CS151
Complexity Theory

Lecture 7

April 24, 2017

Randomness

- 3 examples of the power of randomness
  - communication complexity
  - polynomial identity testing
  - complexity of finding unique solutions
- randomized complexity classes

1. Communication complexity

two parties: Alice and Bob

function $f:\{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$

Alice holds $x \in \{0,1\}^n$; Bob holds $y \in \{0,1\}^n$

- Goal: compute f(x, y) while communicating as few bits as possible between Alice and Bob

- count number of bits exchanged (computation free)
- at each step: one party sends bits that are a function of held input and received bits so far

Communication complexity

- simple function (equality):

  $EQ(x, y) = 1$ iff $x = y$

- simple protocol:
  - Alice sends x to Bob (n bits)
  - Bob sends EQ(x, y) to Alice (1 bit)
  - total: n + 1 bits
  - (works for any predicate f)

Communication complexity

- Can we do better?
  - deterministic protocol?
  - probabilistic protocol?
    - at each step: one party sends bits that are a function of held input and received bits so far and the result of some coin tosses
    - required to output f(x, y) with high probability over all coin tosses

Communication complexity

Theorem: no deterministic protocol can compute EQ(x, y) while exchanging fewer than n+1 bits.

- Proof:
  - "input matrix":

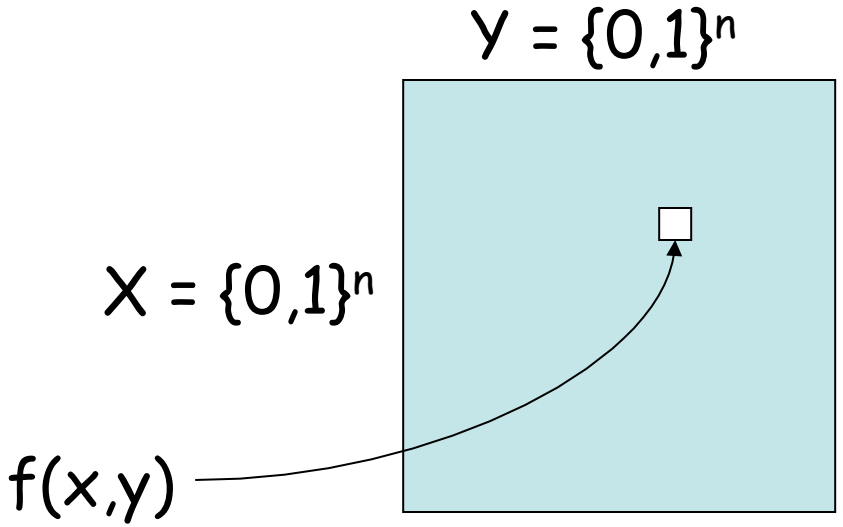

Communication complexity

– assume 1 bit sent at a time, alternating (same proof works in general setting)

– A sends 1 bit depending only on x:

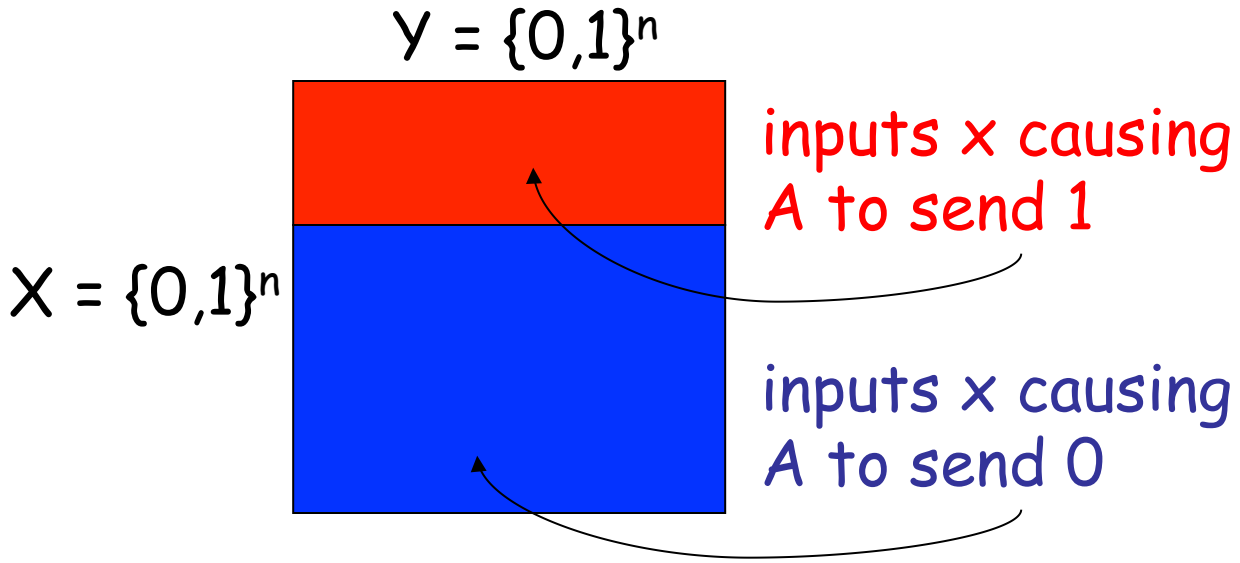

Communication complexity

– B sends 1 bit depending only on y and received bit:

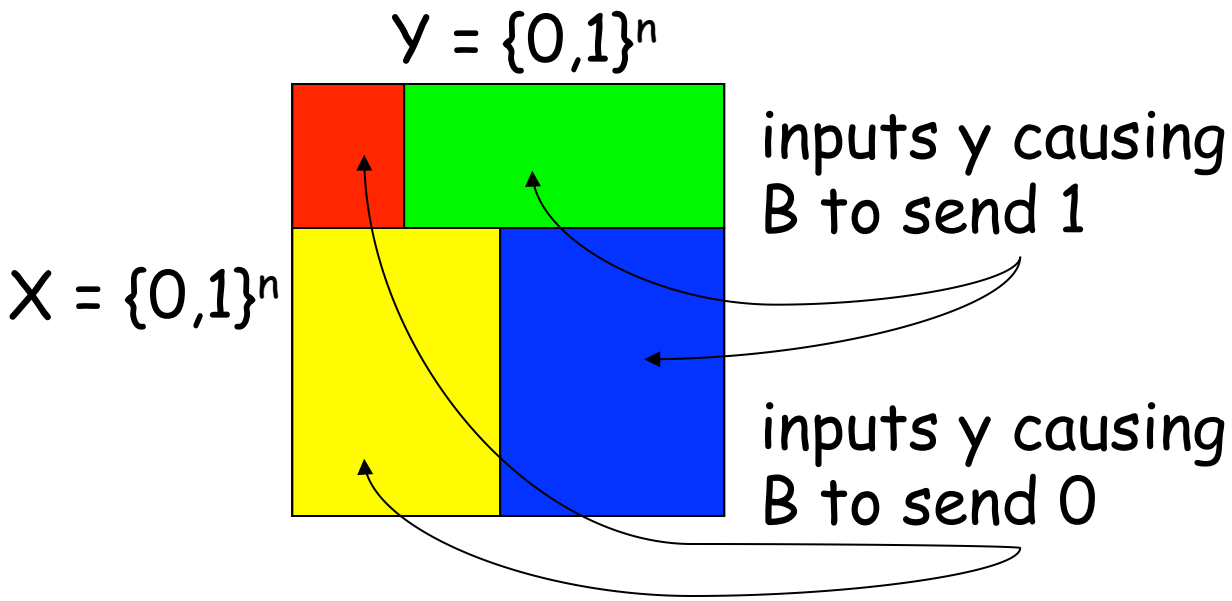

Communication complexity

– at end of protocol involving k bits of communication, matrix is partitioned into at most 2^k combinatorial rectangles

– bits sent in protocol are the same for every input (x, y) in given rectangle

– conclude: f(x,y) must be constant on each rectangle

Communication complexity

Matrix for EQ:

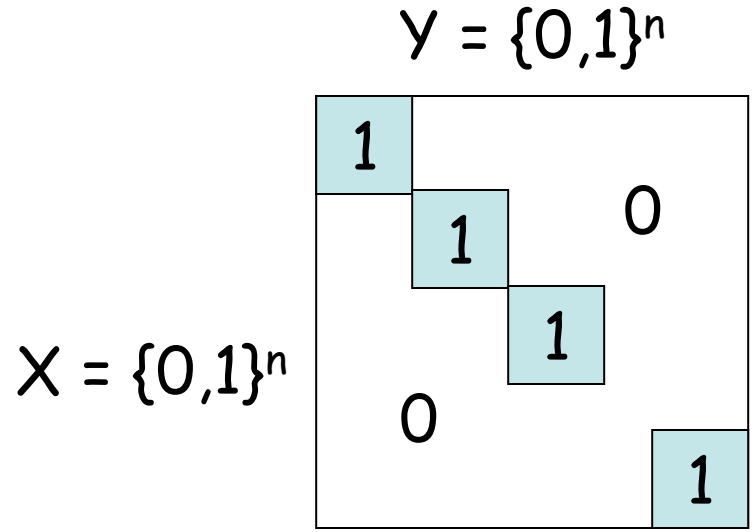

- any partition into combinatorial rectangles with constant f(x,y) must have $2^n + 1$ rectangles
- protocol that exchanges ≤ n bits can only create 2^n rectangles, so must exchange at least n+1 bits.

Communication complexity

- protocol for EQ employing randomness?
  - Alice picks random prime p in $\{1...4n^2\}$, sends:
    - p
    - (x mod p)
  - Bob sends:
    - (y mod p)
  - players output 1 if and only if:

$$(x \bmod p) = (y \bmod p)$$

Communication complexity

– O(log n) bits exchanged
– if x = y, always correct
– if x ≠ y, incorrect if and only if:

p divides |x – y|

– # primes in range is ≥ 2n
– # primes dividing |x – y| is ≤ n
– probability incorrect ≤ 1/2

Randomness gives an exponential advantage!!

2. Polynomial identity testing

- Given: polynomial $p(x_1, x_2, \ldots, x_n)$ as arithmetic formula (fan-out 1):
  - multiplication (fan-in 2)
  - addition (fan-in 2)
  - negation (fan-in 1)

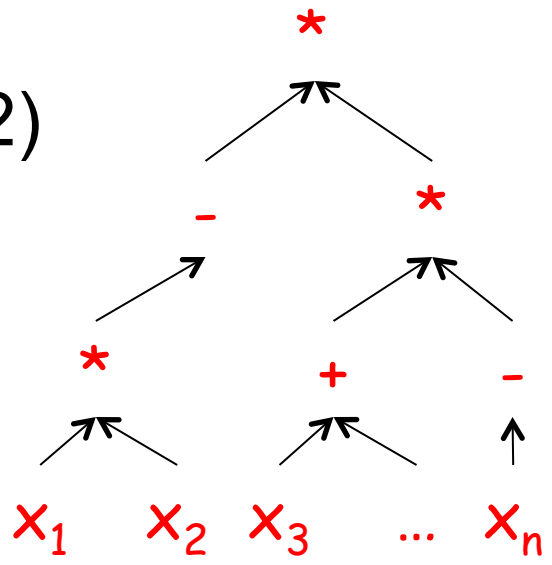

Polynomial identity testing

- Question: Is p identically zero?
  - i.e., is $p(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathbf{F}^n$
  - (assume $|\mathbf{F}|$ larger than degree…)

- "polynomial identity testing" because given two polynomials p, q, we can check the identity $p \equiv q$ by checking if $(p - q) \equiv 0$

Polynomial identity testing

- try all $|\mathbf{F}|^n$ inputs?
  - may be exponentially many
- multiply out symbolically, check that all coefficients are zero?
  - may be exponentially many coefficients

- can randomness help?
  - i.e., flip coins, allow small probability of wrong answer

Polynomial identity testing

Lemma (Schwartz-Zippel): Let

$$p(x_1, x_2, \ldots, x_n)$$

be a total degree d polynomial over a field **F** and let S be any subset of **F**. Then if p is not identically 0,

$$\Pr_{r_1,r_2,\ldots,r_n \in S}[\, p(r_1, r_2, \ldots, r_n) = 0] \leq d/|S|.$$

Polynomial identity testing

- Proof:
  - induction on number of variables n
  - base case: n = 1, p is univariate polynomial of degree at most d
  - at most d roots, so

$$\Pr[\, p(r_1) = 0] \leq d/|S|$$

Polynomial identity testing

– write $p(x_1, x_2, \ldots, x_n)$ as

$$p(x_1, x_2, \ldots, x_n) = \Sigma_i (x_1)^i p_i(x_2, \ldots, x_n)$$

– k = max. i for which $p_i(x_2, \ldots, x_n)$ not id. zero

– by induction hypothesis:

$$\Pr[\, p_k(r_2, \ldots, r_n) = 0] \leq (d\text{-}k)/|S|$$

– whenever $p_k(r_2, \ldots, r_n) \neq 0$, $p(x_1, r_2, \ldots, r_n)$ is a univariate polynomial of degree k

$$\Pr[p(r_1,r_2,\ldots,r_n)=0 \mid p_k(r_2,\ldots,r_n) \neq 0] \leq k/|S|$$

Polynomial identity testing

$$\Pr[\, p_k(r_2, \ldots, r_n) = 0] \le (d-k)/|S|$$

$$\Pr[p(r_1,r_2,\ldots,r_n)=0 \mid p_k(r_2,\ldots,r_n) \neq 0] \le k/|S|$$

– conclude:

$$\Pr[\, p(r_1, \ldots, r_n) = 0] \le (d-k)/|S| + k/|S| = d/|S|$$

– Note: can add these probabilities because

$$\Pr[E_1] = \Pr[E_1|E_2]\Pr[E_2] + \Pr[E_1|\neg E_2]\Pr[\neg\, E_2]$$
$$\le \Pr[E_2] + \Pr[E_1|\neg E_2]$$

Polynomial identity testing

- Given: polynomial $p(x_1, x_2, \ldots, x_n)$
- Is p identically zero?

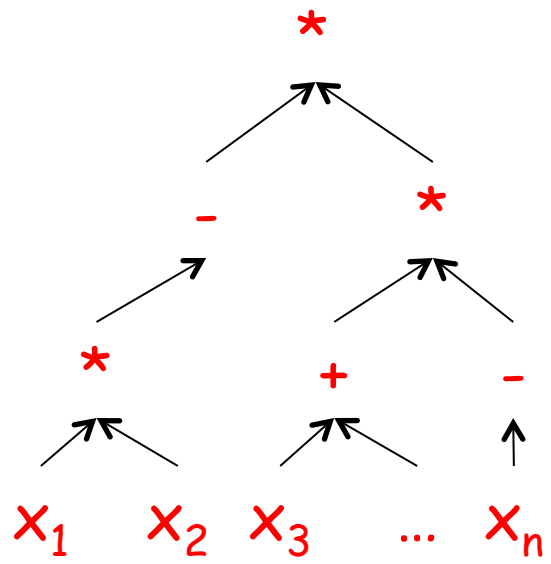

- Note: degree d is at most the size of input

Polynomial identity testing

- randomized algorithm: field **F,** pick a subset S $\subset$ **F** of size 2d
  - pick $r_1, r_2, \ldots, r_n$ from S uniformly at random
  - if $p(r_1, r_2, \ldots, r_n) = 0$, answer "yes"
  - if $p(r_1, r_2, \ldots, r_n) \neq 0$, answer "no"

- if p identically zero, never wrong
- if not, Schwartz-Zippel ensures probability of error at most ½

Polynomial identity testing

- Given: polynomial $p(x_1, x_2, \ldots, x_n)$

- Is p identically zero?

What if polynomial is given as arithmetic circuit?

- max degree?
- does the same strategy work?

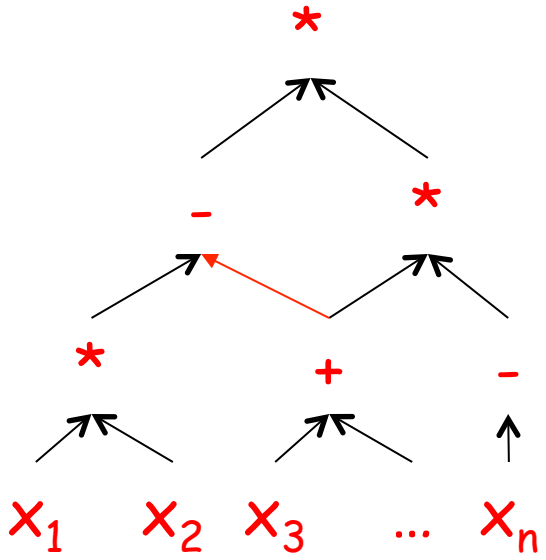

3. Unique solutions

- a positive instance of SAT may have many satisfying assignments
- maybe the difficulty comes from not knowing which to “work on”
- if we knew # satisfying assignments was 1 or 0, could we zoom in on the 1 efficiently?

Unique solutions

Question: given polynomial-time algorithm that works on SAT instances with at most 1 satisfying assignment, can we solve general SAT instances efficiently?

- Answer: yes
  - but (currently) only if "efficiently" allows randomness

Unique solutions

Theorem (Valiant-Vazirani): there is a randomized poly-time procedure that given a 3-CNF formula

$$\varphi(x_1, x_2, \ldots, x_n)$$

outputs a 3-CNF formula φ' such that

- if φ is not satisfiable then φ' is not satisfiable
- if φ is satisfiable then with probability at least $1/(8n)$ φ' has exactly one satisfying assignment

Unique solutions

- Proof:
  - given subset $S \subset \{1, 2, \ldots, n\}$, there exists a 3-CNF formula θ_S on $x_1, x_2, \ldots, x_n$ and additional variables such that:
    - θ_S is satisfiable iff an even number of variables in $\{x_i\}_{i \in S}$ are true
    - for each such setting of the x_i variables, this satisfying assignment is unique
    - $|\theta_S| = O(n)$
    - not difficult; details omitted

Unique solutions

– set $\varphi_0 = \varphi$
– for i = 1, 2, …, n
  • pick random subset S_i
  • set $\varphi_i = \varphi_{i-1} \wedge \theta_{S_i}$
– output random one of the φ_i

– T = set of satisfying assignments for φ
– Claim: if |T| > 0, then

$$\Pr_{k \in \{0,1,2,\ldots,n-1\}}[2^k \leq |T| \leq 2^{k+1}] \geq 1/n$$

Unique solutions

Claim: if $2^k \le |T| \le 2^{k+1}$, then the probability φ_{k+2} has exactly one satisfying assignment is $\ge 1/8$

– fix $t, t' \in T$

S_i contains even # of positions i where $t_i \neq t_i'$

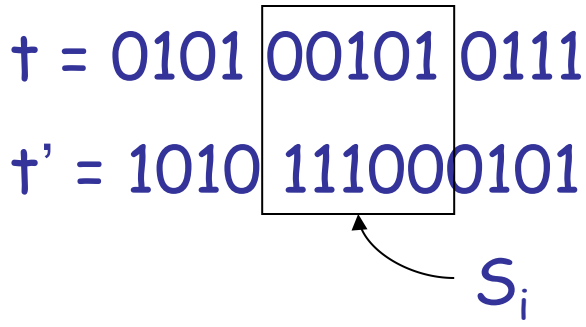

– Pr[t "agrees with" t' on S_i] = ½

– Pr[t agrees with t' on $S_1, S_2, \ldots, S_{k+2}$] = $(½)^{k+2}$

Unique solutions

– Pr[t agrees with *some* t' on $S_1, \ldots, S_{k+2}$]

$\leq (|T|-1)(\frac{1}{2})^{k+2} < \frac{1}{2}$

– Pr[t satisfies $S_1, S_2, \ldots, S_{k+2}$] = $(\frac{1}{2})^{k+2}$

– Pr[t *unique* satisfying assignment of φ_{k+2}]

$> (\frac{1}{2})^{k+3}$

– sum over at least 2^k different $t \in T$ (disjoint events); claim follows.

Randomized complexity classes

- model: probabilistic Turing Machine
  - deterministic TM with additional read-only tape containing "coin flips"
- **BPP** (Bounded-error Probabilistic Poly-time)
  - $L \in$ **BPP** if there is a p.p.t. TM M:

    $x \in L \Rightarrow Pr_y[M(x,y) \text{ accepts}] \geq 2/3$

    $x \notin L \Rightarrow Pr_y[M(x,y) \text{ rejects}] \geq 2/3$
  - "p.p.t" = probabilistic polynomial time

Randomized complexity classes

- **RP** (Random Polynomial-time)
  - $L \in$ **RP** if there is a p.p.t. TM M:

    $x \in L \Rightarrow Pr_y[M(x,y) \text{ accepts}] \geq ½$

    $x \notin L \Rightarrow Pr_y[M(x,y) \text{ rejects}] = 1$
- **coRP** (complement of Random Polynomial-time)
  - $L \in$ **coRP** if there is a p.p.t. TM M:

    $x \in L \Rightarrow Pr_y[M(x,y) \text{ accepts}] = 1$

    $x \notin L \Rightarrow Pr_y[M(x,y) \text{ rejects}] \geq ½$

Randomized complexity classes

One more important class:

- **ZPP** (Zero-error Probabilistic Poly-time)
  - **ZPP = RP ∩ coRP**
  - Pr_y[M(x,y) outputs "fail"] ≤ ½
  - otherwise outputs correct answer

Randomized complexity classes

These classes may capture "efficiently computable" better than **P**.

- "1/2" in **ZPP**, **RP**, **coRP** definition unimportant
  - can replace by 1/poly(n)
- "2/3" in **BPP** definition unimportant
  - can replace by ½ + 1/poly(n)
- Why? error reduction
  - we will see simple error reduction by repetition
  - more sophisticated error reduction later

Error reduction for RP

- given L and p.p.t TM M:

$$x \in L \Rightarrow Pr_y[M(x,y) \text{ accepts}] \geq \varepsilon$$

$$x \notin L \Rightarrow Pr_y[M(x,y) \text{ rejects}] = 1$$

- new p.p.t TM M':
  - simulate M k/ε times, each time with independent coin flips
  - accept if **any** simulation accepts
  - otherwise reject

Error reduction

$$x \in L \Rightarrow Pr_y[M(x,y) \text{ accepts}] \geq \varepsilon$$
$$x \notin L \Rightarrow Pr_y[M(x,y) \text{ rejects}] = 1$$

- if $x \in L$:
  - probability a given simulation "bad" $\leq (1 - \varepsilon)$
  - probability all simulations "bad" $\leq (1-\varepsilon)^{(k/\varepsilon)} \leq e^{-k}$

$$Pr_{y'}[M'(x, y') \text{ accepts}] \geq 1 - e^{-k}$$

- if $x \notin L$:

$$Pr_{y'}[M'(x,y') \text{ rejects}] = 1$$

Error reduction for BPP

- given L, and p.p.t. TM M:

$$x \in L \Rightarrow Pr_y[M(x,y) \text{ accepts}] \geq \tfrac{1}{2} + \varepsilon$$

$$x \notin L \Rightarrow Pr_y[M(x,y) \text{ rejects}] \geq \tfrac{1}{2} + \varepsilon$$

- new p.p.t. TM M':
  - simulate M k/ε^2 times, each time with independent coin flips
  - accept if **majority** of simulations accept
  - otherwise reject

Error reduction for BPP

- X_i random variable indicating "correct" outcome in i-th simulation (out of $m = k/\varepsilon^2$)
  - $\Pr[X_i = 1] \geq \frac{1}{2} + \varepsilon$
  - $\Pr[X_i = 0] \leq \frac{1}{2} - \varepsilon$
- $E[X_i] \geq \frac{1}{2}+\varepsilon$
- $X = \Sigma_i X_i$
- $\mu = E[X] \geq (\frac{1}{2} + \varepsilon)m$

- Chernoff: $\Pr[X \leq m/2] \leq 2^{-\Omega(\varepsilon^2 \mu)}$

Error reduction for BPP

$$x \in L \Rightarrow \Pr_y[M(x,y) \text{ accepts}] \geq \tfrac{1}{2} + \varepsilon$$
$$x \notin L \Rightarrow \Pr_y[M(x,y) \text{ rejects}] \geq \tfrac{1}{2} + \varepsilon$$

– if $x \in L$

$$\Pr_{y'}[M'(x, y') \text{ accepts}] \geq 1 - (\tfrac{1}{2})^{\Omega(k)}$$

– if $x \notin L$

$$\Pr_{y'}[M'(x,y') \text{ rejects}] \geq 1 - (\tfrac{1}{2})^{\Omega(k)}$$

Randomized complexity classes

- We have shown:
  - polynomial identity testing is in **coRP**
  - a poly-time algorithm for detecting unique solutions to SAT implies

NP = RP

Relationship to other classes

- ZPP, RP, coRP, BPP, contain **P**
  - they can simply ignore the tape with coin flips
- all are in **PSPACE**
  - can exhaustively try all strings y
  - count accepts/rejects; compute probability
- **RP** $\subset$ **NP** (and **coRP** $\subset$ **coNP**)
  - multitude of accepting computations
  - **NP** requires only one

Relationship to other classes