Probabilistic computation

Probabilistic computations

Consider the Turing machines M in the following form:

- It has an input tape;
- It has a read-only, no-left-move randomness tape, where each cell (up to infinity) contains either 0 or 1;
- It has internal state, working tapes,...
- It accepts through final states.

We say $\Pr[M(x) = 1] = p$ if $\Pr[M(x, \alpha) = 1 \,|\, \alpha \leftarrow U_\infty] = p$.

- U_∞ is the uniform probability distribution on $\{0, 1\}^\omega$.
  - Each bit will be either 0 or 1 with equal probability, and independently of all other bits.

Classes RP, coRP

The PTM M Monte-Carlo recognizes language L if for all $x \in \{0,1\}^*$:

$$x \in L \Leftrightarrow \Pr[M(x) = 1] \geq 1/2$$
$$x \notin L \Leftrightarrow \Pr[M(x) = 1] = 0$$

Class RP is the set of all languages Monte-Carlo recognizable by polynomial-time PTM-s.

- Running time is polynomial wrt. the length of the first argument of M.

$L \in$ coRP if exists a poly-time PTM M, such that for all $x \in \{0,1\}^*$:

$$x \in L \Leftrightarrow \Pr[M(x) = 1] = 1$$
$$x \notin L \Leftrightarrow \Pr[M(x) = 1] \leq 1/2$$

Example problem in coRP

An arithmetic expression E over variables $x_1, \ldots, x_k$ is one of

- Variable x_i;
- Constant $n \in \mathbb{Z}$;
- Expression $E_1 + E_2$, $E_1 - E_2$, $E_1 \cdot E_2$.

Given expression E. Is it identical to 0?

This is the polynomial identity testing problem. We do not know how to do it in P.

Schwartz-Zippel lemma

Theorem. Let p be a non-zero k-variable $\leq d$-degree polynomial over integers. Let S be a finite subset of $\mathbb{Z}$. If we randomly uniformly pick $x_1, \ldots, x_k$ from S, then $\Pr[p(x_1, \ldots, x_k) = 0] \leq d/|S|$.

Proof. A $\leq d$-degree single-variable polynomial has at most d roots. Continue by induction over the number of variables.

Algorithm for polynomial identity testing. Select a sufficiently large S. Randomly pick $x_1, \ldots, x_k$ from S. Evaluate the arithmetic expression.

Class ZPP

- Let p_i be the probability that $M(x)$ stops in exactly i steps.
- The expected running time of $M(x)$ is $1 \cdot p_1 + 2 \cdot p_2 + \cdots$.
- $L \in$ ZPP if there exists a PTM M, such that
  - M runs in expected polynomial time;
  - If $x \in L$ then $\Pr[M(x) = 1] = 1$.
  - If $x \notin L$ then $\Pr[M(x) = 1] = 0$.
- Such M is called a Las Vegas algorithm for L.

$\mathsf{ZPP} = \mathsf{RP} \cap \mathsf{coRP}$

Theorem. $\mathsf{ZPP} = \mathsf{RP} \cap \mathsf{coRP}$.

Theorem (Obvious) $\mathsf{RP} \subseteq \mathsf{NP}$. $\mathsf{coRP} \subseteq \mathsf{coNP}$. $\mathsf{P} \subseteq \mathsf{ZPP}$.

Handling biased coins

Let the bits in the randomness tape still be independent of each other. Let $0 < p < 1$.

- If bit 1 has probability p, then a tape where bit 1 has probability $1/2$ can still be simulated.
- If bit 1 has probability $1/2$, then a tape where bit 1 has probability p can still be simulated.
  - The bits of p must be computable in polynomial time.

Class BPP

The PTM M recognizes language L with bounded error if for all $x \in \{0,1\}^*$:

$$x \in L \Leftrightarrow \Pr[M(x) = 1] \geq 2/3$$
$$x \notin L \Leftrightarrow \Pr[M(x) = 1] \leq 1/3$$

Class BPP is the set of all languages recognizable by polynomial-time PTM-s with bounded error.

Chernoff bounds

Theorem.

- Let $X_1, \ldots, X_n$ be mutually independent random variables with values from $\{0, 1\}$. Let $X = X_1 + \cdots + X_n$.
- Let $\mu = \mathbf{E}[X] = \sum_{i=1}^{n} \mathbf{E}[X_i]$.
- Let $\delta > 0$. Then

$$\Pr[X \geq (1+\delta)\mu] \leq \left(\frac{e^{\delta}}{(1+\delta)^{(1+\delta)}} \right)^{\mu}$$

$$\Pr[X \leq (1-\delta)\mu] \leq \left(\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}} \right)^{\mu}$$

Some lemmas

- If $X_1, \ldots, X_n$ are mutually independent random variables then $\mathbf{E}[\prod_{i=1}^n X_i] = \prod_{i=1}^n \mathbf{E}[X_i]$.
- If X is a random variable then $\Pr[X \geq k \cdot \mathbf{E}[X]] \leq 1/k$. (**Markov's inequality**)

Proof of Chernoff bound

Let $p_i = \Pr[X_i = 1]$. Let $t = \ln(1+\delta)$. Let $\mathbf{P} = \Pr[X \geq (1+\delta)\mu]$.

$$\begin{aligned}
\mathbf{E}[e^{tX}] &= \prod_{i=1}^{n} \mathbf{E}[e^{tX_i}] = \prod_{i=1}^{n}(1 - p_i + p_i e^t) \leq \prod_{i=1}^{n} \exp(p_i(e^t - 1)) \\
&= e^{\mu\delta} \\
\mathbf{P} &= \Pr[e^{tX} \geq e^{t(1+\delta)\mu}] \leq \frac{\mathbf{E}[e^{tX}]}{e^{t(1+\delta)\mu}} \leq \frac{e^{\mu\delta}}{(1+\delta)^{(1+\delta)\mu}}
\end{aligned}$$

Role of constants in defining RP and BPP

For every $\varepsilon > 0$, define the class BPP_ε as the set of languages L, such that exists poly-time PTM M, such that

$$x \in L \Leftrightarrow \Pr[M(x) = 1] \geq 1 - \varepsilon$$
$$x \notin L \Leftrightarrow \Pr[M(x) = 1] \leq \varepsilon$$

Theorem. If $0 < \varepsilon < 1/2$ then $\mathsf{BPP} = \mathsf{BPP}_\varepsilon$.

More general ε

For any function $e : \mathbb{N} \to \mathbb{R}_+$ define the class BPP_e as the set of languages L, such that exists poly-time PTM M, such that

$$x \in L \Leftrightarrow \Pr[M(x) = 1] \geq 1 - e(|x|)$$
$$x \notin L \Leftrightarrow \Pr[M(x) = 1] \leq e(|x|)$$

Theorem. $\mathsf{BPP}_{\lambda n.1/2-1/poly(n)} = \mathsf{BPP} = \mathsf{BPP}_{\lambda n.2^{-poly(n)}}$

Exercise. What is the corresponding result for RP?

BPP is the model of efficient computation if random choices are allowed.

Strict vs. expected running time

Let M be a PTM that recognizes a language L in expected polynomial time as follows:

$$x \in L \Leftrightarrow \Pr[M(x) = 1] \geq p$$
$$x \notin L \Leftrightarrow \Pr[M(x) = 1] \leq q \; .$$

Then for any ε, there exists a PTM M' that recognizes L in strict polynomial time as follows:

$$x \in L \Leftrightarrow \Pr[M(x) = 1] \geq p - \varepsilon$$
$$x \notin L \Leftrightarrow \Pr[M(x) = 1] \leq q + \varepsilon \; .$$

BPP $\subseteq$ P$/poly$

Theorem. BPP $\subseteq$ P$/poly$.

Proof. Let $L \in$ BPP $=$ BPP$_{\lambda n.2^{-2n}}$. Let M recognize L with error probability 2^{-2n}.

- For each $x \in \{0,1\}^n$, $M(x, \alpha)$ returns whether $x \in L$ for a vast majority of strings α.
- Let E_x be the set of strings α where $M(x, \alpha)$ gives the wrong answer.
- There is a string $\alpha_n \notin \bigcup_{x \in \{0,1\}^n} E_x$.
- The prefix of α_n (bounded by running time of M) is a suitable advice for M.

Corollary. If SAT $\in$ BPP then PH $= \Sigma_2^p$.

BPP $\subseteq \Sigma_2^p \cap \Pi_2^p$

Theorem. BPP $\subseteq \Sigma_2^p \cap \Pi_2^p$

Proof. Let M accept $L \in$ BPP with error probability $\lambda n.2^{-n}$. Let f be running time of M.

- For $x \in \{0,1\}^n$, let $S_x \subseteq \{0,1\}^{f(n)}$ be the set of strings α, such that $M(x,\alpha)$ accepts.
  - $|S_x| \geq (1-2^{-n}) \cdot 2^{f(n)}$ or $|S_x| \leq 2^{f(n)-n}$.
- For $u \in \{0,1\}^{f(n)}$, denote $u \oplus S = \{u \oplus \alpha \,|\, \alpha \in S\}$.
- Let $k = \lceil f(n)/n \rceil + 1$.
- If $\begin{array}{l} |S_x| \leq 2^{f(n)-n} \\ |S_x| \geq (1-2^{-n})2^{f(n)} \end{array}$, then $\begin{array}{l}\text{for all} \\ \text{exists}\end{array}$ $\{u_1, \ldots, u_k\} \subseteq \{0,1\}^{f(n)}$: $\bigcup_{j=1}^{k}(u_j \oplus S_x) \begin{array}{c}\neq \\ =\end{array} \{0,1\}^{f(n)}$.
- $x \in L \Leftrightarrow \exists u_1, \ldots, u_k \forall \alpha : \bigvee M(x, \alpha \oplus u_j)$. A Σ_2-procedure.

The classes RL and BPL

- Same as RP and BPP, but for log-space, not poly-time.
- RL is noteworthy for being known to contain the reachability problem in undirected graphs.
  - A couple of years ago, Omer Reingold showed that undirected reachability is in L.