On the Universality of the Next Bit Test

A.W. Schrift and A. Shamir

Department of Applied Mathematics and Computer Science
The Weizmann Institute of Science
Rehovot 76100, Israel

Abstract

The next bit test was shown by Yao to be a universal test for sources of unbiased independent bits. The aim of this paper is to provide a rigorous methodology of how to test other properties of sources whose output distribution is not necessarily uniform. We prove the surprising result that the natural extension of the next bit test, even in the simplest case of biased independent bits, is no longer universal: We construct a source of biased bits, whose bits are obviously dependent and yet none of these bits can be predicted with probability of success greater than the bias. To overcome this difficulty, we develop new universal tests for arbitrary models of (potentially imperfect) sources of randomness.

1 Introduction

Randomness is an essential resource in many scientific areas, and pseudo-randomness is a good substitute in many applications. In his seminal paper Yao [11] formally defines the notion of perfect pseudo-random bits, i.e. bits that are indistinguishable from truly random bits by any probabilistic polynomial-time observer. He shows that the ability to predict some bit of a given source (the *next bit test*) serves as a universal test for randomness: A natural or pseudo-random source is perfect iff no probabilistic polynomial-time algorithm can, given any prefix of bits, predict the next bit of the source with probability of success significantly greater than 1/2. The next bit test has proved to be a useful tool for constructing perfect pseudo-random bit generators ([3],[2]) and for proving the imperfectness of other generators ([2],[6]).

Several models of natural sources of randomness have been suggested and investigated in [9], [1], [7] and [4]. In all the models the output distribution of natural sources is not uniform: In [9] a natural source outputs biased independent bits, in [1] a source is modeled by a Markov chain and in [7] and [4] the outcomes of the source are controlled by a powerful adversary. Non-uniform distributions appear also in some applications

A.J. Menezes and S.A. Vanstone (Eds.): Advances in Cryptology - CRYPTO '90, LNCS 537, pp. 394-408, 1991.

which require sources of randomness with independent yet biased bits (see for example [10],[5]). Nevertheless no rigorous methodology of how to verify the correspondence between a source of randomness with a non-uniform output distribution and its assumed properties has been given. The aim of our paper is to provide such a formalization.

Consider, for example, the roulette in your favorite Casino, where you are in the habit of placing a variety of bets on 17 with a 1/37 probability of winning each time. However after an unfortunate series of losses you begin to suspect that the roulette has been tampered with. You can easily check that the overall probability of 17 is close to 1/37, but that does not rule out the possibility that the outcomes of the roulette are artificially determined in a way that maintains the overall bias but inhibits 17 from appearing whenever the bets are high. How can you verify that indeed the outcomes of the roulette are independent, and that it is only your bad lack that brought you to the edge of bankruptcy? Clearly the next bit test cannot be employed here since you deal with a biased event.

Using the known notion of polynomial indistiguishability we define the notions of *perfect independence* and in general *perfect simulation* of a source by a mathematical model. We then move to the question of specifying the universal tests for these notions, which will declare a source to be perfect if and only if it passes the universal test. Surprisingly, the natural extension of Yao's work fails, even for the simplest case of independent biased bits. In other words the extended next bit test for biased bits which requires that no observer succeeds in predicting the bits of the source with probability greater than the bias, is no longer a universal test for independence. We introduce the correct test of independence, which we call the *weighted success rate (WSR) test* and prove its universality. We also discuss several alternative tests, and in particular the test we call *the predict or pass (POP) test.*

For general sources of randomness we present the universal test that determines whether a certain mathematical model perfectly simulates a given source. This test is the *comparative version* of the next bit test. The standard next bit test as well as the WSR and POP tests emerge from the comparative next bit test as special cases. Our proof of the universality of the test is the generalization of Yao's original proof, even though the original techniques cannot be implemented directly.

2 Definitions and Notations

Our definitions follow the original definitions of Yao [11]. The notions of a probability distribution, independence etc. are the standard notions from probability theory. All our results are stated in terms of probabilistic polynomial-time algorithms but can be restated in terms of polynomial-size Boolean circuits.

Let Σ^n denote the set of all binary strings of length n. A binary string of length n will be denoted s_1^n. The i-th bit of the string will be denoted by s_i. The substring

starting with the j-th bit and ending with the l-th bit ($1 \leq j < l \leq n$) will be denoted by s_j^l. We use the notation $f < O(\nu(n))$ for any function that vanishes faster than any polynomial, i.e. $\forall$constant k $\exists N$ such that $\forall n > N$: $f < 1/n^k$.

Definition: A *source ensemble* S is a sequence $\{S_n\}$ where S_n is a probability distribution on Σ^n.

Definition: A source S is *biased towards 1* with a fixed bias $1/2 < b < 1$ if for every i: $\Pr_S(s_i = 1) = b$.

Note that by our restriction on the bias the output bits of a biased source have a non-zero probability of being both 0 and 1. This ensures that the definitions of dependency, conditional probabilities etc. remain meaningful.

Let $B = \{B_n\}$ denote the *independent biased ensemble* where the source is biased towards 1 and all the bits are independent. $R = \{R_n\}$ denotes the *truly random ensemble* producing independent unbiased bits. Dealing with arbitrary models, we use $M = \{M_n\}$ for the mathematical model ensemble. We denote by $\Pr_S(E)$ the probability of an event E taking place when the probability distribution is defined by the source ensemble S. Whenever we refer to events that involve a probabilistic algorithm, we explicitly denote only the source ensemble, S, and implicitly assume the probability of the event to be induced by S and by the independent unbiased coin flips of the algorithm.

Definition: A probabilistic polynomial-time algorithm is *constant* if for some value v, $\Pr(\text{algorithm} = v) > 1 - O(\nu(n))$.

To simplify the presentation of our results we require algorithms to be non-constant even when it suffices to require only that for every value v: $\Pr(\text{algorithm} = v) \neq 0$.

Definition: A *distinguisher* is a probabilistic polynomial-time algorithm $D : \{0,1\}^n \rightarrow \{0,1\}$.

Definition: A biased source S outputs *perfect independent* bits if for every distinguisher D: $|\Pr_S(D = 1) - \Pr_B(D = 1)| < O(\nu(n))$.

Definition: A model M is *a perfect simulation* of a source S if for every distinguisher D: $|\Pr_S(D = 1) - \Pr_M(D = 1)| < O(\nu(n))$.

3 Universal Tests of Independence

In this section we construct what seems to be the natural extension of Yao's next bit test. We then show that there exist imperfect sources of randomness that pass the extended next bit test, thus disproving its universality. Our proof is based on the following intuition: Dependencies between the bits of an imperfect source will result in 1 having in some cases probability greater than the bias and in other cases probability smaller than the bias. It is possible, however, for the biased source to be imperfect with 1 remaining always more probable than 0. Hence deterministically predicting 1 is the optimal prediction strategy but has a poor probability of success. Following the same intuition we

suggest the weighted success rate test that is better suited to detect deviations from the bias. In the WSR test we separately compute the probabilities of success in predicting the 0 and 1 values of a next bit, and compose the two terms with appropriate weights into a single measure.

In the following we assume without loss of generality that all our sources to be biased towards 1 with some fixed bias b. It is easy to extend our results to the case where each bit has a different bias. It is worthwhile to emphasize that since we are interested in detecting dependencies among bits that have a particular bias, our basic WSR test may fail to detect imperfectness that results simply from a different overall bias. Testing the condition that the bits of a source have a certain bias can be done easily in polynomial time and with high accuracy using the law of large numbers. We give an alternative universal test, the POP test, with the additional feature that any deviation from the a-priori known bias is automatically detected.

3.1 The Extended Next Bit Test

Trying to extend the definition of Yao's next bit test to biased sources we must take into consideration the fact that the bits of an independent biased source can be trivially predicted with probability of success b, simply by always predicting 1.

Definition: A biased source S *passes the extended next bit test* if for every $1 \leq i \leq n$ and for every probabilistic polynomial-time algorithm A:

$$\Pr_S(A = s_i) < b + O(\nu(n))$$

Theorem 1:

The extended next bit test is not a universal test for independence.

Proof:

Fix a bias b: $1/2 < b \leq 1 - 1/n^t$, for some constant t. We construct a source which is biased towards 1 with bias b. We show that it is imperfect and yet it passes the extended next bit test. The source is the following:

$$\Pr_S(s_i = 1) = \begin{cases} b & \text{for } 1 \leq i \leq n-1 \\ b+\delta & \text{for } i = n \text{ and } s_1^2 = 01 \\ b-\delta & \text{for } i = n \text{ and } s_1^2 = 10 \\ b & \text{for } i = n \text{ and } s_1^2 = 00 \text{ or } 11 \end{cases}$$

Where $\frac{1}{n^q} \leq \delta < \min(b - \frac{1}{2}, 1 - b)$, for some constant q.

Let the distinguisher D be defined by: D=1 iff $s_1^2 = 01$ and $s_n = 1$. Clearly, $\Pr_S(D = 1) = c \cdot (b + \delta)$, while $\Pr_B(D = 1) = c \cdot b$, where $c = b \cdot (1 - b) \geq 1/2n^t$. Therefore, $\Pr_S(D = 1) - \Pr_B(D = 1) = c \cdot \delta \geq 1/2n^{q+t}$, and by definition the source is imperfect. Nevertheless the source passes the extended next bit test: The n-th bit is always biased

towards 1, so the best prediction strategy is to deterministically predict 1 regardless of the known values of the first two bits. It is easy to check that the probability of success of this optimal strategy remains b. □

3.2 The Weighted Success Rate Test

Definition: Fix $1 \leq i \leq n$. The *weighted success rate* of any non-constant probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1\}$ in predicting the i-th bit of a biased source S is:

$$\begin{aligned} ws(A,S,i) &= \frac{\Pr_S(A = s_i | s_i = 1)}{\Pr_S(A = 1)} + \frac{\Pr_S(A = s_i | s_i = 0)}{\Pr_S(A = 0)} = \\ &= \frac{1}{b} \cdot \Pr_S(A = s_i | A = 1) + \frac{1}{1-b} \cdot \Pr_S(A = s_i | A = 0) \end{aligned}$$

Definition: A biased source S *passes the weighted success rate test* if for every $1 \leq i \leq n$ and every non-constant probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1\}$:

$$ws(A,S,i) < 2 + O(\nu(n))$$

Remark: The above definitions do not allow constant prediction algorithms. Remember that we assume that indeed all the tested sources of randomness have a bias b. Since constant algorithms can only detect that the overall bias is other than b, which is not the case, it is possible without loss of the generality to ignore them.

Theorem 2:

A biased source produces perfect independent bits iff it passes the weighted success rate test.

Sketch of Proof:

If a given source fails the weighted success rate test, it is easy to construct a distinguisher which tells the source apart from a truly independent biased source by using the predictions of the WSR test.

To prove the other direction, we show how to construct a weighted success rate test using any distinguisher D for an imperfect source. Following our intuition that imperfect sources can be recognized using the events where the probability of an i-th bit being 1 significantly differs from the bias, we use the distinguisher to single out these events. We prove that one of the following cases always exists:

1. The probability that the distinguisher outputs 1 on the input sequence $s_1^{i-1} 1 s_{i+1}^n$ with $s_1^{i-1} \in S$ and $s_{i+1}^n \in B$ is significantly changed when $s_i = 1$ is taken out of S or B.

2. The probability that the distinguisher outputs 1 on the input sequence $s_1^{i-1}0s_{i+1}^n$ with $s_1^{i-1} \in S$ and $s_{i+1}^n \in B$ is significantly changed when $s_i = 0$ is taken out of S or B.

We construct a different test for each case, and use the corresponding condition on D to prove that in each case one of the terms in the weighted success rate of the corresponding test is significantly greater than 1. We conclude the proof with the following useful lemma:

Prediction Lemma: For any biased source and any non-constant probabilistic polynomial-time algorithm $A : \{0,1\}^{i-1} \rightarrow \{0,1\}$: $\Pr_S(A = s_i | A = 1) \geq b + 1/n^{k_1}$ iff $\Pr_S(A = s_i | A = 0) \geq 1 - b + 1/n^{k_2}$, for some constants k_1 and k_2.

The full proofs of the theorem and the lemma are given in the appendix.

3.3 Alternative Versions

In the following we present three tests, that are all equivalent to the WSR test, but emphasize different aspects of detecting dependencies. The first two definitions are closely related to the WSR test. The last test presents an entirely different approach, which stems from the fact that if a source is imperfect it is possible to detect the event in which 1 is more probable than the given bias and ignore all other events.

Definition: A biased source S *passes the modified WSR test* if for every $1 \leq i \leq n$ and every non-constant probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1\}$:

$$\max\left\{\frac{\Pr_S(A = s_i | s_i = 1)}{\Pr_S(A = 1)}, \frac{\Pr_S(A = s_i | s_i = 0)}{\Pr_S(A = 0)}\right\} =$$

$$= \max\left\{\frac{1}{b}\Pr_S(A = s_i | A = 1), \frac{1}{1-b}\Pr_S(A = s_i | A = 0)\right\} < 1 + O(\nu(n))$$

Definition: A biased source S *passes the behavior test* if for every $1 \leq i \leq n$ and every non-constant probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1\}$:

$$|\Pr_S(A = 1 | s_i = 1) - \Pr_S(A = 1 | s_i = 0)| < O(\nu(n))$$

Definition: A biased source S *passes the predict or pass (POP) test* if for every $1 \leq i \leq n$ and every probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1,*\}$ the following condition holds:

If $\Pr_S(A \neq *) \geq 1/n^l$ for some constant l, then $|\Pr_S(A = s_i | A \neq *) - b| < O(\nu(n))$.

Using similar techniques to those introduced in Theorem 2 we can prove that the above defined tests are universal. The proofs appear in the appendix.

Theorem 3:

The following conditions are equivalent:
1. A biased source outputs perfect independent bits.
2. A biased source passes the modified WSR test.
3. A biased source passes the behavior test.
4. A biased source passes the POP test.

Remark: The above equivalence holds only for biased sources that were a-priori tested to have a certain bias. Otherwise, the POP test behaves differently from the other tests. Its definition allows constant as well as non-constant prediction algorithms. More important is the fact that unlike the WSR test, the POP test succeeds in detecting imperfectness that results merely from a different overall bias.

The POP test introduces what seems to be a new notion of allowing a predictor to be successful only on some non-negligible fraction of its inputs. Despite the fact that this formal definition is novel (as far as we know), known constructions of pseudo-random bit generators often prove their perfectness by showing that they pass what is essentially a POP test (i.e. it is impossible to predict the output bits of the generator even on a non-negligible fraction of the output strings). Indeed, the POP test is particularly useful for constructing perfect generators for biased independent bits.

3.4 Comparison with the Next Bit Test

For unbiased ($b = 1/2$) independent bits the WSR test and its variations all serve as alternative universal tests to the next bit test. We can, however, show an even stronger equivalence between the tests, namely that the same algorithm that succeeds in the prediction of a certain bit with probability significantly greater than $1/2$ (thus proving the source of the bits to be imperfect by the next bit test) have a weighted success rate that is significantly greater than 2 (thus proving the source to be imperfect by the WSR test).

Proposition 4:

For any unbiased source ($b = 1/2$) and any non-constant probabilistic polynomial-time algorithm $A : \{0,1\}^{i-1} \rightarrow \{0,1\}$: $\Pr_S(A = s_i) \geq 1/2 + 1/n^{k_1}$ iff $ws(A, S, i) \geq 2 + 1/n^{k_2}$, for some constants k_1 and k_2.

In terms of the probability of successful prediction for unbiased sources our new definitions are superior:

Proposition 5:

For any unbiased source and any next bit test T there exists a POP test A, such that for every $1 \leq i \leq n$:

$$\Pr_S(A = s_i \mid A \neq *) \geq \Pr_S(T = s_i)$$

The proofs of the two propositions are given in the appendix.

4 Perfectness with Respect to Arbitrary Models

In this section we consider an arbitrary source S, which we believe to have a certain distribution described by a mathematical model M. As for randomness and independence we search for a convenient universal test, based on the probability of correct predictions:

Definition: A source S *passes the comparative next bit test* with respect to a model M if for every $1 \leq i \leq n$ and every probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1\}$:

$$|\Pr_S(A = s_i) - \Pr_M(A = s_i)| < O(\nu(n))$$

Note that the comparative next bit test enables us to avoid performing any a-priori tests on either sources. The test is easiest to implement when the model is described in such a way that the probability of correct bit predictions for the model can be efficiently computed. Yet we can perform the test even when the model is completely unknown and given to us as a black box. In that case the test simply involves a comparison between two boxes: one containing the tested source and the other containing the model black box.

It is instructive to examine simple examples of the comparative next bit test, where the model source is explicitly known:

1. $M = R$, i.e. the model is a source of unbiased independent bits. In that case we know that no matter which algorithm is used $\Pr_R(A = s_i) = 1/2$ and we can immediately derive the next bit test.

2. $M = B$, i.e. the model is a source of biased independent bits. Here we know that for any non-constant algorithm $\Pr_B(A = s_i | A = 1) = b$ and that $\Pr_B(A = s_i | A = 0) = 1 - b$ so that the predictions must be evaluated according to the value that is being predicted. This gives rise to the WSR test.

3. M=a source with a one-bit memory, in which the probability of the i-th bit is determined according to the outcome of the $(i-1)$-th bit. Let $b_i(0) = \Pr(s_i = 1 | s_{i-1} = 0)$ and $b_i(1) = \Pr(s_i = 1 | s_{i-1} = 1)$. Then it is easy to see that the performance of any algorithm must be evaluated not only according to the value of s_i but also according to the value of s_{i-1}. We therefore get that M is a perfect simulation of a source S if for every $1 \leq i \leq n$ and every probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1,*\}$ such that $\Pr_S(A \neq * | s_{i-1} = 0) \geq 1/n^{l_1}$ and $\Pr_S(A \neq * | s_{i-1} = 1) \geq 1/n^{l_2}$ for some constants l_1, l_2:

$$\max \left\{ \begin{array}{l} |\Pr_S(A = s_i | A \neq *, s_{i-1} = 0) - b_i(0)|, \\ |\Pr_S(A = s_i | A \neq *, s_{i-1} = 1) - b_i(1)| \end{array} \right\} < O(\nu(n))$$

It is easy to see that similar analysis holds for any M=Markov chain [1], where predictions must be evaluated according to the output value and to the state (which determines the bias).

Theorem 6:

A model M is a perfect simulation of a source S iff S passes the comparative next bit test with respect to M.

Proof:

It is easy to see that if a source S fails the comparative next bit test it is distinguishable from the model source M. Assume now that we are given that S and M are distinguishable and need to prove that S fails the comparative next bit test w.r.t. M. We cannot implement previous proof techniques directly since they inherently assume independence in concatenating a random prefix of bits taken out of the tested source with a random suffix of bits generated according to the desired distribution. We overcome the problem by using an additional truly random source for the concatenation.

Let $D : \{0,1\}^n \rightarrow \{0,1\}$ be the distinguisher for which $|\Pr_S(D=1) - \Pr_M(D=1)| \geq 1/n^k$ for some constant k. Let p_i^S (p_i^M) denote the probability that D outputs 1 when the first i bits of its input are taken out of S (M) and the rest are independent unbiased coin flips. Note that $p_n^S = \Pr_S(D=1)$, $p_n^M = \Pr_M(D=1)$ and $p_0^S = p_0^M = \Pr_R(D=1)$. Since $|p_n^S - p_n^M| \geq 1/n^k$, by the pigeonhole principle there exists an i which is the first for which p_i^S and p_i^M significantly differ, i.e.:

1. $|p_i^S - p_i^M| \geq 1/n^{k+1}$, and
2. For all $j \leq i-1$, $|p_j^S - p_j^M| < O(\nu(n))$.

We can assume w.l.o.g. $p_i^S - p_i^M > 0$. The comparative next bit test A submits to D the string $s = s_1^{i-1}s_i^n$, where $s_1^{i-1} \in S$ or M and $s_i^n \in R$. If $D(s) = 1$ then A outputs s_i, else A outputs $1 - s_i$. It is easy to see that:

$$\Pr_S(A = s_i) = \frac{1}{2} + p_i^S - p_{i-1}^S$$

While:

$$\Pr_M(A = s_i) = \frac{1}{2} + p_i^M - p_{i-1}^M$$

Hence:

$$\Pr_S(A = s_i) - \Pr_M(A = s_i) \geq \frac{1}{n^{k+1}}$$

□

5 Discussion

In this paper we develop a formal theory for the universal testing of non-uniform probability distributions. Our definitions rely on Yao's pioneering work, but evolve from it in a non-obvious way. In addition to its theoretical significance our results have several practical applications:

1. An important property of one-way functions is the existence of hard bits in the argument that are as hard to compute from a given random value of the function as the entire argument. In a recent work [8] the individual security of every bit of the discrete logarithm modulo a composite was proven. The known definitions of unpredictability could not be applied to the most significant bits, since they are biased towards 0 by definition. For those bits it was necessary to use our new definitions in order to define and prove their security.

2. It is possible to apply the universal test of independence to every biased predicate and use a hard biased predicate to construct a generator of independent biased bits. Consider for example the following construction that is based on the intractability of the discrete logarithm modulo a prime. Let $f_{g,p}(x) = g^x \pmod{p}$, where p is a randomly chosen n-bit prime and g is a random generator of $\mathbf{Z}_p^*$. Let b be the desired bias towards 1 and k the desired output length.

   $\forall i \geq 1: \quad f_{p,g}^i(z) = f_{p,g}\left(f_{p,g}^{i-1}(z)\right) \quad \text{with} \quad f_{p,g}^0(z) = z.$

   Let $G(z) = \begin{cases} 0 & \text{if } z < \lfloor (1-b)\cdot(p-1) \rfloor \\ 1 & \text{otherwise} \end{cases}$

   For a randomly chosen seed x we shall produce k bits $\{s_1^k\}$ by: $s_m = G\left(f_{p,g}^{m-1}(x)\right)$, and use the universal tests of independence to prove that $\{s_1^k\}$ are perfect independent bits with bias b. For constant output lengths the above is a more efficient construction of independent biased bits than the obvious construction of biased bits from pseudo-random (unbiased) bits.

Acknowledgements

We would like to thank Uriel Feige and Rafi Heiman for many stimulating discussions.

References

[1] Blum, M., "Independent Coin Flips From a Correlated Biased Source: a Finite State Markov Chain", *Proc. 25th FOCS, 1984, pp. 425-433.*

[2] Blum, L., Blum, M., Shub, M., "A Simple Secure Pseudo-Random Generator",*SIAM J. of Computing, Vol. 15, No. 2, 1986, pp. 364-383.*

[3] Blum, M., Micali, S., "How to Generate Cryptographically Strong Sequences of Pseudo-Random Bits", *Proc. 26th FOCS, 1982, pp. 112-117.*

[4] Chor, B., Goldreich, O., "Unbiased Bits from Sources of Weak Randomness and Probabilistic Communication Complexity", *Proc. 26th FOCS, 1985, pp. 429-442.*

[5] Feldman, D., Impagliazzo, R., Naor, M., Nisan, N., Rudich, S., Shamir, A., "On Dice and Coins", *ICALP 1989.*

[6] Plumstead, J., "Inferring a Sequence Generated by a Linear Congruence", *Proc. 23rd FOCS, 1982, pp. 153-159.*

[7] Santha, M., Vazirani, U.V., "Generating Quasi-Random Sequences from Slightly-Random Sources", *Proc. 25th FOCS, 1984, pp. 434-440.*

[8] Schrift, A.W., Shamir, A., "The Discrete Log is Very Discreet", *Proc. 22nd STOC, 1990, pp. 405-415.*

[9] von Neumann, J., "Various Techniques Used in Connection with Random Digits", *Notes by G.E. Forsythe, 1951, Reprinted in von Neumann's Collected Works, Vol. 5, Pergamon Press, 1963, pp. 768-770.*

[10] Vazirani, U.V., Vazirani, V.V., "Trapdoor Pseudo-Random Number Generator with Applications to Protocol Design", *Proc. 24th FOCS, 1983, pp. 23-30.*

[11] Yao, A.C., "Theory and Applications of Trapdoor Functions", *Proc. 23rd FOCS, 1982, pp. 80-91.*

Appendix: Full Proofs

Proof of Theorem 2

Given that a source fails the weighted success rate test, it is easy to construct a distinguisher between the source S and a truly independent biased source B by examining the predictions of the test. Formally assume that we are given a non-constant probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1\}$ for the i-th bit of a source S such that $ws(A,S,i) \geq 2+1/n^k$ for some constant k. We shall use A to construct two possible distinguishers and show that for one of them $|\Pr_S(D=1) - \Pr_B(D=1)| \geq 1/n^{k'}$ for some constant k'. Given s_1^n, both D's submits s_1^{i-1} to A and examine A's output. $D_1 = 1$ iff $A = s_i = 1$. $D_2 = 1$ iff $A = 1$. If the overall behavior of A is the same for S and for B, i.e. $|\Pr_S(A=1) - \Pr_B(A=1)| < O(\nu(n))$, then D_1 distinguishes the source. Otherwise D_2 distinguishes. Hence S is imperfect.

To prove the other direction, we show how to construct a weighted success rate test using any distinguisher D for an imperfect source. Let p_i denote the probability that D=1 when the first i input bits are taken out of S and the rest are independent biased coin flips. Note that $p_n = \Pr_S(D=1)$, while $p_0 = \Pr_B(D=1)$. Since D distinguishes between the source and a biased coin, $|p_0 - p_n| \geq 1/n^k$ for some k. By the pigeonhole principle there exists a bit i for which: $|p_i - p_{i-1}| \geq 1/n^{k+1}$. We shall assume w.l.o.g.

that $p_i - p_{i-1} > 0$.
Explicitly:

$$p_i = \sum_{s_1^n} \Pr\left(D(s_1^n) = 1\right) \cdot \Pr_S(s_1^i) \cdot \Pr_B(s_{i+1}^n) =$$

$$= \sum_{s_1^{i-1}, s_{i+1}^n} \Big[\Pr(D(s_1^{i-1} 1 s_{i+1}^n) = 1) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_S(s_i = 1 | s_1^{i-1}) \cdot \Pr_B(s_{i+1}^n) +$$

$$+ \Pr(D(s_1^{i-1} 0 s_{i+1}^n) = 1) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_S(s_i = 0 | s_1^{i-1}) \cdot \Pr_B(s_{i+1}^n)\Big]$$

$$p_{i-1} = \sum_{s_1^n} \Pr\left(D(s_1^n) = 1\right) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_B(s_i^n) =$$

$$= \sum_{s_1^{i-1}, s_{i+1}^n} \Big[\Pr\left(D(s_1^{i-1} 1 s_{i+1}^n) = 1\right) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_B(s_i = 1) \cdot \Pr_B(s_{i+1}^n) +$$

$$+ \Pr\left(D(s_1^{i-1} 0 s_{i+1}^n) = 1\right) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_B(s_i = 0) \cdot \Pr_B(s_{i+1}^n)\Big]$$

Since $p_i - p_{i-1} \geq \frac{1}{n^{k+1}}$, then one of the following two equations hold.

$$(1) \quad \sum_{s_1^{i-1}, s_{i+1}^n} \Big[\Pr\left(D(s_1^{i-1} 1 s_{i+1}^n) = 1\right) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_S(s_i = 1 | s_1^{i-1}) \cdot \Pr_B(s_{i+1}^n) -$$

$$- \Pr\left(D(s_1^{i-1} 1 s_{i+1}^n) = 1\right) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_B(s_i = 1) \cdot \Pr_B(s_{i+1}^n)\Big] \geq \frac{1}{2n^{k+1}}$$

$$(2) \quad \sum_{s_1^{i-1}, s_{i+1}^n} \Big[\Pr\left(D(s_1^{i-1} 0 s_{i+1}^n) = 1\right) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_S(s_i = 0 | s_1^{i-1}) \cdot \Pr_B(s_{i+1}^n) -$$

$$- \Pr\left(D(s_1^{i-1} 0 s_{i+1}^n) = 1\right) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_B(s_i = 0) \cdot \Pr_B(s_{i+1}^n)\Big] \geq \frac{1}{2n^{k+1}}$$

By examining D it is possible to decide which of the two holds and construct a WSR test A_i accordingly. Otherwise by constructing the following two tests A_1 and A_2 you are guaranteed that one of them will be successful.

A_1 submits as input to D the string $s_1^{i-1} 1 s_{i+1}^n$, where $s_1^{i-1} \in S$ and $s_{i+1}^n \in B$. If D=1 then $A_1 = 1$; else $A_1 = 0$. A_2 submits as input to D the string $s_1^{i-1} 0 s_{i+1}^n$, where $s_1^{i-1} \in S$ and $s_{i+1}^n \in B$. If D=1 then $A_2 = 0$; else $A_2 = 1$. We shall now analyze separately the two terms of $ws(A_1, S, i)$ and $ws(A_2, S, i)$. To make the analysis simple we use the second alternative in the definition of the weighted success rate, which compares the probabilities of successful predictions to b or $1 - b$.

$$\Pr_S(A_1 = s_i \mid A_1 = 1) =$$

$$= \frac{\sum_{s_1^{i-1}, s_{i+1}^n} \Pr_S(s_i = 1 \mid s_1^{i-1}) \cdot \Pr\left(D(s_1^{i-1} 1 s_{i+1}^n) = 1\right) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_B(s_{i+1}^n)}{\sum_{s_1^{i-1}, s_{i+1}^n} \Pr\left(D(s_1^{i-1} 1 s_{i+1}^n) = 1\right) \cdot \Pr_S(s_1^{i-1}) \cdot \Pr_B(s_{i+1}^n)} \geq$$

(by equation 1)

$$\geq \frac{\sum_{s_1^{i-1},s_{i+1}^n} \Pr_B(s_i=1)\cdot \Pr\left(D(s_1^{i-1}1s_{i+1}^n)=1\right)\cdot \Pr_S(s_1^{i-1})\cdot \Pr_B(s_{i+1}^n)+\frac{1}{2n^{k+1}}}{\sum_{s_1^{i-1},s_{i+1}^n} \Pr\left(D(s_1^{i-1}1s_{i+1}^n)=1\right)\cdot \Pr_S(s_1^{i-1})\cdot \Pr_B(s_{i+1}^n)} =$$

$(\Pr_B(s_i=1)=b)$

$$= b+\frac{\frac{1}{2n^{k+1}}}{\sum_{s_1^{i-1},s_{i+1}^n} \Pr\left(D(s_1^{i-1}1s_{i+1}^n)=1\right)\cdot \Pr_S(s_1^{i-1})\cdot \Pr_B(s_{i+1}^n)} \geq$$

(the denominator < 1)

$$\geq b+\frac{1}{2n^{k+1}}$$

Similarly for A_2:

$$\Pr_S(A_2=s_i|A_2=0) =$$

$$= \frac{\sum_{s_1^{i-1},s_{i+1}^n} \Pr_S(s_i=0\mid s_1^{i-1})\cdot \Pr\left(D(s_1^{i-1}0s_{i+1}^n)=1\right)\cdot \Pr_S(s_1^{i-1})\cdot \Pr_B(s_{i+1}^n)}{\sum_{s_1^{i-1},s_{i+1}^n} \Pr\left(D(s_1^{i-1}0s_{i+1}^n)=1\right)\cdot \Pr_S(s_1^{i-1})\cdot \Pr_B(s_{i+1}^n)} \geq$$

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

To complete the proof we show that for each of $ws(A_1,S,i)$ and $ws(A_2,S,i)$ the remaining term (that does not appear above) is also significantly greater than 1.

Prediction Lemma: For any biased source and any non-constant probabilistic polynomial time algorithm $A:\{0,1\}^{i-1}\rightarrow\{0,1\}$: $\Pr_S(A=s_i|A=1)\geq b+1/n^{k_1}$ iff $\Pr_S(A=s_i|A=0)\geq 1-b+1/n^{k_2}$, for some constants k_1 and k_2.

Proof:

Assume that $\Pr_S(A=s_i|A=1)\geq b+\varepsilon_1$, where $\varepsilon_1=1/n^{k_1}$. Note that $\Pr_S(A=s_i|A=1)=\Pr_S(s_i=1|A=1)$. Since the overall bias of the source is known to be b,

$$\Pr_S(s_i=1|A=1)\cdot\Pr_S(A=1)+\Pr_S(s_i=1|A=0)\cdot\Pr_S(A=0)=b.$$

Therefore:

$$\Pr_S(s_i=1|A=0)\leq\frac{b-(b+\varepsilon_1)\Pr_S(A=1)}{\Pr_S(A=0)}$$

Simple manipulations give:

$$\begin{aligned}\Pr_S(A=s_i|A=0) &= \Pr_S(s_i=0|A=0)=1-\Pr_S(s_i=1|A=0)\geq\\ &\geq 1-b+\varepsilon_1\cdot\frac{\Pr_S(A=1)}{\Pr_S(A=0)}\end{aligned}$$

Similarly when $\Pr_S(A=s_i|A=0)\geq 1-b+\varepsilon_2$, where $\varepsilon_2=1/n^{k_2}$, we get using the same manipulations that:

$$\Pr_S(A=s_i|A=1)\geq b+\varepsilon_2\cdot\frac{\Pr_S(A=0)}{\Pr_S(A=1)}$$

□

Proof of Theorem 3

The theorem contains three claims, all follow easily from the proof of Theorem 2.

Claim 1: A biased source outputs perfect independent bits iff it passes the modified WSR test.

Proof: If a biased source fails the modified WSR test it is easy to construct a distinguisher between the source and a truly independent biased source in a similar way to the construction in Theorem 2.

If a biased source is imperfect, then by the proof of Theorem 2 there exists a non-constant probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1\}$ such that $\Pr(A = s_i | A = 1) \geq b + 1/n^{k_1}$ and $\Pr(A = s_i | A = 0) \geq 1 - b + 1/n^{k_2}$, for some constants k_1 and k_2. By definition this source fails the modified WSR test. □

Claim 2: A biased source outputs perfect independent bits iff it passes the behavior test.

Proof: For notational simplicity let P_1 denote $\Pr(A = 1 | s_i = 1)$ and P_0 denote $\Pr(A = 1 | s_i = 0)$. We prove that a biased source passes the behavior test iff it passes the WSR test. This follows from the close relation between the two measures: For any biased source and any non-constant probabilistic polynomial-time algorithm A: $\{0,1\}^{i-1} \rightarrow \{0,1\}$:

$$ws(A, S, i) = \frac{P_1}{b \cdot P_1 + (1-b) \cdot P_0} + \frac{1 - P_0}{b \cdot (1 - P_1) + (1-b) \cdot (1 - P_0)}$$

Clearly if S passes the behavior test then it also passes the WSR test. If S fails the behavior test, then for some non-negligible ε: $|P_1 - P_0| \geq \varepsilon$. Assume w.l.o.g. that $P_1 > P_0$. Using the relation between the tests we then get:

$$ws(A, S, i) \geq 2 + \varepsilon \cdot \left\{ \frac{1-b}{P_1 - \varepsilon \cdot (1-b)} + \frac{b}{1 - P_1 + \varepsilon \cdot (1-b)} \right\}$$

Finally note that $\varepsilon \cdot b \leq P_1 - \varepsilon \cdot (1-b) \leq 1 - \varepsilon \cdot (1-b)$, so that the term that is added to 2 is indeed non-negligible. □

Calim 3: A source outputs perfect independent bits iff it passes the POP test.

Proof: Given that a source fails the POP test, it is easy to construct a distinguisher between the source and a truly independent biased source by examining the predictions of the test, as is done in the proof of Theorem 2.

To prove the other direction, assume that S is imperfect and there exists a distinguisher D between S and a truly independent biased source B. Then by the proof of Theorem 2 there exists a non-constant probabilistic polynomial-time prediction algorithm T: $\{0,1\}^{i-1} \rightarrow \{0,1\}$ for the i-th bit of S such that $\Pr_S(T = s_i | T = 1) \geq b + 1/n^k$, for some constant k. From T we construct the following POP test A: $\{0,1\}^{i-1} \rightarrow \{0,1,*\}$: $A = 1$ iff $T = 1$ and $A = *$ iff $T = 0$. Since T is non-constant, then $\Pr_S(T = 1) = \Pr_S(A \neq *) \geq 1/n^l$ for some constant l. We then get that by definition, S fails the POP test A. □

Proof of Proposition 4

Let A be any non-constant probabilistic polynomial-time algorithm: $\{0,1\}^{i-1} \to \{0,1\}$. Clearly:

(1) $\Pr_S(A = s_i) = \Pr_S(A = 1) \cdot \Pr_S(A = s_i | A = 1) + \Pr_S(A = 0) \cdot \Pr_S(A = s_i | A =)$.

The proposition results from the following easily proved two equivalences. We sketch their proofs in brackets:

1. $\Pr_S(A = s_i) \geq 1/2 + 1/n^{k_1}$, for some constant k_1, iff $\Pr_S(A = s_i | A = 1) \geq 1/2 + 1/n^{l_1}$ and $\Pr_S(A = s_i | A = 0) \geq 1/2 + 1/n^{l_2}$ for some constants l_1 and l_2. (If $\Pr_S(A = s_i) \geq 1/2 + 1/n^{k_1}$, then by (1) there exists a value $\alpha \in \{0,1\}$ such that $\Pr_S(A = s_i | A = \alpha) \geq 1/2 + 1/n^{k_1}$. This in turn implies the equivalence according to the Prediction Lemma. The other direction is an immediate consequence of (1).)

2. $ws(A, S, i) \geq 2 + 1/n^{k_2}$ iff $\Pr_S(A = s_i | A = 1) \geq 1/2 + 1/n^{l_1}$ and $\Pr_S(A = s_i | A = 0) \geq 1/2 + 1/n^{l_2}$ for some constants l_1 and l_2. (A direct result from the proof of Theorem 2). □

Proof of Proposition 5

It is obvious that a POP test can always simulate a next bit test (without ever outputting $*$) and therefore for any unbiased source S and any next bit test T there exists a POP test A, such that for every $1 \leq i \leq n$:

$$\Pr_S(A = s_i \mid A \neq *) = \Pr_S(T = s_i)$$

It remains to prove that inequality is also possible. To do so we shall construct an imperfect source S and demonstrate a POP test that does better than any next bit test. The source is the following:

1. The first $i - 1$ bits are independent unbiased coin flips.

2. Fix any $0 \leq \delta \leq \frac{1}{2}$.

$$\Pr_S(s_n = 1) = \begin{cases} \frac{1}{2} + \delta & \text{if } s_1^2 = 00 \\ \frac{1}{2} - \delta & \text{if } s_1^2 = 01 \\ \frac{1}{2} & \text{if } s_1 = 1 \end{cases}$$

Since the next bit test is a global test, for any next bit test T:

$$\Pr_S(T = s_n) \leq \frac{1}{4} \cdot (\frac{1}{2} + \delta) + \frac{1}{4} \cdot (\frac{1}{2} + \delta) + \left(\frac{1}{2}\right)^2 = \frac{1}{2} + \frac{\delta}{2}$$

The POP test A we shall use is: A=1 iff $s_1^2 = 00$; else A=$*$. Clearly

$$\Pr_S(A = s_n \mid A \neq *) = \frac{1}{2} + \delta$$

□