Foundations of Probabilistic Proofs

A course by **Alessandro Chiesa**

Lecture 23

Limitations of IPs

IPs with Bounded Communication

Let $\mathrm{IP}[pc=1]$ = "languages decidable via IPs where the prover sends 1 bit".

Q: Is $\mathrm{IP}[pc=1]$ trivial (contained in BPP)?

Unlikely, because $GNI \in \mathrm{IP}[pc=1]$ and GNI is not known to be in BPP.

→ Even IPs with small communication complexity can decide non-trivial languages.

Q: Can we hope for $SAT \in \mathrm{IP}[pc = o(n)]$? (pc is sublinear in the number of variables n)

Note that $SAT \in NP \subseteq \mathrm{IP}$ so we are asking if there is an IP for SAT that provides an improvement in communication over the trivial IP (send the candidate assignment).

Today we study IPs with BOUNDED COMMUNICATION:

- $\mathrm{IP}[pc, vc, vr]$ = "languages decidable via IPs where { prover sends pc bits; verifier sends vc bits; verifier uses vr random bits }"
- $\mathrm{AM}[pc, vr]$ = "same but via public-coin IPs (where $vc = vr$)"

We see ALGORITHMS FOR IPs, exposing a connection between the communication complexity of an IP and the time complexity of the language it decides.

Bonus: algorithms for IPs often serve as tools to prove limitations of PCPs and IOPs.

Warmup: Short Proof → Easy Language [1/2]

Define $NP[pc]$ = "NP languages where the proof string (i.e. witness) is pc bits".

lemma: $NP[pc] \subseteq DTIME(2^{O(pc)} \cdot poly(n))$

proof: Try every possible proof string.

$A(x)$:= 1. For every NP proof $\tilde{\pi} \in \{0,1\}^{pc}$: if $V_{NP}(x,\tilde{\pi})=1$ then output 1.
2. Output 0.

■

Next we consider proof strings checked with randomness.

Define $MA[\varepsilon_c, \varepsilon_s, pc, vr]$ = "languages decidable via pc-bit proof strings with completeness error ε_c and soundness error ε_s, using vr bits of randomness"

lemma: ① $MA[\varepsilon_c, \varepsilon_s, pc, vr] \subseteq DTIME(2^{O(pc+vr)} \cdot poly(n))$

② $MA[\varepsilon_c, \varepsilon_s, pc, vr] \subseteq BPTIME\left(2^{O(pc)} \cdot poly\left(\frac{1}{1-\varepsilon_c-\varepsilon_s}, n\right)\right)$

Proof of ①: try all possible proof strings and randomness strings.

Proof of ②: we need a new idea because we CANNOT afford trying all randomness strings.

Warmup: Short Proof → Easy Language [2/2]

$\mathsf{MA}[\varepsilon_c, \varepsilon_s, pc, vr] \subseteq \mathsf{BPTIME}\left(2^{O(pc)} \cdot \mathrm{poly}\left(\frac{1}{1-\varepsilon_c-\varepsilon_s}, n\right)\right)$

Idea: APPROXIMATE the acceptance probability for every possible MA proof.

$A(x) :=$
1. Sample $\rho_1, \dots, \rho_t \in \{0,1\}^{vr}$.
2. For every MA proof $\tilde{\pi} \in \{0,1\}^{pc}$:
   - Compute $N(\tilde{\pi}) := \left|\{ i \in [t] : V_{MA}(x, \tilde{\pi}; \rho_i) = 1 \}\right|$.
   - If $\frac{N(\tilde{\pi})}{t} > \frac{1-\varepsilon_c+\varepsilon_s}{2}$ then output 1.
3. Output 0.

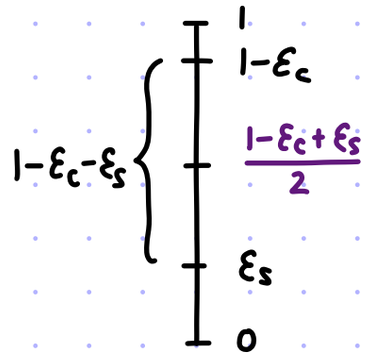

Define $Z(\tilde{\pi}, \rho) :=$ "indicator that $V_{MA}(x, \tilde{\pi}; \rho) = 1$".

$Z(\tilde{\pi}, \rho_1), \dots, Z(\tilde{\pi}, \rho_t)$ are i.i.d. samples from the Bernoulli distribution with bias $\delta(\tilde{\pi}) := \Pr_{\rho}[V_{MA}(x, \tilde{\pi}; \rho) = 1]$.

By an additive Chernoff bound, $\Pr_{\rho_1, \dots, \rho_t}\left[\left|\frac{1}{t}\sum_{i=1}^{t} Z(\tilde{\pi}, \rho_i) - \delta(\tilde{\pi})\right| > \alpha\right] \leq \exp(-t \cdot \alpha^2)$.

$\left\{\begin{array}{l} x \in L \to \exists\, \pi \;\; \delta(\pi) \geq 1-\varepsilon_c \\ x \notin L \to \forall\, \tilde{\pi} \;\; \delta(\tilde{\pi}) \leq \varepsilon_s \end{array}\right\} \to$ We need $\alpha < \frac{1}{2} \cdot ((1-\varepsilon_c)-\varepsilon_s)$ to distinguish between these.

For such α, $\Pr[A(x) \text{ errs}] \leq \Pr\left[\exists\, \tilde{\pi} \in \{0,1\}^{pc} : \left|\frac{N(\tilde{\pi})}{t} - \delta(\tilde{\pi})\right| > \alpha\right] \leq \sum_{\tilde{\pi} \in \{0,1\}^{pc}} \Pr\left[\left|\frac{N(\tilde{\pi})}{t} - \delta(\tilde{\pi})\right| > \alpha\right] \leq 2^{pc} \cdot \exp(-t \cdot \alpha^2)$.

Setting $t := O\left(\frac{pc}{\alpha^2}\right) = O\left(\frac{pc}{(1-\varepsilon_c-\varepsilon_s)^2}\right)$ ensures that A has constant two-sided error. ■

The Case of Interactive Proofs

A similar statement holds for any IP.

theorem: ① $IP[\varepsilon_c, \varepsilon_s, pc, vc, vr] \subseteq DTIME\left(2^{O(pc+vc+vr)} \cdot poly(n)\right)$

② $IP[\varepsilon_c, \varepsilon_s, pc, vc, vr] \subseteq BPTIME\left(2^{O(pc+vc)} \cdot poly\left(\frac{1}{1-\varepsilon_c-\varepsilon_s}, n\right)\right)$

① If we bound (two-way) communication and randomness by B then we can decide the language in deterministic $\approx 2^{O(B)}$ time.

② If we bound (two-way) communication by B (and not bound randomness) then we can decide the language in probabilistic $\approx 2^{O(B)}$ time.

What about ONE-WAY communication (pc only)? We discuss this later.

Example for 3SAT: it is unlikely that $3SAT \in IP[pc = o(n), vc = o(n)]$.

It would imply that $3SAT \in BPTIME(2^{o(n)})$, contradicting RETH.

Randomized Exponential-Time Hypothesis: $\exists c>0$ $3SAT \notin BPTIME(2^{c \cdot n})$

SUCCINCTNESS REQUIRES COMPUTATIONAL SOUNDNESS:

We have seen how using cryptography (e.g. collision-resistant hash functions) we can achieve succinct interactive arguments for NP. For example, an interactive protocol for 3SAT where $pc + vc = poly(\lambda, \log n)$ whose soundness holds against $poly(\lambda)$-size malicious provers.

The above theorem implies that computational soundness is necessary for succinctness.

Game Tree for an IP

Fix an IP verifier V and instance x.

The **IP game tree** $T := T(V,x)$ of $V(x)$ is the tree of all possible augmented transcripts.

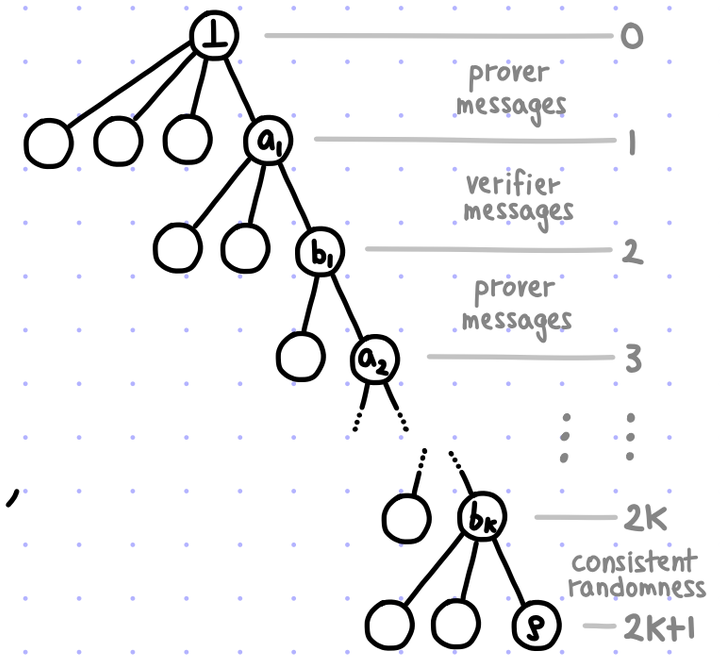

The tree root at level 0 denotes the empty transcript.

Every vertex in the tree represents:

- a transcript $(a_1,b_1,\ldots,a_i,b_i)$ where the prover is about to move,
- a transcript $(a_1,b_1,\ldots,a_i,b_i,a_{i+1})$ where the verifier is about to move,
- a full transcript $(a_1,b_1,\ldots,a_k,b_k)$,
- an augmented transcript $(a_1,b_1,\ldots,a_k,b_k,\rho)$.

The randomness strings consistent with $tr=(a_1,b_1,\ldots,a_i,b_i)$ are $R[x,tr] := \left\{ \rho\in\{0,1\}^{vr} : \begin{array}{l} b_1 = V(x,a_1;\rho) \\ b_2 = V(x,a_1,a_2;\rho) \\ \vdots \\ b_i = V(x,a_1,\ldots,a_i;\rho) \end{array} \right\}$.

Define $R[x,(a_1,b_1,\ldots,a_i,b_i,a_{i+1})] := R[x,(a_1,b_1,\ldots,a_i,b_i)]$.

The possible prover moves for $tr=(a_1,b_1,\ldots,a_i,b_i)$ are $a_{i+1}\in\{0,1\}^{pc_{i+1}}$.

The possible verifier moves for $tr=(a_1,b_1,\ldots,a_i,b_i,a_{i+1})$ are $\{ b_{i+1} : R[x,(a_1,b_1,\ldots,a_i,b_i,a_{i+1},b_{i+1})] \neq \emptyset \}$.

The possible augmentations for $tr=(a_1,b_1,\ldots,a_k,b_k)$ are $R[x,tr]$.

Approximating the Value Suffices

The value of the tree is $\mathrm{val}(T) := \mathrm{val}(\text{root})$.

The value of the vertex is recursively defined:

- $\mathrm{val}\big(\text{leaf vertex } tr = (a_1, b_1, \ldots, a_k, b_k, \rho)\big) := V(x, a_1, \ldots, a_k; \rho)$.
- $\mathrm{val}\big(\text{pre-leaf vertex } tr = (a_1, b_1, \ldots, a_k, b_k) \text{ at level } 2k\big)$

  $:= \mathbb{E}_{\rho \leftarrow R[x,tr]}\big[\mathrm{val}(\text{child vertex } (a_1, b_1, \ldots, a_k, b_k, \rho))\big]$
- $\mathrm{val}\big(\text{internal vertex } tr = (a_1, b_1, \ldots, a_i, b_i) \text{ at level } 2i\big)$

  $:= \max_{a_{i+1} \in \{0,1\}^{pc_{i+1}}} \mathrm{val}\big(\text{child vertex } (a_1, b_1, \ldots, a_i, b_i, a_{i+1}) \text{ at level } 2i+1\big)$
- $\mathrm{val}\big(\text{internal vertex } tr = (a_1, b_1, \ldots, a_i, b_i, a_{i+1}) \text{ at level } 2i+1\big)$

  $:= \mathbb{E}_{\rho \leftarrow R[x,tr]} \mathrm{val}\big(\text{child vertex } (a_1, b_1, \ldots, a_{i+1}, b_{i+1} := V(x, a_1, \ldots, a_i, a_{i+1}; \rho)) \text{ at level } 2i+2\big)$.

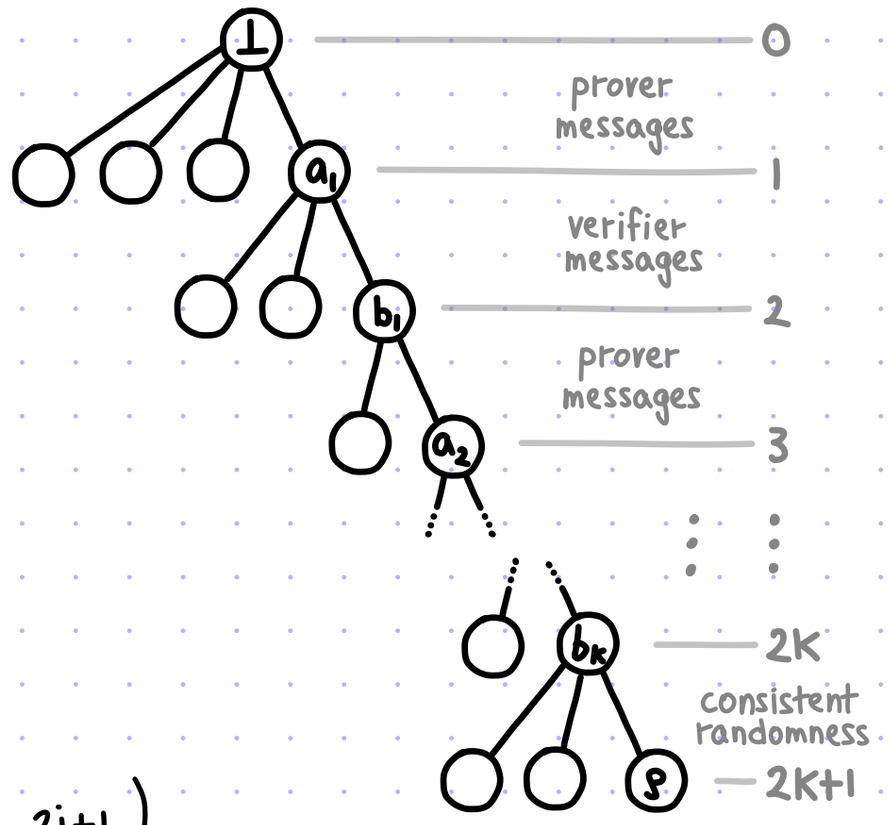

We showed that val(T) is computable in space poly(n) (and thus time $2^{\mathrm{poly}(n)}$).

TODAY: it suffices to approximate val(T), which we can do much faster than time $2^{\mathrm{poly}(n)}$.

↑ within a small-enough additive error with constant probability of error

Bounded Randomness & Two-Way Communication

theorem: $\mathrm{IP}[\varepsilon_c, \varepsilon_s, pc, vc, vr] \subseteq \mathrm{DTIME}\left(2^{O(pc+vc+vr)} \cdot \mathrm{poly}(n)\right)$

Let $c := pc + vc + vr$ be a bound on communication plus randomness.

The number of vertices in the tree $T(V,x)$ is $2^{O(c)}$ because:

- the number of possible transcripts is $\leq 2^{pc+vc}$,
- each transcript has $\leq 2^{vr}$ possible augmentations.

Hence we can compute $\mathrm{val}(T(V,x))$ exactly in time $2^{O(c)} \cdot \mathrm{poly}(n)$, by writing down the tree $T(V,x)$ and recursively computing its value.

Q: How to compute the probabilities of verifier messages?

Associate to each verifier vertex the set of random strings consistent with the transcript so far.

Iterating over this set partitions it by the next verifier message.

No partitioning occurs at prover nodes, so the same randomness may appear in multiple leaves.

NOTE: can set $c = pc + vr$ because the number of augmented transcripts is $\leq 2^{pc+vr}$.

Bounded Two-Way Communication

theorem: $\mathrm{IP}[\varepsilon_c, \varepsilon_s, pc, vc, vr] \subseteq \mathrm{BPTIME}\left(2^{O(pc+vc)} \cdot \mathrm{poly}\left(\frac{1}{1-\varepsilon_c-\varepsilon_s}, n\right)\right)$

Let (P,V) be an IP for $L \in \mathrm{IP}[\varepsilon_c, \varepsilon_s, pc, vc, vr]$.

Let $c := pc + vc$ be a bound on communication ONLY.

There are $\leq 2^c$ possible transcripts, hence $\leq 2^{O(c)}$ internal vertices in the tree $T(V,x)$.

PROBLEM: each transcript may have $2^{vr} = 2^{\mathrm{poly}(n)}$ augmentations,

so we cannot construct the tree $T(V,x)$ in time $2^{O(c)} \cdot \mathrm{poly}\left(\frac{1}{1-\varepsilon_c-\varepsilon_s}, n\right)$

nor compute the probabilities of verifier messages in $T(V,x)$.

IDEA: use randomness to APPROXIMATE $\mathrm{val}(T)$ with bounded probability of error.

We design a $2^{O(c)} \cdot \mathrm{poly}\left(\frac{1}{\varepsilon}, n\right)$-time probabilistic algorithm A s.t.

$$\Pr\left[\, |A(x) - \mathrm{val}(T(V,x))| > \varepsilon \,\right] \leq \frac{1}{100}.$$

Setting $\varepsilon < \frac{1-\varepsilon_c-\varepsilon_s}{2}$ suffices because $\begin{cases} x \in L \to \mathrm{val}(T(V,x)) \geq 1-\varepsilon_c \\ x \notin L \to \mathrm{val}(T(V,x)) \leq \varepsilon_s \end{cases}$.

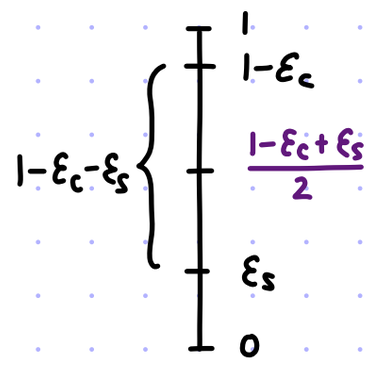

Proof [1/4]

$A(x)$

1. Set $t := \Theta\left(\frac{2^c \cdot c}{\varepsilon^2}\right)$.
2. Sample $s_1, \ldots, s_t \in \{0,1\}^{vr}$ (independently and at random).
3. Construct $T(V_R, x)$, the game tree for the verifier $V(x)$ restricted to sample randomness in $R := \{s_1, \ldots, s_t\}$ rather than $\{0,1\}^{vr}$.
4. Compute and output $\mathrm{val}(T(V,x)[R])$.

Recall:
$c := pc + vc$
$\varepsilon := \frac{1 - \varepsilon_c - \varepsilon_s}{2}$

The algorithm runs in time $2^{O(c)} \cdot \mathrm{poly}(\frac{1}{\varepsilon}, n)$ because:
- $T(V_R, x)$ has size $2^{O(c)} \cdot |R| = 2^{O(c)} \cdot t = 2^{O(c)} \cdot \Theta\left(\frac{2^c \cdot c}{\varepsilon^2}\right) = 2^{O(c)} \cdot \frac{1}{\varepsilon^2}$.
- the running time is $\mathrm{poly}(|T(V_R, x)|, n)$.

We are left to argue CORRECTNESS. The following lemma suffices.

lemma: $\Pr_R\left[\, |\mathrm{val}(T(V_R,x)) - \mathrm{val}(T(V,x))| > \varepsilon \,\right] \leq \exp\left(c \cdot 2^c - O(\varepsilon^2 t)\right) \leq \frac{1}{100}$.

Henceforth we write $T := T(V,x)$ and $T_R := T(V_R, x)$.

Proof [2/4]

lemma: $\Pr_R\left[\,|val(T_R) - val(T)| > \varepsilon\,\right] \leq \frac{1}{100}$

Define V_R to be the verifier V restricted to sample randomness in R rather than $\{0,1\}^{vr}$.

Observe that $val(T_R) = \max_{\tilde{P}} \Pr\left[\langle \tilde{P}, V_R(x)\rangle = 1\right]$.

Fix a prover strategy $\tilde{P}$ and define:

$$\Delta(\tilde{P}, R) := \Pr_{s\leftarrow R}\left[\langle \tilde{P}, V(x;s)\rangle = 1\right] - \Pr_{s\leftarrow\{0,1\}^{vr}}\left[\langle \tilde{P}, V(x;s)\rangle = 1\right]$$

$$= \underbrace{\Pr\left[\langle \tilde{P}, V_R(x)\rangle = 1\right]}_{\text{depends on } R} - \underbrace{\Pr\left[\langle \tilde{P}, V(x)\rangle = 1\right]}_{\text{independent of } R}$$

Observe that $\Pr_R\left[\,|val(T_R) - val(T)| > \varepsilon\right] \leq \Pr_R\left[\exists\, \tilde{P} : |\Delta(\tilde{P}, R)| > \varepsilon\right]$.

Indeed, for every choice of R, the event on the left implies the event on the right:

- $val(T_R) > val(T) + \varepsilon \rightarrow \Pr\left[\langle P_R^*, V_R(x)\rangle = 1\right] > \Pr\left[\langle P^*, V(x) = 1\right] + \varepsilon \geq \Pr\left[\langle P_R^*, V(x)\rangle = 1\right] + \varepsilon$
- $val(T) > val(T_R) + \varepsilon \rightarrow \Pr\left[\langle P^*, V(x) = 1\right] > \Pr\left[\langle P_R^*, V_R(x)\rangle = 1\right] + \varepsilon \geq \Pr\left[\langle P^*, V_R(x) = 1\right] + \varepsilon$

We are left to prove **claim:** $\Pr_R\left[\exists\, \tilde{P} : |\Delta(\tilde{P}, R)| > \varepsilon\right] \leq \frac{1}{100}$. ■

Proof [3/4]

<u>claim:</u> $\Pr_R\left[\exists\, \tilde{P} : |\Delta(\tilde{P},R)| > \varepsilon\right] \leq \frac{1}{100}$.

① We use a concentration argument to show that $\Delta(\tilde{P},R)$ is small w.h.p. over the choice of R:

$$\forall\, \tilde{P},\ \Pr_R\left[|\Delta(\tilde{P},R)| > \varepsilon\right] \leq 2\cdot e^{-2\cdot\varepsilon^2\cdot t}.$$

Define $z_i := \langle \tilde{P}, V(x;\rho_i)\rangle$ where ρ_i is the i-th random string in R.

The random variables $z_1,\dots,z_t$ are i.i.d. because $\rho_1,\dots,\rho_t$ are i.i.d..

Observe that:
- $\mathbb{E}[z_i] = \Pr\left[\langle \tilde{P}, V(x)\rangle = 1\right]$ because each ρ_i is random in $\{0,1\}^{r}$
- $\frac{z_1+\cdots+z_t}{t} = \Pr\left[\langle \tilde{P}, V_R(x)\rangle = 1\right]$

Hence:

$$\Pr_R\left[|\Delta(\tilde{P},R)| > \varepsilon\right]$$

$$= \Pr_R\left[\left|\Pr\left[\langle \tilde{P}, V_R(x)\rangle = 1\right] - \Pr\left[\langle \tilde{P}, V(x)\rangle = 1\right]\right| > \varepsilon\right]$$

$$= \Pr\left[\left|\frac{z_1+\cdots+z_t}{t} - \mathbb{E}[z_1]\right| > \varepsilon\right] \leq 2\cdot e^{-2\cdot\varepsilon^2\cdot t}$$

↑ additive Chernoff bound (for $z_1,\dots,z_t$ i.i.d. in $[0,1]$)

Proof [4/4]

<u>claim:</u> $\Pr_R\left[\exists\,\tilde{P} : |\Delta(\tilde{P},R)| > \varepsilon\right] \leq \frac{1}{100}$.

① We use a concentration argument to show that $\Delta(\tilde{P},R)$ is small w.h.p. over the choice of R:

$$\forall\,\tilde{P},\ \Pr_R\left[|\Delta(\tilde{P},R)| > \varepsilon\right] \leq 2\cdot e^{-2\cdot\varepsilon^2\cdot t}. \quad ✓$$

② We use a union bound to conclude the claim's proof.

Any prover $\tilde{P}$ is a function from "transcript so far" to "next message".

There are at most $(2^c)^{2^c} = 2^{c\cdot 2^c}$ provers (the input and the output is at most c bits).

By a union bound on all such provers and taking $t = \Theta\left(\frac{c\cdot 2^c}{\varepsilon^2}\right)$ large enough,

$$\Pr_R\left[\exists\,\tilde{P} : |\Delta(\tilde{P},R)| > \varepsilon\right] \leq \sum_{\tilde{P}} \Pr_R\left[|\Delta(\tilde{P},R)| > \varepsilon\right] \leq 2^{c\cdot 2^c}\cdot 2\cdot e^{-2\cdot\varepsilon^2\cdot t} \leq \frac{1}{100}.$$

Bounded One-Way Communication [1/2]

The case where we bound ONLY prover communication is more delicate.

With perfect completess, we can non-deterministically decide the complement of the language.

theorem: $IP[\varepsilon_c = 0, \varepsilon_s < 1, pc] \subseteq coNTIME(2^{O(pc)} \cdot poly(n))$

The proof takes a different approach from others in this lecture:

the nondeterministic decider receives as witness an allegedly winning strategy (of size $2^{O(pc)} \cdot poly(n)$)

for a player in a perfect-information 2-player game that has a winning strategy iff $x \notin L$.

The existence of the winning strategy is shown via Zermelo's theorem (a result in game theory).

NOTE: we CANNOT expect an upper bound such as $BPTIME(\exp(pc) \cdot poly(\frac{1}{1-\varepsilon_s}, n))$.

Indeed, $GNI \in IP[\varepsilon_c = 0, \varepsilon_s = \frac{1}{2}, pc = 1, vc = n^2]$ and GNI is not known to be in BPP.

The theorem tells us that $GNI \in coNTIME(2^{O(1)} \cdot poly(n)) = coNP$, which is indeed the case.

Moreover, prior theorems on pc+vc do not yield unexpected conclusions because $vc = n^2$ is large.

Bounded One-Way Communication [2/2]

Next we discuss the case where perfect completeness is NOT required.

The value-approximation strategy can be generalized to work with any IP when given a random continuation sampler (RCS) for the IP.

Public-coin IPs have trivial RCSs, while general IPs may not have an efficient RCS.

Nevertheless an efficient RCS can be constructed given an NP oracle.

theorem: ① $AM[\varepsilon_c, \varepsilon_s, k, pc] \subseteq BPTIME\left(2^{O(pc)} \cdot poly\left(\left(\frac{k}{1-\varepsilon_c-\varepsilon_s}\right)^k, n\right)\right)$

② $IP[\varepsilon_c, \varepsilon_s, k, pc] \subseteq BPTIME\left(2^{O(pc)} \cdot poly\left(\left(\frac{k}{1-\varepsilon_c-\varepsilon_s}\right)^k, n\right)\right)^{NP}$

We CANNOT expect to remove this, since $GNI \in IP[\varepsilon_c=0, \varepsilon_s=\frac{1}{2}, k=1, pc=1, vc=n^2]$.

NOTE: The public-coin IP for GNI has large pc, so we do not obtain unexpected conclusions.

Indeed, $pc = O(n^2)$ because the prover sends (H, π, φ) where $H \in \{0,1\}^{n^2}$, $\pi \in aut(H)$, and $\varphi : [n] \to [n]$.

From ① we infer that $GNI \notin AM\left[\varepsilon_c=\frac{1}{3}, \varepsilon_s=\frac{1}{3}, k=O\left(\frac{\log n}{\log\log n}\right), pc=O(\log n)\right]$, or else $GNI \in P$.

The state of the art improves ②:

theorem: $IP[k, pc] \subseteq coAM[k'=O(k), pc'=2^{pc} \cdot poly(k^k, n)]$

On Interactive Proofs with a Laconic Prover

Oded Goldreich	Salid Vadhan	Avi Wigderson
Weizmann Institute	Harvard	IAS


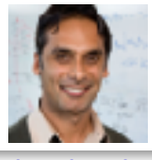
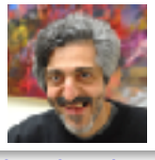

Bibliography

Limitations of IPs

- [BHZ 1987]: Does co-NP have short interactive proofs?, by Ravi Boppana, Johan Håstad, Stathis Zachos.
- [GH 1998]: On the complexity of interactive proofs with bounded communication, by Oded Goldreich, Johan Håstad.
- [GVW 2003]: On interactive proofs with a laconic prover, by Oded Goldreich, Salil Vadhan, Avi Wigderson.
- [CY 2020]: Barriers for succinct arguments in the random oracle model, by Alessandro Chiesa, Eylon Yogev.