Lower Bounds for Bounded Depth Frege

Ivan Monakhov

Academical Physico-Technical University

Plan of the Presentation

General notions and notations

Buss-Pudlák Games

Strategy for Sam using Partial functions

Covering Partial Functions and k-transformations

Sam's Strategy provided by k-transformations

Switching Lemma and Existence of k-transformations

Lower Bound for Frege Proofs

Logical Language

Definition

Our logical language will be restricted to

- Constants 0 (false) and 1 (true).
- Connectives $\{\vee, \neg\}$, $\vee$ is allowed to have unbounded fan-in.

$\wedge$ is a shorthand for $\neg \vee \neg$, and $A \Rightarrow B$ for $\neg A \vee B$.

Definition

The allowable formulas are defined inductively:

1. A literal (either a variable or its negation) is a formula.
2. If A is a formula, then so is $\neg A$.
3. If Γ is a finite set of formulas, then so is $\vee\Gamma$.

We use $A \vee B$ to mean $\vee\{A, B\}$.

Frege System

Definition

Frege system **H** is complete proof system over the basis $\{\vee, \neg\}$

1. Excluded Middle axiom: $\frac{}{A \vee \neg A}$
2. Weakening Rule: $\frac{A}{A \vee B}$
3. Merging Rule: $\frac{\vee(\{\vee\Gamma\} \cup \Delta)}{\vee(\Gamma \cup \Delta)}$
4. Unmerging Rule: $\frac{\vee(\Gamma \cup \Delta)}{\vee(\{\vee\Gamma\} \cup \Delta)}$
5. Cut Rule: $\frac{(A \vee B), (\neg A \vee C)}{B \vee C}$

By $\frac{\phi_1 \ldots \phi_k}{\psi}$ we denote that ψ can be derived from $\{\phi_1, \ldots \phi_k\}$.

Depth of the Formula and Proof

Definition

The *depth* of a literal is 0, the *depth* of a formula ϕ is the maximal number of alternations of connectives in it and the *size* of the formula is the number of occurences of connectives.

We denote by $d(\phi)$ the depth of formula ϕ.

Definition

A Frege proof of a formula ϕ is a sequence of depth d formulas $\pi = \{\phi_1, \ldots \phi_s, \phi\}$, where each formula is either an excluded middle axiom, or is derived from previous lines by other rule.
The *size* of a proof is the sum of the sizes of formulas in it.
The *depth* of the proof is the maximal depth of formulas.

The Pigeonhole Principle

Fix sets D, R: $D \cap R = \emptyset$, $|D| = n + 1$, $|R| = n$, and denote $S = D \cup R$.
Our set of connectives is $\{\vee, \neg\}$, so we use a notation $\wedge(\phi_1, \dots, \phi_k)$ as a shorthand for $\neg(\vee(\neg\phi_1, \dots, \neg\phi_k))$.

Definition

The pigeonhole principle of size n, denoted PHP_n, is the disjunction of four sets of formulas:

$$\begin{array}{ll} \neg \bigvee_{j \in R} p_{ij},\ i \in D & p_{ik} \wedge p_{jk},\ i \neq j \in D,\ k \in R \\ \neg \bigvee_{i \in D} p_{ij},\ j \in R & p_{ij} \wedge p_{ik},\ i \in D,\ j \neq k \in R \end{array}$$

over the variable set p_{ij}, $i \in D$, $j \in R$. Each variable p_{ij} states whether pigeon i occupies pigeonhole j.

Proofs as Games

Under the definition, introduced by Pudlák and Buss,

Definition

The Frege proof of a tautology Φ is a two player game.

- Pavel claimes that Φ is a tautology.
- Sam says that he knows an assignment α setting Φ to 0.
- In round t Pavel presents Sam a Boolean formula ϕ_t.
- Sam answers with a bit b_t, wich is the “value” of $\phi_t(\alpha)$.
- Pavel needs to present an *immediate contradiction*.

Immediate Contradiction

Let B be a set of Boolean gates. In our case $B = \{\neg, \vee\}$.

Definition

An *immediate contradiction* with respect to B is a set of formulas $\psi, \phi_1, \ldots, \phi_k$ and a set of bits $a, b_1, \ldots, b_k$:

1. ψ is $g(\phi_1, \ldots, \phi_k)$, where $g \in B$.
2. Sam was asked formulas $\psi, \phi_1, \ldots, \phi_k$,
   and gave answers $a, b_1, \ldots, b_k$.
3. $a \neq g(b_1, \ldots, b_k)$.

If a set of answers $b_1, \ldots, b_S$ to a set of queries $\phi_1, \ldots, \phi_S$ includes no immediate contradiction as a subset, we call these answers *locally consistent*.

Game Tree

- Frege Proof as the game is a binary tree, called *game tree*. Nodes are labeled by queries and edges by Sam's answers. The root is labeled Φ and has a single edge labeled 0.
- We say that game tree *covicts* Sam if every leaf is labeled by an immediate contradiction.
- A proof has depth d if all queries are depth d formulas.
- *Height* of the proof is the length of longest path from the root to a leaf. The *size* of the proof is the number of nodes.

Theorem

For any Frege system $\mathcal{F}$ there exist integer c:
If Φ has a standard $\mathcal{F}$-proof of size S and maximal depth d, then Φ has a Buss-Pudlák proof of height $\log(S) + O(1)$ and depth $d + c$ and each query is of size at most S.

Partial Functions

Definition

Let S be a set, $D \subseteq S$ and $f : D \rightarrow \{0,1\}$ a function on D.
The ordered pair (D,f) is called a partial Boolean function on S.
The set D is the domain of f, denoted by $\mathrm{Dom}(f)$.
For any set S, let

$$\Delta^S = \{(D,f) | D \subseteq S, f : D \rightarrow \{0,1\}\}$$

For any (D,f) and $b \in \{0,1\}$, $f^{-1}(b) = \{x \in D | f(x) = b\}$.

Transformation of Formulas

Let $\mathcal{T}$ be the game-tree for tautology Φ, proposed by Pavel. Sam applies a transformation, mapping each formula $\phi \in \Sigma_{\mathcal{T}}$ to partial function (D_ϕ, f_ϕ), that satisfies the conditions:

1. $\forall x \in D_\Phi, f_\Phi(x) = 0$.
2. There exists a branch $((\phi_1, b_1), \ldots, (\phi_s, b_s))$ in the game-tree $\mathcal{T}$:

$$\bigcap_{i=1}^{s} (f_{\phi_i})^{-1}(b_i) \neq 0$$

3. For any $\Omega \subseteq \Sigma_{\mathcal{T}}$, if there exists $x \in \cap_{\phi \in \Omega} D_\Phi$, then the answers $(f_\phi(x))_{\phi \in \Omega}$ to the queries $(\phi)_{\phi \in \Omega}$ are locally consistent.

Sam's Strategy

Theorem

Let Φ be a formula and $\mathcal{T}$ a game-tree for Φ. If there exists a set S and a transformation $\phi \overset{\Gamma}{\mapsto} (D_\phi, f_\phi)$: conditions 1,2 and 3 are satisfied, then the game-tree does not convict Sam.

Proof.

- Consider a branch $((\phi_1, b_1), \ldots, (\phi_s, b_s))$ of $\mathcal{T}$ provided by 2.
- Choose any $x \in \bigcap_{i=1}^{s} (f_{\phi_i})^{-1}(b_i)$. Sam answers Pavel's queries $\phi_1, \ldots \phi_s$ along this branch with $b_1, \ldots, b_s$ respectively.
- By 1 Sam answers Pavel's first query $\phi_1 = \Phi$ with $b_1 = 0$.
- Since $x \in \bigcap_{i=1}^{s} \mathrm{Dom}(f_{\phi_i})$, Sam's responses to Pavel's queries along this branch are locally consistent by 3.

Matching and Minimal Matching

- Let D, R be sets: $D \cap R = \emptyset$, $|D| = n+1$, $|R| = n$, and denote $S = D \cup R$. A *matching* between D and R is set of mutually disjoint unordered pairs $\{i, j\}$.
- π cover a vertex i if $\{i, j\} \in \pi$ for some $j \in S$.
  $V(\pi)$ is the set of vertices covered by π.
- For any set $I \subset S$, if π is a matching that covers I but does not cover I on the removal of an edge from it, then π is called *minimal matching* that covers I.
- M^S denotes the set of matchings between D and R.
  For any $I \subseteq S$: $D \not\subseteq I$, define

$$\text{Cover}(I) = \{\pi \in M^S \mid \pi \text{ covers all vertices in } I\}$$
$$\text{MinCover}(I) = \{\pi \in M^S \mid \pi \text{ is a minimal matching that covers } I\}$$

Covering Partial Functions

Note that for all $\pi \in$ MinCover(I), $|\pi| \leq |I|$.

Theorem

Let $S = D \cup R$, where $|D| = n + 1$, $|R| = n$ and $D \cap R = \emptyset$.
Let $I \subseteq S$ and ρ be a matching in M^S: $|\rho| + |I| \leq n$.
Then there exists $\pi \in$ MinCover(I): $\pi \cup \rho \in M^S$.

Definition

A covering partial function over S is an ordered pair (I, f):

- $(\text{Cover}(I), f)$ is a partial function on M^S.
- If $\pi, \pi' \in \text{Cover}(I)$: $\pi \subseteq \pi'$, then $f(\pi') = f(\pi)$.

Merged Form of Formula

Definition

Let ϕ be a disjunction, and ϕ_i are subformulas of ϕ that are not disjunctions, but every subformula of ϕ properly containing them is a disjunction, then the *merged form* of ϕ is defined as the unbounded disjunction $\bigvee_{i \in I} \phi_i$.

Definition

Let (I, f) and $(I_j, f_j), j \in J$ be covering partial functions over S. We say that (I, f) *satisfies Disj*$[\cup_{j \in J}\{(I_j, f_j)\}]$ if for all $\pi \in$ Cover(I)

- $f(\pi) = 1 \Rightarrow \exists j \in J, \pi \in$ Cover(I_j) and $f_j(\pi) = 1$.
- $f(\pi) = 0 \Rightarrow \forall j \in J$, either $\pi \in$ Cover(I_j) and $f_j(\pi) = 0$ or $\pi \notin$ Cover(I_j). (f_j is not defined on π)

k-transformations

Let Σ be closed under taking subformula.

Definition

A k-transformation T is a mapping of formulas $\phi \in \Sigma$ to covering partial functions (I_ϕ, f_ϕ) over S:

1. For all $\phi, |I_\phi| \leq k$.
2. $I_0 = I_1 = \emptyset$ (if $I = \emptyset$, then $\text{Cover}(I) = M^S$), $\forall \pi \in \text{Cover}(I_0), f_0(\pi) = 0$, $\forall \pi \in \text{Cover}(I_1), f_1(\pi) = 1$.
3. $I_{p_{ij}} = \{i, j\}, f_{p_{ij}}(\pi) = 1$ if $\{i, j\} \in \pi$ and $f_{p_{ij}}(\pi) = 0$ otherwise.
4. *[Negation]* $I_{\neg\phi} = I_\phi; f_{\neg\phi}(\pi) = \neg f_\phi(\pi), \forall \pi \in \text{Cover}(I_\phi)$.
5. *[Disjunction]* If ϕ is a disjunction and $\vee_{j\in J}\phi_j$ is the merged form of ϕ, then (I_ϕ, f_ϕ) *satisfies Disj* $\left[\bigcup_{j\in J}\{(I_{\phi_j}, f_{\phi_j})\}\right]$

Proposition 1

Theorem

Let Σ be a set of formulas closed under the operation of taking subformula. Let T be a k-transformation mapping formulas $\phi \in \Sigma$, to covering partial funcitons (I_ϕ, f_ϕ) over S. If for $\Omega \subset \Sigma$, there exists a $\pi \in \bigcap_{\phi \in \Omega} Dom(f_\phi)$, then the answers $(f_\phi(\pi))_{\phi \in I}$ to the queries $(\phi)_{\phi \in I}$ are locally consistent.

Proof.

Let Σ, T, and π be as stated in the lemma.

Since $B = \{\neg, \wedge\}$, it suffices to consider two cases.

[Negation] Let $\phi, \neg\phi \in \Sigma$. By definition of a k-transformation, $f_{\neg\phi}(\pi) = \neg f_\phi(\pi)$ for all $\pi \in \mathrm{Dom}(f_\phi) = \mathrm{Cover}(I_\phi)$. Thus, no immediate contradiction at $\neg$ gate.

Proposition 1. Proof for Disjunction

[Disjunction] Let $\phi = \bigvee_{i \in I} \phi_i$.

- (true case) Let for some $j \in I$, $f_{\phi_j}(\pi) = 1$ and $f_\phi(\pi) = 0$. By definition of a k-transformation, $f_\phi(\pi) = 0$ implies for all $i \in I$, either $\pi \in \text{Cover}(I_{\phi_i})$ and $f_{\phi_i}(\pi) = 0$ or $\pi \notin \text{Cover}(I_{\phi_i})$. This contradicts $f_{\phi_j}(\pi) = 1$. Thus, there is no immediate contradiction in this case.
- (false case) Let for all $j \in I$, $f_{\phi_j}(\pi) = 0$ and $f_\phi(\pi) = 1$. By definition of a k-transformation, $f_\phi(\pi) = 1$ implies there exists $i \in I$: $f_{\phi_i}(\pi) = 1$. This contradicts $f_{\phi_j}(\pi) = 0$. Thus, there is no immediate contradiction in this case too.

Proposition 2

Theorem

If T is k-transformation for a set of formulas containing PHP_n, $k < n-1$, then $f_{PHP_n}(\pi) = 0$ for all $\pi \in Cover(I_{PHP_n})$.

Proof.

PHP_n is the disjunction of formulas of the form $\neg\phi$,
where ϕ ranges over

$$\begin{array}{ll} \bigvee_{j \in R} p_{ij},\ i \in D & \neg p_{ik} \vee \neg p_{jk},\ i \neq j \in D,\ k \in R \\ \bigvee_{i \in D} p_{ij},\ j \in R & \neg p_{ij} \vee \neg p_{ik},\ i \in D,\ j \neq k \in R \end{array}$$

From the definition of a k-transformation, it suffices to show that $f_{\phi}(\pi) = 1, \forall \pi \in \mathrm{Cover}(I_{\phi})$ for each of the above ϕ.

Proposition 2. Proof (1)

Let $i \in D$. Let $\phi = \bigvee_{j \in R} p_{ij}$.
Suppose $f_\phi(\pi) = 0$ for some $\pi \in \mathrm{Cover}(I_\phi)$.
$|I_\phi| \leq k$,$\pi \in \mathrm{MinCover}(I_\phi)$ and $k < n - 1$, imply $|\pi| < n - 1$.
Hence, there exists a $\pi' \in M^S$: $\pi \subseteq \pi'$ and π' covers i.
Let $\{i, j\} \in \pi'$ for some $j \in R$. But then $f_{p_{ij}}(\pi') = 1$
while $f_\phi(\pi') = f_\phi(\pi) = 0$ contradicts the definition of a k-transformation.
Hence, $f_\phi(\pi) = 1, \forall \pi \in \mathrm{Cover}(I_\phi)$ for ϕ of the specified type.

Proposition 2. Proof (2)

Let $i \neq j \in D, k \in R$. Let $\phi = \neg p_{ik} \vee \neg p_{jk}$.
Suppose $f_\phi(\pi) = 0$ for some $\pi \in \text{Cover}(I_\phi)$.
As before, we have $|\pi| < n - 1$.
Since π is a matching, either $\{i, k\} \notin \pi$ or $\{j, k\} \notin \pi$.
Assume $\{i, k\} \notin \pi$. Since $|\pi| < n - 1$, there exists a $\pi' \in M^S$:
$\pi \subseteq \pi'$ and $\{i, r\}, \{s, k\} \in \pi'$ for some $r \neq k \in R$ and $s \neq i \in D$.
We have $\pi' \in \text{Cover}(I_{p_{ik}})$ and $f_{p_{ik}}(\pi') = 0$. Hence, $f_{\neg p_{ik}}(\pi') = 1$.
But $f_\phi(\pi') = f_\phi(\pi) = 0$ again contradicts definition.
The other two types of formulas are proved similarly.

Proposition 3.

Definition

We define $I|_\rho = I \setminus V(\rho)$ for any $I \subseteq S$. For (I, f) a covering partial function over S, we define $f|_\rho : \text{Cover}(I|_\rho) \to \{0, 1\}$ as $f|_\rho(\pi) = f(\pi \cup \rho)$ for all $\pi \in \text{Cover}(I|_\rho)$.

Theorem

Let $\mathcal{T}$ be a game-tree of height r for PHP_n. Let T be a k-transformation mapping formulas ϕ to covering partial functions (I_ϕ, f_ϕ) over $S|_\rho$ for some matching $\rho \in M^S$ of size $n - m$. If $kr \leq m$, then there exists a branch $((\phi_1, b_1), \ldots, (\phi_s, b_s))$ in the game-three $\mathcal{T}$:

$$\bigcap_{i=1}^{s} (f_{\phi_i})^{-1}(b_i) \neq 0$$

Proposition 3. Proof (1)

Consider the following procedure $Walk(\mathcal{T})$, outputing branch of $\mathcal{T}$

1. Set $\pi \leftarrow \emptyset$ and $i \leftarrow 1$.
2. Walk along $\mathcal{T}$ from the root till a leaf reached:
   - (a) Set $\phi_i \leftarrow$ label of current node.
   - (b) Choose a $\pi_i \in \mathrm{MinCover}(I_{\phi_i})$: $\pi \cup \pi_i \in M^{S|_\rho}$.
   - (c) Set $b_i \leftarrow f_{\phi_i}(\pi_i)$ and $\pi \leftarrow \pi \cup \pi_i$.
   - (d) Walk along edge labeled b_i leading out of current node.
   - (e) Increment i.
3. Output $((\phi_1, b_1), \ldots, (\phi_s, b_s))$.

Proposition 3. Proof (2)

- Since $\mathcal{T}$ is a game-tree for PHP_n, we have $\phi_1 = PHP_n$ and $b_1 = 0$ for any branch.
- By Proposition 1, $f_{PHP_n}(\pi) = 0$ for all $\pi \in \mathrm{Cover}(PHP_n)$.
- *Walk* algorithm choose some matching $\pi \in \mathrm{MinCover}(I_{PHP_n})$.
- A matching π_i can be chosen in the loop at Step 2b as long as $|\pi| + k \leq m$.
- $|\pi|$ is extended at most r times by at most k, and $rk \leq m$. Hence, the condition $|\pi| + k \leq m$ is true.

Let π be the matching at the final step of *Walk*.
The branch $((\phi_1, b_1), \ldots, (\phi_s, b_s))$ satisfies $b_i = f_{\phi_i}(\pi)$.
Hence, $\pi \in \bigcap_{i=1}^{s}(f_{\phi_i})^{-1}(b_i)$. Thus, $\bigcap_{i=1}^{s}(f_{\phi_i})^{-1}(b_i) \neq \emptyset$.

Existence of k-transformations

Theorem

(Switching Lemma) *Let* (I_j, f_j) *be covering partial functions over* $S, |I_j| \leq r$ *for all* $j \in J$. *Let* $\ell \geq 10$ *and* $p = \ell/n$. *If* $r \leq \ell$ *and* $p^4 n^3 \leq 1/10$, *then for random* $\rho \in M^S$, $|\rho| = n - \ell$,
Pr{ *"There exists a covering partial function* (I, f) *over* $S|_\rho$*:* (I, f) *satisfies Disj* $\left[\bigcup_{j \in J}\{(I_j|_\rho, f_j|_\rho)\}\right]$ *and* $|I| < 2s$*"*} $\geq 1 - (11p^4 n^3 r)^s$.

Theorem

Let d *be an integer,* $0 < \epsilon < 1/5, 0 < \delta < \epsilon^d$ *and* Σ *a set of formulas of depth* d. *If* $|\Sigma| < 2^{n^\delta}, q = n^{\epsilon^\delta}$ *and* n *is sufficiently large, then there exists a matching* $\rho \in M^S$ *of size* $n - n^{\epsilon^\delta}$*: there is a* $2n^\delta$*-transformation* T *mapping formulas* $\phi \in \Sigma$, *to covering partial functions* (I_ϕ, f_ϕ) *over* $S|_\rho$.

Main Theorem

Theorem

Let $\mathcal{F}$ be a Frege system and let c be the constant that occurs in theorem about Buss-Pudlák Games. Then for sufficiently large n, every depth d proof in $\mathcal{F}$ of PHP_n must have size at least 2^{n^μ}, for $\mu < \frac{1}{2}(\frac{1}{5})^{d+c}$.

Proof.

Let $0 < \epsilon < \frac{1}{5}$ and $0 < \mu < \epsilon^{d+c}/2$. Suppose PHP_n has a depth d proof in $\mathcal{F}$ of size 2^{n^μ}. By the theorem, there exists Buss-Pudlák game-tree $\mathcal{T}$ of height n^μ consisting of formulas of size at most 2^{n^μ} and depth at most $d + c$ convicting Sam on PHP_n.
Let Σ be the set of all formulas in $\mathcal{T}$. Clearly, $|\Sigma| \leq 2^{2n^\mu}$.

Main Theorem. Proof (continue)

- Choose δ: $\mu < \delta < \epsilon^d/2$. For sufficiently large n, $|\Sigma| < 2^{n^\delta}$.
- By the previous theorem, there exists a partial matching ρ of size $n - n^{\epsilon^d}$: Σ has a $2n^\delta$-transformation T mapping formulas $\phi \in \Sigma$ to covering partial functions, (I_ϕ, f_ϕ) over $S|_\rho$.
- By Proposition 2, we have that condition 1 is satisfied since $2n^\delta < n^{\epsilon^d} - 1$ for sufficiently large n.
- Also $2n^\delta \cdot n^\mu \leq n^{\epsilon^d}$ for sufficiently large n, the conditions of Proposition 3 are satisfied.
- Hence, $2n^\delta$-transformation satisfies condition 2.
- By Proposition 1, we have that condition 3 is also satisfied.
- Thus, by the theorem for transformations and strategy, game-tree $\mathcal{T}$ does not convict Sam.
- There is no depth d proof of PHP_n in $\mathcal{F}$ of size less then 2^{n^μ}.

References

Alasdair Urquhart
The complexity of propositional proofs
Bulletin of Symbolic Logic 1(4): 425-467 (1995)

Eli Ben-Sasson, Prahladh Harsha
Lower Bounds for Bounded Depth Frege Proofs via Buss-Pudlák Games
ECCC, Report No. 4 (2003)