Oded: August 12, 2014

On Gentry's ECCC TR14-105

As defined in Sec. I.B (and more generally in Def. 2 of Sec. II.C), a product program (over a multiplicative group G), is a sequence of triples $P = (\mathtt{inp}(i), a_{i,0}, a_{i,1})$, where $\mathtt{inp} : [n] \to [\ell]$ and $a_{i,0}, a_{i,1} \in G$, and its value at $x \in \{0,1\}^\ell$ is defined as $P(x) = \prod_{i=1}^{n} a_{i,x_{\mathtt{inp}(i)}}$. For such product programs, one may consider three natural computational problems:

1. Evaluation: Given such a program P and an assignment x to its variables, compute the value $P(x)$.

2. Satisfiability: Given such a program P and a value v, determine whether there exists an assignment x such that $P(x) = v$. (The search problem version is to find such an x.)

3. Summation (reminiscent of counting): Given such a program P, compute the value $\sum_{x\in\{0,1\}^\ell} P(x)$.

Product problems over different algebras have different expressive power, and this expressive power lower bounds the complexity of the foregoing computational problems. Barrington's celebrated result says that, for every non-solvable group G, the computation of depth d Boolean circuits (of bounded fan-in) can be expressed by G-programs of length exponential in d. The current paper observes that the corresponding counting problem (i.e., counting the number of satisfiable assignment to such a circuit) reduces to the foregoing summation problem (see Fact 1 in Sec. II.D); but this is not the main point of the paper.

The main point of the paper is that the summation problem can be reduced to the computation of the determinant (or the permanent) of matrices over the same algebra. This is what Thm. 5 (of Sec. III.B) says. Actually, it says it only for the determinant, but as we shall see, for the matrices used in the reduction the determinant equals the permanent.

I think it is beneficial to somewhat simplify and rephrase the exposition of the matrix M_P considered in the beginning of Sec. III.B (where also C_k and π_1 are defined). Specifically, first, assume without loss of generality, that all C_k's are of odd size (greater than one); this can be justified just as the author justifies that they are all of size greater than one (i.e., by possibly introducing "dummy" or "neutral" program lines i such that $\mathtt{inp}(i) = k$ and $a_{i,0} = a_{i,1}$ equal the identity element). More importantly, note that M_P has exactly two non-zero entries per each line $i \in [n]$ – the entry $M_P[i,i] = a_{i,0}$ and the entry $M_P[i,\pi_1(i)] = a_{i,1}$. (In the paper, for each $k \in [\ell]$, one of the rows in the π_1-cycle $(i_{k,1}, ..., i_{k,|C_k|})$ carries the sign $(-1)^{|C_k|-1}$, but by the first simplification this is no longer needed.)

Next, recall that each permutation $\sigma : [n] \to [n]$ consists of a set of disjoint cycles. Hence, each permutation σ that corresponds to a non-zero term in the determinant (or permanent) must consist of singleton cycles and cycles that are generated by π_1. The reason is that the term $\prod_{i=1}^{n} M_P[i,\sigma(i)]$ is non-zero only if $\sigma(i) \in \{i, \pi_1(i)\}$ for all i's, whereas $\sigma(i) = \pi_1(i)$ implies that σ agrees with π_1 on the cycle $C_{\mathtt{inp}(i)} = (i, \pi_1(i), ..., \pi_1^{|C_{\mathtt{inp}(i)}|-1}(i))$. This means that such σ must be consistent with some input x; that is, its singleton cycles correspond to the zero-entries of x, whereas its π_1-generated cycles correspond to one-entries of x.

Hence, the value of the permanent of M_P equals the value $\sum_{x\in\{0,1\}^\ell} P(x)$. The same holds for the determinant of M_P, because each σ that is consistent with some x has sign 1 (since it is a collection of odd-length cycles).