Communication Complexity

Lydia Zakynthinou

Corelab
NTUA

June 26, 2014

Outline

Outline

What is it about?

How "much" communication do we need to perform a computational task for which information is distributed among different entities?

Why communication complexity?

- Simple enough so we can prove lower bounds, general enough so we can obtain important applications of these lower bounds.
- Some applications:
  - Lower bounds for Data Structures
  - Lower bounds for parallel and VLSI computations
  - Auctions (cost for prefenences)
  - Polyhedral Theory
  - Time-space tradeoff for Turing Machines

Why communication complexity?

- Simple enough so we can prove lower bounds, general enough so we can obtain important applications of these lower bounds.
- Some applications:
  - Lower bounds for Data Structures
  - Lower bounds for parallel and VLSI computations
  - Auctions (cost for prefenences)
  - Polyhedral Theory
  - Time-space tradeoff for Turing Machines

Yao's Model ['79]

- Two parties (Alice and Bob) with unlimited computational power
- Each holds an n-bit input x, y
- They want to compute $f(x, y)$ where $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ is known to both.
- They have agreed upon a *protocol* of communication P.

$COST(P)$: the number of bits communicated by the players for the worst-case choice of x, y
Communication Complexity of f , $C(f)$: the minimum $COST(P)$ over all *valid* protocols P

Yao's Model ['79]

- Two parties (Alice and Bob) with unlimited computational power
- Each holds an n-bit input x, y
- They want to compute $f(x, y)$ where $f : \{0, 1\}^n \times \{0, 1\}^n \to \{0, 1\}$ is known to both.
- They have agreed upon a *protocol* of communication P.

$COST(P)$: the number of bits communicated by the players for the worst-case choice of x, y
Communication Complexity of f , $C(f)$: the minimum $COST(P)$ over all *valid* protocols P

Yao's Model ['79]

- Two parties (Alice and Bob) with unlimited computational power
- Each holds an n-bit input x, y
- They want to compute $f(x,y)$ where $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ is known to both.
- They have agreed upon a *protocol* of communication P.

$COST(P)$: the number of bits communicated by the players for the worst-case choice of x, y
Communication Complexity of f , $C(f)$: the minimum $COST(P)$ over all *valid* protocols P

Yao's Model ['79]

- Two parties (Alice and Bob) with unlimited computational power
- Each holds an n-bit input x, y
- They want to compute $f(x, y)$ where $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ is known to both.
- They have agreed upon a *protocol* of communication P.

$COST(P)$: the number of bits communicated by the players for the worst-case choice of x, y
Communication Complexity of f , $C(f)$: the minimum $COST(P)$ over all *valid* protocols P

Yao's Model ['79]

- Two parties (Alice and Bob) with unlimited computational power
- Each holds an n-bit input x, y
- They want to compute $f(x, y)$ where $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ is known to both.
- They have agreed upon a *protocol* of communication P.

$COST(P)$: the number of bits communicated by the players for the worst-case choice of x, y
Communication Complexity of f , $C(f)$: the minimum $COST(P)$ over all *valid* protocols P

Example: Parity

$f(x, y)$:the parity of all bits in x, y
It holds that $C(f) = 2$!

- $C(f) \geq 2$ because f depends on both x and y
- $C(f) \leq 2$ because there is a protocol P with $COST(P) = 2$ (Alice sends the parity of x and Bob XORs it with the parity of y)

For every function $C(f) \leq n + 1$

Example: Parity

$f(x, y)$:the parity of all bits in x, y
It holds that $C(f) = 2$!

- $C(f) \geq 2$ because f depends on both x and y
- $C(f) \leq 2$ because there is a protocol P with $COST(P) = 2$ (Alice sends the parity of x and Bob XORs it with the parity of y)

For every function $C(f) \leq n + 1$

Outline

The Fooling Set Method

Observation: If the communication pattern is the same for (x, x) and (x', x') then the output of the protocol is the same for all $(x, x), (x, x'), (x', x), (x', x')$.

- Say there is a protocol P with $COST(P) \leq n - 1$.
- Then there are at most 2^{n-1} communication patterns.
- But there are 2^n input pairs of the form (x, x).
- There exist two distinct pairs with the same communication pattern.

Fooling Set S

- $\forall (x, y) \in S : f(x, y) = b$
- $\forall (x, y'), (x', y) \in S : f(x, y') \neq b or f(x', y) \neq b$

Theorem

If f has a size-M fooling set then $C(f) \leq \log M$

The Fooling Set Method

Observation: If the communication pattern is the same for (x, x) and (x', x') then the output of the protocol is the same for all $(x, x), (x, x'), (x', x), (x', x')$.

- Say there is a protocol P with $COST(P) \leq n - 1$.
- Then there are at most 2^{n-1} communication patterns.
- But there are 2^n input pairs of the form (x, x).
- There exist two distinct pairs with the same communication pattern.

Fooling Set S

- $\forall (x, y) \in S : f(x, y) = b$
- $\forall (x, y'), (x', y) \in S : f(x, y') \neq b or f(x', y) \neq b$

Theorem

If f has a size-M fooling set then $C(f) \leq \log M$

The Fooling Set Method

Observation: If the communication pattern is the same for (x, x) and (x', x') then the output of the protocol is the same for all $(x, x), (x, x'), (x', x), (x', x')$.

- Say there is a protocol P with $COST(P) \leq n - 1$.
- Then there are at most 2^{n-1} communication patterns.
- But there are 2^n input pairs of the form (x, x).
- There exist two distinct pairs with the same communication pattern.

Fooling Set S

- $\forall (x, y) \in S : f(x, y) = b$
- $\forall (x, y'), (x', y) \in S : f(x, y') \neq b or f(x', y) \neq b$

Theorem

If f has a size-M fooling set then $C(f) \leq \log M$

The Fooling Set Method

Observation: If the communication pattern is the same for (x, x) and (x', x') then the output of the protocol is the same for all $(x, x), (x, x'), (x', x), (x', x')$.

- Say there is a protocol P with $COST(P) \leq n - 1$.
- Then there are at most 2^{n-1} communication patterns.
- But there are 2^n input pairs of the form (x, x).
- There exist two distinct pairs with the same communication pattern.

Fooling Set S

- $\forall (x, y) \in S : f(x, y) = b$
- $\forall (x, y'), (x', y) \in S : f(x, y') \neq borf(x', y) \neq b$

Theorem

If f has a size-M fooling set then $C(f) \leq \log M$

The Tiling Method

- $M(f)$ the matrix of f
- Is partitioned into rectangles depending on the protocol bits (that turn out to be monochromatic - why?)
- $\chi(f)$ is the minimum number of rectangles in any monochromatic tiling

Theorem (AUY'83)

$\log \chi(f) \leq C(f) \leq 16(\log \chi(f))^2$

The Rank Method

The rank of a matrix, $rank(M)$, can be expressed as the minimum l s.t.:

$$M = \sum_{i=1}^{l} B_i \text{ , where } rank(B_i) = 1$$

Theorem

For every function f, $\chi(f) \geq rank(M(f))$

The Discrepancy Method

Discrepancy of $M(f)$: $Disc(f) = \max \frac{1}{2^(2n)} |\sum_{x \in A, y \in B} M_{x,y}|$ over all rectangles $A \times B$

Theorem

$\chi(f) \geq \frac{1}{Disc(f)}$

Theorem (Eigenvalue Bound)

$Disc(A \times B) \leq \frac{1}{2^{2n}} \lambda_{max}(M) \sqrt{|A||B|}$

Comparison

- The tiling argument is the strongest lower bound
- $\log \chi(f)$ fully characterizes $C(f)$ within a constant factor
- The rank and fooling set methods are incomparable

Conjecture (log rank conjecture): There is a constant $c > 1$ such that $C(f) = O(\log(rank(M(f)))^c)$ for all f and input sizes n.

Comparison

- The tiling argument is the strongest lower bound
- $\log \chi(f)$ fully characterizes $C(f)$ within a constant factor
- The rank and fooling set methods are incomparable

Conjecture (log rank conjecture): There is a constant $c > 1$ such that $C(f) = O(\log(rank(M(f)))^c)$ for all f and input sizes n.

Outline

The Model

Most interesting model: "Number on the forehead"

Example: $C_3(f) = 3$ where $f(x_1, x_2, x_3) = \oplus maj(x_{1i}, x_{2i}, x_{3i})$

Best known lower bound for the communication complexity of an explicit function (GIP) is of the form $n/2^{-\Omega(k)}$

The Model

Most interesting model: "Number on the forehead"
Example: $C_3(f) = 3$ where $f(x_1, x_2, x_3) = \oplus maj(x_{1i}, x_{2i}, x_{3i})$
Best known lower bound for the communication complexity of an explicit function (GIP) is of the form $n/2^{-\Omega(k)}$

The Model

Most interesting model: "Number on the forehead"
Example: $C_3(f) = 3$ where $f(x_1, x_2, x_3) = \oplus maj(x_{1i}, x_{2i}, x_{3i})$
Best known lower bound for the communication complexity of an explicit function (GIP) is of the form $n/2^{-\Omega(k)}$

Outline

- In a non-deterministic protocol P, the players are both provided an additional input z of some length m that may depend on x, y. We require that $f(x, y) = 1$ iff there exists z that makes the players output 1.
- $COST(P) = |z|+$number of bits communicated
- Inequality and Intersection are in NP.

Outline

- All players have access to a random string r and we define $R(f)$ to be the expected number of bits communicated by the protocol.
- For example, Equality has a randomized protocol with $O(\log n)$ complexity.

Outline

$$P^{CC} = (NP^{CC} \bigcap coNP^{CC}) \subset BPP^{CC}$$

- P^{CC}: deterministic *polylog* time
- RP^{CC}: *polylog* time - error at "no" instances only at most $1/4$ (1-sided error)
- BPP^{CC}: *polylog* time - correct with probability $3/4$ (2-sided error)
- NP^{CC}: non-deterministic *polylog* time

Outline

- Computational Complexity: A Modern Approach [Arora, Barak] : Chapter 13
- Communication Complexity [Kushilevitz, Nisan]
- An Invitation to Mathematics: from Competitions to Research : Chapter 8 by A.Razborov

Thank you!