CAS CS 591 B: Communication Complexity

Prof. Mark Bun | Fall 2019

Lecture Notes 16 & 17:

Deterministic Lifting

Reading.

- Rao-Yehudayoff, Chapter 8
- Chattopadhyay-Kouchký-Loff-Mukhopadhyay, Simulation Theorems via Pseudo-random Properties

We begin a proof of a quite general deterministic lifting theorem. This lifting theorem works for any gadget g satisfying a certain pseudorandom property which we'll call the h-hitting property. The hitting property guarantees that there are distributions over the monochromatic rectangles of g such that any large enough rectangle will intersect a random rectangle from either of these distributions with high probability.

Definition 1. A gadget $g : X \times Y$ has the h-hitting property if there exist distributions σ_0 and σ_1 over the sets of 0-monochromatic rectangles and 1-monochromatic rectangles of g, respectively, such that for every rectangle $A \times B$ with $|A|/|X|, |B|/|Y| \geq 2^{-h}$ we have

$$\Pr_{R \sim \sigma_c}[R \cap (A \times B) \neq \emptyset] \geq 0.99$$

for $c \in \{0, 1\}$.

Theorem 2 (CKLM17)**.** *Let $g : X \times Y$ be a gadget with the h-hitting property and let $n \leq 2^{h/2}$. Then for every $f : \{0,1\}^n \to \{0,1\}$,*

$$\mathcal{C}^{\mathbf{cc}}(f \circ g^n) \geq \frac{1}{10} \cdot \mathcal{C}^{\mathbf{dt}}(f) \cdot h.$$

Examples of gadgets with the hitting property include

- IP_m has the $(m/5)$-hitting property.
- GH_m has the $(m/5)$-hitting property.
- IND_m has the $(3 \log m/20)$-hitting property.

We now begin the proof of Theorem 2. The proof is constructive and uses a communication protocol Π for $f \circ g^n$ to produce a query algorithm (decision tree) computing $f(z)$. This query algorithm attempts to simulate the behavior of Π on inputs that are consistent with the queries to z made so far.

The query algorithm maintains a rectangle $A \times B \subseteq X^n \times Y^n$. In order for the simulation to be successful, we maintain several invariants of the rectangle:

- The rectangle should be consistent with all coordinates of z that have been queries so far. Namely, for every $(x, y) \in A \times B$ we have that $g_i(x, y) = z_i$ for every queried coordinate i.
- $A \times B$ is "sufficiently rich" that it is always possible to maintain the consistency requirement when further coordinates of z are queried. The technical property that captures this is the notion of "thickness": The rectangle $A \times B$ is thick in coordinate i if for every $(x, y) \in A \times B$, the value of $g(x_i, y_i)$ can be arbitrary even given the rest of the coordinates x_{-i}, y_{-i}.

  Formally, a set $A \in X^n$ is τ-thick with respect to S if for every $i \in S$, and all $a \in A_{S-i}$,

  $$d_{\min}(A_S, i) := |x \in A_S : x_{S-i} = a| \geq \tau |X|.$$

  A rectangle $A \times B$ is τ-thick if both A and B are τ-thick.
- We also work with a relaxed notion of thickness called "average thickness" which is enough to guarantee that simulating a step of the protocol is safe. A set A is φ-average-thick with respect to S if for every $i \in S$,

  $$d_{\text{avg}}(A_S, i) := \frac{|A_S|}{|A_{S-i}|} \geq \varphi |X|.$$

  One way to think about thickness and average thickness is in terms of a bipartite graph with left vertices A_i, right vertices A_{S-i}, and edges (a_i, a_{S-i}) corresponding to elements in A. Then $d_{\min}$ is the minimum right-degree of this graph and d_{avg} is the average right-degree of this graph.

The analysis of Algorithm relies on two main lemmas. The first lemma says that if the current rectangle is average-thick, then it can be pruned to restore the thickness invariant while losing only a constant factor if its density.

Lemma 3 (Pruning). *If $A \subseteq X^n$ is φ-average-thick, then there is a $\left(\frac{1}{2n}\varphi\right)$-thick subset $A' \subseteq A$ with $|A'| \geq |A|/2$.*

Proof. We repeatedly remove elements $x \in A$ which violate the thickness condition. Let $\tau = \frac{1}{2n}\varphi$.

Initialize $A' = A$. While there exists an i for which $d_{\min}(A', i) < \tau|X|$, let $a \in A'_{-i}$ such that

$$|E| := |x \in A' : x_{-i} = a| < \tau |X|.$$

Update $A' = A' \setminus E$.

The total number of iterations of this process is at most

$$\sum_{i=1}^{n} |A_{-i}| = \sum_{i=1}^{n} \frac{|A|}{d_{\text{avg}}(A, i)} \leq \frac{n|A|}{\varphi|X|}$$

The algorithm removes at most $\tau|X|$ elements in each iteration, so the total number of elements removed is $|A|/2$. Hence $|A'| \geq |A|/2$. □

The second lemma allows us to find a useful update of the current rectangle when we query a coordinate with low average degree. The update satisfies two crucial properties. First, it maintains the consistency and τ-thickness invariants of the current rectangle. Second, it increases its density. The second property allows us to argue that the number of query steps over the course of the algorithm is not too large, since Simulate and Prune steps each decrease density by only a factor of 2, and density cannot be larger than 1.

Algorithm 1 Decision tree algorithm $T(z)$

```
• Set parameters $\tau = 2^{-h}, \varphi = 4 \cdot 2^{-h/4}$
• Initialize $A \times B = X^n \times Y^n$, $S = [n]$, $v =$ root of the protocol tree
• Until $v$ is a leaf:
  If $A$ and $B$ are both $\varphi$-average-thick w.r.t. $S$:
      1. Simulate the next step of the protocol moving to a new vertex $v$ with consistent rectangle $A \times B$.
         Namely, let $u_0$ and $u_1$ be the children of $v$ in the protocol tree. If it is Alice's turn to speak, choose $c \in \{0,1\}$ so that $|(A_c)_S| \geq |A_S|/2$. Set the new rectangle to $A_c \times B$ and $v = u_c$
      2. Prune $A \times B$ using Lemma 3 to restore $\tau$-thickness.
  Else:
      1. Query a coordinate $i \in S$ for which $d_{\text{avg}}(A_S, i) < \varphi|X|$ or $d_{\text{avg}}(B_S, i) < \varphi|X|$
      2. Use the Projection Lemma 5 to update $A \times B$, set $S = S \setminus \{i\}$
• Output the value at $v$
```

Definition 4. Let $A \times B$ be a rectangle and let $S \subset [n]$. The density of $A \times B$ with respect to S is defined to be

$$\text{dens}_S(A \times B) := \frac{|A_S| \cdot |B_S|}{|X|^{|S|} \cdot |Y|^{|S|}}.$$

Lemma 5 (Projection)**.** *Let $g : X \times Y \to \{0,1\}$ have the h-hitting property. Suppose $A \times B$ is τ-thick w.r.t. S for $\tau \geq 2^{-h}$ and suppose*

$$d_{\text{avg}}(A_S, i) \leq \varphi|X|.$$

Then for $c \in \{0,1\}$ there exists a subrectangle $A' \times B' \subseteq A \times B$ such that

1. *$g(x_i, y_i) = c$ for every $(x,y) \in A' \times B'$*
2. *$A' \times B'$ is τ-thick with respect to $S - i$*
3. *$\text{dens}_{S-i}(A' \times B') \geq \frac{1}{2\varphi} \cdot \text{dens}_S(A \times B)$.*

Proof. For convenience, assume $S = [n]$. Fix $(a,b) \in A_{-i} \times B_{-i}$. Let

$$U_a = \{x_i \in X : (x_i, a) \in A\}, \quad V_b = \{y_i \in Y : (y_i, b) \in B\}.$$

$(\tau = 2^{-h})$-thickness implies that $|U_a|/|X|, |V_b|/|Y| \geq 2^{-h}$. Then by the hitting property of g, there is a distribution σ_c over c-monochromatic rectangles of g for which

$$\Pr_{R \sim \sigma_c}[R \cap (U_a \times V_b) \neq \emptyset] \geq 0.99.$$

Hence there exists a choice of rectangle R such that

$$|\{(a,b) : R \cap (U_a \times V_b) \neq \emptyset\}| \geq |A_{-i} \times B_{-i}|/2.$$

Let $A' \times B' = \{(a,b) \in A \times B : (a_i, b_i) \in R\}$. Then $|A'_{-i} \times B'_{-i}| \geq |A_{-i} \times B_{-i}|/2$.
The rectangle $A' \times B'$ so chosen satisfies

$$\begin{aligned} \mathrm{dens}_{-i}(A' \times B') &= \frac{|A'_{-i}| \cdot |B'_{-i}|}{|X|^{n-1} \cdot |Y|^{n-1}} \\ &= \frac{9}{10} \cdot \frac{|A_{-i}|}{|X|^{n-1}} \cdot \frac{9}{10} \cdot \frac{|B_{-i}|}{|Y|^{n-1}} \\ &\geq \frac{9}{10} \cdot \frac{|A|/(\varphi|X|)}{|X|^{n-1}} \cdot \frac{9}{10} \cdot \frac{|B|/|Y|}{|Y|^{n-1}} \\ &\geq \frac{1}{2\varphi} \cdot \mathrm{dens}(A \times B). \end{aligned}$$

Here, the second-to-last inequality follows because $d_{\mathrm{avg}}(A, i) = |A|/|A_{-i}| \leq \varphi|X|$.

Finally, thickness is inherited from thickness of the original rectangles $A \times B$. (Exercise.) □

We are now ready to prove Theorem 2. The proof is broken into two claims: One that the number of queries is small, and one that the protocol is correct.

Claim 6. *Let ℓ be the length of the protocol Π computing $f \circ g^n$. Then Algorithm makes at most $10\ell/h$ queries to z.*

Proof. By Lemma 5, every Query step increases $\mathrm{dens}_S(A \times B)$ by a factor of $1/2\varphi = 2^{h/4-2}$. Meanwhile, every Simulate and Prune step decreases $\mathrm{dens}_S(A \times B)$ by a factor of at most 2. Let q be the total number of queries. Since density is always at most 1 and the number of Simulate and Prune steps is at most 2ℓ, we have

$$(2^{h/4-2})^q \cdot 2^{2\ell} \leq 1$$

so $q \leq 2\ell/(h/4 - 2) \leq 10\ell/h$. □

Claim 7. *$T(z) = f(z)$ for every $z \in \{0,1\}^n$.*

Proof. Consider the following thought experiment. Augment the original protocol Π to a new protocol Π' where Alice and Bob send each other their inputs x, y at the very end. Let $A \times B$ be the rectangle reached at the end of the simulation of Π, and imagine continuing the simulation of Π'. Then the simulation reaches a rectangle $A' \times B'$ which is a subrectangle of $A \times B$. Since the simulation always enforces consistency, $z_j = g(x_j, y_j)$ for all $j \in [n]$ and $(x, y) \in A' \times B'$. Therefore, there must exist some x, y in the original rectangle $A \times B$ which is consistent with z, so $\Pi(x, y) = f(G(x, y)) = f(z)$. □

1 Hitting Distributions

The last ingredient is to exhibit gadgets which satisfy the hitting property. We'll only do this for the inner product gadget, but see [CKLM17] for the analysis of other gadgets.

Theorem 8. *The gadget IP_m has the h-hitting property for $h = m/4$.*

Recall that the goal is to exhibit distributions σ_0 and σ_1 over the monochromatic rectangles of IP_m which are likely to hit any sufficiently large rectangle. For simplicity, let us only describe the distribution σ_0. (The ideas for σ_1 are similar.)

To sample σ_0, let V be a random subspace of $\mathbb{F}_2^m$ of dimension $m/2$. Let $V^\perp$ be its orthogonal complement, namely $V^\perp = \{u \in \mathbb{F}_2^m : \langle u, v\rangle = 0, \forall v \in V\}$. Let the sampled rectangle R be $V \times V^\perp$ which by construction is 0-monochoromatic with respect to IP_m.

We need to show that R sampled this way is likely to intersect any rectangle $A \times B$ with $|A|, |B| \geq 2^{3m/4}$. Intuitively, a random subspace V of dimension $m/2$ looks like $2^{m/2}$ random points. So the expected size of $V \cap A$ should be $|A| \cdot 2^{-m/2} \geq 2^{m/4}$ and moreover, concentrate around this with high probability. (And similarly for B.)

Let's do the calculation. Assume $0 \notin A$; otherwise we are done since $0 \in V$. Let $W = |V \cap A|$. Then we can express W as a sum of indicator random variables $W = \sum_{x \in A} W_x$ where $W_x = 1$ iff $x \in V$. For each $x \in A$ we have

$$\Pr[W_x = 1] = \frac{2^{d/2} - 1}{2^d - 1}.$$

Moreover, for each $x \neq x' \in A$,

$$\Pr[W_x = 1 \wedge W_{x'} = 1] = \frac{\binom{2^{d/2}-1}{2}}{\binom{2^d-1}{2}} \leq \left(\frac{2^{d/2}-1}{2^d-1}\right)^2.$$

In other words, the random variables W_x are negatively correlated:

$$\mathbb{E}[W_x W_{x'}] \leq \mathbb{E}[W_x]\mathbb{E}[W_{x'}].$$

This implies that

$$\mathrm{V}ar[W] = \mathbb{E}[W^2] - \mathbb{E}[W]^2 \leq \mathbb{E}\left[\sum_{x \in A} W_x^2 - \mathbb{E}[W_x]^2\right] \leq \mathbb{E}[W]$$

By Chebyshev's inequality ($\Pr[|X - \mathbb{E}[X]| \geq k\sqrt{\mathrm{V}ar[X]}] \leq 1/k^2$),

$$\begin{aligned} \Pr[A \cap V = \emptyset] &\leq \Pr[|W - \mathbb{E}[W] \geq \mathbb{E}[W]] \\ &\leq \frac{1}{\mathbb{E}[W]} \\ &\leq 2^{-m/4}. \end{aligned}$$

Similarly, we can argue that $\Pr[B \cap V^\perp = \emptyset] \leq 2^{-m/4}$. Taking a union bound concludes the proof.