1 RL = NL

We denote the class of languages that are accepted by logarithmically space bounded PTMs with one sided error by RL. Just as RP $\subseteq$ NP, it is straightforward to show that RL $\subseteq$ NL. It turns out that RL $=$ NL. While this may seem surprising at first, it follows as a result of the fact that probabilistic log space bounded classes can fruitfully use more than a polynomial number of steps. In contrast, the class NL is the same whether or not we require that the underlying machines halt in polynomial time.

Theorem 1 (Gill) NL $=$ RL.

Proof We show the harder direction, that NL $\subseteq$ RL. Let L be accepted by a NTM M that is $O(\log n)$ space bounded. Without loss of generality, assume that M has at most two possible transitions from any configuration. Let c be a constant such that for all $x \in L$, there is an accepting computation of M on x that halts in at most $t(|x|) = c^{\lceil \log |x| \rceil}$ steps. We construct a PTM M' that accepts L as follows. On input x, the machine iterates the following steps until it halts.

1. Simulate M on x from the initial configuration for $t(|x|)$ steps, where at each step, one of the possible (nondeterministic) transitions of M is chosen uniformly at random. If M accepts, then halt and accept.

2. Generate $t(|x|) + 1$ random bits. If all are heads, reject.

Clearly, M' rejects all inputs not in L with probability 1. If $x \in L$, then any single simulation of M on x in step 1 will halt and accept with probability at least $2^{-t(|x|)}$. Therefore,

$$\begin{aligned}
& \text{Prob}[M' \text{ rejects } x] \\
= \;& \sum_{i=1}^{\infty} \text{Prob}[M' \text{ rejects after exactly } i \text{ iterations}] \\
\leq \;& \sum_{i=1}^{\infty} \text{Prob}[i \text{ iterations of step 1 do not accept}] \cdot \text{Prob}[i\text{th iteration of step 2 rejects}] \\
\leq \;& \sum_{i=1}^{\infty} (1 - 2^{-t(|x|)})^i 2^{-t(|x|)+1} \\
= \;& 2^{-t(|x|)+1} \sum_{i=1}^{\infty} (1 - 2^{-t(|x|)})^i \\
< \;& 2^{-t(|x|)+1} 2^{t(|x|)} = 1/2.
\end{aligned}$$

Hence, $\text{Prob}[M' \text{ accepts } x] > 1/2$, as required. $\square$

2 UPATH is in RLP

Let UPATH (Undirected Graph Reachability) $= \{\langle G, s, t \rangle \mid$ there is a path from s to t in undirected graph $G\}$. We have seen that UPATH is in NL. Here, we show a stronger result. Let RLP be the class of languages that are accepted by a PTM M that is both logarithmically space bounded and polynomial time bounded and has 1-sided error.

Theorem 2 UPATH $\in$ RLP.

The randomized algorithm for UPATH is quite simple. On input $\langle G = (V, E), s, t\rangle$, a random walk is performed on the graph, starting at s. That is, nodes of the graph are visited according to the following random process: initially s is visited, and when node i is visited, an adjacent node j is chosen uniformly at random and node j is visited next. If t is reached at some step, the algorithm halts and accepts. If t is not reached within $6e(n-1)$ steps of the walk, where $n = |V|$ and $e = |E|$, then the algorithm halts and rejects.

Clearly, this algorithm runs in polynomial time and rejects input graphs in which there is no path from s to t. Therefore, let $\langle G, s, t\rangle$ be in UPATH. Let $T(G, s, t)$ be the expected number of edges traversed by a random walk starting at s, until t is reached. We would like to bound $T(G, s, t)$. For an edge $\{i, j\}$ of G, let T_{ij} be the expected time to reach j on a random walk starting at node i. Let σ be any simple path from s to t. Then, $T(G, s, t) \leq \sum_{(i,j)\in\sigma} T_{ij}$. We will show in Section 4 that $T_{ij} \leq 2e$. Then

$$T(G, s, t) = \sum_{(i,j)\in\sigma} T_{ij} \leq 2e \sum_{(i,j)\in\sigma} 1 = 2e|\sigma| \leq 2e(n-1).$$

We now apply Markov's inequality to bound the probability that the algorithm rejects the input $\langle G, s, t\rangle$ in UPATH.

Markov's Inequality. *If X is a nonnegative random variable and k is a positive real then*

$$\text{Prob}[X \geq kE[X]] \leq 1/k.$$

Proof Let $f(X)$ be an indicator variable that is 1 if $X \geq t$ and is 0 otherwise. Then,

$$\text{Prob}[X \geq t] = E[f(X)].$$

Since $f(X) \leq X/t$, we have that $E[f(X)] \leq E[X/t] = E[X]/t$. Choosing $t = kE[X]$, we have that

$$\text{Prob}[X \geq kE[X]] \leq E[X]/(kE[X]) = 1/k.$$

□

The number of edges traversed on a random walk from s to t is a random variable. A direct application of Markov's Inequality with $k = 3$ to this random variable shows that the probability that more than $6e(n-1)$ edges are traversed on a random walk from s to t is at most 1/3. We conclude that the algorithm accepts $\langle G, s, t\rangle$ with probability $> 1/2$.

3 Background: Markov Chains

To bound T_{ij}, we will apply some theory of Markov chains. A finite Markov Chain (with discrete time and stationary transition probabilities) is an infinite sequence $X_0, X_1, X_2, \ldots,$ of random variables, each over state space S such that for all $i, j, a_0, a_1, \ldots, a_{k-2} \in S$,

$$\text{Prob}[X_k = j \mid X_0 = a_0, X_1 = a_1, \ldots X_{k-2} = a_{k-2}, X_{k-1} = i] = \text{Prob}[X_k = j \mid X_{k-1} = i] = P_{ij},$$

where P_{ij} may depend on i and j but not on $k, a_0, a_1, \ldots, a_{k-2}$. For example, the matrix

$$P = \left(\begin{array}{ll} 1-p & p \\ q & 1-q \end{array} \right)$$

together with the initial condition that$X_0 = 1$ with probability 1 defines a Markov chain over state space $\{1, 2\}$. In what follows, we assume that the state space is $\{1, 2, \ldots, n\}$. For any $l \in \mathbb{N}$, P^l_{ij} is the probability of going from state i to state j in exactly l steps; this can be proved by induction on l.

Definition 1 *A Markov Chain is* irreducible *if and only if for all i, j, there exists t such that*

$$\mathit{Prob}[X_t = j \mid X_0 = i] > 0.$$

Theorem 3 *Let P be the transition probability matrix of an irreducible finite Markov chain. Then $\pi P = \pi$ has a unique solution π up to constant multiplicative factors.*

A proof of this can be found in "A First Course in Stochastic Processes" by Karlin and Taylor, 1975, Chapter 3, Theorem 1.3.
Example. For the P of our example above, $\pi = (q/(p+q), p/(p+q))$.

Definition 2 *A* random walk *on an undirected connected graph $G = (V, E)$ is an irreducible Markov chain with state space equal to V and*

$$P_{uv} = \begin{cases} 1/d(u), & \textit{if } \{u, v\} \in E, \\ 0, & \textit{otherwise}. \end{cases}$$

Let P be the transition probability matrix of a random walk on G. Let π be the unique vector such that $\pi P = \pi$ and $\sum_i \pi_i = 1$. Let $N(u)$ be the set of neighbours of a node u of G and let $d(u) = |N(u)|$ be the degree of node u. Then, $\pi_i = d(i)/(2e)$. To see this, note that for all $j \in V$,

$$[\pi P]_j = \sum_{i \in V} \pi_i P_{ij} = \sum_{i \in N(j)} \frac{d(i)}{2e} \frac{1}{d(i)} = \frac{1}{2e} \sum_{i \in N(j)} 1 = \frac{d(j)}{2e}$$

and

$$\sum_{i \in V} \pi_i = \sum_{i \in V} d(i)/(2e) = 1.$$

Example. Consider the line graph $G = (\{1, 2, 3, 4, 5\}, \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 5\}\}$. For this graph, if π satisfies Theorem 3 and its elements sum to 1, then $\pi_1 = \pi_5 = 1/8$, and $\pi_2 = \pi_3 = \pi_4 = 1/4$.

Definition 3 *A* random commute *from i to j is a random walk starting at i that ends the first time it returns to i after having at some point visited j.*

For $i, j, u, v \in V$ with $\{u, v\} \in E$, let θ_{ijuv} be the expected number of times edge $\{u, v\}$ is visited from u to v on a random commute from i to j.
Example. Continuing informally with our line graph example, 2, 3, 4, 3, 4, 5, 4, 3, 2 is an example of a commute from 2 to 3.

Also for this example, it is not hard to see that $\theta_{1212} = 1$. Consider θ_{1223}. When a walk moves from node 1 to node 2, half of the time the walk returns to node 1 without visiting node 3 at all. The other half of the time, node 3 is visited, in which case we expect to visit node 2 two more times before completing our commute at node 1. Therefore,

$$\theta_{1223} = (1/2)0 + (1/2)2 = 1.$$

In fact, for all edges $\{u, v\}$, $\theta_{12uv} = 1$.

We will next show that for any graph G, θ_{ijuv} is independent of u and v. We first consider v.

Claim 1 *θ_{ijuv} is independent of v. That is, for all $v' \in N(u)$, $\theta_{ijuv} = \theta_{ijuv'}$.*

This claim follows from the fact that each time u is visited on a random commute from i to j, v and v' are visited after u with equal probability. Let θ_{iju} equal θ_{ijuv} for any v.

Claim 2 *θ_{iju} is independent of u.*

Proof For all $u \in V$,

$$d(u)\theta_{iju} = \sum_{v \in N(u)} \theta_{ijv} = \sum_{v \in N(u)} d(v)\theta_{ijv} 1/d(v) = \sum_{v \in V} d(v)\theta_{ijv} P_{vu}. \tag{1}$$

Here, the term on the left is the expected number of times the random commute leaves u and the term on the right is the expected the number of times that the random commute enters u. Let $\pi'_u = d(u)\theta_{iju}$. Then from Equation (1), $\pi' = \pi' P$. Since π is unique up to constant multiplicative factors,

$$d(u)\theta_{iju} = \lambda_{ij} \frac{d(u)}{2e},$$

where λ_{ij} is a constant independent of u. Therefore, $\theta_{iju} = \lambda_{ij}/(2e)$. □

4 Back to T_{ij}

Recall that T_{ij} is the expected number of steps to reach j on a random walk starting at node i, where $\{i, j\}$ is an edge of G.

Theorem 4 $T_{ij} \leq 2e$.

Proof

$$T_{ij} \leq \sum_{\{u,v\} \in E} (\theta_{ijuv} + \theta_{ijvu}) = \sum_{\{u,v\} \in E} 2\theta_{ijij} = \theta_{ijij}(\sum_{\{u,v\} \in E} 2) = \theta_{ijij}(2e) \leq 2e.$$

□