# Cryptanalysis of stream ciphers with linear masking 

Don Coppersmith Shai Halevi Charanjit Jutla

IBM T. J. Watson Research Center, NY, USA<br>\{copper,shaih, csjutla@watson.ibm.com\}

May 31, 2002


#### Abstract

We describe a cryptanalytical technique for distinguishing some stream ciphers from a truly random process. Roughly, the ciphers to which this method applies consist of a "non-linear process" (say, akin to a round function in block ciphers), and a "linear process" such as an LFSR (or even fixed tables). The output of the cipher can be the linear sum of both processes. To attack such ciphers, we look for any property of the "non-linear process" that can be distinguished from random. In addition, we look for a linear combination of the linear process that vanishes. We then consider the same linear combination applied to the cipher's output, and try to find traces of the distinguishing property.

In this report we analyze two specific "distinguishing properties". One is a linear approximation of the non-linear process, which we demonstrate on the stream cipher SNOW. This attack needs roughly $2^{95}$ words of output, with work-load of about $2^{100}$. The other is a "low-diffusion" attack, that we apply to the cipher Scream-0. The latter attack needs only about $2^{43}$ bytes of output, using roughly $2^{50}$ space and $2^{80}$ time.


Key words: Hypothesis testing, Linear cryptanalysis, Linear masking, Low-Diffusion attacks, Stream ciphers.

## 1 Introduction

A stream cipher (or pseudorandom generator) is an algorithm that takes a short random string, and expands it into a much longer string, that still "looks random" to adversaries with limited resources. The short input string is called the seed (or key) of the cipher, and the long output string is called the output stream (or key-stream). Although one could get a pseudorandom generator simply by iterating a block cipher (say, in counter mode), it is believed that one could get higher speeds by using a "special purpose" stream cipher.

One approach for designing such fast ciphers, is to use some "non-linear process" that may resemble block cipher design, and to hide this process using linear masking. A plausible rationale behind this design, is that the non-linear process behaves roughly like a block cipher, so we expect its state at two "far away" points in time to be essentially uncorrelated. For close points, on the other hand, it can be argued they are masked by independent parts of the linear process, and so again they should not be correlated.

Some examples of ciphers that use this approach include SEAL [18] and Scream [11], where the non-linear process is very much like a block cipher, and the output from each step is obtained by adding together the current state of the non-linear process and some entries from fixed (or slowly
modified) secret tables. Other examples are PANAMA [3] and MUGI [21], where the linear process (called buffer) is an LFSR (Linear Feedback Shift Register), which is used as input to the non-linear process, rather than to hide the output. Yet another example is SNOW [4], where the linear LFSR is used both as input to the non-linear finite state machine, and also to hide its output.

In this work we describe a technique that can be used to distinguish such ciphers from random. The basic idea is very simple. We first concentrate on the non-linear process, looking for a characteristic that can be distinguished from random. For example, a linear approximation that has noticeable bias. We then look at the linear process, and find some linear combination of it that vanishes. If we now take the same linear combination of the output stream, then the linear process would vanish, and we are left with a sum of linear approximations, which is itself a linear approximation. As we show below, this technique is not limited to linear approximations. In some sense, it can be used with "any distinguishing characteristic" of the non-linear process. In this report we analyze in details two types of "distinguishing characteristics", and show some examples of its use for specific ciphers.

Perhaps the most obvious use of this technique, is to devise linear attacks (and indeed, many such attacks are known in the literature). This is also the easiest case to analyze. In Section 4 we characterize the statistical distance between the cipher and random as a function of the bias of the original approximation of the non-linear process, and the weight distribution of a linear code related to the linear process of the cipher.

Another type of attacks uses the low diffusion in the non-linear process. Namely, some input/output bits of this process depend only on very few other input/output bits. For this type of attacks, we again analyze the statistical distance, as a function of the number of bits in the low-diffusion characteristic. This analysis is harder than for the linear attacks. Indeed, here we do not have a complete characterization of the possible attacks of this sort, but only an analysis for the most basic such attack.

We demonstrate the usefulness of our technique by analyzing two specific ciphers. One is the cipher SNOW [4], for which we demonstrate a linear attack, and the other is the variant Scream-0 of the stream cipher Scream [11], for which we demonstrate a low-diffusion attack.

### 1.1 Relation to prior work

Linear analyses of various types are the most common tool for cryptanalyzing stream ciphers. Much work was done on LFSR-based ciphers, trying to discover the state of the LFSRs using correlation attacks (starting from Meier and Staffelbach [17], see also, e.g., [14, 13]). Golić [8, 9] devised linear models (quite similar to our model of linear attacks) that can be applied in principle to any stream cipher. He then used them to analyze many types of ciphers (including, for example, a linear distinguisher for RC4 [10]). Some examples of linear distinguishers for LFSR based ciphers, very similar to our analysis of SNOW, are [1,5], among others. Few works used also different cryptanalytical tools. Among them are the distinguishers for SEAL [12, 6] and for RC4 [7].

The main contribution of the current work is in presenting a simple framework for distinguishing attacks. This framework can be applied to many ciphers, and for those ciphers it incorporates linear analysis as a special case, but can be used to devise many other attacks, such as our "low-diffusion attacks". (Also, the attacks on SEAL due to [12] and [6] can be viewed as special cases of this framework.) For linear attacks, we believe that our explicit characterization of the statistical distance (Theorem 6) is new and useful. In addition to the cryptanalytical technique, the explicit formulation of attacks on stream ciphers, as done in Section 3, is a further contribution of this work.

Organization. In Section 2 we briefly review some background material on statistical distance and hypothesis testing. In Section 3 we formally define the framework in which our techniques apply. In Section 4 we describe how these techniques apply to linear attacks, and in Section 5 we show how they apply to low-diffusion attacks.

## 2 Elements of statistical hypothesis testing

If $\mathcal{D}$ is a distribution over some finite domain $X$ and $x$ is an element of $X$, then by $\mathcal{D}(x)$ we denote probability mass of $x$ according to $\mathcal{D}$. For notational convenience, we sometimes denote the same probability mass by $\operatorname{Pr}_{\mathcal{D}}[x]$. Similarly, if $S \subseteq X$ then $\mathcal{D}(S)=\operatorname{Pr}_{\mathcal{D}}[S]=\sum_{x \in S} \mathcal{D}(x)$.

Definition 1 (Statistical distance) Let $\mathcal{D}_{1}, \mathcal{D}_{2}$ be two distributions over some finite domain $X$. The statistical distance between $\mathcal{D}_{1}, \mathcal{D}_{2}$, is defined as

$$
\left|\mathcal{D}_{1}-\mathcal{D}_{2}\right| \stackrel{\text { def }}{=} \sum_{x \in X}\left|\mathcal{D}_{1}(x)-\mathcal{D}_{2}(x)\right|=2 \cdot \max _{S \subseteq X} \mathcal{D}_{1}(S)-\mathcal{D}_{2}(S)
$$

(We note that the statistical distance is always between 0 and 2.) In our analysis, we often view the statistical distance $\left|\mathcal{D}_{1}-\mathcal{D}_{2}\right|$ as (a scaling of) the expected value of $\left|\mathcal{D}_{1}(x)-\mathcal{D}_{2}(x)\right|$, where $x$ is chosen according to the uniform distribution. Namely, we can write

$$
\left|\mathcal{D}_{1}-\mathcal{D}_{2}\right|=|X| \cdot \sum_{x} \frac{1}{|X|} \cdot\left|\mathcal{D}_{1}(x)-\mathcal{D}_{2}(x)\right|=|X| \cdot E_{x}\left[\left|\mathcal{D}_{1}(x)-\mathcal{D}_{2}(x)\right|\right]
$$

Below are two useful facts about this measure:

- Denote by $\mathcal{D}^{N}$ the distribution which is obtained by picking independently $N$ elements $x_{1}, \ldots, x_{n} \in$ $X$ according to $\mathcal{D}$. If $\left|\mathcal{D}_{1}-\mathcal{D}_{2}\right|=\epsilon$, then to get $\left|\mathcal{D}_{1}^{N}-\mathcal{D}_{2}^{N}\right|=1$, the number $N$ needs to be between $\Omega(1 / \epsilon)$ and $O\left(1 / \epsilon^{2}\right)$. (A proof can be found, for example, in [20, Lemma 3.1.15].) In this work we sometimes make the heuristic assumption that the distributions that we consider are "smooth enough", so that we really need to set $N \approx 1 / \epsilon^{2}$.
- If $\mathcal{D}_{1}, \ldots, \mathcal{D}_{N}$ are distributions over $n$-bit strings, we denote by $\sum \mathcal{D}_{i}$ the distribution over the sum (exclusive-or), $\sum_{i=1}^{N} x_{i}$, where each $x_{i}$ is chosen according to $\mathcal{D}_{i}$, independently of all the other $x_{j}$ 's. Denote by $\mathcal{U}$ the uniform distribution over $\{0,1\}^{n}$. If for all $i,\left|\mathcal{U}-\mathcal{D}_{i}\right|=\epsilon_{i}$, then $\left|\mathcal{U}-\sum \mathcal{D}_{i}\right| \leq \prod_{i} \epsilon_{i}$. (We include a proof of this simple "xor lemma" in Section 2.1 below.) In the analysis in this paper, we sometimes assume that the distributions $\mathcal{D}_{i}$ are "smooth enough", so that we can use the approximation $\left|\mathcal{U}-\sum \mathcal{D}_{i}\right| \approx \prod_{i} \epsilon_{i}$.

Hypothesis testing. We provide a brief overview of (binary) hypothesis testing. This material is covered in many statistics and engineering textbooks (e.g., [16, Ch.5]). In a binary hypothesis testing problem, there are two distributions $\mathcal{D}_{1}, \mathcal{D}_{2}$, defined over the same domain $X$. We are given an element $x \in X$, which was drawn according to either $\mathcal{D}_{1}$ or $\mathcal{D}_{2}$, and we need to guess which is the case. A decision rule for such hypothesis testing problem is a function $D R: X \rightarrow\{1,2\}$, that tells us what should be our guess for each element $x \in X$. Perhaps the simplest notion of success for a decision rule $D R$, is the statistical advantage that it gives (over a random coin-toss), in the case that the distributions $\mathcal{D}_{1}, \mathcal{D}_{2}$ are equally likely a-priori. Namely,

$$
\operatorname{adv}(\mathrm{DR})=\frac{1}{2}\left(\underset{\mathcal{D}_{1}}{\operatorname{Pr}}[D R(x)=1]+\underset{\mathcal{D}_{2}}{\operatorname{Pr}}[D R(x)=2]\right)-\frac{1}{2}
$$

Proposition 2 For any hypothesis-testing problem $\left\langle\mathcal{D}_{1}, \mathcal{D}_{2}\right\rangle$, the decision rule with the largest advantage is the maximum-likelihood rule,

$$
M L(x)= \begin{cases}1 & \text { if } \mathcal{D}_{1}(x)>\mathcal{D}_{2}(x) \\ 2 & \text { otherwise }\end{cases}
$$

The advantage of the ML decision rule equals a quarter of the statistical distance, adv $(M L)=$ ${ }_{\frac{1}{4}}\left|\mathcal{D}_{1}-\mathcal{D}_{2}\right|$.

### 2.1 Proof of the xor-lemma for statistical distance

Lemma 3 Let $\mathcal{D}_{1}, \mathcal{D}_{2}$ be two distributions over $\{0,1\}^{k}$, let $\mathcal{D}_{3}=\mathcal{D}_{1}+\mathcal{D}_{2}$, and denote by $\mathcal{U}$ the uniform distribution over $\{0,1\}^{k}$, and $\epsilon_{i}=\left|\mathcal{U}-\mathcal{D}_{i}\right|$. Then $\epsilon_{3} \leq \epsilon_{1} \epsilon_{2}$.

Proof: For each $r, s \in\{0,1\}^{k}$, denote $e_{r}=\mathcal{D}_{1}(r)-2^{-k}$ and $f_{s}=\mathcal{D}_{2}(s)-2^{-k}$. By definition of statistical distance, we have $\epsilon_{1}=\left|\mathcal{U}-\mathcal{D}_{1}\right|=\sum_{r}\left|e_{r}\right|$ and similarly $\epsilon_{2}=\sum_{s}\left|f_{s}\right|$. For each $t \in\{0,1\}^{k}$, we then have

$$
\begin{aligned}
\mathcal{D}_{3}(t) & =\sum_{r+s=t}\left(2^{-k}+e_{r}\right)\left(2^{-k}+f_{s}\right) \\
& =2^{k} \cdot 2^{-2 k}+\sum_{r+s=t} 2^{-k}\left(e_{r}+f_{s}\right)+\sum_{r+s=t} e_{r} f_{s}=2^{-k}+\sum_{r+s=t} e_{r} f_{s}
\end{aligned}
$$

(where the last equality holds since $\sum_{r} e_{r}=\sum_{s} f_{s}=0$ ). Therefore we have

$$
\begin{aligned}
\left|\mathcal{U}-\mathcal{D}_{3}\right| & =\sum_{t}\left|\mathcal{D}_{3}(t)-2^{-k}\right|=\sum_{t}\left|\sum_{r+s=t} e_{r} f_{s}\right| \\
& \leq \sum_{t} \sum_{r+s=t}\left|e_{r} f_{s}\right|=\sum_{r, s}\left|e_{r} f_{s}\right|=\left(\sum_{r}\left|e_{r}\right|\right)\left(\sum_{s}\left|f_{s}\right|\right)=\epsilon_{1} \epsilon_{2}
\end{aligned}
$$

Corollary 4 If $\mathcal{D}_{i}, i=1 \ldots N$ are distributions with $\left|\mathcal{U}-\mathcal{D}_{i}\right|=\epsilon_{i}$, then $\left|\mathcal{U}-\sum_{i} \mathcal{D}_{i}\right| \leq \prod_{i} \epsilon_{i}$.

## 3 Formal framework

We consider ciphers that are built around two repeating functions (processes). One is a non-linear function $N F(x)$ and the other is a linear function $L F(w)$. The non-linear function $N F$ is usually a permutation on $n$-bit blocks (typically, $n \approx 100$ ). The linear function $L F$ is either an LFSR, or just fixed tables of size between a few hundred and a few thousand bits. The state of such a cipher consists of the "non-linear state" $x$ and the "linear state" $w$. In each step, we apply the function $N F$ to $x$ and the function $L F$ to $w$, and we may also "mix" these states by xor-ing some bits of $w$ into $x$ and vice versa. The output of the current state is also computed as an xor of bits from $x$ and $w$. To simplify the presentation of this report, we concentrate on a special case, similar to Scream. ${ }^{1}$ A pictorial decription of this case is shown in Figure 1. In each step $i$ we do the following:

1. Set $w_{i}:=L F\left(w_{i-1}\right)$
2. Set $y_{i}:=L 1\left(w_{i}\right), \quad z_{i}=L 2\left(w_{i}\right) \quad / / L 1, L 2$ are some linear functions
3. Set $x_{i}:=N F\left(x_{i-1}+y_{i}\right)+z_{i} \quad / /{ }^{\prime}+$ ' denotes exclusive-or
4. Output $x_{i}$

[^0]

Figure 1: A style of cipher to which our techniques apply

### 3.1 The linear process

The only property of the linear process that we care about, is that the string $y_{1} z_{1} y_{2} z_{2} \ldots$ can be modeled as a random element in some known linear subspace of $\{0,1\}^{\star}$. Perhaps the most popular linear process is to view the "linear state" $w$ as the contents of an LFSR. The linear modification function $L F$ clocks the LFSR some fixed number of times (e.g., 32 times), and the functions $L 1, L 2$ just pick some bits from the LFSR. If we denote the LFSR polynomial by $p$, then the relevant linear subspace is the subspace orthogonal to $p \cdot Z_{2}[x]$.

A different approach is taken in Scream. There, the "linear state" resides in some tables, that are "almost fixed". In particular, in Scream, each entry in these tables is used 16 times before it is modified (via the non-linear function $N F$ ). For our purposes, we model this scheme by assuming that whenever an entry is modified, it is actually being replaced by a new random value. The masking scheme in Scream can be thought of as a "two-dimensional" scheme, where there are two tables, which are used in lexicographical order. ${ }^{2}$ Namely, we have a "row table" $R[\cdot]$ and a "column table" $C[\cdot]$, each with 16 entries of $2 n$-bit string. The steps of the cipher are partitioned into batches of 256 steps each. At the beginning of a batch, all the entries in the tables are "chosen at random". Then, in step $i=j+16 k$ in a batch, we set $\left(y_{i} \mid z_{i}\right):=R[j]+C[k]$.

### 3.2 Attacks on stream ciphers

We consider an attacker that just watches the output stream and tries to distinguish it from a truly random stream. The relevant parameters in an attack are the amount of text that the attacker must see before it can reliably distinguish the cipher from random, and the time and space complexity of the distinguishing procedure. The attacks that we analyze in this report exploit the fact that for a (small) subset of the bits of $x$ and $N F(x)$, the joint distribution of these bits differs from the uniform distribution by some noticeable amount. Intuitively, such attacks never try to exploit correlations between "far away" points in time. The only correlations that are considered, are the ones between the input and output of a single application of the non-linear function. ${ }^{3}$

Formally, we view the non-linear process not as one continuous process, but rather as a sequence of uncorrelated steps. That is, for the purpose of the attack, one can view the non-linear state $x$ at the beginning of each step as a new random value, independent of anything else. Under this view,

[^1]the attacker sees a collection of pairs $\left\langle x_{j}+y_{j}, N F\left(x_{j}\right)+z_{j}\right\rangle$, where the $x_{j}$ 's are chosen uniformly at random and independently of each other, and the $y_{j}, z_{j}$ 's are taken from the linear process.

One example of attacks that fits in this model are linear attacks. In linear cryptanalysis, the attacker exploits the fact that a one-bit linear combination of $\langle x, N F(x)\rangle$ is more likely to be zero than one (or vice versa). In these attack, it is always assumed that the bias in one step is independent of the bias in all the other steps. Somewhat surprisingly, differential cryptanalysis too fits into this framework (under our attack model). Since the attacker in our model is not given chosen-input capabilities, it exploits differential properties of the round function by waiting for the difference $x_{i}+x_{j}=\Delta$ to happen "by chance", and then using the fact that $N F\left(x_{i}\right)+N F\left(x_{j}\right)=\Delta^{\prime}$ is more likely than you would expect from a random process. It is clear that this attack too is just as effective against pairs of uncorrelated steps, as when given the output from the real cipher.

We are now ready to define formally what we mean by "an attack on the cipher". The attacks that we consider, observe some (linear combinations of) input and output bits from each step of the cipher, and try to decide if these indeed come from the cipher, or from a random source. This can be framed as a hypothesis testing problem. According to one hypothesis (Random), the observed bits in each step are random and independent. According to the other (Cipher), they are generated by the cipher.

Definition 5 (Attacks on stream ciphers with linear masking) An attack is specified by a linear function $\ell$, and by a decision rule for the following hypothesis-testing problem: The two distributions that we want to distinguish are

Cipher. The Cipher distribution is $\mathcal{D}_{c}=\left\langle\ell\left(x_{j}+y_{j}, N F\left(x_{j}\right)+z_{j}\right)\right\rangle_{j=1,2, \ldots}$, where the $y_{j} z_{j}$ 's are chosen at random from the appropriate linear subspace (defined by the linear process of the cipher), and the $x_{j}$ 's are random and independent.

Random. Using the same notations, the "random process" distribution is $\mathcal{D}_{r} \stackrel{\text { def }}{=}\left\langle\ell\left(x_{j}, x_{j}^{\prime}\right)\right\rangle_{j=1,2, \ldots}$, where the $x_{j}$ 's and $x_{j}^{\prime}$ 's are random and independent.
We call the function $\ell$, the distinguishing characteristic used by attack.
The amount of text needed for the attack is the smallest number of steps for which the decision rule has a constant advantage (e.g., advantage of $1 / 4$ ) in distinguishing the cipher from random. Other relevant parameters of the attack are the time and space complexity of the decision rule. An obvious lower bound on the amount of text is provided by the statistical distance between the Cipher and Random distributions after $N$ steps.

## 4 Linear attacks

A linear attack [15] exploits the fact that some linear combination of the input and output bits of the non-linear function is more likely to be zero than one (or vice versa). Namely, we have a (non-trivial) linear function $\ell:\{0,1\}^{2 n} \rightarrow\{0,1\}$, such that for a randomly selected $n$ bit string $x$, $\operatorname{Pr}[\ell(x, N F(x))=0]=(1+\epsilon) / 2$. The function $\ell$ is called a linear approximation (or characteristic) of the non-linear function, and the quantity $\epsilon$ is called the bias of the approximation.

When trying to exploit one such linear approximation, the attacker observes for each step $j$ of the cipher, a bit $\sigma_{j}=\ell\left(x_{j}+y_{j}, N F\left(x_{j}\right)+z_{j}\right)$. Note that $\sigma_{j}$ by itself is likely to be unbiased, but the $\sigma$ 's are correlated. In particular, since the $y, z$ 's come from a linear subspace, it is possible
to find some linear combination of steps for which they vanish. Let $J$ be a set of steps such that $\sum_{j \in J} y_{j}=\sum_{j \in J} z_{j}=0$. Then we have

$$
\sum_{j \in J} \sigma_{j}=\sum_{j \in J} \ell\left(x_{j}, N F\left(x_{j}\right)\right)+\sum_{j \in J} \ell\left(y_{j}, z_{j}\right)=\sum_{j \in J} \ell\left(x_{j}, N F\left(x_{j}\right)\right)
$$

(where the equalities follow since $\ell$ is linear). Therefore, the bit $\xi_{J}=\sum_{j \in J} \sigma_{j}$ has bias of $\epsilon^{|J|}$. If the attacker can observe "sufficiently many" such sets $J$, it can reliably distinguish the cipher from random.

This section is organized as follows: We first bound the effectiveness of linear attacks in terms of the bias $\epsilon$ and the weight distribution of some linear subspace. As we explain below, this bound suggests that looking at sets of steps as above is essentially "the only way to exploit linear correlations". Then we show how to devise a linear attack on SNOW, and analyze its effectiveness.

### 4.1 The statistical distance

Recall that we model an attack in which the attacker observes a single bit per step, namely $\sigma_{j}=$ $\ell\left(x_{j}+y_{j}, N F\left(x_{j}\right)+z_{j}\right)$. Below we denote $\tau_{j}=\ell\left(x_{j}, N F\left(x_{j}\right)\right)$ and $\rho_{j}=\ell\left(y_{j}, z_{j}\right)$. We can re-write the Cipher and Random distributions as

Cipher. $\mathcal{D}_{c} \stackrel{\text { def }}{=}\left\langle\tau_{j}+\rho_{j}\right\rangle_{j=1,2, \ldots}$, where the $\tau_{j}$ 's are independent but biased, $\operatorname{Pr}\left[\tau_{j}=0\right]=(1+\epsilon) / 2$, and the string $\rho_{1} \rho_{2} \ldots$ is chosen at random from the appropriate linear subspace (i.e., the image under $\ell$ of the linear subspace of the $y_{j} z_{j}$ 's).

Random. $\mathcal{D}_{r} \stackrel{\text { def }}{=}\left\langle\sigma_{j}\right\rangle_{j=1,2, \ldots}$, where the $\sigma_{j}$ 's are independent and unbiased.
Below we analyze the statistical distance between the Cipher and Random distributions, after observing $N$ bits $\sigma_{1} \ldots \sigma_{N}$. Denote the linear subspace of the $\rho^{\prime}$ 's by $L \subseteq\{0,1\}^{N}$, and let $L^{\perp} \subseteq\{0,1\}^{N}$ be the orthogonal subspace. The weight distribution of the space $L^{\perp}$ plays an important role in our analysis. For $r \in\{0,1, \ldots, N\}$, let $\mathcal{A}_{N}(r)$ be the set of strings $\vec{\chi} \in L^{\perp}$ of Hamming weight $r$, and let $A_{N}(r)$ denote the cardinality of $\mathcal{A}_{N}(r)$. We prove the following theorem:

Theorem 6 The statistical distance between the Cipher and Random distributions from above, is bounded by $\sqrt{\sum_{r=1}^{N} A_{N}(r) \epsilon^{2 r}}$.
Proof: Recall that the statistical distance |Cipher-Random| (for $N$ observed bits) can be expressed in terms of the expected value of $\left|\operatorname{Pr}_{\text {Cipher }}[\vec{\sigma}]-\operatorname{Pr}_{\text {Random }}[\vec{\sigma}]\right|$, where $\vec{\sigma}$ is chosen uniformly at random from $\{0,1\}^{N}$. Fix a string $\vec{\sigma} \in\{0,1\}^{N}$, and we want to analyze the probability $\operatorname{Pr}_{\text {Cipher }}[\vec{\sigma}]$. That probability is

$$
\underset{\text { Cipher }}{\operatorname{Pr}}[\vec{\sigma}]=\sum_{\vec{\rho} \in L} \frac{1}{|L|} \cdot \prod_{j=1}^{N}\left(\frac{1}{2}+\frac{\epsilon}{2} \cdot \operatorname{sign}\left(\rho_{i}+\sigma_{i}\right)\right)
$$

where the sign indicator is taken to be $(+1)$ if $\rho_{i}=\sigma_{i}$, and $(-1)$ otherwise. In other words, $\operatorname{sign}(x) \stackrel{\text { def }}{=}(-1)^{x}$. We can break the expression above into a power series in $\epsilon$. In this power series, the constant term is $2^{-N}$, and the series looks as follows $\operatorname{Pr}_{\text {Cipher }}[\vec{\sigma}]=2^{-N}\left(1+\sum_{r=1}^{N} \epsilon^{r}\right.$ coef $\left._{r}\right)$, where the coefficients coef ${ }_{r}$ are defined as

$$
\operatorname{coef}_{r} \stackrel{\text { def }}{=} \sum_{\vec{p} \in L} \frac{1}{|L|} \cdot \sum_{\left\{j_{1} \ldots j_{r}\right\}} \prod_{t=1}^{r} \operatorname{sign}\left(\sigma_{j_{t}}+\rho_{j_{t}}\right)=\sum_{\left\{j_{1} \ldots j_{r}\right\}} \frac{1}{|L|} \cdot \sum_{\vec{p} \in L} \operatorname{sign}\left(\sum_{t=1}^{r} \sigma_{j_{t}}+\rho_{j_{t}}\right)
$$

The summation over $\left\{j_{1} \ldots j_{r}\right\}$ in the expression above ranges over all ordered sets of cardinality $r$ in $[1, N]$ (i.e., $1 \leq j_{1}<j_{2} \cdots<j_{r} \leq N$.) Consider one such $r$-set $J=\left\{j_{1} \ldots j_{r}\right\}$, and we analyze its contribution to the total sum. Let $\chi(J)$ be the characteristic vector of this set. That is, $\chi(J)$ is an $N$-bit string, which is 1 in bit positions $\left\{j_{1} \ldots j_{r}\right\}$, and 0 everywhere else.

Proposition 7 Let $J=\left\{j_{1} \ldots j_{r}\right\}$ be a set of cardinality $r$. If $\chi(J) \notin L^{\perp}$, then the total contribution to $\operatorname{coef}_{r}$ due to the set $J$ is zero. If $\chi(J) \in L^{\perp}$ then the total contribution to $\operatorname{coef}_{r}$ due to the set $J$ is $\operatorname{sign}\left(\sum_{j \in J} \sigma_{j}\right)$.

Proof: If $\vec{\chi}=\chi(J)$ is not in $L^{\perp}$, then for exactly half of the strings $\vec{\rho} \in L$ it holds that $\sum_{j \in J} \rho_{j}=$ $\langle\vec{\chi}, \vec{\rho}\rangle=0$. Thus, for exactly half of the strings $\vec{\rho} \in L$ we have $\operatorname{sign}\left(\sum_{t=1}^{r} \sigma_{j_{t}}+\rho_{j_{t}}\right)=+1$, and for the other half we have $\operatorname{sign}\left(\sum_{t=1}^{r} \sigma_{j_{t}}+\rho_{j_{t}}\right)=-1$, so $\sum_{\vec{\rho} \in L} \operatorname{sign}\left(\sum_{t=1}^{r} \sigma_{j_{t}}+\rho_{j_{t}}\right)=0$. If $\chi(J) \in L^{\perp}$, then for all $\vec{\rho} \in L$ we have $\sum_{t=1}^{r} \rho_{j_{t}}=0$, and therefore $\operatorname{sign}\left(\sum_{t=1}^{r} \sigma_{j_{t}}+\rho_{j_{t}}\right)=\operatorname{sign}\left(\sum_{t=1}^{r} \sigma_{j_{t}}\right)$. Thus, we get $\frac{1}{|L|} \cdot \sum_{\vec{\rho} \in L} \operatorname{sign}\left(\sum_{t=1}^{r} \sigma_{j_{t}}+\rho_{j_{t}}\right)=\operatorname{sign}\left(\sum_{t=1}^{r} \sigma_{j_{t}}\right)$.

We now view the terms in the power series above as random variables. For any set $J$ with $\chi(J) \in L^{\perp}$, denote $\xi_{J}(\vec{\sigma}) \stackrel{\text { def }}{=} \operatorname{sign}\left(\sum_{j \in J} \sigma_{j}\right)$, and we view the $\xi_{J}$ 'es as random variables, which are defined over the choice of $\vec{\sigma}$ uniformly at random in $\{0,1\}^{N}$. Then, we define the normalized probability difference

$$
\Delta(\vec{\sigma}) \stackrel{\text { def }}{=} 2^{N} \cdot\left(\operatorname{Pr}_{\text {Cipher }}[\vec{\sigma}]-\operatorname{Pr}_{\text {Random }}[\vec{\sigma}]\right)=\sum_{r=1}^{N} \epsilon^{r} \sum_{\chi(J) \in \mathcal{A}_{N}(r)} \xi_{J}(\vec{\sigma})
$$

Again, we stress that we view $\Delta(\vec{\sigma})$ as a random variable over the uniform choice of $\vec{\sigma} \in\{0,1\}^{N}$. It is easy to see that for any non-empty $J$, we have $E\left[\xi_{J}\right]=0$ and $\operatorname{VAR}\left[\xi_{J}\right]=1$. Also, if $J_{1} \neq J_{2}$, then $\xi_{J_{1}}, \xi_{J_{2}}$ are independent. Therefore, the variable $\Delta$ has zero mean, and its variance equals the weighted sum of the $\xi_{J}$ variances. Namely, $\operatorname{VAR}[\Delta]=\sum_{r=1}^{N} A_{N}(r) \epsilon^{2 r}$. We can now write the statistical distance between the Cipher and Random distributions as

$$
\mid \text { Cipher - Random }\left|=\sum_{\vec{\sigma}}\right| \operatorname{Pr}_{\text {Cipher }}[\vec{\sigma}]-\operatorname{Pr}_{\text {Random }}[\vec{\sigma}]\left|=\sum_{\vec{\sigma}} 2^{-N}\right| \Delta(\vec{\sigma}) \mid=E_{\vec{\sigma}}[|\Delta|]
$$

By the convexity of the squaring function, $E[|\Delta|] \leq \sqrt{\operatorname{VAR}[\Delta]}$, and therefore

$$
\begin{equation*}
\mid \text { Cipher - Random } \mid=E_{\vec{\sigma}}[|\Delta|] \leq \sqrt{\operatorname{VAR}[\Delta]}=\sqrt{\sum_{r=1}^{N} A_{N}(r) \epsilon^{2 r}} \tag{1}
\end{equation*}
$$

Remark. Heuristically, this bound is nearly tight. In the proof we analyzed the random variable $\Delta$ and used the bound $E[|\Delta-E[\Delta]|] \leq \sqrt{\operatorname{VAR}[\Delta]}$. One can argue heuristically that as long as the statistical distance is sufficiently small, " $\Delta$ should behave much like a Gaussian random variable". If it were a Gaussian, we would have $E[|\Delta|]=\sqrt{\operatorname{VAR}[\Delta]} \cdot \sqrt{2 / \pi}$. Thus, we expect the bound from Theorem 6 to be tight up to a constant factor $\sqrt{2 / \pi} \approx 0.8$.

### 4.2 Interpretations of Theorem 6

There are a few ways to view Theorem 6. The obvious way is to use it in order to argue that a certain cipher is resilient to linear attacks. For example, in [11] we use Theorem 6 to deduce a lower-bound on the amount of text needed for any linear attack on Scream-0.

Also, one could notice that the form of Theorem 6 exactly matches the common practice (and intuition) of devising linear attacks. Namely, we always look at sets where the linear process vanishes, and view each such set $J$ as providing "statistical evidence of weight $\epsilon^{2|J| "}$ for distinguishing the cipher from random. Linear attacks work by collecting enough of these sets, until the weights sum up of to one. One can therefore view Theorem 6 as asserting that this is indeed the best you can do.

Finally, we could think of devising linear attacks, using the heuristic argument about this bound being tight. However, the way Theorem 6 is stated above, it usually does not imply efficient attacks. For example, when the linear space $L$ has relatively small dimension (as is usually the case with LFSR-based ciphers, where the dimension of $L$ is at most a few hundreds), the statistical distance is likely to approach one for relatively small $N$. But it is likely that most of the "mass" in the power series in Eq. (1) comes from terms with a large power of $\epsilon$ (and therefore very small "weight"). Therefore, if we want to use a small $N$, we would need to collect very many samples, and this attack is likely to be more expensive than an exhaustive search for the key.

Alternatively, one can try and use an efficient sub-optimal decision rule. For a given bound on the work-load $W$ and the amount of text $N$, we only consider the first few terms in the power series. That is, we observe the $N$ bits $\vec{\sigma}=\sigma_{1} \ldots \sigma_{N}$, but only consider the $W$ smallest sets $J$ for which $\vec{\chi}(J) \in L^{\perp}$. For each such set $J$, the sum of steps $\sum_{j \in J} \sigma_{j}$ has bias $\epsilon^{|J|}$, and these can be used to distinguish the cipher from random. If we take all the sets of size at most $R$, we expect the advantage of such a decision rule to be roughly $\frac{1}{4} \sqrt{\sum_{r=1}^{R} A_{N}(r) \epsilon^{2 r}}$. The simplest form of this attack (which is almost always the most useful), is to consider only the minimum-weight terms. If the minimum-weight of $L^{\perp}$ is $r_{0}$, then we need to make $N$ big enough so that $\frac{1}{4} \sqrt{A_{N}\left(r_{0}\right)}=\epsilon^{-r_{0}}$.

### 4.3 The attack on SNOW

The stream cipher SNOW was submitted to NESSIE in 2000, by Ekdahl and Johansson. A detailed description of SNOW is available from [4]. Here we outline a linear attack on SNOW along the lines above, that can reliably distinguish it from random after observing roughly $2^{95}$ steps of the cipher, with work-load of roughly $2^{100}$.

SNOW consists of a non-linear process (called there a Finite-State Machine, or FSM), and a linear process which is implemented by an LFSR. The LFSR of SNOW consists of sixteen 32-bit words, and the LFSR polynomial, defined over $G F\left(2^{32}\right)$, is $p(z)=z^{16}+z^{13}+z^{7}+\alpha$, where $\alpha$ is a primitive element of $G F\left(2^{32}\right)$. (The orthogonal subspace $L^{\perp}$ is therefore the space of (bitwise reversal of) polynomials over $Z_{2}$ of degree $\leq N$, which are divisible by the LFSR polynomial $p$.) At a given step $j$, we denote the content of the LFSR by $L_{j}[0 . .15]$, so we have $L_{j+1}[i]=L_{j}[i-1]$ for $i>0$ and $L_{j+1}[0]=\alpha \cdot\left(L_{j}[15]+L_{j}[12]+L_{j}[6]\right)$.

The "FSM state" of SNOW in step $j$ consists of only two 32 -bit words, denoted $R 1_{j}, R 2_{j}$. The FSM update function modifies these two values, using one word from the LFSR, and also outputs one word. The output word is then added to another word from the LFSR, to form the step output. We denote the "input word" from the LFSR to the FSM update function by $f_{j}$, and the "output word" from the FSM by $F_{j}$. The FSM uses a " $32 \times 32 \mathrm{~S}$-box" $S[\cdot]$ (which is built internally as an SP-network, from four identical $8 \times 8$ boxes and some bit permutation). A complete step of SNOW is described in Figure 2. In this figure, we deviate from the notations in the rest of the paper, and denote exclusive-or by $\oplus$ and integer addition $\bmod 2^{32}$ by + . We also denote 32 -bit cyclic rotation to the left by $\ll$.

To devise an attack we need to find a good linear approximation of the non-linear FSM process, and low-weight combinations of steps where the $L_{j}[\cdot]$ values vanish (i.e., low-weight polynomials

1. $f_{j}:=L_{j}[0]$
2. $F_{j}:=\left(f_{j}+R 1_{j}\right) \oplus R 2_{j}$
3. output $F_{j} \oplus L_{j}[15]$
4. $R 1_{j+1}:=R 1_{j} \oplus\left(\left(R 2_{j}+F_{j}\right) \lll 7\right)$
5. $R 2_{j+1}:=S\left[R 1_{j}\right]$
6. update the LFSR

Figure 2: One step of SNOW: $\oplus$ is xor and + is addition $\bmod 2^{32}$.
which are divisible by the LFSR polynomial $p$ ). The best linear approximation that we found for the FSM process, uses six bits from two consecutive inputs and outputs, $f_{j}, f_{j+1}, F_{j}, F_{j+1}$. Specifically, for each step $j$, the bit

$$
\sigma_{j} \xlongequal{\text { def }}\left(f_{j}\right)_{15}+\left(f_{j}\right)_{16}+\left(f_{j+1}\right)_{22}+\left(f_{j+1}\right)_{23}+\left(F_{j}\right)_{15}+\left(F_{j+1}\right)_{23}
$$

is biased. (Of these six bits, the bits $\left(f_{j}\right)_{15},\left(F_{j}\right)_{15}$ and $\left(F_{j+1}\right)_{22}$ are meant to approximate carry bits.) We measured the bias experimentally, and it appears to be at least $2^{-8.3}$.

At first glance, one may hope to find weight-4 polynomials that are divisible by the LFSR polynomial $p$. After all, $p$ itself has only four non-zero terms. Unfortunately, one of these terms is the element $\alpha \in G F\left(2^{32}\right)$, whereas we need a low-weight polynomial with $0-1$ coefficients. What we can show, however, is the existence of 0-1 polynomials of weight-six that are divisible by $p$.

Proposition 8 The polynomial $q(z)=z^{16 \times 2^{32}-7}+z^{13 \times 2^{32}-7}+z^{7 \times 2^{32}-7}+z^{9}+z^{6}+1$ is divisible by the LFSR polynomial $p(z)=z^{16}+z^{13}+z^{7}+\alpha$.

Proof: Since $\alpha \in G F\left(2^{32}\right)$, then the polynomial $t+\alpha$ divides $t^{232}+t$. That is, there is a polynomial $r(\cdot)$ (with coefficients in $G F\left(2^{32}\right)$ ) such that $r(t) \cdot(t+\alpha)=t^{2^{32}}+t$, as formal polynomials over $G F\left(2^{32}\right)$. It follows that for any polynomial $t(z)$ over $G F\left(2^{32}\right)$, we have $r(t(z)) \cdot(t(z)+\alpha)=$ $t(z)^{2^{32}}+t(z)$, again, as formal polynomials over $G F\left(2^{32}\right)$. Specifically, if we take $t(z)=z^{16}+z^{13}+z^{7}$, we get

$$
r(t(z)) \cdot\left(z^{16}+z^{13}+z^{7}+\alpha\right)=z^{16 \times 2^{32}}+z^{13 \times 2^{32}}+z^{7 \times 2^{32}}+z^{16}+z^{13}+z^{7}
$$

so the polynomial on the right hand side is divisible by $p(z)$. Since $p(z)$ is co-prime with the polynomial $z$, we can divide the right-hand-side polynomial by $z^{7}$ and still get a polynomial divisible by $p(z)$.

Corollary 9 For all $m, n$, the polynomial

$$
\begin{aligned}
q_{m, n}(z) \stackrel{\text { def }}{=} q(z)^{2^{m}} \cdot z^{n}= & z^{16 \times 2^{32+m}-7 \times 2^{m}+n}+z^{13 \times 2^{32+m}-7 \times 2^{m}+n} \\
& +z^{7 \times 2^{32+m}-7 \times 2^{m}+n}+z^{9 \times 2^{m}+n}+z^{6 \times 2^{m}+n}+z^{n}
\end{aligned}
$$

is divisible by $p(z)$.
If we take, say, $m=0,1, \ldots 58$ and $n=0,1, \ldots 2^{94}$, we get about $2^{100}$ different $0-1$ polynomials, all with weight 6 and degree less than $N=2^{95}$, and all divisible by $p(z)$. Each such polynomial yields a sequence of six steps, $J_{m, n}$, such that the sum of the $L_{j}[\cdot]$ values in these steps vanishes. Specifically, the polynomial $q_{m, n}(z)$ corresponds to the sequence of steps

$$
J_{m, n}=\left\{\begin{array}{ll} 
& N-n-16 \cdot 2^{32+m}+7 \cdot 2^{m}, \\
& N-n-9 \cdot 2^{m}, \\
& N-n-13 \cdot 2^{32+m}+7 \cdot 2^{m}, \\
& N-n-6 \cdot 2^{m}, \\
& N-n-7 \cdot 2^{32+m}+7 \cdot 2^{m},
\end{array} \quad N-n, ~ \$\right.
$$

with the property that for all $m, n, \quad \sum_{j \in J_{m, n}} L_{j}[0 . .15]=\sum_{j \in J_{m, n}} L_{j+1}[0 . .15]=[0,0, \ldots, 0]$. Therefore, if we denote the output word of SNOW at step $j$ by $S_{j}$, then for all $m, n$ we have,

$$
\tau_{m, n} \stackrel{\text { def }}{=} \sum_{j \in J_{m, n}}\left(S_{j}\right)_{15}+\left(S_{j+1}\right)_{23}=\sum_{j \in J_{m, n}} \sigma_{j}
$$

and therefore each $\tau_{m, n}$ has bias of $2^{-8.3 \times 6}=2^{-49.8}$. Since we have roughly $2^{100}$ of them, we can reliably distinguish them from random.

## 5 Low-diffusion attacks

In low-diffusion attacks, the attacker looks for a small set of (linear combinations of) input and output bits of the non-linear function $N F$, whose values completely determine the values of some other (linear combinations of) input and output bits. The attacker tries to guess the first set of bits, computes the values of the other bits, and uses the computed value to verify the guess against the cipher's output. The complexity of such attacks is exponential in the number of bits that the attacker needs to guess.

We introduce some notations in order to put such attacks in the context of our framework. To simplify the notations, we assume that the guessed bits are always input bits, and the determined bits are always output bits. (Eliminating this assumption is usually quite straightforward.) As usual, let $N F:\{0,1\}^{n} \rightarrow\{0,1\}^{n}$ be the non-linear function. The attack exploits the fact that some input bits $\ell_{\text {in }}(x)$ are related to some output bits $\ell_{\text {out }}(N F(x))$ via a known deterministic function $f$. That is, we have

$$
\ell_{\text {out }}(N F(x))=f\left(\ell_{\text {in }}(x)\right)
$$

Here, $\ell_{\text {in }}, \ell_{\text {out }}$ are linear functions, and $f$ is an arbitrary function, all known to the attacker. We denote the output size of $\ell_{\text {in }}, \ell_{\text {out }}$ by $m, m^{\prime}$, respectively. That is, $\ell_{\text {in }}:\{0,1\}^{n} \rightarrow\{0,1\}^{m}$, $\ell_{\text {out }}:$ $\{0,1\}^{n} \rightarrow\{0,1\}^{m^{\prime}}$, and $f:\{0,1\}^{m} \rightarrow\{0,1\}^{m^{\prime}}$.

In each step $j$, the attacker observes the bits $\ell_{\text {in }}\left(x_{j}+y_{j}\right)$ and $\ell_{\text {out }}\left(N F\left(x_{j}\right)+z_{j}\right.$ ) (where $y_{j}, z_{j}$ are from the linear process, as in Section 3.1). Below we denote $u_{j}=\ell_{\text {in }}\left(x_{j}\right), u_{j}^{\prime}=\ell_{\text {out }}\left(N F\left(x_{j}\right)\right)$, $v_{j}=\ell_{\text {in }}\left(y_{j}\right), v_{j}^{\prime}=\ell_{\text {out }}\left(z_{j}\right)$, and $w_{j}=u_{j}+v_{j}, w_{j}^{\prime}=u_{j}^{\prime}+v_{j}^{\prime}$. We can re-write the Cipher and Random distributions for this case as

Cipher. $\mathcal{D}_{c} \stackrel{\text { def }}{=}\left\langle\left(w_{j}=u_{j}+v_{j}, w_{j}^{\prime}=u_{j}^{\prime}+v_{j}^{\prime}\right)\right\rangle_{j=1,2, \ldots . .}$, where the $u_{j}$ 's are uniform and independent, $u_{j}^{\prime}=f\left(u_{j}\right)$, and the string $v_{1} v_{1}^{\prime} v_{2} v_{2}^{\prime} \ldots$ is chosen at random from the appropriate linear subspace (i.e., the image under $\ell_{\text {in }}, \ell_{\text {out }}$ of the linear subspace of the $y, z$ 's).
Random. $\mathcal{D}_{r} \stackrel{\text { def }}{=}\left\langle\left(w_{j}, w_{j}^{\prime}\right)\right\rangle_{j=1,2, \ldots}$, all uniform and independent.
It is not hard to see that there may be enough information there to distinguish these two distributions after only a moderate number of steps of the cipher. Suppose that the dimension of the linear subspace of the $v_{j}$ 's and $v_{j}^{\prime}$ 's is $a$, and the attacker observes $N$ steps such that $m^{\prime} N>a$. Then, the attacker can (in principle) go over all the $2^{a}$ possibilities for the $v_{j}$ 's and $v_{j}^{\prime}$ 's. For each guess, the attacker can compute the $u_{j}$ 's and $u_{j}^{\prime}$ 's, and verify the guess by checking that $u_{j}^{\prime}=f\left(u_{j}\right)$ for all $j$. This way, the attacker guesses $a$ bits and gets $m^{\prime} N$ bits of consistency checks. Since $m^{\prime} N>a$ we expect only the "right guess" to pass the consistency checks.

This attack, however, is clearly not efficient. To devise an efficient attack, we can again concentrate on sets of steps where the linear process vanishes: Suppose that we have a set of steps $J$,
such that $\sum_{j \in J}\left[v_{j}, v_{j}^{\prime}\right]=[0,0]$. Then we get

$$
\sum_{j \in J}\left(w_{j}, w_{j}^{\prime}\right)=\sum_{j \in J}\left(u_{j}, u_{j}^{\prime}\right)=\sum_{j \in J}\left(u_{j}, f\left(u_{j}\right)\right)
$$

and the distribution over such pairs may differ from the uniform distribution by a noticeable amount. The distance between this distribution and the uniform one, depends on the specific function $f$, and on the cardinality of the set $J .{ }^{4}$ Below we analyze in details perhaps the simplest cases, where $f$ is a random function. Later we explain how this analysis can be extended for other settings, and in particular for the case of the functions in Scream.

### 5.1 Analysis for random functions

For a given function, $f:\{0,1\}^{m} \rightarrow\{0,1\}^{m^{\prime}}$, and an integer $n$, we denote

$$
\mathcal{D}_{f}^{n} \stackrel{\text { def }}{=}\left\langle d=\sum_{j=1}^{n} u_{j}, d^{\prime}=\sum_{j=1}^{n} f\left(u_{j}\right)\right\rangle
$$

where the $u_{j}$ 's are uniform in $\{0,1\}^{m}$ and independent. We assume that the attacker knows $f$, and it sees many instances of $\left\langle d, d^{\prime}\right\rangle$. The attacker needs to decide if these instances come from $\mathcal{D}_{f}^{n}$ or from the uniform distribution on $\{0,1\}^{m+m^{\prime}}$. Below we denote the uniform distribution by $\mathcal{R}$. If the function $f$ "does not have any clear structure", it makes sense to analyze it as if it was a random function. Here we prove the following:

Theorem 10 Let $n, m, m^{\prime}$ be integers with $n^{2} \ll 2^{m}$. ${ }^{5}$ For a uniformly selected function $f$ : $\{0,1\}^{m} \rightarrow\{0,1\}^{m^{\prime}}, E_{f}\left[\left|\mathcal{D}_{f}^{n}-\mathcal{R}\right|\right] \leq c(n) \cdot 2^{\frac{m^{\prime}-(n-1) m}{2}}$, where

$$
c(n)= \begin{cases}\sqrt{(2 n)!/\left(n!2^{n}\right)} & \text { if } n \text { is odd } \\ (1+o(1)) \sqrt{\frac{(2 n)!}{n!2^{n}}-\left(\frac{n!}{(n / 2)!2^{n / 2}}\right)^{2}} & \text { if } n \text { is even }\end{cases}
$$

Proof: Fix $n, m, m^{\prime}$. For the rest of the proof, these integers will always be implicit (for example, we write $\mathcal{D}_{f}$ instead of $\mathcal{D}_{f}^{n}$, etc.). Recall that we denote the probability mass of ( $d, d^{\prime}$ ) according to $\mathcal{D}_{f}$ by $\mathcal{D}_{f}\left(d, d^{\prime}\right)$. We can express the expected value of $\left|\mathcal{D}_{f}-\mathcal{R}\right|$, where $f$ is chosen at random, as:

$$
\begin{align*}
& E_{f}\left[\left|\mathcal{D}_{f}-\mathcal{R}\right|\right]  \tag{2}\\
& =E_{f}\left[\sum_{d, d^{\prime}}\left[\left|\mathcal{D}_{f}\left(d, d^{\prime}\right)-2^{-m-m^{\prime}}\right|\right]\right]=2^{m^{\prime}} \sum_{d} E_{f, d^{\prime}}\left[\left|\mathcal{D}_{f}\left(d, d^{\prime}\right)-2^{-m-m^{\prime}}\right|\right]
\end{align*}
$$

In the last term, we view the $\mathcal{D}_{f}\left(d, d^{\prime}\right)$ 's as random variables over the choice of $f, d^{\prime}$, and we have $2^{m}$ such variables, one for each $d$. Some properties of these random variables are summarized in the following proposition (which is proved later).

[^2]Proposition 11 For any fixed $d$, $E_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)\right]=2^{-m-m^{\prime}}$. Also, for odd $n$ and any fixed $d$, $V A R_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)\right] \leq 2^{-m(n+1)-m^{\prime}} \cdot \frac{(2 n)!}{n!2^{n}}$, and for even $n$ and fixed $d$, $V A R_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)\right] \leq$

$$
\begin{cases}2^{-m n-m^{\prime}}\left(\left(\frac{n!}{(n / 2)!2^{n / 2}}\right)^{2}+o(1)\left(\frac{(2 n)!}{n!2^{n}}\right)\right) & \text { for } d=0 \\ 2^{-m(n+1)-m^{\prime}}(1+o(1))\left(\frac{(2 n)!}{n!2^{n}}-\left(\frac{n!}{(n / 2)!2^{n / 2}}\right)^{2}\right) & \text { for } d \neq 0\end{cases}
$$

We can now continue Eq. (2) as follows:

$$
\begin{align*}
E_{f}\left[\left|\mathcal{D}_{f}-\mathcal{R}\right|\right] & =2^{m^{\prime}} \sum_{d} E_{f, d^{\prime}}\left[\left|\mathcal{D}_{f}\left(d, d^{\prime}\right)-2^{-m-m^{\prime}}\right|\right] \\
& =2^{m^{\prime}} \sum_{d} E_{f, d^{\prime}}\left[\left|\mathcal{D}_{f}\left(d, d^{\prime}\right)-E_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)\right]\right|\right] \\
& \stackrel{(a)}{\leq} 2^{m^{\prime}} \sum_{d} \sqrt{\operatorname{VAR}_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)\right]} \tag{3}
\end{align*}
$$

where (a) follows since for any random variable $X, E[|X-E[X]|] \leq \sqrt{\operatorname{VAR}[X]}$. Plugging the variance bounds from Proposition 11 completes the proof of Theorem 10.

How tight is this bound? Here too we can argue heuristically that the random variables in the proof "should behave like Gaussian random variables", and again we expect the ratio between $E[|X-E[X]|]$ and $\sqrt{\operatorname{VAR}[X]}$ to be roughly $\sqrt{2 / \pi}$. Therefore, we expect the constant $c(n)$ to be replaced by $\sqrt{2 / \pi} \cdot c(n) \approx 0.8 c(n)$. Indeed we ran some experiments to measure the statistical distance $\left|\mathcal{D}_{f}^{n}-\mathcal{R}\right|$, for random functions with $n=4$ and a few values of $m, m^{\prime}$. (Note that $c(4)=(1+o(1)) \sqrt{96} \approx 9.8$ and $\sqrt{2 / \pi} \cdot c(4) \approx 7.8)$. These experiments are described in Appendix A. The results confirm that the distance between these distributions is just under $7.8 \cdot 2^{\left(m^{\prime}-3 m\right) / 2}$.
Proof: (of Proposition 11) We recall that for a given $f, d, d^{\prime}$, the term $\mathcal{D}_{f}\left(d, d^{\prime}\right)$ is defined as $\mathcal{D}_{f}\left(d, d^{\prime}\right)=\operatorname{Pr}_{\vec{u}}\left[\sum u_{i}=d, \quad \sum f\left(u_{i}\right)=d^{\prime}\right]$, where the probability is taken over the choice of $\vec{u}=u_{1} \ldots u_{n}$, uniformly at random in $\{0,1\}^{m n}$. Analyzing the expected value of $\mathcal{D}_{f}\left(d, d^{\prime}\right)$ is straightforward. For any fixed $d$, we have

$$
\begin{aligned}
E_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)\right] & =E_{f, d^{\prime}}\left[\underset{\vec{u}}{\operatorname{Pr}}\left[\sum u_{i}=d, \sum f\left(u_{i}\right)=d^{\prime}\right]\right] \\
& =\operatorname{Pr}_{f, d^{\prime}, \vec{u}}\left[\sum u_{i}=d, \sum f\left(u_{i}\right)=d^{\prime}\right] \\
& =\underset{\vec{u}}{\operatorname{Pr}}\left[\sum u_{i}=d\right] \cdot \underset{d^{\prime} \ldots}{\operatorname{Pr}}\left[d^{\prime}=\sum f\left(u_{i}\right) \mid \cdots\right]=2^{-m} \cdot 2^{-m^{\prime}}
\end{aligned}
$$

To analyze the variance, we need to introduce some more notations. For a vector $\vec{u}=u_{1} \ldots u_{n}$ of $m$-bit strings, denote by $[\vec{u}]_{2}$ the set of strings that appear in $\vec{u}$ odd number of times, and notice that $\sum_{i=1}^{n} u_{i}=\sum_{u \in[\vec{u}]_{2}} u$ and also $\sum_{i=1}^{n} f\left(u_{i}\right)=\sum_{u \in[\vec{u}]_{2}} f(u)$. (In particular, it follows that when $[\vec{u}]_{2}=\emptyset$, then $\sum u_{i}=0$ and $\sum f\left(u_{i}\right)=0$ for all $f$.) With these notations, we can express the second moment as

$$
\begin{aligned}
& E_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)^{2}\right] \\
& =\operatorname{Pr}_{f, d^{\prime}, \vec{u}, \vec{v}}^{\operatorname{Pr}}\left[\sum u_{i}=\sum v_{i}=d, \quad \sum f\left(u_{i}\right)=\sum f\left(v_{i}\right)=d^{\prime}\right] \\
& =\underset{f, d^{\prime}, \vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}, \sum u_{i}=\sum v_{i}=d, \quad \sum f\left(u_{i}\right)=\sum f\left(v_{i}\right)=d^{\prime}\right] \\
& \quad+\underset{f, d^{\prime}, \vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2} \neq[\vec{v}]_{2}, \sum u_{i}=\sum v_{i}=d, \quad \sum f\left(u_{i}\right)=\sum f\left(v_{i}\right)=d^{\prime}\right]
\end{aligned}
$$

where the last term in this equation is bounded by

$$
\begin{aligned}
& \quad \operatorname{Pr}_{f, d^{\prime}, \vec{u}, \vec{v}}\left[[\vec{u}]_{2} \neq[\vec{v}]_{2}, \sum u_{i}=\sum v_{i}=d, \quad \sum f\left(u_{i}\right)=\sum f\left(v_{i}\right)=d^{\prime}\right] \\
& =\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2} \neq[\vec{v}]_{2}, \sum u_{i}=\sum v_{i}=d\right] \\
& \quad \cdot \underset{f \ldots}{\operatorname{Pr}}\left[\sum f\left(u_{i}\right)=\sum f\left(v_{i}\right) \mid[\vec{u}]_{2} \neq[\vec{v}]_{2}, \cdots\right] \cdot \underset{d^{\prime} \ldots}{\operatorname{Pr}}\left[d^{\prime}=\sum f\left(u_{i}\right) \mid \ldots\right] \\
& \leq 2^{-2 m} \cdot 2^{-m^{\prime}} \cdot 2^{-m^{\prime}}
\end{aligned}
$$

Therefore, for any fixed $d$, the variance is bounded by

$$
\begin{aligned}
& \operatorname{VAR}_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)\right]=E_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)^{2}\right]-E_{f, d^{\prime}}\left[\mathcal{D}_{f}\left(d, d^{\prime}\right)\right]^{2} \\
& \begin{aligned}
& \leq\left(\underset{f, d^{\prime}, \vec{u}, \vec{v}}{ }\left[[\vec{u}]_{2}=[\vec{v}]_{2}, \sum u_{i}=\sum v_{i}=d, \sum f\left(u_{i}\right)=\sum f\left(v_{i}\right)=d^{\prime}\right]\right. \\
&\left.+2^{-2 m-2 m^{\prime}}\right)-2^{-2 m-2 m^{\prime}} \\
& \stackrel{(a)}{=} \operatorname{Pr}_{\vec{u}, \vec{v}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}, \sum u_{i}=d\right] \cdot \operatorname{Pr}\left[d^{\prime}=\sum f\left(u_{i}\right)\right] \\
&= 2^{-m^{\prime}} \cdot \operatorname{Prr}_{\vec{u}, \vec{v}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}, \sum u_{i}=d\right]
\end{aligned}
\end{aligned}
$$

where Equality ( $a$ ) holds because $[\vec{u}]_{2}=[\vec{v}]_{2}$ implies both $\sum u_{i}=\sum v_{i}$ and $\sum f\left(u_{i}\right)=\sum f\left(v_{i}\right)$. Bounding the last term, $\operatorname{Pr}_{\vec{u}, \vec{v}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}, \sum u_{i}=d\right]$, is where we need to distinguish between odd and even $n$ and between $d=0$ and $d \neq 0$. In the case analysis below, we make use of the following proposition, which is proved later.

Proposition 12 (i) For two vectors $\vec{u}, \vec{v}$, we have $[\vec{u}]_{2}=[\vec{v}]_{2}$ if and only if $[\vec{u} \mid \vec{v}]_{2}=\emptyset$ (where $\vec{u} \mid \vec{v}$ is the concatenation of the two vectors).
(ii) If $n$ is odd, then for any $n$-vector $\vec{u},[\vec{u}]_{2} \neq \emptyset$. If $n$ is even, then when we pick a random $n$-vector $\vec{u}$ we have,

$$
B(m, n) \cdot\left(1-\frac{n^{2}}{2^{m+3}}\right) \leq \operatorname{Pr}\left[[\vec{u}]_{2}=\emptyset\right] \leq B(m, n)
$$

where

$$
B(m, n) \stackrel{\text { def }}{=} \frac{n!}{(n / 2)!2^{n / 2}} \cdot 2^{-n m / 2}
$$

(iii) When $n$ is odd and $\vec{u}=\left(u_{1} \ldots u_{n}\right), \vec{v}=\left(v_{1} \ldots v_{n}\right)$ are chosen at random, we have for any fixed $d$,

$$
\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[\sum u_{i}=d \mid[\vec{u}]_{2}=[\vec{v}]_{2}\right]=2^{-m}
$$

(iv) When $n$ is even and $\vec{u}=\left(u_{1} \ldots u_{n}\right), \vec{v}=\left(v_{1} \ldots v_{n}\right)$ are chosen at random, we have

$$
\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[\sum u_{i} \neq 0 \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right] \geq\left(1-\frac{n^{2}}{2^{m+3}}\right) \cdot\left(1-\frac{1}{2^{m}-n}\right)
$$

and moreover, for any fixed $d \neq 0$,

$$
\operatorname{Pr}_{\vec{u}, \vec{v}}\left[\sum u_{i}=d \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right]=\frac{\operatorname{Pr}_{\vec{u}, \vec{v}}\left[\sum u_{i} \neq 0 \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right]}{2^{m}-1}
$$

Using Proposition 12, we now analyze three cases: when $n$ is odd, when $n$ is even and $d=0$, and when $n$ is even and $d \neq 0$.

Case 1: Odd $n$. This is the simplest case. Here, for every fixed $d$, we bound

$$
\begin{align*}
\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}, \sum u_{i}=d\right] & \stackrel{(a)}{=} \underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u} \mid \vec{v}]_{2}=\emptyset\right] \cdot \operatorname{Pr}_{\vec{u}, \vec{v}}\left[\sum v_{i}=d \mid[\vec{u}]_{2}=[\vec{v}]_{2}\right] \\
& \stackrel{(b)}{\leq} B(m, 2 n) \cdot 2^{-m}=2^{-m(n+1)} \cdot\left(\frac{(2 n)!}{n!2^{n}}\right) \tag{5}
\end{align*}
$$

where equality $(a)$ follows from part (i) of Proposition 12, and inequality $(b)$ follows from the bounds in parts (ii) and (iii).

Case 2: Even $n$ and $d=0$. Here we have

$$
\begin{align*}
& \underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}, \sum u_{i}=0\right]  \tag{6}\\
& =\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}=\emptyset\right]+\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right] \cdot \operatorname{Pr}_{\vec{u}, \vec{v}}\left[\sum u_{i}=0 \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right] \\
& \leq \underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}=\emptyset\right]+\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}\right] \cdot \underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[\sum u_{i}=0 \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right] \\
& \left(\frac{(a)}{\leq} 2^{-m n}\left(\frac{n!}{(n / 2)!2^{n / 2}}\right)^{2}+2^{-m n}\left(\frac{(2 n)!}{n!2^{n}}\right) \cdot\left(1-\left(1-\frac{n^{2}}{2^{m+3}}\right)\left(1-\frac{1}{2^{m}-n}\right)\right)\right. \\
& =2^{-m n}\left(\left(\frac{n!}{(n / 2)!2^{n / 2}}\right)^{2}+\left(\frac{(2 n)!}{n!2^{n}}\right) \cdot o(1)\right)
\end{align*}
$$

Inequality (a) follows from the bounds in parts (ii) and (iv) of Proposition 12.

Case 3: Even $n$ and $d \neq 0$. For any fixed $d \neq 0$, we have

$$
\begin{align*}
& \operatorname{Pr}_{\vec{u}, \vec{v}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}, \sum u_{i}=d\right] \stackrel{(a)}{=} \underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset, \sum u_{i}=d\right]  \tag{7}\\
& =\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right] \cdot \underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[\sum u_{i}=d \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right] \\
& \stackrel{(b)}{=}\left(\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}\right]-\underset{\vec{u}, \vec{v}}{\operatorname{Pr}}\left[[\vec{u}]_{2}=[\vec{v}]_{2}=\emptyset\right]\right) \cdot \frac{\operatorname{Pr}_{\vec{u}, \vec{v}}\left[\sum u_{i} \neq 0 \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right]}{2^{m}-1} \\
& \stackrel{(c)}{\leq}\left(2^{-m n} \frac{(2 n)!}{n!2^{n}}-\left(2^{-m n / 2} \frac{n!}{(n / 2)!2^{n / 2}} \cdot\left(1-\frac{n^{2}}{2^{m+3}}\right)\right)^{2}\right) \cdot \frac{1}{2^{m}-1} \\
& =2^{-m(n+1)}\left(\frac{(2 n)!}{n!2^{n}}-\left(\frac{n!}{(n / 2)!2^{n / 2}}\right)^{2}\right)(1+o(1))
\end{align*}
$$

Equality ( $a$ ) holds because $\sum u_{i}=d \neq 0$ implies $[\vec{u}]_{2} \neq \emptyset$. Inequality (b) follows from part (iv) of Proposition 12, and inequality ( $c$ ) follows from the bounds in part (ii).

Plugging the bounds from Equations (5), (6) and (7) into Eq. (4) completes the proof of Proposition 11.

Proof:(of Proposition 12) Part (i) is obvious. To prove the bounds in part (ii), we count the number of vectors with $[\vec{u}]_{2}=\emptyset$ as follows: First we partition the $n$ entries in $\vec{u}$ into pairs, then we pick a sequence of $n / 2$ strings, and finally we assign each string to one pair. Hence the number of such vectors is at most

$$
\frac{\binom{n}{2}\binom{n-2}{2} \cdots\binom{4}{2}}{(n / 2)!} \cdot \frac{2^{m n / 2}}{(n / 2)!} \cdot(n / 2)!=2^{m n / 2} \frac{n!}{(n / 2)!2^{n / 2}}
$$

On the other hand, as long as the sequence of strings that we pick in this process does not contain duplicates, which happen with probability at least $1-\binom{(n / 2)}{2} / 2^{m}$, then each of the choices above yields a different vector $\vec{u}$, hence the lower bound.
(iii) When $n$ is odd, the sum $\sum u_{i}$ is independent of the event $[\vec{u}]_{2}=[\vec{v}]_{2}$. Indeed, for any fixed $d 1, d 2$, denote $\Delta=d 1+d 2$, and consider the mapping

$$
\psi(\vec{u})=\psi\left(u_{1} \ldots u_{n}\right)=\left(u_{1}+\Delta \ldots u_{n}+\Delta\right)
$$

The mapping $\psi$ is a permutation over $\{0,1\}^{m n}$, that preserves the relation $[\vec{u}]_{2}=[\vec{v}]_{2}$, and it satisfies $\sum \psi(u)_{i}=d 2$ if and only if $\sum u_{i}=d 1$.
(iv) The "moreover" part is proved similarly to (iii) above. For any fixed non-zero $d 1, d 2$, denote $\Delta=\frac{d 2}{d 1}$ (where the operations are in $G F\left(2^{n}\right)$ ), and consider the mapping

$$
\phi(\vec{u})=\psi\left(u_{1} \ldots u_{n}\right)=\left(u_{1} \cdot \Delta \ldots u_{n} \cdot \Delta\right)
$$

(again, everything in $G F\left(2^{n}\right)$ ). As before, the mapping $\phi$ is a permutation over $\{0,1\}^{m n}$, that preserves the relation $[\vec{u}]_{2}=[\vec{v}]_{2}$, and it satisfies $\sum \phi(u)_{i}=d 2$ if and only if $\sum u_{i}=d 1$. It follows that conditioned on $\sum u_{i} \neq 0$, the sum $\sum u_{i}$ is independent of the event $[\vec{u}]_{2}=[\vec{v}]_{2}$. We note also that conditioned on $\sum u_{i} \neq 0$, the events $[\vec{u}]_{2}=[\vec{v}]_{2}$ and $[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset$ coincide.

Proving the bound on $\operatorname{Pr}\left[\sum u_{i} \neq 0 \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right]$ is a bit harder. Denote by $S$ the space of pairs of vectors $\vec{w}=(\vec{u} \mid \vec{v})$, with $[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset$, restricted so that each string that appears in $\vec{w}$, appears in it exactly twice. In terms of the process for selecting such $\vec{w}$ 's, as described in part (ii) above, this means that there are no duplicates in the sequence of strings that we select. As before, the space $S$ captures at least a fraction $1-\frac{n^{2}}{2^{m+3}}$ of the entire space of $[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset$. Consider now the following random process, for picking an element out of $S$ :

1. Pick a uniformly selected element $\vec{w}=(\vec{u}, \vec{v}) \in S$.
2. Pick a pair of indexes $(i, j)$, so that $w_{i}=w_{j}$ and the index $i$ is in the $\vec{u}$ part and the index $j$ is in the $\vec{v}$ part. (At least one such pair must exist, since $[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset$. If there is more than one, then just pick the first one.)
3. Pick at random a string that does not appear anywhere else in $\vec{w}$, and replace entries $i, j$ in $\vec{w}$ by this string.
4. Return the modified vector $\vec{w}^{\prime}$.

It is easy to see that this procedure returns a uniformly selected element in $S$. On the other hand, since the string in step 3 is chosen at random from a set of $2^{m}-n+1$ strings, then the probability
of $\sum u_{i}=0$ is at most $1 /\left(2^{m}-n+1\right)$. We therefore conclude that

$$
\begin{aligned}
& \operatorname{Pr}\left[\sum u_{i} \neq 0 \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right] \\
& \geq \operatorname{Pr}\left[(\vec{u}, \vec{v}) \in S \mid[\vec{u}]_{2}=[\vec{v}]_{2} \neq \emptyset\right] \cdot \operatorname{Pr}\left[\sum u_{i} \neq 0 \mid(\vec{u}, \vec{v}) \in S\right] \\
& \geq\left(1-\frac{n^{2}}{2^{m+3}}\right)\left(1-\frac{1}{2^{m}-n}\right)
\end{aligned}
$$

### 5.2 Variations and extensions

Here we briefly discuss a few possible extensions to the analysis from above.

Using different $f$ 's for different steps. Instead of using the same $f$ everywhere, we may have different $f$ 's for different steps. I.e., in step $j$ we have $\ell_{\text {out }}\left(N F\left(x_{j}\right)\right)=f_{j}\left(\ell_{\text {in }}\left(x_{j}\right)\right)$, and we assume that the $f_{j}$ 's are random and independent. The distribution that we want to analyze is therefore $\left\langle d=\sum u_{j}, d^{\prime}=\sum f_{j}\left(u_{j}\right)\right\rangle$. The analysis from above still works for the most part (as long as $\ell_{\text {in }}, \ell_{\text {out }}$ are the same in all the steps). The main difference is that the factor $c(n)$ is replaced by a smaller one (call it $c^{\prime}(n)$ ).

For example, if we use $n$ independent functions, we get $c^{\prime}(n)=1$, since all the symmetries in the proof of Proposition 11 disappear. Another example (which is used in the attack on Scream-0) is when we have just two independent functions, $f_{1}=f_{3}=\cdots$ and $f_{2}=f_{4}=\cdots$. In this case (and when $n$ is divisible by four), we get $c^{\prime}(n)=(1+o(1)) \sqrt{\left(\frac{n!}{(n / 2)!2^{n / 2}}\right)^{2}-\left(\frac{(n / 2)!}{(n / 4)!2^{n / 4}}\right)^{4}}$.

When $f$ is a sum of a few functions. An important special case, is when $f$ is a sum of a few functions. For example, in the functions that are used in the attack on Scream-0, the $m$-bit input to $f$ can be broken into three disjoint parts, each with $m / 3$ bits, so that $f(x)=$ $f^{1}\left(x^{1}\right)+f^{2}\left(x^{2}\right)+f^{3}\left(x^{3}\right)$. (Here we have $\left|x^{1}\right|=\left|x^{2}\right|=\left|x^{3}\right|=m / 3$ and $x=x^{1} x^{2} x^{3}$.) If $f^{1}, f^{2}, f^{3}$ themselves do not have any clear structure, then we can apply the analysis from above to each of them. That analysis tells us that each of the distributions $\mathcal{D}^{i} \stackrel{\text { def }}{=}\left(\sum_{j} u_{j}^{i}, \sum_{j} f^{i}\left(u_{j}^{i}\right)\right)$ is likely to be roughly $c(n) \cdot 2^{\left(m^{\prime}-(n-1) m / 3\right) / 2}$ away from the uniform distribution.

It is not hard to see that the distribution $\mathcal{D}_{f}^{n}$ that we want to analyze can be cast as $\mathcal{D}^{1}+\mathcal{D}^{2}+\mathcal{D}^{3}$, so we expect to get $\left|\mathcal{D}_{f}^{n}-\mathcal{R}\right| \approx \prod\left|\mathcal{D}^{i}-\mathcal{R}\right| \approx\left(c(n) \cdot 2^{\left(m^{\prime}-(n-1) m / 3\right) / 2}\right)^{3}=c(n)^{3} 2^{\left(3 m^{\prime}-(n-1) m\right) / 2}$. More generally, suppose we can write $f$ as a sum of $r$ functions over disjoint arguments of the same length. Namely, $f(x)=\sum_{i=1}^{r} f^{i}\left(x^{i}\right)$, where $\left|x^{1}\right|=\ldots=\left|x^{r}\right|=m / r$ and $x=x^{1} \ldots x^{r}$. Repeating the argument from above, we get that the expected distance $\left|\mathcal{D}_{f}^{n}-\mathcal{R}\right|$ is about $c(n)^{r} 2^{\left(r m^{\prime}-(n-1) m\right) / 2}$ (assuming that this is still smaller than one). As before, one could use the "Gaussian heuristics" to argue that for the "actual distance" we should replace $c(n)^{r}$ by $(c(n) \cdot \sqrt{2 / \pi})^{r}$. (And if we have different functions for different steps, as above, then we would get $\left(c^{\prime}(n) \cdot \sqrt{2 / \pi}\right)^{r}$.)

Linear masking over different groups. Another variation is when we do linear masking over different groups. For example, instead of xor-ing the masks, we add them modulo some prime $q$, or modulo a power of two. Again, the analysis stays more or less the same, but the constants change. If we work modulo a prime $q>n$, we get a constant of $c^{\prime}(n)=\sqrt{n!}$, since the only symmetry that
is left is between all the orderings of $\left\{u_{1}, \ldots, u_{n}\right\}$. When we work modulo a power of two, the constant will be somewhere between $c^{\prime}(n)$ and $c(n)$, probably closer to the former.

### 5.3 Efficiency considerations

The analysis from above says nothing about the computational cost of distinguishing between $\mathcal{D}_{f}^{n}$ and $\mathcal{R}$. It should be noted that in a "real life" attack, the attacker may have access to many different relations (with different values of $m, m^{\prime}$ ), all for the same non-linear function $N F$. To minimize the amount of needed text, the attacker may choose to work with the relation for which the quantity $(n-1) m-m^{\prime}$ is minimized. However, the choice of relations is limited by the attacker's computational resources. Indeed, for large values of $m, m^{\prime}$, computing the maximum-likelihood decision rule may be prohibitively expensive in terms of space and time. Below we review some strategies for computing the maximum-likelihood decision rule.

Using one big table. Perhaps the simplest strategy, is for the attacker to prepare off-line a table of all possible pairs $\left\langle d, d^{\prime}\right\rangle$ with $d \in\{0,1\}^{m}, d^{\prime} \in\{0,1\}^{m^{\prime}}$. For each pair $\left\langle d, d^{\prime}\right\rangle$ the table contains the probability of this pair under the distribution $\mathcal{D}_{f}^{n}$ (or perhaps just one bit that says whether this probability is more than $2^{-m-m^{\prime}}$ ).

Given such a table, the on-line part of the attack is trivial: for each set of steps $J$, compute $\left(d, d^{\prime}\right)=\sum_{j \in J}\left(w_{j}, w_{j}^{\prime}\right)$, and look into the table to see if this pair is more likely to come from $\mathcal{D}_{f}^{n}$ or from $\mathcal{R}$. After observing roughly $2^{(n-1) m-m^{\prime}} / c(n)^{2}$ such sets $J$, a simple majority vote can be used to determine if this is the cipher or a random process. Thus, the on-line phase is linear in the amount of text that has to be observed, and the space requirement is $2^{m+m^{\prime}}$.

As for the off-line part (in which the table is computed), the naive way is to go over all possible values of $u_{1} \ldots u_{n} \in\{0,1\}^{m}$, for each value computing $d=\sum u_{i}$ and $d^{\prime}=\sum f\left(u_{i}\right)$ and increasing the corresponding entry $\left\langle d, d^{\prime}\right\rangle$ by one. This takes $2^{m n}$ time. However, in the (typical) case where $m^{\prime} \ll$ $(n-1) m$, one can use a much better strategy, whose running time is only $O\left(\log n\left(m+m^{\prime}\right) 2^{m+m^{\prime}}\right)$.

First, we represent the function $f$ by a $2^{m} \times 2^{m^{\prime}}$ table, with $F[x, y]=1$ if $f(x)=y$, and $F[x, y]=0$ otherwise. Then, we compute the convolution of $F$ with itself, ${ }^{6}$

$$
E[s, t] \stackrel{\text { def }}{=}(F \star F)[s, t]=\sum_{x+x^{\prime}=s} \sum_{y+y^{\prime}=t} F[x, y] \cdot F\left[x^{\prime}, y^{\prime}\right]=|\{x: f(x)+f(x+s)=t\}|
$$

(Note that $E$ represents the distribution $\mathcal{D}_{f}^{2}$.) One can use the Walsh-Hadamard transform to perform this step in time $O\left(\left(m+m^{\prime}\right) 2^{m+m^{\prime}}\right)$ (see, e.g., [19]). Then, we again use the WalshHadamard transform to compute the convolution of $E$ with itself,

$$
\begin{aligned}
D\left[d, d^{\prime}\right] & \stackrel{\text { def }}{=}(E \star E)\left[d, d^{\prime}\right]=\sum_{s+s^{\prime}=d} \sum_{t+t^{\prime}=d^{\prime}} E(s, t) \cdot E\left(s^{\prime}, t^{\prime}\right) \\
& =\left|\left\{\langle x, s, z\rangle: f(x)+f(x+s)+f(z)+f(z+s+d)=d^{\prime}\right\}\right| \\
& =\left|\left\{\langle x, y, z\rangle: f(x)+f(y)+f(z)+f(x+y+z+d)=d^{\prime}\right\}\right|
\end{aligned}
$$

thus getting the distribution $\mathcal{D}_{f}^{4}$, etc. After $\log n$ such steps, we get the distribution of $\mathcal{D}_{f}^{n}$.

[^3]When $f$ is a sum of functions. We can get additional flexibility when $f$ is a sum of functions on disjoint arguments, $f(x)=f^{1}\left(x^{1}\right)+\cdots+f^{r}\left(x^{r}\right)$ (with $x=x^{1} \ldots x^{r}$ ). In this case, one can use the procedure from above to compute the tables $D^{i}\left[d, d^{\prime}\right]$ for the individual $f^{i}$,s. If all the $x^{i}$,s are of the same size, then each of the $D^{i}$, s takes up $2^{m^{\prime}+(m / r)}$ space, and can be computed in time $O\left(\log n\left(m^{\prime}+(m / r)\right) 2^{m^{\prime}+(m / r)}\right)$. Then, the "global" $D$ table can again be computed using convolutions. Specifically, for any fixed $d=d^{1} \ldots d^{r}$, the $2^{m^{\prime}}$-vector of entries $D[d, \cdot]$ can be computed as the convolutions of the $2^{m^{\prime}}$-vectors $D^{1}\left[d^{1}, \cdot\right], D^{2}\left[d^{2}, \cdot\right], \ldots, D^{r}\left[d^{r}, \cdot\right]$,

$$
D[d, \cdot]=D^{1}\left[d^{1}, \cdot\right] \star D^{2}\left[d^{2}, \cdot\right] \star \cdots \star D^{r}\left[d^{r}, \cdot\right]
$$

At first glance, this does not seem to help much: Computing each convolution takes time $O(r$. $m^{\prime} 2^{m^{\prime}}$ ), and we need to repeat this for each $d \in\{0,1\}^{m}$, so the total time is $O\left(r m^{\prime} 2^{m+m^{\prime}}\right)$. However, we can do much better than that.

Instead of storing the vectors $D^{i}\left[d^{i}, \cdot\right]$ themselves, we store their image under the WalshHadamard transform, $\Delta^{i}\left[d^{i}, \cdot\right] \stackrel{\text { def }}{=} \mathcal{H}\left(D^{i}\left[d^{i}, \cdot\right]\right)$. Then, to compute the vector $D\left[\left\langle d^{1} \ldots d^{r}\right\rangle, \cdot\right]$, all we need is to multiply (point-wise) the corresponding $\Delta^{i}\left[d^{i}, \cdot \cdot\right.$ 's, and then apply the inverse WalshHadamard transform to the result. Thus, once we have the tables $D^{i}[\cdot, \cdot]$, we need to compute $r \cdot 2^{m / r}$ "forward transforms" (one for each vector $D^{i}\left[d^{i}, \cdot\right]$ ), and $2^{m}$ inverse transforms (one for each $\left\langle d^{1} \ldots d^{r}\right\rangle$. Computing each transform (or inverse) takes $O\left(m^{\prime} 2^{m^{\prime}}\right)$ time. Hence, the total time (including the initial computation of the $D^{i}$ 's) is $O\left(\log n\left(r m^{\prime}+m\right) 2^{m^{\prime}+(m / r)}+m^{\prime} 2^{m+m^{\prime}}\right)$, and the total space that is needed is $O\left(2^{m+m^{\prime}}\right)$.

If the amount of text that is needed is less than $2^{m}$, then we can optimize even further. In this case the attacker need not store the entire table $D$ in memory. Instead, it is possible to store only the $D^{i}$ tables (or rather, the $\Delta^{i}[\cdot, \cdot]$ vectors), and compute the entries of $D$ during the on-line part, as they are needed. Using this method, the off-line phase takes $O\left(\log n\left(\mathrm{rm}^{\prime}+m\right) 2^{m^{\prime}+(m / r)}\right)$ time and $O\left(r 2^{m^{\prime}+m / r}\right)$ space to compute and store the vectors $\Delta_{i}[\cdot, \cdot]$, and the on-line phase takes $O\left(m^{\prime} 2^{m^{\prime}}\right)$ time per sample. Thus the total time complexity here is $O\left(\log n\left(r m^{\prime}+m\right) 2^{m^{\prime}+(m / r)}+S m^{\prime} 2^{m^{\prime}}\right)$, where $S$ is the number of samples needed to distinguish $\mathcal{D}$ from $\mathcal{R}$.

### 5.4 An attack on Scream-0

The stream cipher Scream (with its variants Scream-0 and Scream-F) was proposed very recently by Coppersmith, Halevi and Jutla. A detailed description of Scream is available in [11]. Below we only give a partial description of Scream-0, which suffices for the purpose of our attack.

Scream-0 maintains a 128 -bit "non-linear state" $x$, two 128 -bit "column masks" $c 1, c 2$ (which are modified every sixteen steps), and a table of sixteen "row masks" $R[0 . .15]$. It uses a non-linear function $N F$, somewhat similar to a round of Rijndael. Roughly speaking, the steps of Scream-0 are partitioned to chunks of sixteen steps. A description of one such chunk is found in Figure 3.

Here we outline a low-diffusion attack on the variant Scream-0, along the lines above, that can reliably distinguish it from random after observing merely $2^{43}$ bytes of output, with memory requirement of about $2^{50}$ and work-load of about $2^{80}$. This attack is described in more details in the long version of [11].

As usual, we need to find a "distinguishing characteristic" of the non-linear function (in this case, a low-diffusion characteristic), and a combination of steps in which the linear process vanishes. The linear process consists of the $c_{i}$ 's and the $R[i]$ 's. Since each entry $R[i]$ is used sixteen times before it is modified, we can cancel it out by adding two steps were the same entry is used. Similarly, we can cancel $c_{2}$ by adding two steps within the same "chunk" of sixteen steps. However, since $c 1$

1. for $i=0$ to 15 do
$x:=N F(x+c 1)+c 2$
output $x+R[i]$
if $i$ is even, rotate $c 1$ by 64 bits
if $i$ is odd, rotate $c 1$ by some other amount
. end-for
2. modify $c 1, c 2$, and one entry of $R$, using the function $N F(\cdot)$

Figure 3: sixteen steps of Scream-0.
is rotated after each use, we need to look for two different characteristics of the $N F$ function, such that the pattern of input bits in one characteristic is a rotated version of the pattern in the other.

The best such pair of "distinguishing characteristics" that we found for Scream-0, uses a lowdiffusion characteristic for $N F$ in which the input bits pattern is 2-periodic (and the fact that $c 1$ is rotated every other step by 64 bits). Specifically, the four input bytes $x_{0}, x_{5}, x_{8}, x_{13}$, together with two bytes of linear combinations of the output $N F(x)$, yield the two input bytes $x_{2}, x_{10}$, and two other bytes of linear combinations of the output $N F(x)$. In terms of the parameters that we used above, we have $m=48$ input and output bits, which completely determine $m^{\prime}=32$ other input and output bits.

To use this relation, we can observe these ten bytes from each of four steps, (i.e., $j, j+1, j+$ $16 k, j+1+16 k$ for even $j$ and $k<16$ ). We can then add them up (with the proper rotation of the input bytes in steps $j+1, j+17$ ), to cancel both the "row masks" $R[i]$ and the "column masks" $c 1, c 2$. This gives us the following distribution $\mathcal{D}=\left\langle u_{1}+u_{2}+u_{3}+u_{4}, f_{1}\left(u_{1}\right)+f_{2}\left(u_{2}\right)+f_{1}\left(u_{3}\right)+f_{2}\left(u_{4}\right)\right\rangle$, where the $u_{i}$ 's are modeled as independent, uniformly selected, 48-bit strings, and $f_{1}, f_{2}$ are two known functions $f_{j}:\{0,1\}^{48} \rightarrow\{0,1\}^{32}$. (The reason that we have two different functions is that the order of the input bytes is different between the even and odd steps.) Moreover, each of the two $f_{j}$ 's can be written as a sum of three functions over disjoint parts, $f_{j}(x)=f_{j}^{1}\left(x^{1}\right)+f_{j}^{2}\left(x^{2}\right)+f_{j}^{3}\left(x^{3}\right)$ where $\left|x^{1}\right|=\left|x^{2}\right|=\left|x^{3}\right|=16$.

This is one of the "extensions" that were discussed in Section 5.2. Here we have $n=4, m=48$, $m^{\prime}=32, r=3$, and two different functions. Therefore, we expect to get statistical distance of $c^{\prime}(n)^{3} \cdot 2^{\left(3 m^{\prime}-(n-1) m\right) / 2}$, with

$$
c^{\prime}(n) \approx \sqrt{2 / \pi} \cdot \sqrt{\left(\frac{n!}{(n / 2)!2^{n / 2}}\right)^{2}-\left(\frac{(n / 2)!}{(n / 4)!2^{n / 4}}\right)^{4}}
$$

Plugging in the parameters, we have $c^{\prime}(4) \approx \sqrt{2 / \pi} \cdot \sqrt{8}$, and the expected statistical distance is roughly $(16 / \pi)^{3 / 2} \cdot 2^{-24} \approx 2^{-20.5}$. We therefore expect to be able to reliably distinguish $\mathcal{D}$ from random after about $2^{41}$ samples. Roughly speaking, we can get $8 \cdot\binom{14}{2} \approx 2^{10}$ samples from 256 steps of Scream-0. (We have 8 choices for an even step in a chunk of 16 steps, and we can choose two such chunks from a collection of 14 in which the three row masks in use remain unchanged.) So we need about $2^{31} \cdot 256=2^{39}$ steps, or $2^{43}$ bytes of output.

Also, in Section 5.3 we show how one could efficiently implement the maximum-likelihood decision rule to distinguish $\mathcal{D}$ from $\mathcal{R}$, using Walsh-Hadamard transforms. Plugging the parameters of the attack on Scream-0 into the general techniques that are described there, we have space complexity of $O\left(r 2^{m^{\prime}+m / r}\right)$, which is about $2^{50}$. The time complexity is $O\left(\log n\left(r m^{\prime}+m\right) 2^{m^{\prime}+(m / r)}+\right.$ $S m^{\prime} 2^{m^{\prime}}$ ), where in our case $S=2^{41}$, so we need roughly $2^{80}$ time.

## 6 Conclusions

In this work we described a general cryptanalytical technique that can be used to attack ciphers that employ a combination of a "non-linear" process and a "linear process". We analyze in details the effectiveness of this technique for two special cases. One is when we exploit linear approximations of the non-linear process, and the other is when we exploit the low diffusion of (one step of) the non-linear process. We also show how these two special cases are useful in attacking the ciphers SNOW [4] and Scream-0 [11].

It remains an interesting open problem to extend the analysis that we have here to more general "distinguishing characteristics" of the non-linear process. For example, extending the analysis of the low-diffusion attack from Section 5.1 to the case where the functions $f$ is key-dependent (and thus not known to the adversary) may yield an effective attack on Scream [11].

In addition to the cryptanalytical technique, we believe that another contribution of this work is our formulation of attacks on stream ciphers. We believe that explicitly formalizing an attack as considering sequence of uncorrelated steps (as opposed to one continuous process) can be used to shed light on the strength of many ciphers.

## References

[1] A. Canteaut and E. Filiol. Ciphertext only reconstruction of stream ciphers based on combination generators. In Fast Software Encryption, volume 1978 of Lecture Notes in Computer Science, pages 165-180. Springer-Verlag, 2000.
[2] D. Coppersmith, S. Halevi, and C. Jutla. Cryptanalysis of stream ciphers with linear masking. In CRYPTO'02, Lecture Notes in Computer Science. Springer-Verlag, 2002. to appear. A longer version is available on-line from http://eprint.iacr.org/2002/020/.
[3] J. Daemen and C. S. K. Clapp. Fast hashing and stream encryption with Panama. In S. Vaudenay, editor, Fast Software Encryption: 5th International Workshop, volume 1372 of Lecture Notes in Computer Science, pages 23-25. Springer-Verlag, 1998.
[4] P. Ekdahl and T. Johansson. SNOW - a new stream cipher. Submitted to NESSIE. Available on-line from http://www.it.lth.se/cryptology/snow/.
[5] P. Ekdahl and T. Johansson. Distinguishing attacks on SOBER-t16 and t32. In Fast Software Encryption, Lecture Notes in Computer Science. Springer-Verlag, 2002. to appear.
[6] S. Fluhrer. Cryptanalysis of the SEAL 3.0 pseudorandom function family. In Proceedings of the Fast Software Encryption Workshop (FSE'01), 2001.
[7] S. R. Fluhrer and D. A. McGraw. Statistical analysis of the alleged RC4 keystream generator. In Proceedings of the 7th Annual Workshop on Fast Software Encryption, (FSE'2000), volume 1978 of Lecture Notes in Computer Science, pages 19-30. Springer-Verlag, 2000.
[8] J. D. Golić. Correlation properties of a general binary combiner with memory. Journal of Cryptology, 9(2):111-126, 1996.
[9] J. D. Golić. Linear models for keystream generators. IEEE Trans. on Computers, 45(1):41-49, Jan 1996.
[10] J. D. Golić. Linear statistical weakness of alleged RC4 keystream generator. In W. Fumy, editor, Advances in Cryptology - Eurocrypt'97, volume 1233 of Lecture Notes in Computer Science, pages 226-238. Springer-Verlag, 1997.
[11] S. Halevi, D. Copersmith, and C. Jutla. Scream: a software-efficient stream cipher. In Fast Software Encryption, Lecture Notes in Computer Science. Springer-Verlag, 2002. to appear. A longer version is available on-line from http://eprint.iacr.org/2002/019/.
[12] H. Handschuh and H. Gilbert. $\chi^{2}$ cryptanalysis of the SEAL encryption algorithm. In Proceedings of the 4th Workshop on Fast Software Encryption, volume 1267 of Lecture Notes in Computer Science, pages 1-12. Springer-Verlag, 1997.
[13] T. Johansson and F. Jönsson. Fast correlation attacks based on turbo code techniques. In Advances in Cryptology - CRYPTO '99, volume 1666 of Lecture Notes in Computer Science, pages 181-197. Springer-Verlag, 1999.
[14] T. Johansson and F. Jönsson. Improved fast correlation attacks on stream ciphers via convolution codes. In Advances in Cryptology - Eurocrypt '99, volume 1592 of Lecture Notes in Computer Science, pages 347-362. Springer-Verlag, 1999.
[15] M. Matsui. Linear cryptanalysis method for DES cipher. In Advances in Cryptology, EUROCRYPT'93, volume 765 of Lecture Notes in Computer Science, pages 386-397. Springer-Verlag, 1993.
[16] R. N. McDonough and A. D. Whalen. Detection of Signals in Noise. Academic Press, Inc., 2nd edition, 1995.
[17] W. Meier and O. Staffelbach. Fast correlation attacks on stream ciphers. Journal of Cryptology, 1(3):159-176, 1989.
[18] P. Rogaway and D. Coppersmith. A software optimized encryption algorithm. Journal of Cryptology, 11(4):273-287, 1998.
[19] D. Sundararajan. The Discrete Fourier Transform: Theory, Algorithms and Applications. World Scientific Pub Co., 2001.
[20] S. P. Vadhan. A Study of Statistical Zero-Knowledge Proofs. PhD thesis, MIT Department of Mathematics, August 1999.
[21] D. Watanabe, S. Furuya, H. Yoshida, and B. Preneel. A new keystream generator MUGI. In Fast Software Encryption, Lecture Notes in Computer Science. Springer-Verlag, 2002. Description available on-line from http://www.sdl.hitachi.co.jp/crypto/mugi/index-e.html.

## A Experimental results

We tested our analysis from Section 5.1, by choosing a few random functions $f:\{0,1\}^{m} \rightarrow\{0,1\}^{m^{\prime}}$ (for several settings of $m, m^{\prime}$ ), and evaluating the distance $\left|\mathcal{D}_{f}^{4}-\mathcal{R}\right|$. For each function $f$, we used the techniques from Section 5.3 (based on the Welsh-Hadamard transform) to compute the statistical distance. We used the SPIRAL implementation of the Welsh-Hadamard transform, due to Markus Pueschel, Bryan Singer, and Adrian Sox (see http://www.ece.cmu.edu/ $\sim$ spiral).

For each setting of $m, m^{\prime}$, we chose sixteen random functions, and computed the average distance for these functions. The results are presented below. One can see that the only deviation from the expected values in our analysis, is in the cases where $m$ is significantly smaller than $m^{\prime}$. In these cases, the distance is less than what we expect from the analysis. We speculate that the reason for this deviation, is that for such settings the variables in the proof are "not as smooth", and therefore, there is a larger gap between the quantities $E[|X-E[X]|]$ and $\sqrt{\operatorname{VAR}[X]}$.

```
m=6, m'=6: average distance is 1.174e-01 = 7.514 * 2^{(m'-3m)/2}
m=8, m'=8: average distance is 3.022e-02 = 7.736 * 2^{(m'-3m)/2}
m=10, m'=10: average distance is 7.569e-03 = 7.750 * 2^{(m'-3m)/2}
m=6, m'=12: average distance is 5.700e-01 = 4.560 * 2^{(m'-3m)/2}
m=8, m'=12: average distance is 8.417e-02 = 5.387* 2^{(m'-3m)/2}
m=10, m'=12: average distance is 1.310e-02 = 6.706 * 2^{(m'-3m)/2}
m=12, m'=6: average distance is 2.380e-04 = 7.799 * 2^{(m'-3m)/2}
m=12, m'=8: average distance is 4.767e-04 = 7.811 * 2^{(m'-3m)/2}
m=12, m'=10: average distance is 9.520e-04 = 7.799 * 2^{(m'-3m)/2}
```


[^0]:    ${ }^{1}$ We show how our techniques can handle other variants when we describe the attack on SNOW, but we do not attempt to characterize all the variants where such techniques apply.

[^1]:    ${ }^{2}$ The scheme in Scream is actually slightly different than the one described here, but this difference does not effect the analysis in any significant way.
    ${ }^{3}$ When only a part of $x$ is used as output, we may be forced to look at a few consecutive applications of NF. This is the case in SNOW, for example.

[^2]:    ${ }^{4}$ When $|J|=2$, this is just a differential attack, which uses the fact that for some values of $\Delta=u_{1}+u_{2}$, a corresponding $\Delta^{\prime}=f\left(u_{1}\right)+f\left(u_{2}\right)$ is more likely than in the random process.
    ${ }^{5}$ It can be shown that the same bounds hold also for larger $n$ 's, but assuming $n^{2} \ll 2^{m}$ makes some proofs a bit easier.

[^3]:    ${ }^{6}$ Recall that the convolution operator is defined on one-dimensional vectors, not on matrices. Indeed, in this expression we view the table $F$ as a one-dimensional vector, whose indexes are $m+m^{\prime}$-bits long.

