

Undirected s - t connectivity in deterministic logspace

Following Vadhan, Chapter 4

Gil Cohen

December 21, 2020

Overview

1 Space-bounded computation

2 s - t connectivity

3 $SL = L$

The model

We consider a Turing machine with four tapes:

- Input tape (R, \leftrightarrow): Read-only, can move left and right.
- Output tape (W, \rightarrow): Write-only, left to right
- Randomness tape (R, \rightarrow): Read-only left to right
- Work tape: (RW, \leftrightarrow), Read-write, can move left and right

Definitions and remarks

- The **space complexity**, s , is the number of cells used in the work tape.
- n usually denotes the input length. Typically $n \gg s$.
- The **randomness complexity** is the number of bits read from the randomness tape.

The model

Technicalities

- Every tape has a **head** - a pointer to the current location on the tape. The machine does not “remember” the head location. If the programmer wishes to do so (and she almost always does) then that should be done as part of the program. In particular, the space required will be accounted for in the space complexity.
- n is typically the largest parameter (we don't care about the head of the randomness tape). By paying an additional $O(\log n)$ in space, the above technicality can be ignored.
- For that price, we can also maintain any constant number of additional work tapes.

BPL vs. L

- A huge open problem in complexity theory: Simulate any space s (one, or better yet, two sided error) randomized algorithm deterministically in space $O(s)$.
- The regime $s = \Omega(\log n)$ is most interesting, and $s = \Theta(\log n)$ is complete for that regime.
- **L** stands for the class of all languages computable in deterministic logarithmic space. **BPL** is the class of all languages computable by a two-sided error randomized algorithm in logarithmic space.
- It is conjectured that **BPL** = **L**. The best known result due to Saks-Zhou (from the mid 90s) gives **BPL** \subseteq **L**^{3/2}.

s - t connectivity

Given a graph on n vertices and two vertices s, t , decide whether there exists a path from s to t .

- Solving s - t connectivity on directed graphs in space $O(\log n)$ would imply $\mathbf{NL} = \mathbf{L}$. Savitch's theorem (1970) gives a solution in space $O(\log^2 n)$, hence $\mathbf{NL} \subseteq \mathbf{L}^2$.
- Reingold (2005) solved s - t connectivity on undirected graphs, deterministically, in logarithmic space, proving $\mathbf{SL} = \mathbf{L}$.
- Approximating the probability a random walk on directed graphs starting at s reaches t to within a constant additive error in deterministic logarithmic space would imply $\mathbf{BPL} = \mathbf{L}$.

Overview

1 Space-bounded computation

2 s - t connectivity

3 $SL = L$

Reingold's idea

- Observe that the problem is easy if G happens to be a constant degree γ -spectral expander with $\gamma > 0$ constant. Indeed, the diameter of G is then logarithmic. Hence, a simple search will do.
- Transform the given graph G to a constant degree expander G' , while respecting connectivity, and use the above on G' .
- To obtain G' we repeatedly apply squaring (to improve expansion) and the Zig-Zag product (to reduce degree).

The algorithm

Input. An undirected graph G on n vertices, and two vertices s, t .

Ingredient. H a d -regular $\frac{3}{4}$ -spectral expander on d^4 vertices with self loops on every vertex.

- 1 Reduce (G, s, t) to a graph (G_0, s_0, t_0) which is d^2 -regular such that every connected component is nonbipartite. Moreover, s, t are connected in G if and only if s_0, t_0 are connected in G_0 .
- 2 For $k = 1, \dots, \ell = O(\log n)$,
 - 1 Let $G_k = G_{k-1}^2 \otimes H$
 - 2 Set s_k, t_k to be any two vertices in the clouds of G_k corresponding to s_{k-1}, t_{k-1} .
- 3 Exhaustively search all paths of length $O(\log n)$ in G_ℓ from s_ℓ and accept if one of them reach t_ℓ .

Correctness

- Let C_k denote the connected component of G_k containing s_k . Observe that $C_k = C_{k-1}^2 \circledast H$.
- C_0 being connected, undirected and nonbipartite implies that $\gamma(C_0) \geq \frac{1}{\text{poly}(n)}$. Now,

$$\gamma(C_{k-1}^2) \geq 2\gamma(C_{k-1}) - \gamma(C_{k-1})^2,$$

and so

$$\begin{aligned} \gamma(C_k) &= \gamma(C_{k-1}^2 \circledast H) \\ &\geq \left(\frac{3}{4}\right)^2 (2\gamma(C_{k-1}) - \gamma(C_{k-1})^2) \\ &\geq \min\left(\frac{35}{32} \cdot \gamma(C_{k-1}), \frac{1}{18}\right). \end{aligned}$$


Space analysis

The space analysis is delicate. We want to show that computing π_{G_k} can be done in space that is only a **constant** larger than the space required for computing $\pi_{G_{k-1}}$.


When analyzing constant space computation (or sub logarithmic space) the statement is model dependent, and so we next specify the exact model.

The end result is not model dependent as it is about logarithmic space.

Space analysis



$$\Downarrow \mathcal{T}_G(v, i) = (w, j)$$



Space analysis

Denote by $\text{space}(G)$ the amount of space required on the third tape for evaluating π_G .

Claim

If H is a graph of constant size then

$$\begin{aligned}\text{space}(G^2) &= \text{space}(G) + O(1) \\ \text{space}(G \otimes H) &= \text{space}(G) + O(1).\end{aligned}$$

Extra space for the proof

Extra space for the proof