Undirected *s*-*t* connectivity in deterministic logspace

Gil Cohen

February 8, 2024

Overview

1. Space-bounded computation
2. s-t connectivity
3. **SL** = **L**

Adding two numbers

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

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

- On input x, the space complexity, $s(x)$, is the number of cells used in the work tape (taken worst case over the randomness).
- n usually denotes the input length, and $s(n) = \max\{s(x) \mid x \in \{0,1\}^n\}$. Typically $n \gg s$.
- The randomness complexity $r(x)$ (and $r(n)$) is the number of bits read from the randomness tape on input x (and on any length-n input, resp.).

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 fundamental 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.

BPL vs. L

- A fundamental 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. **RL** captures one-sided error.
- It is conjectured that **BPL** = **L**. The best known result due to Saks-Zhou (from the mid 90s) gives $\mathbf{BPL} \subseteq \mathbf{L}^{3/2}$.

Overview

s-t connectivity

Given a graph $G = (V, E)$ on n vertices and two vertices s, t, decide whether there exists a path from s to t in G.

- Solving s-t connectivity for 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 $\mathbf{NL} \subseteq \mathbf{L}^2$.

s-t connectivity

Given a graph $G = (V, E)$ on n vertices and two vertices s, t, decide whether there exists a path from s to t in G.

- Solving s-t connectivity for 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 $\mathbf{NL} \subseteq \mathbf{L}^2$.
- Deterministically 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 $\mathbf{BPL} = \mathbf{L}$.

s-t connectivity

Given a graph $G = (V, E)$ on n vertices and two vertices s, t, decide whether there exists a path from s to t in G.

- Solving s-t connectivity for 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 $\mathbf{NL} \subseteq \mathbf{L}^2$.
- Deterministically 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 $\mathbf{BPL} = \mathbf{L}$.
- s-t connectivity for undirected graphs in logarithmic space using randomness is easy.

s-t connectivity

Given a graph $G = (V, E)$ on n vertices and two vertices s, t, decide whether there exists a path from s to t in G.

- Solving s-t connectivity for 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 $\mathbf{NL} \subseteq \mathbf{L}^2$.
- Deterministically 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 $\mathbf{BPL} = \mathbf{L}$.
- s-t connectivity for undirected graphs in logarithmic space using randomness is easy.
- Reingold (2005) solved s-t connectivity for undirected graphs, deterministically, in logarithmic space, proving $\mathbf{SL} = \mathbf{L}$.

Overview

Reingold's idea

- Note that the problem is easy for G a constant degree ω-spectral expander, where $\omega < 1$ is a constant.

Reingold's idea

- Note that the problem is easy for G a constant degree ω-spectral expander, where $\omega < 1$ is a constant.
- Transform the given graph G to a constant degree expander G', while respecting connectivity, and use the above on G'.

Reingold's idea

- Note that the problem is easy for G a constant degree ω-spectral expander, where $\omega < 1$ is a constant.
- 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 Zig-Zag (to reduce the degree).

The algorithm

Input. An undirected graph $G = (V, E)$ on n vertices, and $s, t \in V$.

Ingredient. H a d-regular $\frac{1}{4}$-spectral expander on d^4 vertices with a self loop on every vertex.

1. Transform (G, s, t) to a graph (G_0, s_0, t_0) which is d^2-regular such that every connected component is non-bipartite. Moreover, s, t are connected in $G \iff s_0, t_0$ are connected in G_0.
2. For $k = 1, \ldots, \ell = O(\log n)$,
    1. Let $G_k = G_{k-1}^2 \text{ⓩ} H$.
    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 $C_k = C_{k-1}^2 \, ⓩ \, H$.
- C_0 being connected, undirected and non-bipartite implies that $\gamma(C_0) \geq \frac{1}{\text{poly}(n)}$ (problem set). Now, $\omega(C_{k-1}^2) = \omega(C_{k-1})^2$ and so
$$\gamma(C_{k-1}^2) = 2\gamma(C_{k-1}) - \gamma(C_{k-1})^2.$$

Thus,

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

Space analysis

The space analysis is delicate. We want to show that computing

$$\pi_{G_k} : [V_k] \times [d^2] \to [V_k] \times [d^2]$$

can be done in space that is only constant larger compared to the space required for computing

$$\pi_{G_{k-1}} : [V_{k-1}] \times [d^2] \to [V_{k-1}] \times [d^2].$$

Space analysis

The space analysis is delicate. We want to show that computing

$$\pi_{G_k} : [V_k] \times [d^2] \to [V_k] \times [d^2]$$

can be done in space that is only constant larger compared to the space required for computing

$$\pi_{G_{k-1}} : [V_{k-1}] \times [d^2] \to [V_{k-1}] \times [d^2].$$

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

Denote by 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

$$\text{space}(G^2) = \text{space}(G) + O(1)$$
$$\text{space}(G \textcircled{z} H) = \text{space}(G) + O(1).$$

Space analysis

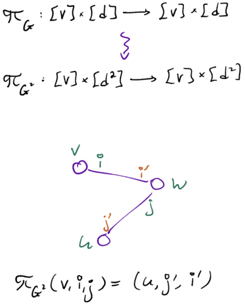

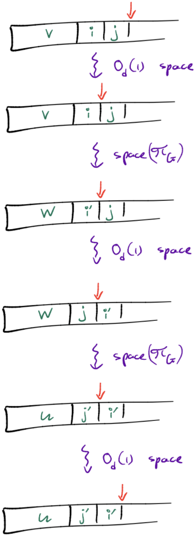

Space analysis

$$\pi_G : [V]\times[d] \longrightarrow [V]\times[d]$$

$$\pi_H : [d]\times[c] \longrightarrow [d]\times[c]$$

$$\rightsquigarrow$$

$$\pi_{G ⓩ H} : \big([V]\times[d]\big)\times\big([d]\times[c]\big) \circlearrowleft$$

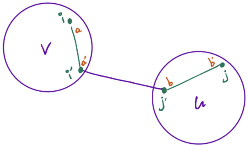

$$\pi_{G ⓩ H}(v, i, a, b) = (u, j, b', a')$$

Extra space for the proof

Extra space for the proof