Lecture 9: Geiger's Theorem

Instructor: Jin-Yi Cai Scribe: Aaron Gorenstein

Last time we proved the following two theorems:

Theorem 1 If $\mathcal{F}$ is closed (under composition, but technically also projection), then Pol Inv $\mathcal{F} = \mathcal{F}$. Recall we assume all this is on a finite domain set.

Theorem 2 (Extending Partial Polymorphisms) Every partial polymorphism of a clone ($\wedge$, $\exists$, append, repeat, and permutation) can be extended to a full polymorphism. Being partial means that the it has to hold only when the function is defined for every row in the table.

1 Geiger's Theorem

Today we will focus on the following theorem:

Theorem 3 (Geiger's Theorem)

$$\langle \Gamma \rangle = \text{Inv Pol } (\Gamma) \tag{1}$$

Recall the (straightforward) fact that Pol Γ = Pol $\langle \Gamma \rangle$. This completes the foundation of the Bulatov et al. theory of polymorphisms in this context.

Proof. Proof What do we want? We want, for any relation $Q \notin \langle \Gamma \rangle$, a polymorphism f of Γ such that $Q \notin$ Inv f. Note that this is the main reason we wanted theorem 2, because it suffices now to have a *partial* polymorphism of Γ that does not commute with Q. This is basically providing a witness, a table, that it does not commute. Also note that it is enough to show that the right-hand side of theorem 3 is no bigger—we already know that it must be at least as big.

Now we introduce a crucial construction:

$$P = \bigcap_{P' \in \langle \Gamma \rangle, P' \supseteq Q} P' \tag{2}$$

which is nonempty (since A^n is there) and a finite intersection (because it is over relations of arity n over a finite domain set). As $P \in \langle \Gamma \rangle$ (because $\langle \Gamma \rangle$ is closed under intersection), we essentially have a minimal containment. Hence $P \neq Q$, and therefore $\exists t \in P - Q$. We list all tuples of Q of arity n in a table. Consider the table:

$$N = \left[\begin{pmatrix} \vdots \end{pmatrix} \in Q \ldots \begin{pmatrix} \vdots \end{pmatrix} \in Q \right] \xrightarrow{f} t \notin Q. \tag{3}$$

If f were a partial polymorphism of Γ, this would be a witness we want for Q. So if we can define f just on this (a partial polymorphism) then we're done.

Here is a problem: what if there are 2 identical rows in the table N? We can address this by defining the set of tuples $Q' = Q$ with "one row erased". The formal meaning of erasing a row is as follows: Say i, j are the indices of the two identical rows. We define Q' on $(x_1, x_2, \ldots, x_i, \ldots, \hat{x_j}, \ldots, x_n)$, so that is on $n - 1$ arity.

$$Q'(x_1, x_2, \ldots, x_i, \ldots, \hat{x_j}, \ldots, x_n) = \exists x_j [Q(x_1, \ldots, x_i, \ldots, x_j, \ldots, x_n) \wedge x_i = x_j] \notin \langle \Gamma \rangle \tag{4}$$

So we establish that $Q' \notin \langle \Gamma \rangle$, otherwise we could obtain Q by repeat and permutation operators on Q'.

Now we proceed through induction: We claim that since $Q' \notin \langle \Gamma \rangle \exists$ a polymorphism g of $\langle \Gamma \rangle$ such that $Q' \notin$ Inv g. That means there's a table such that all the columns are in Q', but it produces a tuple not in Q'. What does that table look like? The table $\exists N'$ consists of a sequence of columns of N after the j^{th} row was removed and $g_{\rightarrow}(N') \notin Q'$ (by induction). Therefore $g_{\rightarrow}(N)$ is defined! We just put in that one row with repetition, that being one of the operations we're closed under. But wait: what is our base case?

Our induction is on arity n. At the base case, $n = 1$, there are no repeated rows. Our base case is this: suppose N has no identical rows. Now we define our f as on every row to produce t. So $f_{\rightarrow}(N) = t$. That's the *only* definition of f. It is a partial function! All that remains now is to show that f is a partial *polymorphism*. We shall do so with contradiction.

Let P_1 be a relation in $\langle\Gamma\rangle$ such that $f \notin \text{Inv }(P_1)$. Say P_1 is of minimal arity—clearly there is such a minimal. Then $\exists N_1$ such that $f_{\rightarrow}(N_1)$ is defined and the columns of N_1 are in P_1 and $f_{\rightarrow}(N_1) \notin P_1$. We first argue that there are no repeated rows: if there are, then let $P_2 = \exists x_j(P_1 \wedge x_i = x_j)$. Then we say that N_2 is a table for P_2, where N_2 is just N_1 with the j^{th} row removed. So then $f_{\rightarrow}(N_2)$ is defined, $P_2 \in \langle\Gamma\rangle$, $N_2 \subseteq P_2$, but $f_{\rightarrow}(N_2) \notin P_2$. That contradicts our claim of minimal arity, so we know that any minimal example has no repeated rows.

We've defined f on N with no repeated rows, and we have a counter-example. So our witness is some rows of N, maybe under permutation or deletion. So N_1 has no repeated rows and $f_{\rightarrow}(N_1)$ is defined, but f is only defined on rows of N. We conclude: N_1 consists of some subsets of rows of N, possibly permuted.

Take those rows, and say they're the first l rows, $1 \leq l \leq n$. We define the following relation:

$$P_1 \times A^{n-l}. \tag{5}$$

In other words, just those l entries, and the anything else we'd like. So we have an n-arity relation in $\langle\Gamma\rangle$. Certainly this contains Q, we are more permissive than before. So $P_1 \times A^{n-l} \supseteq Q$. So this must contain the minimal such relation, namely P, so $P_1 \times A^{n-l} \supset P \ni t$. Contradiction! We know t cannot be in $P_1 \times A^{n-l}$. □

2 Roadmap

Our next goal is to show that $\#CSP(\Gamma) \equiv_T \#CSP(\langle\Gamma\rangle)$. We can implement conjunction ($\wedge$) really easily, but simulating the existential quantifier will be complex. There we try to say, for $Q(\overrightarrow{y}) = \exists x P(x, \overrightarrow{y})$.

One direction of the equivalence is obvious, we're really trying to show that $\#CSP(\Gamma) \geq_T \#CSP(\langle\Gamma\rangle)$. We will do this with replacement by induction, one at a time.

This will ultimately show:

$$\#CSP(\Gamma) \equiv_T \#CSP(\langle\Gamma\rangle) \leftrightarrow \text{Pol }\langle\Gamma\rangle = \Gamma. \tag{6}$$

This raises the question: what qualities in the polymorphism world characterize tractability? We will get the partial result that the *Malt'sev Polymorphism* is a necessary condition for tractability, it must be that $m(x, y, z) \in \text{Pol }\Gamma$.

Such a polymorphism is equivalent to saying it is "congruence permutable", there's an equivalence relation in $\langle\Gamma\rangle$ that becomes a predicate you can use to express equivalence in some way.

The idea is that it defines the congruence pair of α and β:

$$\alpha \circ \beta = \{(a, c) | \exists b(a, b) \in \alpha, (b, c) \in \beta\} = \beta \circ \alpha. \tag{7}$$

As on exercise, show that saying $\alpha \circ \beta = \beta \circ \alpha$ is the same as saying $\alpha \circ \beta \subseteq \beta \circ \alpha$.

3 The Incompressible Cleverness of M.P.

Recall: we define a ternary relation:

$$P = \{(x, y, z) | (x, y) \in \alpha, (y, z) \in \beta\} \in \langle\Gamma\rangle \tag{8}$$

And given

$$\begin{array}{ccc} a & & c \\ b & \rightarrow & \star \\ c & & a \end{array} \tag{9}$$

we want to find a value for $\star$ such that $(a,b) \in \alpha, (c,\star) \in \alpha$ and $(b,c) \in \beta, (\star, a) \in \beta$.

We exploit the equivalence relation from the Malt'sev polymorphism!

$$\begin{array}{cccc} a & \underline{a} & c \rightarrow c \\ a & \underline{b} & c \rightarrow d \\ a & \underline{c} & c \rightarrow a \end{array} \tag{10}$$

By the Malt'sev polynomial, all 4 of these columns must be in P. Critically, there must exist that useful d.

We will show that it is iff and equivalent to a quality called "rectangularity" (Dyer-Richerby).

3.1 A Glimpse of Rectangularity

Consider huge tuples, and sub-tuples (a,c) (so combined they are a single valid 100-tuple, for example), and another (b,d). We want to say that if:

$$\begin{aligned} (a,c) &\in R \\ (a,d) &\in R \\ (b,d) &\in R \end{aligned}$$

then we can conclude that $(b,c) \in R$ as well. If you draw these out, they look like a rectangle, hence the name. If all relations in the clone are rectangular, we call it *strongly* rectangular.

This property has a strong affinity with congruence. Think in terms of a graph: the quality is if they are a path that shares a tail. Then we have an equivalence relation! We can define α as

$$\alpha = \{(x,y),(x,z) | (x,y) \in R, (x,z) \in R\} \tag{11}$$

which means "share tail", and we can similarly define β for "share head". See that $(a,c),(b,d)$ are in the *composition* of α, β.

4 Using Malt'sev

The main theorem of Bulatov is that the existence of the Malt'sev polymorphism is necessary for tractability. (Technically we say for not-#P-hardness, because if P=#P the first statement becomes nonsense.)

Bulatov conjectured that it was also sufficient, but found a counter-example.

4.1 Using Malt'sev

We will sketch how the lack of a Malt'sev polymorphism causes hardness. If $\langle \Gamma \rangle$ does not have a Malt'sev polymorphism, then there must exist two congruences α, β such that $\alpha \circ \beta \neq \beta \circ \alpha$. In turn that means $\exists (a,b,c) s.t.\ (a,b) \in \alpha,\ (b,c) \in \beta$, but there is no b' such that $(a,b') \in \beta, (b',c) \in \alpha$.

We then know that $(a,c) \in \alpha \circ \beta$, but $(c,a) \notin \alpha \circ \beta$. *That* relation, $[\exists (a,c) \in R, (c,a) \notin R]$ is reflexive but not symmetric. What's the simplest embodiment of such a relation? Expressing the relation as a look-up matrix:

$$T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}. \tag{12}$$

That is a graph homomorphism! Not only that, but it is for partial order on $\{0, 1\}$. Flipping the diagonal gets independent set. Consider adding to each edge this situation: $u \to v \leftarrow w$, so we put T and T^T. Their product is

$$\begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}. \tag{13}$$

By "thickening" (adding identical, repeated edges) we power each element in the matrix. This causes the 2 to blow up to such a degree that we "notice" and can mod out that value. Thus we can count vertex covers! Or independent sets, as the case may be.