CS 880: Complexity of Counting Problems	May 1, 2012

Lecture 27: Dyer and Richerby

Instructor: Jin-Yi Cai Scribe: Aaron Gorenstein

This is a quick summary of the lecture. We are reviewing the main ideas and results in Dyer and Richerby's paper "An Effective Dichotomy for the Counting Constraints Satisfaction Problem."

Recall the Bulatov and Dalmua result, that for tractability you must have th Malt'sev polymorphism in $\langle\Gamma\rangle$. Recall that for a relation $R \subset D^n$ the map $\phi : D^k \to D$ is a polymorphism of R if

$$\forall (a_{1,n} \dots a_{1,n}) \dots (a_{k,1} \dots a_{k,n}) \in R \implies \phi(a_{1,1} \dots a_{k,1}) \in R \tag{1}$$

Recall that a relation $R \subseteq A_1 \times A_2$ is called rectangular if $(a, c)(a, d)(b, c) \in R \implies (b, d) \in R$.

We define $R \subseteq D^n$ for $n \geq 2$ is rectangular if every expression of R as a binary relation in $D^k \times D^{n-k}, 1 \leq k < n$ is rectangular. Note that our real goal is just that any partition of $[n]$ as 2 disjoint parts is our goal. When we go to $\langle\Gamma\rangle$, permutation grants us the rest. If every $R \in \langle\Gamma\rangle$ is rectangular than we say Γ is strongly rectangular, but we may use that term interchangeably.

We now spend some time showing the equivalence between $\langle\Gamma\rangle$ being rectangular and Γ having a Malt'sev polymorphism. Dyer and Richerby really capture the notion that, if $\alpha 0 \dots, \alpha 1 \dots, \beta 1 \dots$ are in, then $\beta 0 \dots$ must also be in. Note that we do need to consider the tail (the $\dots$ following) in the sense that the contents of the tail are dictated by the prefix α, β.

Now we move on to the technical stuff: We can decide if Γ is strongly rectangular by checking if one of the maps $\phi : D^3 \to D$ is a Malt'sev polymorphism. The runtime is clearly bounded by q^{q^3}, which (q being the size of the domain) is a constant. In the paper they say it is $O(\|\Gamma\|^4)$, where $\|\Gamma\|$ is the "table form": if all S relations are of arity R, the table would be size $S \times R$ (times the bit size of each domain element). We already know that a Malt'sev polymorphism is *required* (otherwise the problem is #P-hard, so let's assume from now on it exists.

We define $\sim_i$ as an equivelance relation on domain D. It is the case that $a \sim_i b$ if and only if a, b share a common prefix—this is foreshadowing use of the Malt'sev polymorphism. The i is the index at which a, b must appear. So if $a \sim_i b$ then $\mathbf{x}a\mathbf{y}, \mathbf{x}b\mathbf{z} \in R$. Observe that if $i = 1$ then we have an empty prefix. We define projection:

$$\mathrm{Pr}_i R = \{a \in D | \exists \mathbf{x}a\mathbf{y} \in R, \mathbf{x} \in D^{i-1}, \mathbf{y} \in D^{n-1}\} \text{ for } 1 \leq i \leq n. \tag{2}$$

Note that $\sim_i \in \mathrm{Pr}_i R$.

Consider the following lemma: $\sim_i$ is an equivalence relation on $\mathrm{Pr}_i R$ for rectangular R. Using the fact that we have a Malt'sev polymorphism, it is not hard to show that this is true. For the rest of the lecture, we will go through some of this paper, and then transition to the more general result by Cai, Chen, and Lu. But the following numbers refer to the Dyer-Richerby paper.

Corollary 11: $\exists$ a common prefix $\mathbf{u}_{i,k} \in D^{i-1} \forall a \in \varepsilon_{i,k} (k \in [\kappa], i \in [n])$. This allows us to find a common prefix for a when we need it.

Lemma 12: Let H be an n-ary relation. If $I \subseteq [n]$, then $c\ell_\phi \mathrm{Pr}_i H = \mathrm{Pr}_i c\ell_\phi H$, where ϕ is any Malt'sev polymorphism. This is a very useful lemma, and the proof is quite straightforward. Basically, think about it. Any resulting tuple from either side must also be made on the other side, basically. The proof does not even need the Malt'sev property.

Lemma 13: given $S = \{\mathbf{t_1}, \dots, \mathbf{t_s}\}$ set of n-tuples. So for $I \subseteq [n]$ we can compute $T \subseteq c\ell_\phi(S)$ such that $\mathrm{Pr}_I c\ell_\phi(S) = c\ell_\phi \mathrm{Pr}_I(S)$. If $\ell = |\mathrm{Pr}_I c\ell_\phi(S)|$ then T can be computed in $O(n\ell^3 + s\ell^4)$.

Basically, T is a witness to the projection result, a small collection of things that establishes the equivalence classes. We compute it following the algorithm on page 11. We produce a new tuple, and *only if it is unique in the range of elements defined by I* do we keep it. Thus, T does not get too big.

We conclude with a key definition: a *frame*. For rectangular R on D^n, a set $F \subseteq R$ is a frame if:

- $\mathrm{Pr}_i F = \mathrm{Pr}_i R \forall i \in [n]$

- $\exists \mathbf{v}_{i,k} \in D^{n-i}$ for each equivalence class $\varepsilon_{i,k}$ of $\sim_i, i \in [n], k \in [\kappa_i]$ such that $\forall a \in \varepsilon_{i,k} \exists \mathbf{w}_a \in F, \Pr_{[i]} \mathbf{w}_a = \mathbf{v}_{i,k} a$.

What this means is that for any a, b that share a prefix in R (they are equivalent), there is at least one witness to that in F. We really use this table for the purpose of the *witness function*, some function $w : D \times [n] \to D^n \to R \cup \{\perp\}$. Furthermore, if $a \sim_i b$ then $\Pr_{[i-1]} w(a, i) = \Pr_{[i-1]} w(b, i)$, they all have the *same* witness. The image of the witness function is the frame.

We conclude with this homework question: how do you find a good frame for the relation that equals D^n?