Improved Lower Bounds for Pointer Chasing via Gadgetless Lifting

Xinyu Mao, Guangxu Yang, and Jiapeng Zhang

University of Southern California

2024/05/24 @Complexity Network

Round Communication Trade-Off and Pointer Chasing

Round communication trade-off

Do more rounds of interaction allow
two parties to solve problems with less communication?

Example. Parity and constant-depth circuits

Theorem. Any circuit of depth d that computes $\bigoplus_n$ must be of size $\Omega\left(2^{n^{\frac{1}{d-1}}}\right)$.

Karchmer-Wigderson game KW_f.

- Alice holds $x \in f^{-1}(0)$, Bob holds $y \in f^{-1}(1)$.
- They want to find an index i such that $x_i \neq y_i$.

Depth d, size S circuit computing f $\Leftrightarrow$ d round protocol for KW_f with $\log S$ communication

Corollary. Any d-round protocol that computes $KW_{\bigoplus_n}$ must communicate $\Omega\left(n^{\frac{1}{d-1}}\right)$ bits.

The pointer chasing problem

- Alice holds $f_A \in [n]^n$, Bob hold $f_B \in [n]^n$.
- The k step pointer chasing function $PC_k \colon [n]^n \times [n]^n \to \{0,1\}$
  - $pt_0 \coloneqq 1$
  - for odd r's, $pt_r \coloneqq f_A(pt_{r-1})$
  - for even r's, $pt_r \coloneqq f_B(pt_{r-1})$
  - $PC_k(f_A, f_B) \coloneqq pt_k \bmod 2$.

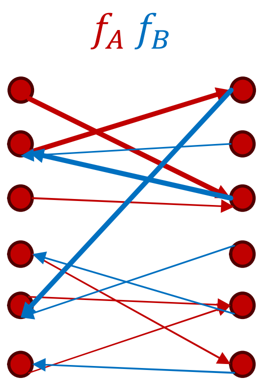

Theorem (Yehudayoff 2016). Any randomized $(k-1)$-round protocol for PC_k that is correct with probability 0.9 requires $\Omega\left(\frac{n}{k} - k \log n\right)$ bits of communication.

This work. $\Omega\left(\frac{n}{k}\right)$ lower bound via a completely different, combinatorial proof.

A simple class of protocols for pointer chasing

- Alice and Bob choose a subset $I \subseteq [n]$ of size $S \coloneqq 10\frac{n}{k}$ uniformly at random, and then send $f_A(I)$ and $f_B(I)$ to the other party.
- Alice and Bob run the naïve (k rounds) protocol, but they can skip one round if the pointer falls into I.
- If the skip round never happens, Alice and Bob simply abort at the last round.
- The skip round event happen with high probability.

Gadgetless Lifting

Gadgetless lifting

- Identify a simple class of protocols $\mathcal{K}$.
- Prove lower bound for these simple protocols.
- Prove that every protocol can be simulated by a combination of simple protocols.

$$CC(f) \coloneqq \min_{\Pi:\Pi \text{ computes } f} CC(\Pi) = \min_{\Pi \in \mathcal{K}} CC(\Pi) =: CC_{\mathcal{K}}(\Pi).$$

- For pointer chasing, $\mathcal{K}$ is the set of protocols where Alice and Bob only send values of some coordinate to each other.

Lifting theorems

- Let $g\colon \{0,1\}^q \times \{0,1\}^q \to \{0,1\}$ be a **gadget** function.
- Consider functions of the form $f \circ g^n$ for some outer function $f\colon \{0,1\}^n \to \{0,1\}$,

$$(f \circ g^n)\big((x_1, y_1), \dots, (x_n y_n)\big) \coloneqq f\big(g(x_1, y_1), \dots, g(x_n, y_n)\big).$$

$CC(f \circ g^n) = \Omega(Q(f) \cdot q)$, where $Q(f)$ denotes the query complexity of f.

- Not all functions can be written as $f \circ g^n$.
- Often need q to be large.
  - Proving lift theorems for constant gadget size q is very hard and has many implications.

Decomposition and Sampling Process

Density restoring partition

Def. For a random variable $\boldsymbol{X}$, its min-entropy is defined as $\mathbf{H}_\infty(\boldsymbol{X}) := \log \frac{1}{\max_x \Pr[\boldsymbol{X}=x]}$.

Def. We say a random variable $\boldsymbol{X}$ over $[n]^J$ is γ-dense if $\mathbf{H}_\infty\big(\boldsymbol{X}(I)\big) \geq \gamma \log n\ |I|$ for all $I \subseteq J$.

For a set X, $\boldsymbol{X} :=$ uniform distribution over X.

Theorem([GPW17]). For any $X \subseteq [n]^J$, there is a partition $X = X^1 \cup \cdots \cup X^r$ and each X^i is associated with a set I_i with the following properties.

- X^i is fixed on I_i: there exists some $\alpha_i \in [n]^{I_i}$ such that $x(I_i) = \alpha_i$ for all $x \in X^i$.
- $\boldsymbol{X}^i(J \setminus I_i)$ is γ-dense.
- $\mathbf{D}_\infty\left(\boldsymbol{X}^i(J \setminus I_i)\right) \leq \mathbf{D}_\infty(X) - (1-\gamma)\log n\,|I_i| + \delta_i$ where $\delta_i = \log \frac{|X|}{|\cup_{j \geq i} X^j|}$.

▶ $\mathbf{D}_\infty(\boldsymbol{X}) := |J| \log n - \mathbf{H}_\infty(\boldsymbol{X})$ if $\boldsymbol{X}$ is supported on $[n]^J$.

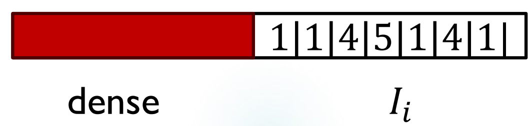

Protocol tree

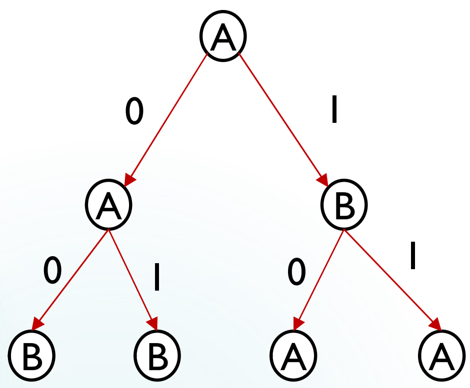

- For each internal vertex v,
  - v is owned by either Alice or Bob
  - v corresponds to a rectangle $\Pi_v = X_v \times Y_v$, the input that leads to v.
  - v has two children u_0, u_1
    - If v is owned by Alice, $X_{u_0} \cup X_{u_1}$ is a partition of X_v and $Y_{u_0} = Y_{u_1} = Y$.
    - If v is owned by Bob, $Y_{u_0} \cup Y_{u_1}$ is a partition of Y_v and $X_{u_0} = X_{u_1} = X$.
- Each leaf specifies an output.

Yao's min-max principle

To prove lower bound for all **randomized** protocols, it suffices to prove lower bound for all **deterministic** protocols under some input distribution μ.

Here we let μ to be the uniform distribution on all inputs $[n]^n \times [n]^n$.

Decomposition and sampling process $DS(\Pi)$

Input: A protocol Π
Output: A rectangle $R = X \times Y \subseteq [n]^n \times [n]^n$, J_A, $J_B \subseteq [n]$.
Initialization: $X := Y := [n]^n, J_A := J_B := [n]$, skip $:= false$, $r := 0, v := root$.

$$\Pr[DS(\Pi) \text{ outputs } R] = \frac{|R|}{|\text{all inputs}|}$$

1. Partition X into $X = X^0 \cup X^1$ according to node v.
2. Sample $\boldsymbol{b} \in \{0,1\}$ such that $\Pr[\boldsymbol{b} = b] = \frac{|X^b|}{|X|}$.
3. Update $X := X^{\boldsymbol{b}}, v := u_{\boldsymbol{b}}$.
4. If u_b is owned by Bob:
   - Further partition X into $X = X^0 \cup X^1$ where $X^b := \{f_A \in X : f_A(z_{r-1}) \bmod 2 = b\}$.
   - Sample $\boldsymbol{b} \in \{0,1\}$ such that $\Pr[\boldsymbol{b} = b] = \frac{|X^b|}{|X|}$.
   - Update $X := X^{\boldsymbol{b}}, r := r + 1$.
5. Let $X = X^1 \cup \cdots \cup X^m$ be density restoring partition of X with associated $I_1, \ldots, I_m$.
6. Sample a random element $\boldsymbol{j} \in [m]$ such that $\Pr[\boldsymbol{j} = j] = \frac{|X^j|}{|X|}$ for $j \in [m]$.
7. Update $X := X^{\boldsymbol{j}}, J_A := J_A \setminus I_{\boldsymbol{j}}$.
8. If u_b is owned by Bob $z_{r-1} \notin J_B$, skip := true.

Suppose Alice owns node v.
Let u_0, u_1 be the children of v.

As a new round begins, we do an extra partition to fix the parity of pt_r.

1|1|4|5|1|4|1|

J_A

X_{I_j} is fixed;
X_{J_A} is dense.

Loop invariant

Input: A protocol Π
Output: A rectangle $R = X \times Y \subseteq [n]^n \times [n]^n,\ \ J_A,\ J_B \subseteq [n]$.
Initialization: $X \coloneqq Y \coloneqq [n]^n, J_A \coloneqq J_B \coloneqq [n], \mathrm{skip} \coloneqq false, r \coloneqq 0, v \ \coloneqq root$.

Lemma. *Set $\gamma \ \coloneqq 1 - \frac{0.1}{\log n}$. Then in the running of $DS(\Pi)$, we have the following loop invariants: After each iteration,*

- $X \times Y \subseteq \Pi_v$.
- $\boldsymbol{X}(J_A), \boldsymbol{Y}(J_B)$ *are γ-dense.*
- *There exists some $\alpha_A \in [n]^{\overline{J_A}}, \alpha_B \in [n]^{\overline{J_B}}$ such that $x(\overline{J_A}) = \alpha_A, y(\overline{J_B}) = \alpha_B$ for all $x \in X, y \in Y$.*
- *There exists some $z_r \in [n]$ such that $pt_r(f_A\ , f_B) = z_r$ for all $f_A \in X,\ \ f_B \in Y$.*

We only fix the party but the density restoring partition helps to fix pt_r.
This is way we save the $k \log n$ factor in the previous result.

Relating accuracy and *average fixed size*

Input: A protocol Π
Output: A rectangle $R = X \times Y \subseteq [n]^n \times [n]^n,\ \ J_A,\ J_B \subseteq [n]$.
Initialization: $X \coloneqq Y \coloneqq [n]^n, J_A \coloneqq J_B \coloneqq [n], \mathrm{skip} \coloneqq false, r \coloneqq 0, v \coloneqq root$.

Lemma. If $DS(\Pi)$ outputs $(R = X \times Y, J_A, J_B)$ and $\mathrm{skip} = false$ in the end, then

$$\Pr_{(f_A, f_B) \leftarrow R}[\Pi(f_A, f_B) = PC_k(f_A, f_B)] \leq \frac{2^{0.1}}{2}.$$

Lemma. $\Pr[\mathrm{skip} = true] \leq \frac{2^{0.1}}{n} \cdot k \cdot \mathbf{E}[|\overline{J_A}| + |\overline{J_B}|]$.

Union bound for k rounds

If we can prove $\mathbf{E}[|\overline{J_A}| + |\overline{J_B}|] = O(c)$, then we have

$$\frac{2^{0.1}}{n} \cdot k \cdot O(c) = \Omega(1) \Rightarrow\ c = \Omega\left(\frac{n}{k}\right).$$

Average fixed size is bounded by communication: A density increment argument

- In the running of $DS(\Pi)$, we track the value of the following value:

$$D_\infty(R) \coloneqq D_\infty\big(X(J_A)\big) + D_\infty\big(Y(J_B)\big).$$

$\mathbf{D}_\infty(\boldsymbol{X}) \coloneqq |J| \log n - \mathbf{H}_\infty(\boldsymbol{X})$

- In the beginning, $D_\infty([n]^n \times [n]^n) = 0$.
- In expectation (over the choice of $\boldsymbol{b}$), each communication bit/new round increase $D_\infty(R)$ by at most 1:

$$\frac{|X^0|}{|X|}\log\frac{|X^0|}{|X|} + \frac{|X^1|}{|X|}\log\frac{|X^1|}{|X|} \le 1.$$

Since X is fixed outside J_A,
$\boldsymbol{X}(J_A)$ is a uniform distribution.

- In expectation (over the choice of $\boldsymbol{j}$), $D_\infty(R)$ decreases by at least $(1-\gamma)\log n\, \mathbf{E}_j\big[|I_j|\big] + 1$.
  - $\mathbf{D}_\infty\left(\boldsymbol{X}^i(J \setminus I_i)\right) \le \mathbf{D}_\infty(X) - (1-\gamma)\log n\,|I_i| + \delta_i$ where $\delta_i = \log\frac{|X|}{|\cup_{j\ge i} X^j|}$.
  - $\mathbf{E}_j[\delta_j] = \sum_j p_j\, \delta_j = \sum_j p_j \log\frac{1}{\sum_{t\ge j} p_t} \le \int_0^1 \frac{1}{1-x}\, dx \le 1.$

$p_j \coloneqq \frac{|X^j|}{|X|}$

- $D_\infty(R) \ge 0 \rightarrow \mathbf{E}[|\bar{J}_A| + |\bar{J}_B|] = \mathbf{E}[|I_1| + |I_2| + \cdots +] \le O\left(\frac{c}{(1-\gamma)\,\mathrm{lo}}\right)$

total increment $\ge$ total decrement.
Not a round-by-round bound!

Recap

- **The decomposition and sampling process**: Use **density restoring partition** to decompose the behavior of Π into the combination of simple protocols (i.e., fixing some coordinates).
- Relating accuracy and **average fixed size**.
- Average fixed size is bounded by communication.

Discussion

- More generic density restoring partition?
- **Open question**: Can we prove parity not in AC0 using a **top-down** approach?
  - [RSS' FOCS 23] gave a proof for depth 4 circuits.
- Round communication trade-off for other problems?

Theorem. Any randomized $(k-1)$-round protocol (where Alice speaks first) for PC_k that is correct with probability 0.9 requires $\Omega\left(\frac{n}{k}\right)$ bits of communication.

Thanks for listening ☺

Appendix: Proof of density restoring partition lemma

A greedy algorithm

- Input: $X \subseteq [n]^J$.
- Output: a partition $X = X^1 \cup \cdots \cup X^m$ and $I_1, \ldots, I_m \subseteq [J]$.
- While $X \neq \emptyset$
  1. Find the maximal $I \subseteq J$ such that X_I is not γ-dense.
     - $\exists \alpha_i \in [n]^I \ s.t. \ \Pr_{x \leftarrow X}[x(I) = \alpha_i] \geq n^{-\gamma|I|}$.
  2. $X^i := \{x \in X : x(I) = \alpha_i\}, I_i := I$.
  3. $X := X \setminus X^i, \ J := J \setminus I_i, \ i := i + 1$.

- X^i is fixed on I_i by construction.
- $X^i(J \setminus I_i)$ is γ-dense: if not, then $\exists\ K \subseteq J \setminus I_i$ that violates the min-entropy condition at the moment I_i is chosen.
  - $\Pr_{x \leftarrow X^i}[x(K) = \beta] \geq n^{-\gamma|K|}$.
  - $I_i \cup K$ violates the maximality of I_i.