Communication Complexity 25 Oct, 2011 (@ IMSc)

18. Monotone depth lower bound for st connectivity

Lecturer: Karteek Sreenivasaiah *Scribe: Nitin Saurabh*

In Lecture 2, we saw how communication complexity lower bounds yield lower bounds for circuit depth. In particular, we showed that for any function f, $D(KW_f) = \text{depth}(f)$ and $D(KW_f^+) = \text{depth}^+(f)$, where KW_f denotes the Karchmer-Wigderson game on f. Using this, we showed that monotone circuits for matching require $\Omega(n)$ depth.

In this lecture we will show that circuits solving directed s-t connectivity require $\Omega(\log^2 n)$ depth. The directed s-t connectivity function DSTCON_n is defined as follows: Given a directed graph G on n nodes, a source vertex s and a target vertex t,

$$\mathsf{DSTCON}_n(G, s, t) = 1 \iff \text{there is a directed path in the graph } G \text{ from } s \text{ to } t$$

Clearly, this function is monotone; adding edges cannot remove an already existing path.

We assume without loss of generality that the vertices are numbered from 1 to n and $s = 1$ and $t = n$. In the KW^+ game on this function, Alice is given a graph G_1 that has an $s - t$ path, while Bob has a graph G_0 that does not have an $s - t$ path. The goal is to find an edge (u, v) that appears in G_1 but not in G_0.

For the purposes of showing a lower bound, we will restrict our attention to special inputs. We will define a FORK relation and show that the communication game for the FORK relation reduces to $\mathsf{KW}^+_{\mathsf{DSTCON}}$ on special inputs. We will then give a lower bound for the FORK relation by repeatedly using round elimination and amplification. The references for today's lecture include Sections 5.3 and 10.3 of [KN97].

18.1 The FORK relation

Let Σ be an alphabet consisting of w letters, say $\{1, \ldots, w\}$. Define a relation $\mathsf{FORK}_{w,l} \subseteq \Sigma^l \times \Sigma^l \times [l]$ where $(x, y, i) \in \mathsf{FORK}_{w,l} \Leftrightarrow (x_i = y_i \text{ and } x_{i+1} \neq y_{i+1})$. If $x = y$, then there does not exist any i such that $(x, y, i) \in \mathsf{FORK}_{w,l}$. Therefore we will implicitly pad x and y with additional 0 and $l + 1$ positions such that $x_0 = y_0 = 1$, $x_{l+1} = w$, and $y_{l+1} = w - 1$. This ensures that $\forall x, y \in \Sigma^l$, $\exists i \in \{0, \ldots, l\}$ such that $(x, y, i) \in \mathsf{FORK}_{w,l}$.

18.2 Reducing the FORK relation to DSTCON

In the communication game on the FORK relation, Alice has the string $x \in \Sigma^l$, Bob has the string $y \in \Sigma^l$, and they want to determine an i such that $(x, y, i) \in \mathsf{FORK}_{w,l}$. We will show that a communication protocol for the DSTCON function can be used to solve this game. We will only need instances of DSTCON that are layered graphs consisting of $l + 2$ layers, with each layer having w vertices. s belongs to layer 0 and t belongs to layer $l + 1$. Each edge connects a vertex in some layer i to a vertex in the next layer $i + 1$. We refer to DSTCON_n, restricted to such instances, as $\mathsf{DSTCON}_{w,(l+2)}$, where $n = w(l + 2)$.

Lemma 18.1.

$$\mathsf{FORK}_{w,l} \leq \mathsf{KW}^{+}_{\mathsf{DSTCON}_{w,(l+2)}}$$

Proof. Let Π be a protocol for $\mathsf{KW}^{+}_{\mathsf{DSTCON}_{w(l+2)}}$. We will show that this protocol can be used to solve $\mathsf{FORK}_{w,l}$. Alice and Bob can solve $\mathsf{FORK}_{w,l}$ as follows:

Alice is given $x \in \Sigma^l$. Alice constructs the layered graph G_1 with $l+2$ layers. Each layer has w vertices, corresponding to the w letters of the alphabet. Alice constructs the path P_x corresponding to $x_0, \ldots, x_{l+1}$ by choosing from each layer i the vertex x_i and connecting it to the vertex x_{i+1} in layer $i+1$. Let $v_{i,j}$ denote vertex j in layer i. So Alice's graph consists of just one path connecting $v_{0,1}$ to $v_{l+1,w}$.

Bob is given a string $y \in \Sigma^l$. Bob constructs the graph G_0 with the same number of layers and vertices as above. Bob's graph contains the path P_y corresponding to $y_0, \ldots, y_{l+1}$. In addition, Bob also adds an edge between a vertex of layer i that is not in the path P_y to all the vertices of layer $i+1$. Observe that from $v_{0,1}$, we can only go along the path P_y. But P_y does not reach $v_{l+1,w}$ (since $x_{l+1} = w$ and $y_{l+1} = w-1$), and hence $v_{0,1}$ is not connected to $v_{l+1,w}$ in G_0.

Choosing s to be $v_{0,1}$ and t to be $v_{l+1,w}$, we see that G_1 is a yes instance and G_0 is a no instance. Now Alice and Bob use the protocol Π on G_1 and G_0 and get as output an edge (u, v) that appears in G_1 but not in G_0. Let u belongs to some layer i and v belongs to layer $i+1$. Note that (u, v) belongs to path P_x since only edges of P_x are present in G_1. Further, u belongs to path P_y but v does not, because these are the only kind of edges missing in G_0. Therefore $x_i = y_i$ but $x_{i+1} \neq y_{i+1}$. So $(x, y, i) \in \mathsf{FORK}$ as desired, and both Alice an Bob know i after running the protocol. □

18.3 Lower bound for the FORK relation

We now show a lower bound for the communication game of the FORK relation. In particular, we show that $\mathrm{D}(\mathsf{FORK}_{w,l}) = \Omega(\log l \log w)$.

For each fixed w, we define the notion of an (α, l) protocol. For $0 \leq \alpha \leq 1$, we say that a protocol is an (α, l) protocol if there exists a set $S \subseteq \Sigma^l$ of size $|S| \geq \alpha \cdot |\Sigma|^l$ such that for all $x, y \in S$, the protocol gives a correct answer for $\mathsf{FORK}_{w,l}(x, y)$.

Lemma 18.2 (Round elimination)**.** *If there exists a c-bit (α, l) protocol for the relation $\mathsf{FORK}_{w,l}$, then there is also a $(c-1)$-bit $(\alpha/2, l)$ protocol for $\mathsf{FORK}_{w,l}$.*

This lemma says that we can eliminate one bit from the message transcript and still be correct on a large fraction of the inputs.

Lemma 18.3 (Amplification)**.** *Let $\alpha \geq \lambda/w$ (for a large enough constant λ). If there exists a c-bit (α, l) protocol for $\mathsf{FORK}_{w,l}$, then there is also a c-bit $(\sqrt{\alpha}/2, l/2)$ protocol for $\mathsf{FORK}_{w,l/2}$.*

This lemma says that a protocol with a "success probability" (fraction of inputs on which correct) in a suitable range ($\sqrt{\alpha}/2 \geq \alpha \geq \lambda/w$) can be converted into a protocol with a larger success probability, though on smaller inputs. (Since we fix w in this argument, l is a good measure of input length.)

Assuming the above lemmas, we can prove the following theorem.

Theorem 18.4.

$$\mathrm{D}(\mathsf{FORK}_{w,l}) = \Omega(\log l \log w)$$

Proof. Let $C(\alpha, l)$ denote the minimum number of bits required by an (α, l) protocol for $\mathsf{FORK}_{w,l}$. Then $\mathrm{D}(\mathsf{FORK}_{w,l}) = C(1, l)$. Since $C(1, l) \geq C(1/w^{1/3}, l)$, it suffices to prove that $C(1/w^{1/3}, l) = \Omega(\log l \log w)$. From Lemma 18.2 we know that $C(\alpha, l) \geq C(\alpha/2, l) + 1$. Applying this $\log(\frac{1}{4}w^{1/3})$ times, we get $C(1/w^{1/3}, l) \geq \Omega(\log w) + C(4/w^{2/3}, l)$. Apply Lemma 18.3 once to get $C(4/w^{2/3}, l) \geq C(1/w^{1/3}, l/2)$. Hence we have $C(1/w^{1/3}, l) \geq \Omega(\log w) + C(1/w^{1/3}, l/2)$. Repeating the above argument $\Theta(\log l)$ times, we get $C(1/w^{1/3}, l) \geq \Omega(\log l \log w) + C(1/w^{1/3}, 1)$. But $C(\alpha, 1) \leq C(1, 1) \leq \log w$. The result follows. □

We will now establish the two lemmas.

Proof of Lemma 18.2. Assume without loss of generality that Alice sends the first bit in the (α, l) protocol Π (the case when Bob sends the first bit is similar). Let $S \subseteq \Sigma^l$ be the good set guaranteed by the (α, l) property. Let $S_0, S_1 \subseteq S$ be the sets of strings for which Alice sends 0 and 1 as the first bit respectively. Let S_b be the larger among S_0 and S_1, then, clearly $|S_b| \geq |S|/2$. Define a new protocol Π' which is exactly like Π except that the first bit is not sent at all; Alice and Bob assume the first bit to be b and then follow Π. Then, Π' is a $(c-1)$-bit protocol with good set S_b. Hence, Π' is a $(c-1)$-bit $(\alpha/2, l)$ protocol for $\mathsf{FORK}_{w,l}$. □

We will need the following claim to prove Lemma 18.3.

Claim 18.5. *Consider an $n \times n$ 0-1 matrix. Let m be the number of 1s in it, and m_i be the number of 1s in the i-th row. Denote by $\alpha = m/n^2$ the fraction of 1-entries in the matrix and by $\alpha_i = m_i/n$ the fraction of the 1-entries in the i-th row. Then, at least one of the following holds:*

(a) There is some row i with $\alpha_i \geq \sqrt{\alpha/2}$.

(b) The number of rows for which $\alpha_i \geq \alpha/2$ is at least $\sqrt{\alpha/2} \cdot n$.

Proof. Say $\sqrt{\alpha/2}$ is high-density, and $\alpha/2$ is moderate density, of 1s. Then the claim says that either there is a high-density row, or there are many moderate-density rows. To see why, observe that $\sum_{i=1}^{n} \alpha_i = \sum_{i=1}^{n} m_i/n = m/n = \alpha \cdot n$. Now suppose neither (a) nor (b) holds. This means that for all rows $\alpha_i < \sqrt{\alpha/2}$, and for less than $\sqrt{\alpha/2} \cdot n$ rows $\alpha_i \geq \alpha/2$. Therefore,

$$\alpha \cdot n = \sum_{i=0}^{n} \alpha_i < (\sqrt{\alpha/2} \cdot n) \cdot \sqrt{\alpha/2} + n \cdot \alpha/2 = \alpha n,$$

a contradiction. □

Proof of Lemma 18.3. Let S be the good set corresponding to the (α, l) protocol Π. Consider a matrix M whose rows and columns correspond to strings in $\Sigma^{l/2}$. An entry corresponding to row u and column v of M is 1 if the string $u \circ v$ is in S and 0 otherwise. Since $|S| \geq \alpha|\Sigma^l|$, the density of 1s in the matrix is at least α. Applying Claim 18.5 to the matrix M, we get that it satisfies either (a) or (b) (or both).

Suppose the matrix satisfies (a). Then there exist a row, corresponding to some string $u \in \Sigma^{l/2}$, with density at least $\sqrt{\alpha/2}$. The new protocol Π' for $\mathsf{FORK}_{w,l/2}$ works as follows: on input $x, y \in \Sigma^{l/2}$, Alice and Bob use the original c-bit (α, l) protocol Π on the strings $x' = u \circ x$ and $y' = u \circ y$. Since we are prefixing both x and y with the same string u, whenever $(x', y', i) \in \mathsf{FORK}$, we know that $i \geq l/2$, and hence $(x, y, i - l/2) \in \mathsf{FORK}$. The protocol Π succeeds whenever $u \circ x$ and $u \circ y$ are in S. Let $S' = \{x | u \circ x \in S\}$. Then Π' succeeds whenever $x, y \in S'$, so S' is good for the protocol Π'. Since (a) holds with respect to S, we know that $|S'| \geq \sqrt{\alpha/2}|\Sigma|^{l/2}$. So Π' is a c-bit $(\sqrt{\alpha}/2, l/2)$ protocol for $\mathsf{FORK}_{w,l/2}$.

Suppose the matrix satisfies (b). Let S' be the set of all rows with density at least $\alpha/2$; then $|S'| \geq \sqrt{\alpha/2} \cdot |\Sigma|^{l/2}$. We will show that there exist functions $f, g : \Sigma^{l/2} \to \Sigma^{l/2}$ and a set $S'' \subseteq S'$ such that the following holds:

1. $\forall x \in S''$, $x \circ f(x) \in S$,
2. $\forall y \in S''$, $y \circ g(y) \in S$,
3. $\forall x, y \in S''$, the strings $f(x)$ and $g(y)$ are different in all coordinates, and
4. S'' contains $\sqrt{\alpha}/2$ fraction of the strings in $\Sigma^{l/2}$.

Assuming that we can show the existence of f, g and S'', the new protocol Π' is as follows: On input $x, y \in \Sigma^{l/2}$ Alice and Bob use the original c-bit (α, l) protocol on $x' = x \circ f(x)$ and $y' = y \circ g(y)$. By properties (1) and (2), for all x and y in S'', $x', y' \in S$, and so Π identifies an i such that $(x', y', i) \in \mathsf{FORK}$. By property (3), $i \leq l/2$, and $(x, y, i) \in \mathsf{FORK}$. By property (4), this is a c-bit $(\sqrt{\alpha}/2, l/2)$ protocol for $\mathsf{FORK}_{w,l/2}$.

Now we prove the existence of f, g and S'' with the desired properties. Let $A_1, \ldots, A_{l/2}$ be subsets of Σ where each A_i is of size $w/2$. If we ensure that $f(x) \in A = A_1 \times \cdots \times A_{l/2}$ and $g(y) \in B = \overline{A}_1 \times \cdots \times \overline{A}_{l/2}$, then property (3) immediately holds. So it remains to show that there exist such sets for which the other properties also hold. We will choose the A_is at random and show that this happens with non-zero probability. We choose A_is as follows: first choose at random $w/2$ strings $v^1, \cdots, v^{w/2}$ each of length $l/2$. Then we define A_i to include the i-th letter in each of these $w/2$ strings and extend it into a set of size $w/2$ randomly. (Note that the resulting sets $A_1, \ldots, A_{l/2}$ are indeed random and independent.) Now, fix $x \in S'$. An extension x' is a good choice for $f(x)$ if $x \cdot x' \in S$. Since $x \in S'$, we know that a random x' is good with probability at least $\alpha/2$. Hence the probability that none of the vectors in A is a good choice for $f(x)$ is less than $(1 - \alpha/2)^{w/2} < e^{-\alpha w/4}$. A similar analysis holds for good choices for $g(y)$ in B. Therefore, the probability that either A or the corresponding B is not good is at most $2e^{-\alpha w/4}$. So, for every $x \in S'$, at least $(1 - 2e^{-\alpha w/4})$ fraction of the partitions (A, B) is good. Hence, there is a partition that is good for at least $1 - 2e^{-\alpha w/4}$ of the elements of S'. Let S'' be this set of elements. The fraction of elements of $|\Sigma|^{l/2}$ in S'' is thus at least $(1 - 2e^{-\alpha w/4}) \cdot \sqrt{\alpha/2}$, which is at least $\sqrt{\alpha}/2$, for $\alpha \geq \lambda/w$ (for some constant λ). □

18.4 Putting it together

Using Lemma 18.1, Theorem 18.4, and choosing $l + 2 = w = \sqrt{n}$ we have

$$\mathrm{D}(\mathsf{KW}^{+}_{\mathsf{DSTCON}_n}) \geq \mathrm{D}(\mathsf{FORK}_{\sqrt{n}, \sqrt{n}-2}) = \Omega(\log^2 n)$$

Now using Theroem 2.14 from Lecture 2, we get the following theorem.

Theorem 18.6.

$$\mathrm{depth}^{+}(\mathsf{DSTCON}_n) = \Omega(\log^2 n)$$

References

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997. doi:10.2277/052102983X.