Ryan Williams: The Orthogonal Vectors Conjecture and Non-Uniform Circuit Lower Bounds

Nikolai Chukhin, Alexander S. Kulikov, Ivan Mihajlin, Arina Smirnova: Deriving Nonuniform Lower Bounds from Uniform Nondeterministic Lower Bounds

presented by Petr Chmel

Theorem 1 (Win-win lower bounds)**.**
At least one of the following lower bounds holds:

1. E^{NP} does not have Valiant series-parallel circuits of size $\mathcal{O}(n)$.
2. There is an $\varepsilon > 0$ such that Boolean Inner Product on n-bit vectors does not have $2^{\varepsilon n}$-size ETHR$\circ$ETHR circuits. Moreover, there is an explicit monotone function $f \in$ coNP that requires monotone circuits of size $2^{\Omega(n/\log n)}$.

Hypothesis 1 (Strong Exponential Time Hypothesis (SETH), Impagliazzo & Paturi'01)**.**
For any $\varepsilon > 0$, there exists a k such that k-SAT cannot be solved in time $\mathcal{O}(2^{(1-\varepsilon)n})$.

Hypothesis 2 (Nondeterministic Strong Exponential Time Hypothesis (NSETH), Carmosino et al.'16)**.**
For any $\varepsilon > 0$, there exists a k such that k-SAT cannot be solved in *co-nondeterministic* time $\mathcal{O}(2^{(1-\varepsilon)n})$.

Definition 1 (Orthogonal Vectors problem (OV))**.**
Given two sets $A, B \subseteq \{0,1\}^d$ with $|A| = |B| = n$, decide whether there exist $a \in A$ and $b \in B$ such that $\sum_{i=1}^{d} a_i b_i = 0$.

Conjecture 1 (Orthogonal Vectors conjecture (OVC))**.**
For every $\varepsilon > 0$ there exists a constant $c \geq 1$ such that the Orthogonal Vectors problem cannot be solved in time $n^{2-\varepsilon}$ on instances with $d = c \log n$.
(Note that the obvious algorithm runs in time $\mathcal{O}(n^2 d)$.)

Theorem 2 (SETH implies OVC, Williams'04)**.**
There exists an algorithm that, given a CNF formula φ with n variables and m clauses outputs two sets $A_\varphi, B_\varphi \subseteq \{0,1\}^m$ such that $|A_\varphi| = |B_\varphi| = 2^{n/2}$ and φ is satisfiable if and only if A_φ, B_φ form a YES instance of OV. The algorithm runs in time $\mathcal{O}(mn \cdot 2^{n/2})$.
In other words, faster algorithms for OV imply faster algorithms for SAT (and breaking SETH).

Theorem 3 (The first lower bound, Carmosino et al.'16 + Jahanjou, Miles & Viola'15)**.**
If NSETH is false, then there exists a function family in E^{NP} that requires Valiant series-parallel circuits of size $\omega(n)$.

Part 1: matrix decompositions (Williams)

Theorem 4 (The second lower bound, part 1)**.**
If the disjointness matrix for d-bit vectors has weak equality rank at most $f(d)$, where $f(d)$ is a subexponential function, then $\forall c \geq 1, \varepsilon > 0$, OV on n vectors in $c \log n$ dimensions can be solved in time $n^{1+\varepsilon}$ deterministically.

Definition 2 (Equality matrix, weak equality rank)**.**
A matrix $A \in \{0,1\}^{m \times n}$ is an equality matrix if there exist $u \in \mathbb{N}^m, v \in \mathbb{N}^n$ such that $A[i,j] = 1 \Leftrightarrow u[i] = v[j]$.
The weak equality rank of a matrix A is the smallest number of equality matrices $M_1, \ldots, M_r$ with constants $\alpha_1, \ldots, \alpha_r$ such that $A[i,j] = 0 \Rightarrow \sum_{k=1}^{r} \alpha_k M_k[i,j] = 0$ and $A[i,j] = 1 \Rightarrow \sum_{k=1}^{r} \alpha_k M_k[i,j] \neq 0$.

Definition 3 (Exact threshold function)**.**
An exact threshold function $f : \{0,1\}^n \to \{0,1\}$ is defined by weights $\alpha_1, \ldots, \alpha_n, t \in \mathbb{R}$ such that $f(x_1, \ldots, x_n) = 1 \Leftrightarrow \sum \alpha_i x_i = t$.

Theorem 5 (Circuits and weak equality rank, Williams'18)**.**
For a function $f : \{0,1\}^{2n} \to \{0,1\}$ and its descriptor matrix $M_f \in \{0,1\}^{n \times n}$ such that $M_f[x,y] = f(xy)$, we define $\neg M_f[x,y] = 1 - M_f[x,y]$. Then, if f has an ETHR $\circ$ ETHR circuit of size s, then $\neg M_f$ has weak equality rank at most $s+1$.

Definition 4 (Satisfying pairs problem)**.**
Let $\mathcal{M}$ be a family of matrices $M_d \in \{0,1\}^{2^d \times 2^d}$. The $\mathcal{M}$-Satisfying-Pairs problem is defined as follows: on input $d \geq 1$ with two sets $L, R \subseteq [2^d]$, are there $i \in L, j \in R$ such that $M_d[i,j] = 1$?

Theorem 6 (Algorithms for Satisfying pairs)**.**
Let $\mathcal{M}$ be a family of matrices $M_d \in \{0,1\}^{2^d \times 2^d}$. Suppose that M_d has a weak equality rank at most r with an explicit rank decomposition. Then, we can solve $\mathcal{M}$-Satisfying-Pairs with $|L| = |R| = n$ in dimension d in randomized time $r \cdot n \cdot \mathrm{poly}(d, \log n)$ and space $n \cdot \mathrm{poly}(d, \log n)$. We can also make it deterministic in time $r^2 \cdot n \cdot \mathrm{poly}(d, \log n)$.

Theorem 7 (Uniformization)**.**
Suppose that for some fixed k, r, the disjointness matrix on k-bit strings has weak equality rank r. Then for all $d \geq k$, the disjointness matrix on d-bit strings has weak equality rank at most $\mathcal{O}(r^{d/k})$ with an explicit rank decomposition.

Part 2: Monotone circuits (Chukhin, Kulikov, Mihajlin, Smirnova)

Theorem 8 (The second lower bound, part 2)**.**
If NSETH is true, then there exists an explicit monotone Boolean function family in coNP with monotone circuit size $2^{\Omega(n/\log n)}$.

Definition 5 (Monotone functions and circuits)**.**
A function $f : \{0,1\}^n \to \{0,1\}$ is monotone if changing a zero in the input to one never decreases the output value.
A Boolean circuit is monotone if it only uses AND and OR gates.

Lemma 1.
Let φ be a CNF formula. Then φ is *unsatisfiable* if and only if $(A_\varphi, \overline{B_\varphi})$ from the Williams reduction can be separated by a monotone function.

More generally, a pair $A, B \subseteq \{0,1\}^d$ has no orthogonal pair if and only if there exists a monotone function $f : \{0,1\}^d \to \{0,1\}$ such that $f[A] = \{0\}, f[\overline{B}] = 1$, where $\overline{B} = \{\overline{b} \colon b \in B\}$.

Lemma 2.
There exists a parameter $\ell \in \mathcal{O}(n \log n)$ and an injective encoding e that encodes all k-CNF formulas with n variables and at most βn clauses into ℓ-bit strings with exactly $\ell/2$ ones and both encoding and decoding can be done in time polynomial in n.

Bonuses

Theorem 9 (Sparsification lemma, Impagliazzo, Paturi & Zane'01)**.**
For any $k \geq 3$ and $\varepsilon > 0$, there exists $\alpha = \alpha(k, \varepsilon)$ and an algorithm that, given a k-CNF formula φ over n variables, outputs $t \leq 2^{\varepsilon n}$ formulas $\varphi_1, \ldots, \varphi_t$ in k-CNF such that all formulas have at most n variables and at most αn clauses, annd φ is satisfiable iff at least one of φ_i is satisfiable. The algorithm runs in time $\mathcal{O}(\mathrm{poly}(n) \cdot 2^{\varepsilon n})$.