CS 880: Complexity of Counting Problems April 26, 2012

Lecture 26: The Giant Gadget

Instructor: Jin-Yi Cai Scribe: Aaron Gorenstein

This is a quick summary of the lecture.

We discuss the effect of replacing each edge with the giant gadget on page 84. Note that the picture depicts $G^{[1]}$, which is just about half the actual gadget. In reality there are 2 more "wings", both attached to the "spine" consisting of x, y, u, v. The u, v in the spine are the original $(u, v) \in E$. Note that there is no longer an edge directly between them! In this lecture we only discussed a quarter, the "bottom" wing present in the figure.

We have $A_{(0,u)(0,v)}$, for $u, v \in \mathbb{Z}_{\mathcal{Q}}$. Note that A is still bipartite: the upper-left and lower-right values are nonzero, while the other corners must be zero.

Going through each vertex in the figure, we account for their degrees. All vertices have degree $0 \mod N$ *except* the nodes a_i, b. They have degree r, relating to the $D^{[r]}$.

Recall that our goal is, with $\Lambda_r = \{x | D^{[r]}_{0,x} \neq 0\}$ (a coset), and $S = \{r | \Lambda_1 \neq 0\}$, we want to say that the value is in an exponential quadratic form. As an aside, the difficulty of verifying this gadget's qualities gives pause for $P \neq NP$: if the P-time verification takes so long, surely the process of making the proof must have been intractable!

We have a tremendous formula specifying the value of a particular $A_{u,v,x,y}$. Now we "harvest" all of the hard work from the previous parts of the paper, and dramatically simplify the sum. The summation over the w, for example, is simplified to an inner product, using the fact that some edges are conjugate-values of the other, and that the rows form a group. The inner products become either 0 (when non-equal) or m, as before. This furthermore places constraints on how x, y may relate to each other. Both w and z sums, for similar reasons, simplify to m terms.

The sum over c (which is in two parts) becomes m^{N-1}. The leftmost term, the product over $D^{[r]}_{(0,a_i)}$ simplifies by using the fact they are all roots of unity. Along the simplifications for w, z, c we concluded $a_i = a_j \forall i, j$, so the product over all a_i for D becomes just powering, and by the root of unity it becomes its converse: $\overline{D^{[r]}_{(0,a)}}$.

With all this we have a greatly simplified term for $A_{u,v,x,y}$. We simplify a bit further, and by using the fact that Λ_r is a coset, show that $u - v$ must be in the linear part of that set. So $u - v = a - b \in \Lambda_r^{lin}$, and we assume they have this property (otherwise the value is zero.

Now we introduce, on the bottom of page 86, the term T. We return to page 26, and spend some time discussing $\mathcal{D}_{1-4}$. We discussed primarily the purpose of $\mathfrak{a}$ (note the distinctive font!) introduced in $\mathcal{L}_3$. It is used to relate the product of D values to a function if $\omega_N^{\alpha} \cdot F$. This is to determine important properties of the term T.

The punchline is that with

$$F_{x,\tilde{b}} = \omega_{q_k}^{x_k^b} \tag{1}$$

we see that the product between two D terms only depends on x_k in particular, and no other x.

Amazingly, only *after* all this reasoning about T are we able to, in the middle of page 87, conclude that A (with our giant gadget) is actually symmetric! Unlike other gadgets, this one did not obviously produce a symmetric A. After this point, there are only 2 more pages left to prove hardness.

However, next week we will return to CSP, starting with Dyer's paper, which is available on the course webpage.