Handout6: $NP \subset PCP(\log n, (\log\log n)^{O(1)})$

May 28, 2003

Using the parallelization techniques of the last lecture, we shall complete the proof that $NP \subset PCP(\log n, (\log\log n)^{O(1)})$, by combining an outer PCP with an inner PCP. The only part of this proof that was not covered in full in the course is the total degree test. This is addressed to some extent in the homework.

Most likely, time will permit to begin describing the second inner PCP. It will probe only a constant number of bits in the inner PCP witness, but will be relatively inefficient in terms of randomness. The encoding scheme E for this inner verifier will no longer be degree-h multinomials over $\mathcal{F}^m$. Instead, it will be linear functions (total degree 1) over $\{0,1\}^n$. The linearity test that is associated with this code is somewhat simpler (both to describe and to analyse) than the total degree test that was used over $\mathcal{F}^m$. The sum check protocol (which appears to inherently require a nonconstant number of probes) will be replaced by a test that reads only a constant number of bits, but uses far more randomness.

Homework.

Recall the total degree test for a table A that supposedly represents the values of a multinomial in $\mathcal{F}^m$ of total degree d. It uses also a table B that for every line in $\mathcal{F}^m$ gives a degree d polynomial. The test picks a random line $l(t) = x + ty$ and a random value $t_r \in \mathcal{F}$, accesses the table B at location (x, y) to get the polynomial $q(t)$, evaluates l and q at t_r, and rejects if $q(t_r) \neq A(l(t_r))$.

Assume that A is quite close to some total degree d polynomial p. That is, $\Delta(A, p) = \delta$ for some small δ (e.g., $\delta = 1/4$). Show that in this case the test rejects with probability $\Omega(\delta) - O(d/|\mathcal{F}|)$.

(Here is a recommended outline to an instructive proof, though other proofs are also possible. Show that points on a random line are pairwise independent. Then using Chebyschev's inequality show that most random lines are "typical", in the sense that the fraction of points where A and p disagree is close to δ. Then use the fact that every univariate degree d polynomial q that disagrees with p on the random line disagrees with p almost everywhere, and then if the random line is typical, one gets that q disagrees with A on $\Omega(\delta)$ fraction of the points, when $|\mathcal{F}|$ is large enough.)