CMSC 28100

Introduction to Complexity Theory

Spring 2024
Instructor: William Hoza

Robustness of P

- The complexity class P is highly robust against modifications to the computational model
  - Multi-tape Turing machines, two-dimensional Turing machines, etc.
- After studying many examples, one begins to suspect that every realistic model of computation can be simulated by Turing machines with only polynomial slowdown

Extended Church-Turing Thesis

- Let L be a language

Extended Church-Turing Thesis:

It is physically possible to build a device that decides L in polynomial time if and only if $L \in \mathrm{P}$.

Extended Church-Turing Thesis

Extended Church-Turing Thesis:

It is physically possible to build a device that decides L in polynomial time if and only if $L \in \mathrm{P}$.

- If it were true, the thesis would justify using P as our model of tractability
- However, it seems increasingly likely that the thesis is false!
- Two key challenges: Randomized Computation and Quantum Computation

Randomized computation

- Sometimes, in our efforts to solve problems and figure things out, we want to make random choices
  - Random sampling for opinion polls
  - Randomized controlled trials in science/medicine
- What happens if we incorporate this ability into our computational model?

Randomized computation

- To properly study the role of randomness in computing, we ought to define and study a randomized variant of the Turing machine model
- However, let's temporarily set Turing machines aside – we will circle back to them later
- To build intuition, let's study the role of randomness in a different situation first

Communication Complexity

Communication complexity

- Goal: Compute $f(x, y)$ using as little communication as possible
- In each round, one party sends a single bit while the other party listens
- At the end, both parties announce $f(x, y)$

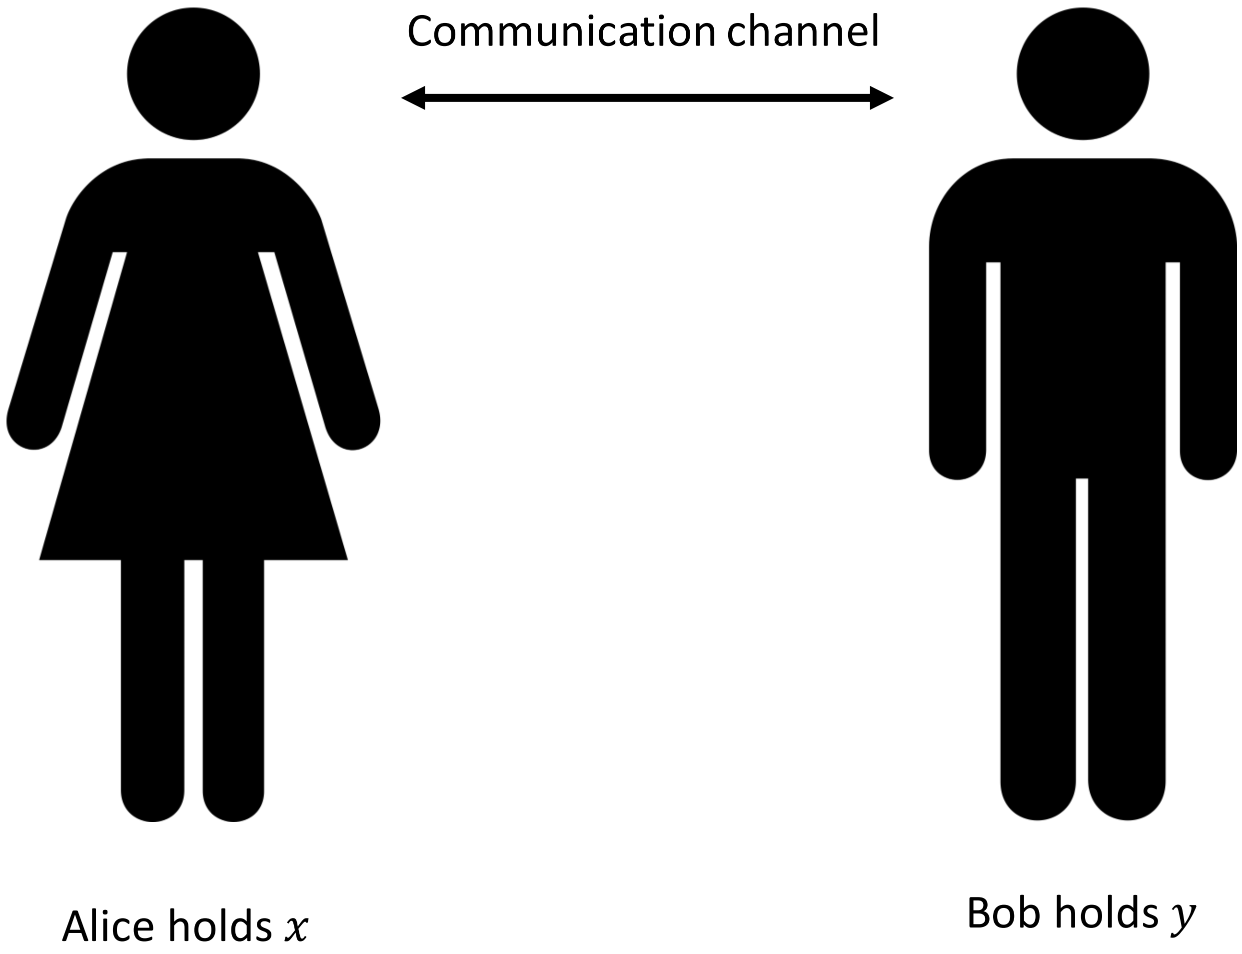

The equality function

- We will focus on the case $f = \mathrm{EQ}_n$
- $\mathrm{EQ}_n \colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$
- Definition:

$$\mathrm{EQ}_n(x,y) = \begin{cases} 1 & \text{if } x = y \\ 0 & \text{if } x \neq y \end{cases}$$

- "Does your copy of the database match my copy?"

Protocols for equality

Protocol A:

1) Alice sends x
2) Bob sends $\mathrm{EQ}_n(x, y)$

$n + 1$ bits of communication

Protocol B:

1) For $i = 1$ to n:
   a) Alice sends x_i
   b) Bob sends a bit indicating whether $x_i = y_i$

$2n$ bits of communication
(in the worst case)

Communication complexity of equality

- Is there a better protocol?

Theorem: Every deterministic communication protocol for EQ_n uses at least $n + 1$ bits of communication in the worst case

- Before we can prove it, we must clarify how we model communication protocols mathematically

Communication protocol model

- Idea: We model a communication protocol as a binary tree
- We start at the root node
- Someone transmits a zero ⇔ We move to the left child
- Someone transmits a one ⇔ We move to the right child
- (Alice and Bob both know where we are in the tree)

Example protocol tree

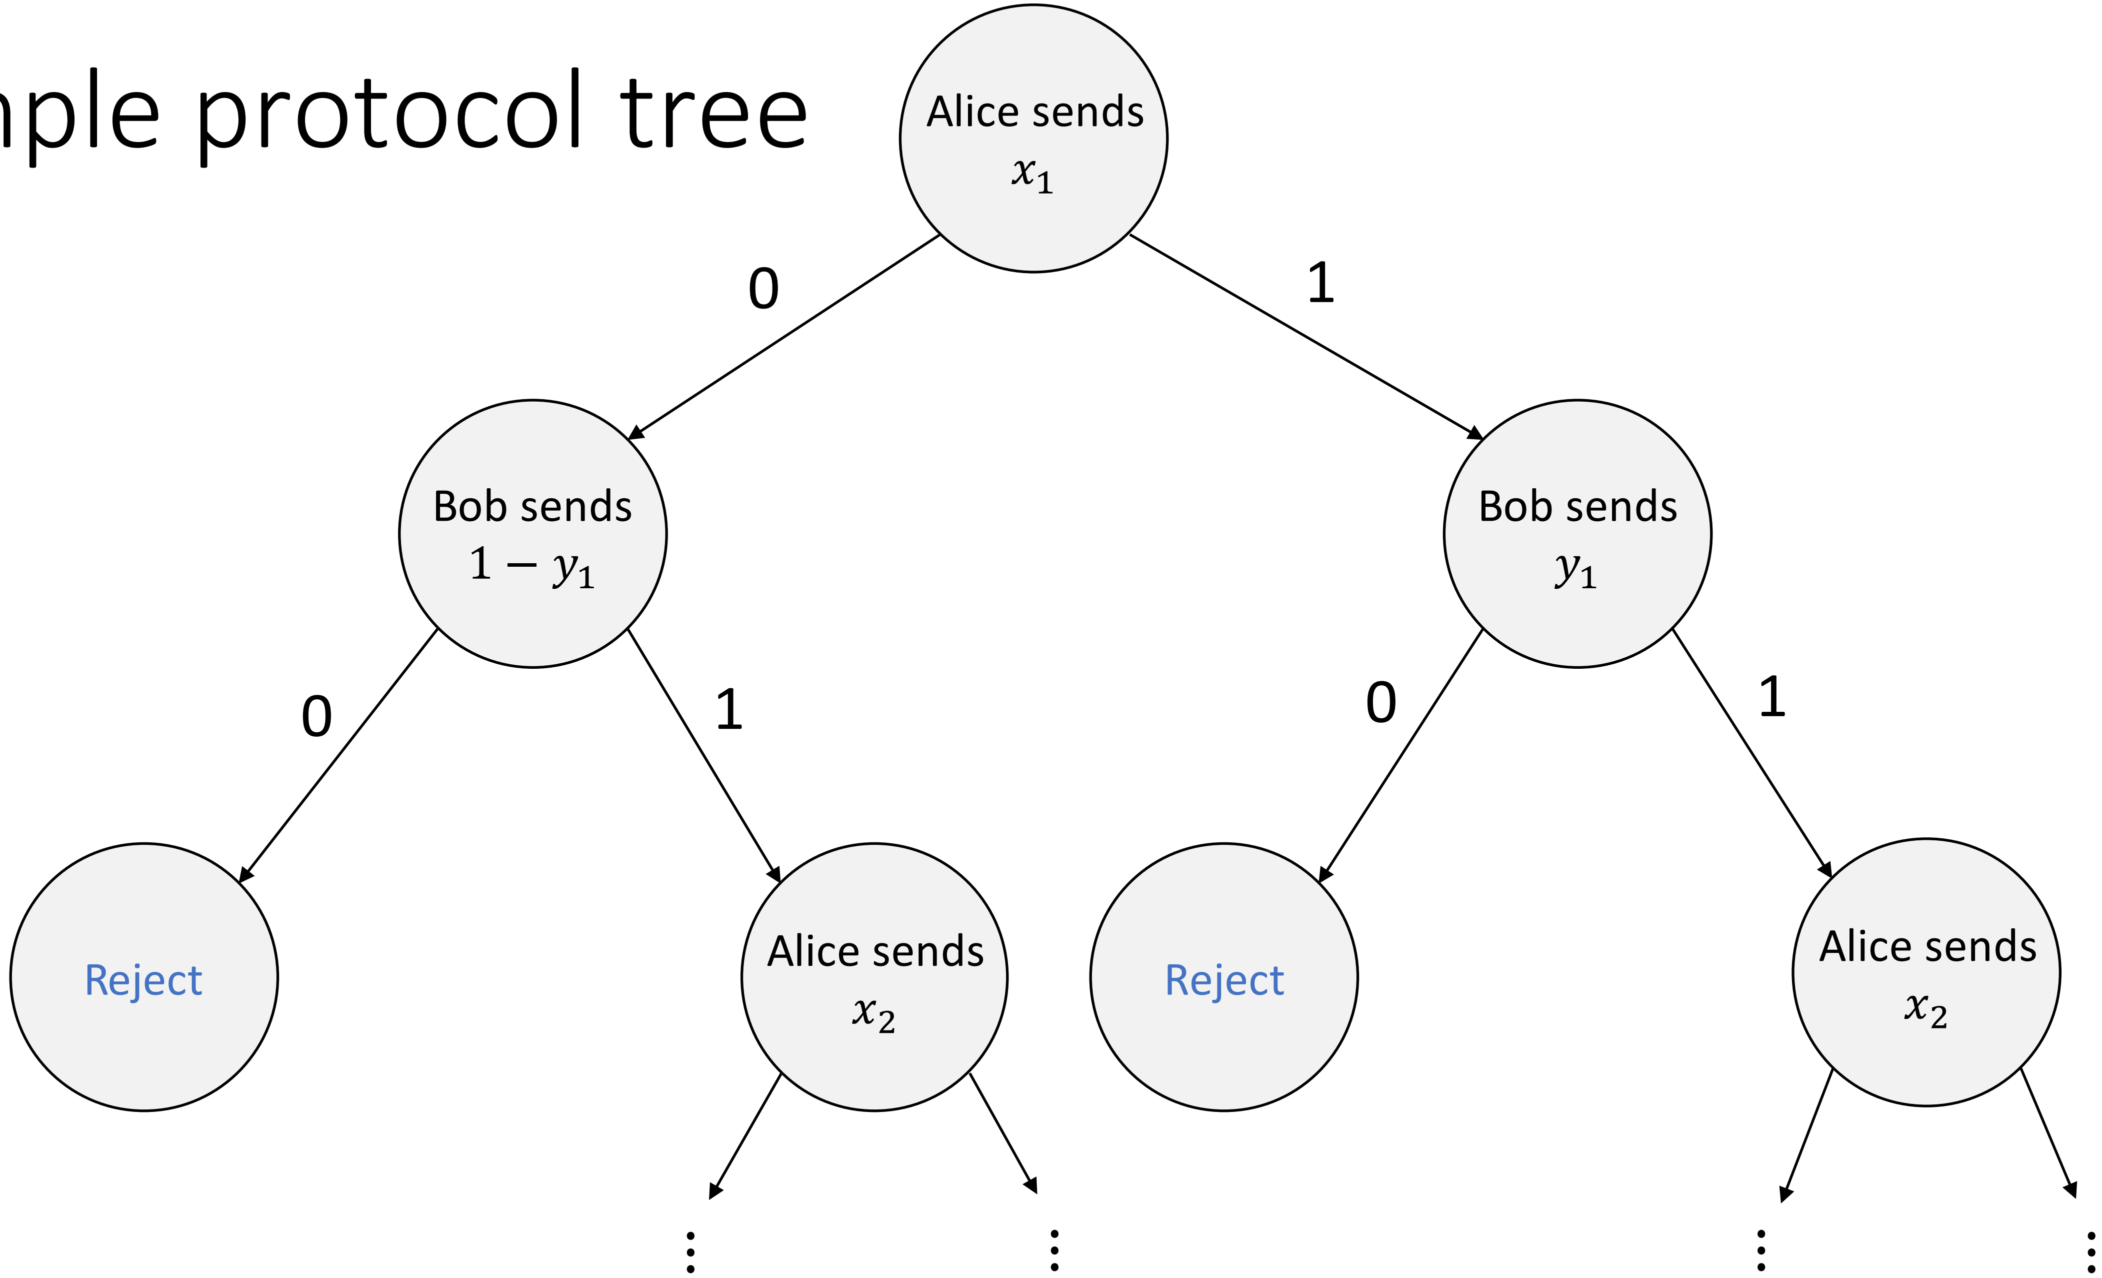

Rigorously defining communication protocols

- A deterministic communication protocol with n-bit inputs is a rooted binary tree π with the following features
  - The vertex set V is partitioned into three parts: $V = V_A \cup V_B \cup V_L$
  - Each vertex $v \in V_A \cup V_B$ has two children (ℓ and r) and is labeled with a function $\delta_v : \{0,1\}^n \to \{\ell, r\}$
  - Each vertex $v \in V_L$ has zero children and is labeled “accept” or “reject”

Rigorously defining communication protocols

- For $x, y \in \{0,1\}^n$, we define a sequence of vertices $v_0, v_1, \dots, v_T$
  - $v_0 =$ the root vertex
  - If $v_i \in V_A$ (Alice speaks next), then $v_{i+1} = \delta_{v_i}(x)$
  - If $v_i \in V_B$ (Bob speaks next), then $v_{i+1} = \delta_{v_i}(y)$
  - If $v_i \in V_L$ (the conversation is over), then $T = i$
- We define $\text{leaf}(x, y) = v_T$
- We define $\pi(x, y) = \begin{cases} 1 \text{ if leaf}(x, y) \text{ is labeled "accept"} \\ 0 \text{ if leaf}(x, y) \text{ is labeled "reject"} \end{cases}$

Communication complexity

- We say that π computes f if for every $x, y \in \{0,1\}^n$, we have

$$\pi(x,y) = f(x,y)$$

- The cost of the communication protocol π is the depth of the tree, i.e., the length of the longest path f

- (Cost = number of rounds =

In this model, what happens if Alice and Bob speak at the same time?

A: Trick question. In this model, they never speak simultaneously

B: Only one of the messages is successfully transmitted

C: Both of the messages are successfully transmitted

D: Neither message is successfully transmitted

Respond at PollEv.com/whoza or text "whoza" to 22333

Rectangle lemma

$$\begin{array}{ccc} (x,y) & \text{--------} & (x',y) \\ \vdots & & \vdots \\ (x,y') & \text{--------} & (x',y') \end{array}$$

- Let π be any communication protocol with n-bit inputs
- Let $x, x', y, y' \in \{0,1\}^n$ and let v be any leaf

Rectangle Lemma: If $\text{leaf}(x,y) = \text{leaf}(x',y') = v$,

then $\text{leaf}(x,y') = \text{leaf}(x',y) = v$

- **Proof (sketch):** Let $v_0, v_1, \dots, v_T$ be the vertices from the root to v
- If $v_i \in V_A$, we must have $\delta_{v_i}(x) = \delta_{v_i}(x') = v_{i+1}$. Similarly if $v_i \in V_B$

Communication complexity of equality

Theorem: Every deterministic communication protocol that computes EQ_n has cost at least $n+1$

- **Proof:** Let π be any communication protocol that computes EQ_n
- Assume WLOG that every leaf is at the same depth m
- Our job is to prove that $m \geq n+1$

$$\begin{tikzcd}
(x,x) \arrow[r, dash, dashed] \arrow[d, dash, dashed] & (x,y) \arrow[d, dash, dashed] \\
(y,x) \arrow[r, dash, dashed] & (y,y)
\end{tikzcd}$$

Communication complexity of $\mathbf{EQ}_n$

- If $x \neq y$, then $\mathrm{leaf}(x, x) \neq \mathrm{leaf}(x, y)$
- By the rectangle lemma, it follows that $\mathrm{leaf}(x, x) \neq \mathrm{leaf}(y, y)$
- Therefore, there are at least 2^n distinct leaves labeled “accept”
- There is also at least one leaf labeled “reject”
- Therefore, there are more than 2^n leaves
- Therefore, $2^m > 2^n$, hence $m \geq n + 1$

Communication complexity of $\mathbf{EQ}_n$

- We just proved that computing EQ_n requires $n + 1$ bits of communication
- However, there is a loophole!
- Our impossibility proof only applies to deterministic protocols!

Randomized communication complexity

- In a randomized communication protocol, Alice and Bob are permitted to make decisions based on coin tosses

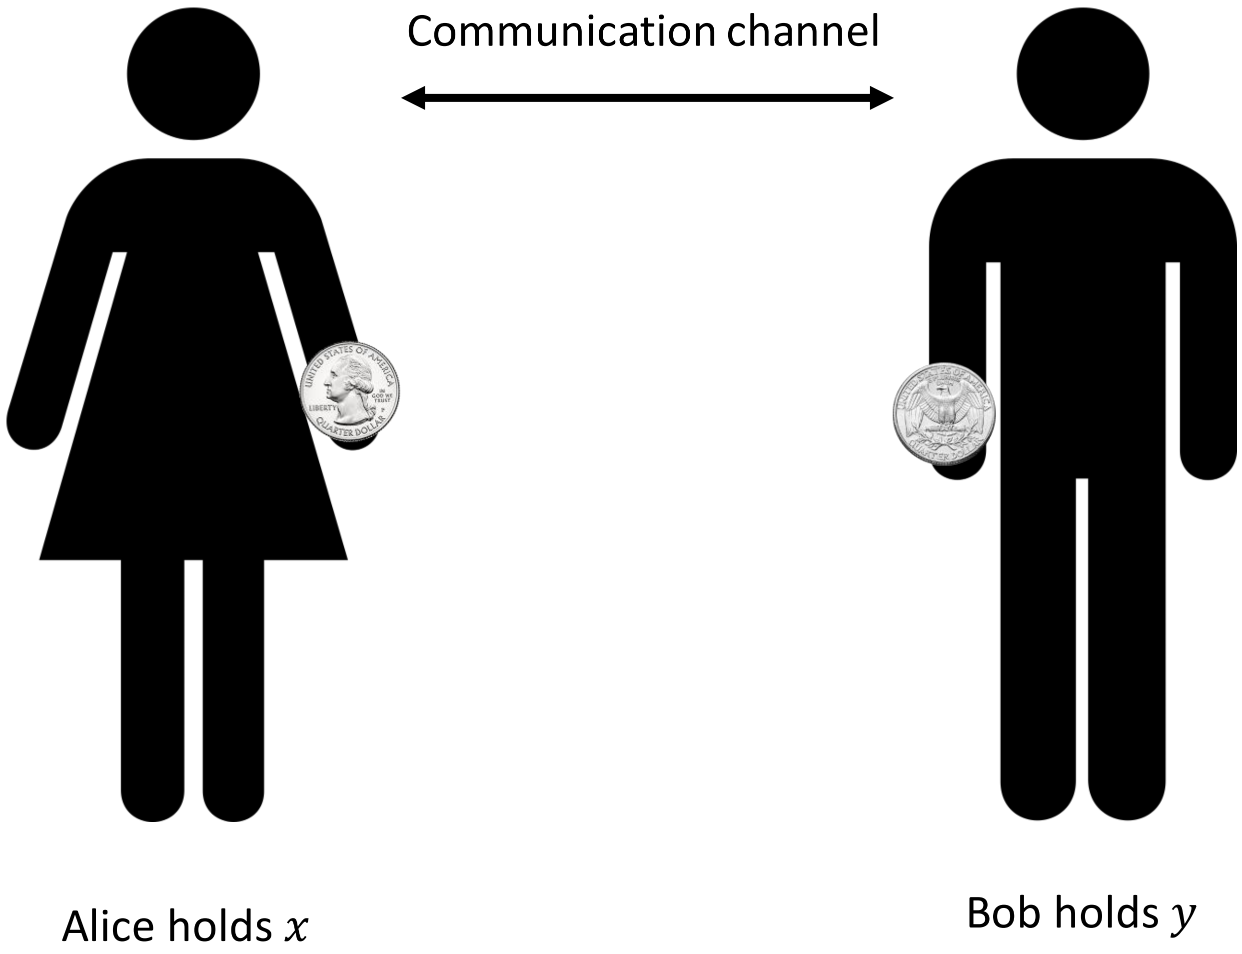

Randomized communication protocols

- Mathematically, we model a randomized communication protocol with n-bit inputs as a deterministic communication protocol with $(n + r)$-bit inputs for some $r \geq 0$
- Alice holds xu, where $x \in \{0, 1\}^n$ and $u \in \{0, 1\}^r$
- Bob holds yw, where $y \in \{0, 1\}^n$ and $w \in \{0, 1\}^r$
- Interpretation: x, y are the "actual inputs," and u, w are the coin tosses