CS395T Pseudorandomness and Combinatorial Constructions
Lecture 23, April 10, 2001
Lecturer: Prof. David Zuckerman, Scribe: Vladimir Trifonov

Trevisan's Extractor

1 Idea

As we saw in Lecture 15, extractors are functions which take two inputs – an n-bit string sampled from a weak random source of min-entropy at least k (e.g. $k = n/2$), a uniform random seed of length d (e.g. $d = \text{poly-log}(n)$), and output an m-bit string (m as close to k as possible) which is close to uniform. On the other hand in Lecture 19 we saw the Nisan-Widgerson PRG (NW), which given a hard function in **EXP** is able to stretch poly-log(n) truly random bits to n bits which look random to every sufficiently small circuit.

The idea of Trevisan is that considering the hard function used in NW as input to the generator, i.e. the generator takes two inputs – the truth table of the hard function and a random seed, we obtain an extractor. Since a random function is hard, the first input will come from a weak random source. The only obstacle (which probably made this idea so hard to conceive) is that the output of a PRG has to look random to small circuits, whereas the output of an extractor has to be statistically close to uniform. Trevisan overcame this obstacle by strengthening the analysis of NW. Namely he noticed that the analysis of NW implies that given any test T (i.e. no restriction on the circuit size) which distinguishes the output of the generator from uniform, we can construct a small circuit which using oracle access to T computes the function used by the generator.

From this, the result of Nisan and Widgerson follows immediatelly. Even better – we are able to analyze the extracting capabilities of NW. Notice that the hard function used by NW has to be hard on average. Using list-decodable code, Trevisan obtains average case hardness and thus constructs an extractor which is able to extract $m = k^{1-\alpha}$ bits, for any k and α, using $d = \log^2 n/\log k$ bits. The additional randomness required by the Trevisan's extractor was the best among the extractors working for all possible min-entropies. On the other hand this extractor extracts a sublinear number of bits and there are extractors which extract a constant fraction of the randomness in the weak source or even all of it.

2 Construction

Our goal will be to construct an extractor producing an m-bit ε-uniform string from an n-bit source of min-entropy at least k using d additional truly random bits. Set $\delta = \varepsilon/m$.

The construction uses a code $C : \{0,1\}^n \to \{0,1\}^{\bar{n}}$, $\bar{n} = \text{poly}(n, 1/\delta)$, that has at most $1/\delta^2$ codewords in every ball of radius $(1/2 - \delta)\bar{n}$. From Lecture 10 we know that binary linear codes with distance at least $(1/2 - \delta^2)\bar{n}$ have this property and that we can construct such code using δ^2-biased distribution.

Let $l = \log \bar{n}$ and $\mathcal{S} = \{S_1, \ldots, S_m\}$ be the design required by the construction of NW, i.e.

- $S_i \subseteq [d]$,
- $|S_i| = l$, and
- $|S_i \cap S_j| \leq a = \log(k/2m)$.

Let $NW_{f,\mathcal{S}}(x) = f(x|_{S_1}) \circ \cdots \circ f(x|_{S_m})$ be the Nisan–Widgerson PRG using the design $\mathcal{S}$ and f as a hard function. Let $\bar{x} = C(x)$. The Trevisan's extractor is defined in the following way:

$$E(x,y) = NW_{\bar{x},\mathcal{S}}(y)$$

The subscript $\bar{x}$ in $NW_{\bar{x},\mathcal{S}}$ should be understood as the function with truth table $\bar{x}$.

3 Analysis

Theorem 1 *For every $T : \{0,1\}^m \to \{0,1\}$, there are at most $2^{m2^a}/\delta^2$ strings $x \in \{0,1\}^n$ such that*

$$|\Pr[T(E(x,U_d)) = 1] - \Pr[T(U_m) = 1]| \geq \varepsilon.$$

To prove this theorem we will use the following lemma.

Lemma 1 *For every $T : \{0,1\}^m \to \{0,1\}$ there exists $W \subseteq \{0,1\}^{\bar{n}}$, $|W| \leq 2^{m2^a}$, such that for every $f : \{0,1\}^l \to \{0,1\}$ satisfying*

$$|\Pr[T(NW_{f,\mathcal{S}}(U_d)) = 1] - \Pr[T(U_m) = 1]| \geq \varepsilon,$$

there is $w \in W$ which agrees with the truth table of f in at least $(1/2+\delta)\bar{n}$ bits.

Proof:(Theorem 1) Let W be the set whose existence is guaranteed by Lemma 1. By the properties of C, for every $w \in W$ there are at most $1/\delta^2$ strings $x \in \{0,1\}^n$ whose encoding agrees with w in at least $(1/2+\delta)\bar{n}$ bits. Therefore there are at most $|W|/\delta^2 \leq 2^{m2^a}/\delta^2$ elements of $\{0,1\}^n$ with the property given in Theorem 1. $\square$

Proof:(Lemma 1) Essentially repeats the proof of Nisan–Widgerson.

Let T be the distinguisher given in the condition of the Lemma. Using a hybrid argument we obtain a predictor that can predict the i-th bit of the output of the generator given the first $i-1$ bits. Fix the bits of the input to the generator outside of S_i so that its bias is preserved. By the requirements of the design $|S_i \cap S_j| \leq a$, for $j \in [i-1]$, i.e. there are 2^{2^a} possible ways in which the j-th bit of the output depends on the bits of the input after the fixing. If we knew the correct dependence on each $S_i \cap S_j$, then we can compute f from S_i. We do not, but there are at most $\left(2^{2^a}\right)^{i-1}$ possibilities. These possibilities give rise to W. Multiplying by 2 because of $|\cdot|$ we obtain that $|W| \leq 2 \cdot 2^{(m-1)2^a} \leq 2^{m2^a}$ as needed. $\square$