= \prod_{j=1}^m (x + \beta_j), \\ h(x) &:= \sum_{i=1}^m \left[z_i (x + \beta_i)^{2^b-1} \cdot \prod_{j \neq i} (x + \beta_j) \right], \\ h_*(x) &:= \sum_{i=1}^m z_i \prod_{j \neq i} (x + \beta_j). \end{aligned}$$

Then we have $f(x) = h(x)/g(x)$ as formal polynomials. Furthermore, for any $y \notin \{\beta_j\}_j$, we have $h(y) = h_*(y)$, using the identity $y^{2^b-2} = 1$ which is valid for any $y \neq 0$. Since our goal is to evaluate $f(\alpha_i)$ for all i , this is now seen to be equivalent to evaluating $h_*(\alpha_i)/g(\alpha_i)$, because $\alpha_i \neq \beta_j$ for all i, j .

Notice that, for each β_j , we have $h_*(\beta_j) = z_j \cdot g'(\beta_j)$, where $g'(x) = \sum_{i \in [m]} \prod_{j \neq i} (x + \beta_j)$ is the derivative of g . So, another way to view $h_*(x)$ is that it is the unique degree $\leq m-1$ polynomial interpolating the points $\{(\beta_j, z_j \cdot g'(\beta_j))\}_{j \in [m]}$. The algorithm is now the following.

- (1) Compute $g(x)$ and $g'(x)$ in coefficient form.
- (2) Evaluate $g'(\beta_j)$ for each β_j .
- (3) Compute all values of $z_j \cdot g'(\beta_j)$.
- (4) Interpolate the points $\{(\beta_j, z_j \cdot g'(\beta_j))\}$ to obtain $h_*(x)$ in coefficient form.
- (5) Evaluate both $g(\alpha_j)$ and $h_*(\alpha_j)$ for each α_j .
- (6) Compute each value of $f(\alpha_j) = h_*(\alpha_j)/g(\alpha_j)$.

We note that steps (1) and (2) do not involve the vector z and thus can be preprocessed, and that steps (3) and (6) can easily be done with m operations over $\text{GF}(2^b)$ each. For the remaining steps, we use the following results which can be found in (for example) von zur Gathen and Gerhard [2003, Ch. 10] and which hold for any commutative ring with unity R .

THEOREM 3.14 ([VON ZUR GATHEN AND GERHARD 2003], COROLLARY 10.8). *Evaluation of a polynomial in $R[x]$ of degree $< m$ at m points can be done with $O(m \cdot \log^2 m \cdot \log \log m)$ operations in R .*

THEOREM 3.15 ([VON ZUR GATHEN AND GERHARD 2003], COROLLARY 10.12). *Given m distinct values $u_1, \dots, u_m \in R$ and m arbitrary values $v_1, \dots, v_m \in R$, the unique polynomial in $R[x]$ of degree $< m$, which interpolates $\{(u_i, v_i)\}_i$, can be computed in coefficient form with $O(m \cdot \log^2 m \cdot \log \log m)$ operations in R .*

As a result, steps (4) and (5), and thus the entire multiplication by C , can be performed with the stated number of operations in $\text{GF}(2^b)$. $\square$

One round of $\mathcal{F}_2$ consists of the following three steps:

- (1) m parallel instances of exponentiation in $\text{GF}(2^b)$ (i.e., $x \mapsto x^{2^b-2}$);
- (2) one instance of multiplication by $M \in \text{GF}(2^b)^{m \times m}$;
- (3) one instance of the round key XOR.

Because finite field arithmetic is computable by polynomial size circuits, step (1) can be computed by a circuit with at most $m \cdot b^{O(1)}$ wires. For step (2), we have size at most $m \cdot \log^3 m \cdot b^{O(1)}$ by Theorem 3.13. Step (3) can clearly be done with $O(mb)$ wires. Thus, given the preceding choices of m , b , and r , the r rounds of $\mathcal{F}_2 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ are computable by a circuit of size $n \cdot \log^{O(1)} n$, and the key size is $|k| = mbr = O(n \log n)$.

3.4. Candidate 3

In this section, we define a candidate PRF $\mathcal{F}_3 : \{0, 1\}^n \rightarrow \{0, 1\}$ parameterized by a key of length $O(n)$ and computable by TC^0 circuits of size $O(n^{1+\epsilon})$ for any $\epsilon > 0$. The construction is again inspired by the SPN structure, and the S-box S is defined identically to that of $\mathcal{F}_2$, but the linear transformation M takes a somewhat different form.

The Linear Transformation M . M is constructed using a good error correcting code as before; specifically, we use codes given by the following theorem, which follows from Gál et al. [2013, Theorem 1].

THEOREM 3.16 ([GÁL ET AL. 2013]). *For any constant $\epsilon > 0$, there exist constants $c, \delta > 0$ such that for sufficiently large ℓ , there exists a linear code $C_\epsilon : \{0, 1\}^{\ell/c} \rightarrow \{0, 1\}^\ell$ which has distance $\geq \delta \cdot \ell$ and is computable by a TC^0 circuit of size $O(\ell^{1+\epsilon})$.*

Rather than using a portion of C_ϵ 's generator matrix as with $\mathcal{F}_2$, M consists of the entire matrix that generates C_ϵ . As a result, the internal state grows by a factor of c during each round, and thus the M used at round i will be a $c^i n \times c^{i-1} n$ matrix.

To see the advantage that this has over the previous choice of M , consider an input vector to M in which all b -bit bundles are nonzero. If we use only the fact that M has maximal branch number (Definition 2.1), then we are only guaranteed that M 's output will have one nonzero bundle. However, if we instead take $M = C_\epsilon$, then we are guaranteed that at least $\delta \cdot m$ of the output bundles will be nonzero (where $n = mb$), even if all input bundles were nonzero.

Definition of $\mathcal{F}_3$. Let $m, b, r \in \mathbb{N}$ be arbitrary for now, and set $n := mb$. Fix any $\epsilon > 0$; let c, δ, C_ϵ be given by Theorem 3.16, and for $1 \leq i \leq r$, let $M^{(i)}$ be the matrix that generates C_ϵ when $\ell := c^i n$ in Theorem 3.16. Let $k = (k_0, \dots, k_{r+1})$ denote the key of $\mathcal{F}_3$, where $|k_i| = c^i n$ for $0 \leq i \leq r$ and $|k_{r+1}| = |k_r| = c^r n$.

$\mathcal{F}_3 : \{0, 1\}^n \rightarrow \{0, 1\}$ is computed over r rounds. On input x , $\mathcal{F}_3(x)$ gives $x \oplus k_0$ as input to the first round; the output of round i becomes the input to round $i+1$ (for $1 \leq i < r$), and $\mathcal{F}_3(x)$ outputs $\langle y, k_{r+1} \rangle \in \{0, 1\}$, where y denotes the output of round r . Round i ($1 \leq i \leq r$) is computed over three steps: (1) $c^{i-1}m$ parallel applications of S ; (2) application of $M^{(i)}$ to the entire state; (3) XOR of the entire state with the round key k_i .

Note that this structure is the same as an SPN, except for two changes: M is no longer a permutation (though it is still injective), and the last step is an *inner product* with the final round key, rather than an XOR.

Efficiency of $\mathcal{F}_3$. We now show that $\mathcal{F}_3$ can be computed by TC^0 circuits of size $O(n^{1+\epsilon})$. For $1 \leq i \leq r$, round i consists of the following:

- (1) $c^{i-1}m$ parallel instances of exponentiation in $\text{GF}(2^b)$ (i.e., $x \rightarrow x^{2^{b-2}}$);
- (2) One instance of multiplication by $M^{(i)}$;
- (3) One instance of the round key xor.

Step (1) is computable by a TC^0 circuit of size $c^{i-1} \cdot m \cdot b^{O(1)}$, using the technique described in Section 3.1. Step (2) is computable with size $O(c^{(1+\epsilon)i} \cdot n^{1+\epsilon})$ by Theorem 3.16. Step (3) can be computed with size $O(c^i n)$. The final inner product with k_{r+1} can be computed with size $O(c^r n)$.

Putting it together, there exists a constant κ such that the entire function can be computed by a threshold circuit of size $O(r \cdot (c^r m b^\kappa + c^{(1+\epsilon)r} n^{1+\epsilon}))$ and depth $O(r)$. Let $b \in \mathbb{N}$ be sufficiently large, and set $m := \lceil b^{(\kappa-\epsilon-1)/\epsilon} \rceil$ and $r := \lceil \kappa/\epsilon \rceil$. With $n := mb$, this ensures that $m b^\kappa \leq n^{1+\epsilon}$ (and also that $b^r \geq n$). Thus, the entire function is indeed computable by a TC^0 circuit of size $O(n^{1+\epsilon})$, where both the depth and the hidden constant depend on ϵ .

Security of $\mathcal{F}_3$. Here we are able to leverage techniques from differential cryptanalysis to prove that $\mathcal{F}_3$ is almost 3-wise independent. Specifically, the following proof uses a technique from Nyberg's proof of Theorem 3.10 [Nyberg 1993].

THEOREM 3.4. $\mathcal{F}_3$ is $(3, 2^{-\Omega(n)})$ -wise independent.

PROOF. We will show that $\mathcal{F}_3$ is a 3-wise $2^{-\Omega(n)}$ -bias generator, that is, that for every $d \leq 3$ and any distinct $x_1, \dots, x_d \in \{0, 1\}^n$, $|\Pr_k[\sum_i \mathcal{F}_3(x_i) = 0] - 1/2| < 2^{-\Omega(n)}$. By a well-known fact (cf. [Alon et al. 1992, Lemma 1]), this implies the theorem.

For any input x , let $\mathcal{F}_3^*(x) \in \{0, 1\}^{c^r n}$ denote the state just before the final inner product; that is, $\mathcal{F}_3(x) = \langle \mathcal{F}_3^*(x), k_{r+1} \rangle$. Then, $\sum_i \mathcal{F}_3(x_i) = \langle \sum_i \mathcal{F}_3^*(x_i), k_{r+1} \rangle$. So, we will show that for any $d \leq 3$ and any distinct $x_1, \dots, x_d \in \{0, 1\}^n$,

$$\Pr_{(k_0, \dots, k_r)} \left[\sum_i \mathcal{F}_3^*(x_i) = 0^{c^r n} \right] < 2^{-\Omega(n)},$$

which will complete the proof because $z \neq 0 \Rightarrow \Pr_{k_{r+1}}[\langle z, k_{r+1} \rangle = 0] = 1/2$.

For $d = 2$, this probability is 0 simply because $x_1 \neq x_2 \Rightarrow \mathcal{F}_3^*(x_1) \neq \mathcal{F}_3^*(x_2)$ due to the fact that each component of $\mathcal{F}_3$ is injective.

Fix distinct $x_1, x_2, x_3 \in \{0, 1\}^n$. Fix any values for $(k_0, \dots, k_{r-2})$, the round keys used prior to round $r - 1$. Let $y_i \in \{0, 1\}^{c^{r-1}n}$ denote the state of the computation of $\mathcal{F}_3(x_i)$ immediately prior to the XOR with round key k_{r-1} in round $r - 1$, and let $\Delta_i := y_1 \oplus y_i$. Let z_1, z_2, z_3 be jointly-distributed random variables, over a uniform choice of k_{r-1} , defined by $z_i := y_i \oplus k_{r-1}$; note that (z_1, z_2, z_3) is uniformly distributed over all tuples with xors $\{\Delta_i\}_i$.

Fix any $j \leq m$, and let $\bar{z}_1$ denote the j th bundle of z_1 and $\bar{\Delta}_i$ denote the j th bundle of Δ_i . We wish to bound

$$\Pr_{k_{r-1}} \left[(\bar{z}_1)^{2^b-2} + (\bar{z}_1 + \bar{\Delta}_2)^{2^b-2} + (\bar{z}_1 + \bar{\Delta}_3)^{2^b-2} = 0 \right], \quad (5)$$

the probability that the outputs of the j th S-box in round r sum to 0. If $\bar{\Delta}_1 = \bar{\Delta}_2 = 0$, then the equation is satisfied if and only if $\bar{z}_1 = 0$, in which case (5) = 2^{-b} . Now assume that at least one of $\bar{\Delta}_1, \bar{\Delta}_2$ are not zero. If we assume that $\bar{z}_1 \notin \{0, \bar{\Delta}_1, \bar{\Delta}_2\}$, then we may multiply both sides of the equation by $\prod_i (\bar{z}_1 + \bar{\Delta}_i) \neq 0$ to get a quadratic polynomial in $\bar{z}_1$. Thus, there are at most five values of $\bar{z}_1$ for which the equation is satisfied (including $\{0, \bar{\Delta}_1, \bar{\Delta}_2\}$), so we can bound (5) $< 6/2^b$.

Finally, because each bundle of k_{r-1} is chosen independently, and because the remaining steps in round r are linear, we have

$$\Pr_{(k_0, \dots, k_r)} \left[\sum_{i \leq 3} \mathcal{F}_3^*(x_i) = 0^{c^r n} \right] < \left(\frac{6}{2^b} \right)^{c^{r-1}m} = 2^{-\Omega(n)}. \quad \square$$

Note that this proof does not use any properties of the code C_ϵ aside from injectivity. We remark why this proof does not show that $\mathcal{F}_3$ is almost d -wise independent for $d \geq 4$. When the number of inputs d is even, Equation (5) is satisfied for all values of $\bar{z}_1$ if and only if $\bar{\Delta}_1, \dots, \bar{\Delta}_d$ can be partitioned into $d/2$ pairs such that the two values in each pair are equal. Indeed, for even $d \in (2, 2^b]$, it is possible to construct a set $\{\Delta_1, \dots, \Delta_d\}$ which admits such a partition for all j and yet satisfies the minimum-distance property of C_ϵ (which guarantees that $\geq \delta c^{r-1}m$ bundles of each Δ_i are nonzero for $i > 1$, and further that $\geq \delta c^{r-1}m$ bundles of $(\Delta_i \oplus \Delta_j)$ are nonzero for all $i \neq j$). However, it seems counterintuitive that the differences at round r would satisfy such a specialized property with noticeable probability, and we believe that this proof can be extended to higher values of d .

3.5. Candidate 4

For our next candidate, we choose the extreme setting of $b = n$ and $r = 1$, which means that the function is computed over one round and essentially consists of just a single S-box. More specifically, the function is indexed by a seed $(k_0, k_1) \in \{0, 1\}^{2n}$ and is computed as

$$\mathcal{F}_4(x) := (x + k_0)^{2^n-2} + k_1.$$

Though $\mathcal{F}_4$ does indeed have resistance to differential and linear cryptanalysis, we note that the seed can be recovered with four known input/output pairs using an attack similar in spirit to the so-called interpolation attack of Jakobsen and Knudsen [2001].

CLAIM 3.17. *Let $\mathcal{F}_4$ be the preceding function indexed by $k_0, k_1 \in \{0, 1\}^n$, and let p_1, p_2, p_3, p_4 be any set of distinct inputs with $p_1 + p_2 \neq p_3 + p_4$. Then, with probability*

$(1 - 1/2^{n-2})$ over k_0 , the values of k_0 and k_1 can be recovered from $\{(p_i, \mathcal{F}_4(p_i))\}_i$ in time $\text{poly}(n)$.

PROOF. The attack is performed by using $\{(p_i, \mathcal{F}_4(p_i))\}_i$ to create two equations over $\text{GF}(2^n)$ that are linear in the seed, as follows. Assume that $k_0 \notin \{p_i\}_i$, which happens with probability $(1 - 1/2^{n-2})$. Denote $c_i := \mathcal{F}_4(p_i)$. Then the equation

$$(c_i + k_1) \cdot (p_i + k_0) = 1 \quad (6)$$

holds for $1 \leq i \leq 4$. We can rewrite these equations as

$$k_0 k_1 + c_i k_0 + p_i k_1 + c_i p_i = 1. \quad (7)$$

If we sum (7) for $i = 1, 2$, the quadratic terms cancel and we obtain

$$(c_1 + c_2)k_0 + (p_1 + p_2)k_1 + (c_1 p_1 + c_2 p_2) = 0.$$

Summing (7) for $i = 3, 4$ gives another linear equation in k_0, k_1 . The attack concludes by solving the two linear equations. In the unlikely event that $c_1 + c_2 = c_3 + c_4$, k_1 can still be uniquely recovered because $p_1 + p_2 \neq p_3 + p_4$, and then k_0 can be recovered from (6). $\square$

The function $\mathcal{F}_4$ can be seen as a concrete instantiation of the Even-Mansour cipher [Even and Mansour 1997], where the random permutation is replaced with (the asymptotic version of) the AES S-box. This cipher is easily breakable as we have just observed, but we now consider a slight modification to $\mathcal{F}_4$ that is not susceptible to this simple attack, and furthermore fools all parity tests that look at $\leq 2^{0.9n}$ outputs. The modified function $\mathcal{F}'_4 : \{0, 1\}^n \rightarrow \{0, 1\}$ is defined as follows:

$$\mathcal{F}'_4(x) := \langle (x + k_0)^{2^n-2}, k_1 \rangle.$$

In other words, we combine the AES S-box with the Goldreich-Levin hardcore predicate [Goldreich and Levin 1989]. Note that we now output only a single bit. This modification—replacing the second XOR with an inner product—can also be applied to the Even-Mansour cipher. We consider it an interesting question to what extent the assumptions necessary for the pseudorandomness of Even-Mansour can be relaxed in this setting. (In their setting, the assumption is that all parties have oracle access to a truly random permutation.)

The next theorem shows that $\mathcal{F}'_4$ fools all parity tests that look at $\leq 2^{0.9n}$ outputs. This result is reminiscent of the “exponentiation” small-bias generator in Alon et al. [1992], where the x th output bit is $\langle k_0^x, k_1 \rangle$. Indeed, our proof is inspired by theirs. However we face the extra difficulty that the polynomials we work with are not of low degree.

THEOREM 3.5. *For any choice of $d \leq 2^n$, $\mathcal{F}'_4$ is a d -wise small-bias generator with error $d/2^n$. That is, for any distinct $a_1, \dots, a_d \in \{0, 1\}^n$,*

$$\left| \Pr_{k_0, k_1} \left[\sum_{i=1}^d \mathcal{F}'_4(a_i) = 0 \right] - \frac{1}{2} \right| < \frac{d}{2^n}.$$

PROOF. Fix any distinct choices of $a_1, \dots, a_d$. Then, identifying elements of $\text{mathrm{GF}}(2^n)$ with elements of $\{0, 1\}^n$, we have

$$\begin{aligned} \sum_{i \leq d} \mathcal{F}'_4(a_i) &= \sum_{i \leq d} \langle (a_i + k_0)^{2^n-2}, k_1 \rangle \\ &= \left\langle p(x) := \sum_{i \leq d} (a_i + k_0)^{2^n-2}, k_1 \right\rangle. \end{aligned}$$

We now show that the polynomial $p(x) = \sum_{i \leq d} (a_i + x)^{2^n - 2}$ has at most $2d - 1$ distinct roots. This will conclude the proof because when k_0 is not a root of $p(x)$, we have $\Pr_{k_1}[\langle p(k_0), k_1 \rangle = 0] = 1/2$. Therefore,

$$\left| \Pr_{k_0, k_1} \left[\sum_{i=1}^d \mathcal{F}'_4(a_i) = 0 \right] - \frac{1}{2} \right| \leq \frac{1}{2} \Pr_{k_0} [p(k_0) = 0] < \frac{d}{2^n}.$$

To show the bound on the number of roots, define the following polynomials:

$$\begin{aligned} \bar{p}(x) &:= p(x) \cdot \prod_{i \leq d} (a_i + x) = \sum_{i \leq d} \left[(a_i + x)^{2^n - 1} \prod_{j \neq i} (a_j + x) \right], \\ \bar{p}_*(x) &:= \sum_{i \leq d} \prod_{j \neq i} (a_j + x). \end{aligned}$$

Observe that any root y of $p(x)$ is also a root of $\bar{p}(x)$. Moreover, note for any $y \notin \{a_j : j \leq d\}$, $\bar{p}(y) = \bar{p}_*(y)$, using the identity $y^{2^b - 2} = 1$ which is valid for any $y \neq 0$.

Also observe that $\bar{p}_*(x)$ is not identically zero. Indeed, by inspection, the constant term of the polynomial $\bar{p}_*(x + a_1)$ is $\prod_{j \neq 1} (a_j + a_1)$, which is nonzero because the a_j are distinct; therefore, $\bar{p}_*(x + a_1)$ is not identically zero, and so neither is $\bar{p}_*(x)$. Since $\bar{p}_*(x)$ is a nonzero polynomial of degree $d - 1$, it has at most $d - 1$ distinct roots.

So, if $p(x)$ has r roots, also $\bar{p}$ has r roots. At least $r - d$ of these do not belong to $\{a_j : j \leq d\}$, and so they are also roots of $\bar{p}_*(x)$. Therefore, $r - d \leq d - 1$, or $r \leq 2d - 1$. $\square$

By a result of Braverman [2009] (cf. [Bazzi 2009; Razborov 2009]), we obtain that $\mathcal{F}'_4$ also fools small-depth AC^0 circuits of any size $w = 2^{n^{o(1)}}$ (that look at only w fixed output bits of the candidate).

Indeed, fix any function $w = 2^{n^{o(1)}}$ and any constant $d = O(1)$; let $N := 2^n$. By Theorem 3.5, any w output bits have bias $< w/N$. By Alon et al. [2003], for any $k \leq w$, the output distribution on those w bits is $w^k w/N$ -close to a k -wise independent distribution. By Braverman [2009], $k = \lg^{O(d^2)} w \leq n^{o(1)}$ is sufficient to fool circuits of depth d with error $1/w$. Hence the overall error will be $1/w = 1/2^{n^{o(1)}}$ plus

$$\frac{w^k w}{N} = \frac{(2^{n^{o(1)}})^{n^{o(1)}}}{N} \leq \frac{1}{\sqrt{N}},$$

for a total of $1/w + 1/\sqrt{N} = O(1/w)$.

Efficiency. As noted in Section 3.1, $\mathcal{F}'_4$ is computable by Boolean circuits of size $O(n \log^2 n \log \log n)$ and TC^0 circuits of size $n^{O(1)}$.

3.6. Candidate 5

Our final candidate $\mathcal{F}_5$ preserves the structure of AES almost exactly. For any n that is a multiple of 32, we set $b = 8$, $m = n/8$, and $r = n$, and we use $S(x) := x^{2^b - 2}$. The linear transformation M is of a slightly different form than that of the previous candidates, which we explain now.

M is computed in two (linear) steps. In the first step, a permutation $\pi : [m] \rightarrow [m]$ is used to shuffle the b -bit bundles of the state; namely, bundle i moves to position $\pi(i)$. The permutation π is computed as follows. First, the m bundles are placed column-wise

into a $4 \times m/4$ matrix. Then row i of the matrix ($0 \leq i < 4$) is shifted circularly to the left by i places, and finally the bundles are extracted column-wise from the new matrix.

In the second step, a maximal-branch-number matrix $\phi \in \text{GF}(2^8)^{4 \times 4}$ is applied in parallel to each consecutive group of 4 bundles.

Efficiency: Small Circuits. In each round, the $O(n)$ instances of S and ϕ each perform computations on a constant number of bits; because permuting the bundles and adding the round key can also be done with $O(n)$ wires, each round of $\mathcal{F}_5$ can be computed by a circuit of depth $d = O(1)$ and size $w = O(n)$. Thus the entire (r -round) circuit for $\mathcal{F}_5$ has depth $d = O(n)$ and size $w = O(n^2)$.

Efficiency: Fast Turing Machines. Similarly, for any fixed seed k , each round of $\mathcal{F}_5$ can be computed in time $O(n)$ on a single-tape Turing machine with $O(n^2)$ states. To do so, we encode the bundles on the tape so that the matrix used by π is written column-wise. As before, the $O(n)$ instances of S and ϕ in a single round can be done in time $O(n)$. To see that π can also be computed in time $O(n)$, note that due to the column-wise representation, each bundle needs to move ≤ 3 places away, except for the six bundles which are shifted circularly to the other end of the tape. Finally, encoding the $O(n^2)$ -bit seed in the TM's state transitions, the addition of each round key also takes time $O(n)$. Therefore, the $r = n$ rounds of $\mathcal{F}_5$ can be computed in time $O(n^2)$.

Alternatively, consider the Turing machine variant with two tapes, in which the first tape is read-only and contains the n -bit input followed by the $n(n+1)$ -bit seed, the second tape is read/write, and the TM has $O(1)$ states. Then $\mathcal{F}_5$ can again be computed in time $O(n^2)$ exactly as just described, because in round i , only bits $in+1, \dots, in+n$ of the seed are used.

4. CONCLUSION AND FUTURE WORK

Two obvious directions for future work are to extend the analysis of $\mathcal{F}_1$ to handle inverse queries (necessarily choosing the S-box as a random *permutation*), and to extend Theorem 3.4 to prove almost d -wise independence of $\mathcal{F}_3$ for $d > 3$. A more foundational question left unanswered is to understand how the degree of each output bit of an SPN (as a polynomial in the input bits) is affected by the degree of the S-box and by the “mixing” properties of the linear transformation.

Exploring other choices of the S-box besides inversion may lead to more efficient constructions, and utilizing other properties of the linear transformation besides maximal-branch-number may allow stronger proofs of security. This could potentially give a (plausibly secure) SPN computable by circuits of size $O(n)$. Recall from Section 1 that a PRF computable with size $O(n)$ and with security 2^n would bring the natural proofs barrier to the current frontier of lower bounds against unbounded-depth circuits.

Abstracting from the SPN structure, one may arrive to the following paradigm for constructing PRF: alternate the application of (1) an error-correcting code and (2) a bundle-wise application of any local map that has high degree over $\text{GF}(2)$ and resists attacks corresponding to linear and differential cryptanalysis. This viewpoint may lead to a PRF candidate computable in ACC^0 , since for stipulation (1) one just needs parity gates, while, say, taking parities of suitable stipulation mod 3 maps one should get a map that satisfies (2). However a good choice for this latter map is not clear to us at this moment.

We believe a good candidate PRF should be the simplest candidate that resists known attacks. As noted in Daemen and Rijmen [2002], some of the choices in the design of AES are not motivated by any known attack but are there as a safeguard (e.g., one can reduce the number of rounds and still no attack is known). While this is comprehensible

when having to choose a standard that is difficult to change or when deploying a system that is to be widely used, one could argue that a better way for the research community to proceed is to put forth the simplest candidate PRF, possibly break it, and iterate until hopefully converging to a secure PRF. We view this article as a step in this direction.

APPENDIXES

A. SECURITY AGAINST LINEAR/DIFFERENTIAL CRYPTANALYSIS

In this section, we fill in the missing details from Section 2 on how the security of an SPN is evaluated against linear and differential cryptanalysis, and we prove Theorem 2.3 via an inductive extension of the results of Kang et al. [2001].

A.1. Linear Cryptanalysis

Recall the following two definitions from Section 2.

$$\text{Cor}_{\Gamma_{\text{in}}, \Gamma_{\text{out}}}(f) := 2 \cdot \Pr_x[\langle \Gamma_{\text{in}}, x \rangle = \langle \Gamma_{\text{out}}, f(x) \rangle] - 1 \in [-1, 1].$$

$$p_{\text{LC}}(C_k) := \max_{\Gamma_{\text{in}}, \Gamma_{\text{out}} \neq 0^n} (\mathbb{E}_k[\text{Cor}_{\Gamma_{\text{in}}, \Gamma_{\text{out}}}(C_k)^2]).$$

To bound $p_{\text{LC}}(C_k)$, the concept of a *linear trail* is used. Let ρ_k^i denote the i th round function of an SPN C_k , that is, $C_k(x) = \rho_k^r(\rho_k^{r-1}(\dots(\rho_k^1(x \oplus k_0))\dots))$. A linear trail is a vector $\Gamma = (\Gamma_0, \dots, \Gamma_r) \in (\{0, 1\}^n)^{r+1}$, and the correlation of C_k with respect to Γ is (cf. Daemen and Rijmen [2002, Eq. 7.59])

$$\text{Cor}_{\Gamma}(C_k) := \prod_{i=1}^r \text{Cor}_{\Gamma_{i-1}, \Gamma_i}(\rho_k^i).$$

This equation is defined for a fixed key k , but in fact for SPNs, only the sign of this product is affected by the value of the key [Daemen and Rijmen 2002, Section 7.9.2]. In particular, $\text{Cor}_{\Gamma}(C_k)^2$ is the same for every key k .

For any pair of input/output parities $\Gamma_{\text{in}}, \Gamma_{\text{out}}$, we have the following theorem.

THEOREM A.1 ([DAEMEN AND RIJMEN 2002], THM. 7.9.1).

$$\mathbb{E}_k[\text{Cor}_{\Gamma_{\text{in}}, \Gamma_{\text{out}}}(C_k)^2] = \sum_{\Gamma: \Gamma_0 = \Gamma_{\text{in}}, \Gamma_r = \Gamma_{\text{out}}} \text{Cor}_{\Gamma}(C_k)^2.$$

A naïve evaluation of this sum would lead to a useless bound on p_{LC} (i.e., a bound ≥ 1) due to the large number of vectors Γ that have the specified first and final elements. Kang et al. [2001] give an exponentially small bound on this sum (Theorem 2.2) in the case where $r = 2$ and the linear transformation M has maximal branch number.

A.2. Differential Cryptanalysis

Recall the following two definitions from Section 2.

$$\text{DPP}_{\Delta_{\text{in}}, \Delta_{\text{out}}}(f_k) := \Pr_{x, k}[f_k(x) \oplus f_k(x \oplus \Delta_{\text{in}}) = \Delta_{\text{out}}].$$

$$p_{\text{DC}}(C_k) := \max_{\Delta_{\text{in}}, \Delta_{\text{out}} \neq 0^n} (\text{DPP}_{\Delta_{\text{in}}, \Delta_{\text{out}}}(C_k)).$$

Similarly to how linear trails were used in the previous section, *differential trails* are used to bound $p_{\text{DC}}(C_k)$. A differential trail is a vector $\Delta = (\Delta_0, \dots, \Delta_r) \in (\{0, 1\}^n)^{r+1}$. For any SPN C_k , again let ρ_k^i denote its i th round function, and let $C_k^{(i)}(x)$ denote

the output of the i th round of $C_k(x)$, with $C_k^{(0)}(x) := x \oplus k_0$. That is, for any $i \leq r$, $C_k^{(i)}(x) = \rho_k^i(\rho_k^{i-1}(\dots(\rho_k^1(x \oplus k_0))\dots))$.

Then for any Δ_0, Δ_r , we have

$$\text{DPP}_{\Delta_0, \Delta_r}(C_k) = \sum_{\Delta_1, \dots, \Delta_{r-1}} \Pr_{x, k} \left[\bigwedge_{i=1}^r [\rho_k^i(C_k^{(i-1)}(x)) \oplus \rho_k^i(C_k^{(i-1)}(x \oplus \Delta_0)) = \Delta_i] \right].$$

This can be seen by noting that for any fixed values of x, k, Δ_0 , and Δ_r , there is at most one tuple $(\Delta_1, \dots, \Delta_{r-1})$ for which the conjunction evaluates to true. To simplify this equation, we use the following two facts.

- The independence of the round keys ensures that, conditioned on two inputs to round i having XOR Δ_{i-1} , the inputs are uniformly distributed over all pairs with XOR Δ_{i-1} and are independent of the inputs to all previous rounds.
- XORing the round key does not affect the DPP of a given round. That is, letting ρ denote the round function without the key XOR, we have $\text{DPP}_{\Delta_{in}, \Delta_{out}}(\rho) = \text{DPP}_{\Delta_{in}, \Delta_{out}}(\rho_k^i)$ for all $i, \Delta_{in}, \Delta_{out}$.

Using these facts and an application of the chain rule, we have

$$\text{DPP}_{\Delta_0, \Delta_r}(C_k) = \sum_{\Delta_1, \dots, \Delta_{r-1}} \prod_{i=1}^r \text{DPP}_{\Delta_{i-1}, \Delta_i}(\rho).$$

Kang et al. [2001] again give an exponentially small bound when $r = 2$ (Theorem 2.2).

A.3. Proof of Theorem 2.3

We now prove Theorem 2.3 via an inductive extension of Theorem 2.2. We restate both theorems for convenience.

THEOREM 2.2 ([KANG ET AL. 2001], THMS. 5 & 6). *Let $C_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an SPN with $r = 2$ rounds and S -box S . Let $q := \max_{\Gamma_{in}, \Gamma_{out} \neq 0^n} (\text{Cor}_{\Gamma_{in}, \Gamma_{out}}(S)^2)$ and $p := \max_{\Delta_{in}, \Delta_{out} \neq 0^n} (\text{DPP}_{\Delta_{in}, \Delta_{out}}(S))$. If $\text{Br}(M) = m + 1$, then $p_{\text{LC}}(C_k) \leq q^m$ and $p_{\text{DC}}(C_k) \leq p^m$.*

THEOREM 2.3. *Let $C_k : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an SPN with $r = 2\ell \geq 2$ rounds and S -box S . Let $q := \max_{\Gamma_{in}, \Gamma_{out} \neq 0^n} (\text{Cor}_{\Gamma_{in}, \Gamma_{out}}(S)^2)$ and $p := \max_{\Delta_{in}, \Delta_{out} \neq 0^n} (\text{DPP}_{\Delta_{in}, \Delta_{out}}(S))$. If $\text{Br}(M) = m + 1$, then $p_{\text{LC}}(C_k) \leq q^{\ell m} \cdot 2^{(\ell-1)n}$ and $p_{\text{DC}}(C_k) \leq p^{\ell m} \cdot 2^{(\ell-1)n}$.*

PROOF. We prove part 1; part 2 is essentially identical.

We proceed inductively on ℓ . The base case $\ell = 1$ is given by Theorem 2.2. Fix $\ell > 1$, and let $\Gamma_0, \Gamma_{2\ell}$ be any nonzero input/output parities. Then,

$$\begin{aligned} p_{\text{LC}}(C_k) &= \sum_{\Gamma=(\Gamma_0, \dots, \Gamma_{2\ell})} \text{Cor}_{\Gamma}(C_k)^2 \\ &= \sum_{\Gamma_1, \dots, \Gamma_{2\ell-1}} \prod_{i=1}^{2\ell} \text{Cor}_{\Gamma_{i-1}, \Gamma_i}(\rho_{k_i})^2 \\ &= \sum_{\Gamma_1, \dots, \Gamma_{2\ell-2}} \prod_{i=1}^{2\ell-2} \text{Cor}_{\Gamma_{i-1}, \Gamma_i}(\rho_{k_i})^2 \sum_{\Gamma_{2\ell-1}} \prod_{i=2\ell-1}^{2\ell} \text{Cor}_{\Gamma_{i-1}, \Gamma_i}(\rho_{k_i})^2 \end{aligned}$$

$$\leq q^m \cdot \sum_{\Gamma_1, \dots, \Gamma_{2\ell-2}} \prod_{i=1}^{2\ell-2} \text{Cor}_{\Gamma_{i-1}, \Gamma_i}(\rho_{k_i})^2 \quad (8)$$

$$\begin{aligned} &= q^m \cdot \sum_{\Gamma_{2\ell-2}} \left(\sum_{\Gamma_1, \dots, \Gamma_{2\ell-3}} \prod_{i=1}^{2\ell-2} \text{Cor}_{\Gamma_{i-1}, \Gamma_i}(\rho_{k_i})^2 \right) \\ &\leq q^m \cdot \sum_{\Gamma_{2\ell-2}} q^{(\ell-1)m} \cdot 2^{(\ell-2)n} \quad (9) \end{aligned}$$

$$= q^{\ell m} \cdot 2^{(\ell-1)n}, \quad (10)$$

where (8) is by Theorem 2.2, (9) is by the inductive hypothesis, and (10) is by the fact that there are 2^n choices for $\Gamma_{2\ell-2}$. $\square$

B. DISTINGUISHING $O(N)$ -DEGREE PRFS

In this section, we show (in Theorem 2.4 in the following) that any PRF $f_k : \{0, 1\}^n \rightarrow \{0, 1\}$ that is computable by an $o(n)$ -degree polynomial over $\text{GF}(2)$ cannot have hardness 2^n . This just follows from the fact that in time 2^n one can write down the polynomial representation of f restricted to $\Omega(n)$ input bits. Details follow.

For simplicity, we instead show that any such PRF can be broken in time $2^{O(n)}$. This implies the desired goal, for if we had a PRF $f_k : \{0, 1\}^n \rightarrow \{0, 1\}$ with hardness 2^n , we could consider it over bn input bits (where b is larger than the hidden constant in the $2^{O(n)}$ runtime of the attack), note that the degree would still be $o(n) = o(bn)$, and obtain a contradiction.

To start, let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be any function, and define the following three values.

- $T_f \in \{0, 1\}^{2^n}$ is the truth table of f ; that is, $(T_f)_i := f(i)$, identifying a natural number with its binary representation.
- $C_f \in \{0, 1\}^{2^n}$ is the coefficient vector of f , defined as follows. Fix some ordering on the 2^n possible multilinear monomials in n variables. Then, $(C_f)_i = 1$ if and only if the i th monomial appears in the polynomial representation of f over $\text{GF}(2)$.
- $A \in \{0, 1\}^{2^n \times 2^n}$ is the matrix with rows indexed by the set $\{0, 1\}^n$ and columns indexed by the set of degree $\leq n$ multilinear monomials (as with C_f), defined by $A_{ij} := 1$ iff monomial j has value 1 under input i .

Note that A is independent of the function f . Furthermore, A is invertible because it has full rank, which follows from the fact that any two distinct linear combinations of A 's columns give the truth tables of two distinct polynomials. We now show how to distinguish a low-degree PRF using the fact that $A \cdot C_f = T_f$ for all f .

THEOREM 2.4. *Let $F = \{f_k : \{0, 1\}^n \rightarrow \{0, 1\}\}_k$ be any set of functions such that, for each key k , the polynomial representation of f_k over $\text{GF}(2)$ has degree $o(n)$. Then there is an adversary that runs in time $2^{O(n)}$ and distinguishes a random $f_k \in F$ from a random function with advantage $1 - 2^{-2^{\Omega(n)}}$.*

PROOF. For any function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, we can use C_f to check if the polynomial representation of f contains a monomial of degree $\geq n/2$. Clearly this will be false for any f_k drawn from the PRF, and for a uniformly random function F , we have

$$\Pr_F[F \text{ has a monomial of degree } \geq n/2] \geq 1 - 2^{-\binom{n}{n/2}} \geq 1 - 2^{-2^{\Omega(n)}},$$

which can be seen by viewing F as being randomly chosen by including each possible monomial independently with probability $1/2$. Finally, note that C_f can be computed from the truth table of f in time $2^{O(n)}$ as $C_f = A^{-1} \cdot T_f$. $\square$

ACKNOWLEDGMENTS

We thank Craig Gentry for pointing out that Candidate 4 is computable in quasilinear size, Guevara Noubir for helpful discussions, and Salil Vadhan for mentioning AES. We also thank the anonymous referees for very helpful feedback, including pointing us to Kang et al. [2001].

REFERENCES

- Scott Aaronson and Avi Wigderson. 2009. Algebrization: A new barrier in complexity theory. *ACM Trans. Comput. Theory* 1, 1 (2009) Article 2. DOI=<http://dx.doi.org/10.1145/1490270.1490272>
- Adi Akavia, Andrej Bogdanov, Siyao Guo, Akshay Kamath, and Alon Rosen. 2014. Candidate weak pseudorandom functions in $AC^0 \circ MOD_2$. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science (ITCS'14)*. ACM, New York, NY, 251–260. DOI=<http://dx.doi.org/10.1145/2554797.2554821>
- Eric Allender and Michal Koucký. 2010. Amplifying lower bounds by means of self-reducibility. *J. ACM* 57, 3, (2010) Article 14. DOI=<http://dx.doi.org/10.1145/1706591.1706594>
- Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. 1992. Simple constructions of almost k -wise independent random variables. *Random Struct. Algorith.* 3, 3 (1992), 289–304. DOI=<http://dx.doi.org/10.1002/rsa.3240030308>
- Noga Alon, Oded Goldreich, and Yishay Mansour. 2003. Almost k -wise independence versus k -wise independence. *Inf. Process. Lett.* 88, 3 (2003), 107–110. DOI=[http://dx.doi.org/10.1016/S0020-0190\(03\)00359-4](http://dx.doi.org/10.1016/S0020-0190(03)00359-4)
- Theodore Baker, John Gill, and Robert Solovay. 1975. Relativizations of the $P=?NP$ question. *SIAM J. Comput.* 4, 4 (1975), 431–442. DOI=<http://dx.doi.org/10.1137/0204037>
- Abhishek Banerjee and Chris Peikert. 2014. New and improved key-homomorphic pseudorandom functions. In *Proceedings of the 34th Annual Cryptology Conference (CRYPTO), Part I*. 353–370. DOI=http://dx.doi.org/10.1007/978-3-662-44371-2_20
- Abhishek Banerjee, Chris Peikert, and Alon Rosen. 2012. Pseudorandom functions and lattices. In *Proceedings of the 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*. 719–737. DOI=http://dx.doi.org/10.1007/978-3-642-29011-4_42
- Louay M. J. Bazzi. 2009. Polylogarithmic independence can fool DNF formulas. *SIAM J. Comput.* 38, 6 (2009), 2220–2272. DOI=<http://dx.doi.org/10.1137/070691954>
- Mihir Bellare, Joe Kilian, and Phillip Rogaway. 2000. The security of the cipher block chaining message authentication code. *J. Comput. Syst. Sci.* 61, 3 (Dec. 2000), 362–399. DOI=<http://dx.doi.org/10.1006/jcss.1999.1694>
- E. Biham and A. Shamir. 1991. Differential cryptanalysis of DES-like cryptosystems. *J. Cryptol.* 4, 1 (1991), 3–72. DOI=<http://10.1007/BF00630563>
- Mark Braverman. 2009. Poly-logarithmic independence fools AC^0 circuits. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity (CCC'09)*. IEEE Computer Society, 3–8. <http://dx.doi.org/10.1109/CCC.2009.35>
- Alex Brodsky and Shlomo Hoory. 2008. Simple permutations mix even better. *Random Struct. Algorith.* 32, 3 (May 2008), 274–289. DOI=<http://dx.doi.org/10.1002/rsa.v32:3>
- Hong-Su Cho, Soo Hak Sung, Daesung Kwon, Jung-Keun Lee, Jung Hwan Song, and Jongin Lim. 2004. New method for bounding the maximum differential probability for SPNs and ARIA. In *Proceedings of the 7th International Conference on Information Security and Cryptology (ICISC'04)*, Choon-sik Park and Seong-taek Chee (Eds.), Springer-Verlag, Berlin, Heidelberg, 21–32. DOI=http://dx.doi.org/10.1007/11496618_4
- Joan Daemen. 1995. Cipher and hash function design strategies based on linear and differential cryptanalysis. Ph.D. Dissertation, K.U. Leuven, Leuven, Flanders, Belgium.
- Joan Daemen and Vincent Rijmen. 2002. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Springer Verlag, New York, Inc., Secaucus, NJ.
- Shimon Even and Yishay Mansour. 1997. A construction of a cipher from a single pseudorandom permutation. *J. Cryptol.* 10, 3 (June 1997), 151–162. DOI=<http://dx.doi.org/10.1007/s001459900025>
- Anna Gál, Kristoffer Arnsfelt Hansen, Michal Koucký, Pavel Pudlák, and Emanuele Viola. 2013. Tight bounds on computing error-correcting codes by bounded-depth circuits with arbitrary gates. *IEEE Trans. Inform. Theory* 59, 10 (2013), 6611–6627. DOI=<http://dx.doi.org/10.1109/TIT.2013.2270275>

- Joseph Naor and Moni Naor. 1993. Small-bias probability spaces: Efficient constructions and applications. *SIAM J. Comput.* 22, 4 (August 1993), 838–856. DOI=<http://dx.doi.org/10.1137/0222053>
- Moni Naor and Omer Reingold. 1999. On the construction of pseudorandom permutations: Luby-Rackoff revisited. *J. Cryptol.* 12, 1 (January 1999), 29–66. DOI=<http://dx.doi.org/10.1007/PL00003817>
- Moni Naor and Omer Reingold. 2004. Number-theoretic constructions of efficient pseudo-random functions. *J. ACM* 51, 2 (March 2004), 231–262. DOI=<http://dx.doi.org/10.1145/972639.972643>
- Moni Naor, Omer Reingold, and Alon Rosen. 2002. Pseudorandom functions and factoring. *SIAM J. Comput.* 31, 5 (May 2002), 1383–1404. DOI=<http://dx.doi.org/10.1137/S0097539701389257>
- Kaisa Nyberg. 1993. Differentially uniform mappings for cryptography. In *Proceedings of the Workshop on the Theory and Application of Cryptographic Techniques on Advances in Cryptology (EUROCRYPT'93)*. Tor Hellesest (Ed.), Springer-Verlag New York, Inc., Secaucus, NJ, 55–64.
- Josef Pieprzyk. 1991. On bent permutations. In *Proceedings of the International Conference on Finite Fields, Coding Theory, and Advances in Communications and Computing*. Lecture Notes in Pure and Applied Math, Vol. 141. Marcel Dekker, Inc.
- Zulfikar Ramzan and Leonid Reyzin. 2000. On the round security of symmetric-key cryptographic primitives. In *Proceedings of the 20th Annual International Cryptology Conference on Advances in Cryptology (CRYPTO'00)*, Mihir Bellare (Ed.), Springer-Verlag, London, 376–393.
- Alexander Razborov and Steven Rudich. 1997. Natural proofs. *J. Comput. Syst. Sci.* 55, 1 (August 1997), 24–35. DOI=<http://dx.doi.org/10.1006/jcss.1997.1494>
- Alexander A. Razborov. 2009. A simple proof of Bazzi's theorem. *ACM Trans. Comput. Theory* 1, 1 (February 2009). Article 3. DOI=<http://dx.doi.org/10.1145/1490270.1490273>
- Ron M. Roth and Gadiel Seroussi. 1985. On generator matrices of MDS codes. *IEEE Trans. Inf. Theor.* 31, 6 (November 1985), 826–830. DOI=<http://dx.doi.org/10.1109/TIT.1985.1057113>
- Claude Shannon. 1949. Communication theory of secrecy systems. *Bell Syst. Tech. J.* 28, 4 (1949), 656–715. DOI=<http://dx.doi.org/10.1002/j.1538-7305.1949.tb00928.x>
- Salil P. Vadhan and Colin Jia Zheng. 2012. Characterizing pseudoentropy and simplifying pseudorandom generator constructions. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC'12)*. ACM, New York, NY, 817–836. DOI=<http://dx.doi.org/10.1145/2213977.2214051>
- Joachim von zur Gathen and Jürgen Gerhard. 2003. *Modern Computer Algebra* (2nd ed.). Cambridge University Press, New York, NY.
- Ryan Williams. 2011. Non-uniform ACC circuit lower bounds. In *Proceedings of the IEEE Conference on Computational Complexity (CCC)*. 115–125. DOI=<http://dx.doi.org/10.1109/CCC.2011.36>
- Hongjun Wu. 2011. The hash function JH. <http://www3.ntu.edu.sg/home/wuhj/research/jh/index.html>.

Received September 2012; revised August 2015; accepted September 2015