Lecture 05

Foundations of Probabilistic Proofs
Fall 2020
Alessandro Chiesa

IPs with Bounded Resources

Let IP[pc=1] be the languages decidable via IPs where prover sends 1 bit only.
Is IP[pc=1] trivial (contained in P)?
Probably no, since GNI ∈ IP[pc=1] and GNI is not known to be in P.

⇒ even IPs with small communication can decide non-trivial languages.

Could we hope for SAT ∈ IP[pc = o(n)] (pc is sublinear in #vars)?
Note that SAT ∈ NP ⊆ IP so the question is about whether there exists an IP for SAT that provides some efficiency benefits over the trivial IP.

To formally study this question we consider:

IP[pc, vc, vr] = "languages decidable by IP where prover sends pc bits, verifier sends vc bits, and verifier uses vr random bits (& any # of rounds)"

AM[pc, vc, vr] = "similar but with public-coin IPs"

Limitations of Bounded Resources

We will learn about several limitations of IPs with bounded resources:

<u>theorem 1</u>: $IP[pc, vc, vr] \subseteq DTIME(2^{O(pc+vc+vr)} poly(n))$

<u>theorem 2</u>: $IP[pc, vc, *] \subseteq BPTIME(2^{O(pc+vc)} poly(n))$

<u>theorem 3</u>: $AM[pc, *, *] \subseteq BPTIME(2^{O(pc \cdot \log pc)} poly(n))$

<u>theorem 4</u>: $IP[pc, *, *] \subseteq BPTIME(2^{O(pc \cdot \log pc)} poly(n))^{NP}$

$\Rightarrow$ there is a relation between communication complexity of IP and the time complexity of the language it decides

<u>Observation</u>:

$GNI \in IP[pc=1]$ falls under theorem 4

$GNI \in AM[pc = O(n^2)]$ falls under theorem 3

<u>but</u> unless $GNI \in P$ we should <u>not</u> expect that $GNI \in AM[pc = o(\log n)]$

prover sends pre-image $H \in \{0,1\}^{n^2}$ & isomorphism $\phi: [n] \to [n]$

Game Tree

A transcript (of interaction) is a tuple $(a_1, b_1, \dots, a_k, b_k)$.

An augmented transcript is $(a_1, b_1, \dots, a_k, b_k, r)$ where r is verifier randomness.

Fix a verifier V and instance x.

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

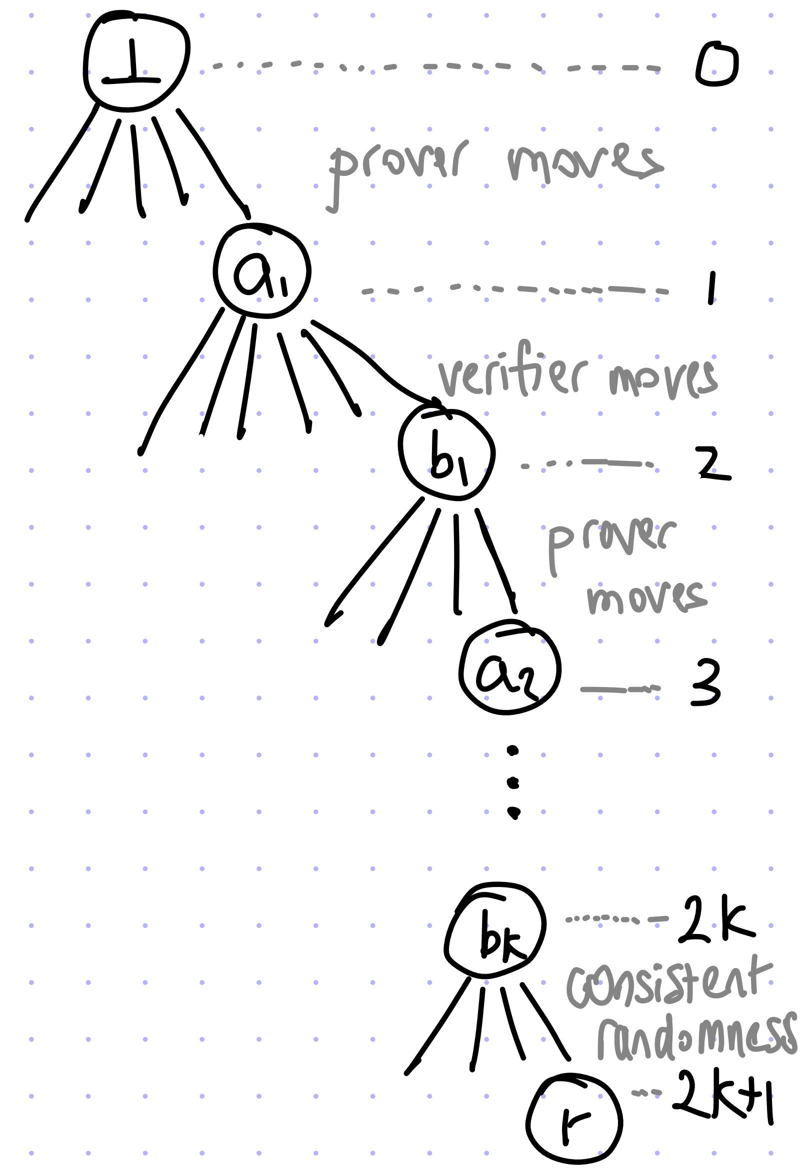

For $i = 0, 1, \dots, k-1$:

- prover moves at level $2i$
- verifier moves at level $2i+1$

Edges from $2i$ to $2i+1$ are possible moves by prover.

Edges from $2i+1$ to $2(i+1)$ are possible moves by verifier.

Edges from $2k$ to $2k+1$ are possible random strings consistent with transcript.

Approximating the Value Suffices

<u>def</u>: val(T) is the value of the root, which is recursively computed as follows:

- value of a leaf node at location $(a_1, b_1, \ldots, a_k, b_k, r)$ is the bit $V(x, a_1, \ldots, a_k; r) \in \{0,1\}$
- value of an internal node at level $2i$ is the maximum of its children's values [prover maximizes]
- value of an internal node at level $2i+1$ is the weighted average of its children's values where the weights are the probabilities of each verifier message [this includes second to last layer where the randomness r can be viewed as a fictitious final verifier message]

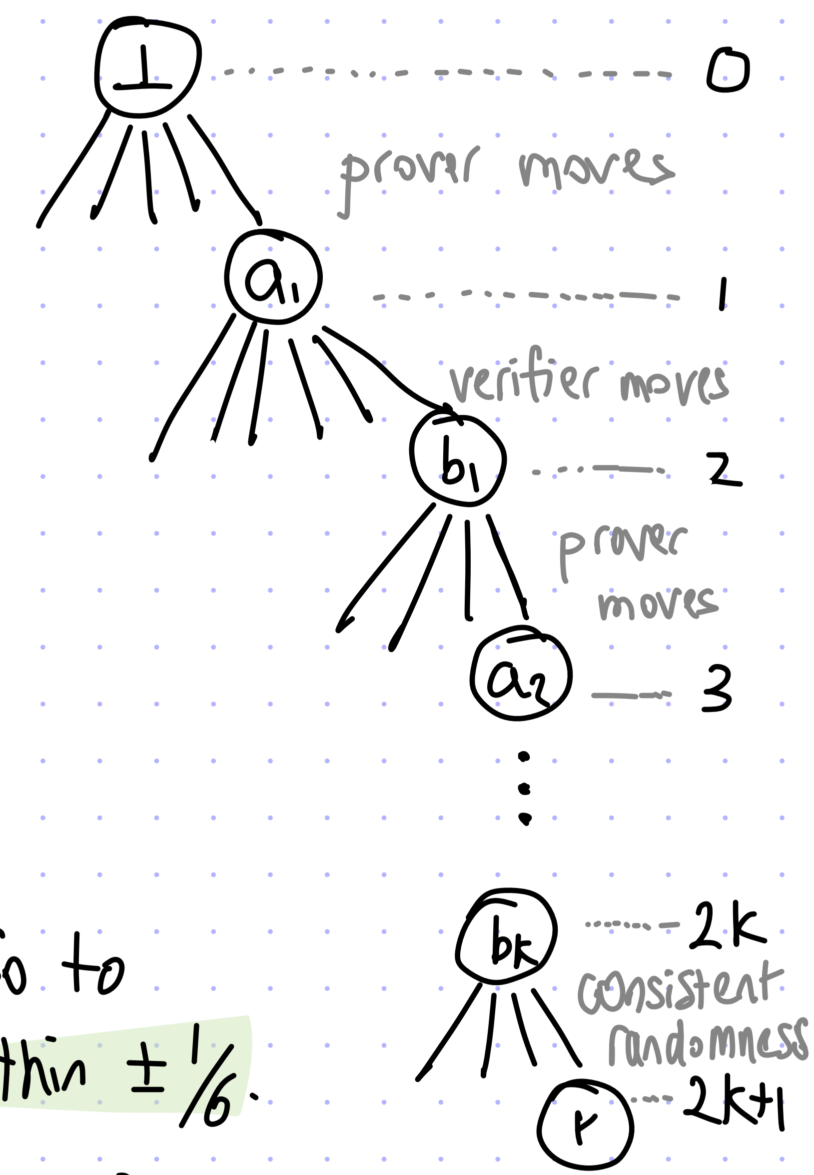

If $x \in L$ then $\mathrm{val}(T) \geq 2/3$, else if $x \notin L$ then $\mathrm{val}(T) \leq 1/3$. So to decide if $x \in L$ or $x \notin L$ it suffices to approximate val(T) to within $\pm 1/6$.

<u>Note</u>: can compute val(T) in poly(n) space and exp(poly(n)) time.

Today we are interested in time complexity to <u>approximate</u> val(T).

Theorem 1: $IP[pc, vc, vr] \subseteq DTIME(2^{O(pc+vc+vr)} poly(n))$

Let $c = pc + vc + vr$ be a bound on communication complexity and randomness.

The number of nodes in T is $2^{O(c)}$ because there are $\leq 2^c$ possible transcripts and each has $\leq 2^c$ possible augmentations, yielding $\leq 2^{2c}$ leaves.

Hence, can compute val(T) (exactly) in $2^{O(c)} poly(n)$ time, by writing out the tree explicitly and following the recursive computation.

Note: we can actually set $c = pc + vr$ since the number of augmented transcripts can be bounded by $2^{pc} \cdot 2^{vr}$.

Note: how do we compute the probabilities of verifier messages?
Associate to each node where verifier moves the set of all random strings consistent with transcript so far. To generate the probabilities iterate over this set, which will partition set according to verifier's move.

[We are not partitioning randomness when prover moves.
Hence the same randomness r may appear in more than 1 leaf.]

Theorem 2: $IP[pc, vc, *] \subseteq BPTIME(2^{O(pc+vc)} poly(n))$

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

There are still $\leq 2^c$ possible transcripts. (Hence $\leq 2^{O(c)}$ internal nodes.)

But now each transcript may have $2^{poly(n)}$ augmentations.

Hence, we cannot construct T in the allotted time $(2^{O(c)} poly(n))$, nor compute the probabilities of verifier messages inside the tree.

Instead: will use randomness to approximate $val(T)$ in $2^{O(c)} poly(n)$ time

Probabilistic algorithm:

1. sample $R = \{r_1, \dots, r_m\}$ independently in $\{0,1\}^{vr}$, with $m = \Theta(2^c \cdot c)$
2. compute $val(T[R])$ where $T[R]$ is the residual game tree obtained by omitting nodes inconsistent with R (and adjusting weights)

The algorithm runs in time $2^{O(c)} poly(n)$ because $|T[R]| = 2^{O(c)} \cdot |R| = 2^{O(c)}$.

We are left to argue correctness.

lemma: $\Pr_{R}\left[\left| \mathrm{val}(T[R]) - \mathrm{val}(T) \right| \leq \frac{1}{10} \right] \geq \frac{99}{100}$.

proof: A concentration argument applied to the right random variables.

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

$$\mathrm{val}(T[R]) = \left[\begin{array}{c} \text{maximum acceptance probability of } V^R(x) \text{ when} \\ \text{interacting with any prover strategy} \end{array} \right].$$

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

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

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

We now argue that $|\Delta(\tilde{P}, R)|$ is small w.h.p. over the choice of R.

<u>claim:</u> $\forall \tilde{P},\ \Pr_R\left[\,|\Delta(\tilde{P},R)| > \frac{1}{10}\right] \leq 2\cdot e^{-2\cdot\left(\frac{1}{10}\right)^2\cdot m}$.

<u>proof:</u>

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

The random variables $z_1,\dots,z_m$ are i.i.d. because $r_1,\dots,r_m$ are.

Moreover: • $\mathbb{E}[z_i] = \Pr[\langle \tilde{P}, V(x)\rangle = 1]$ as each r_i is random in $\{0,1\}^{v_r}$

• $\frac{z_1+\cdots+z_m}{m} = \Pr[\langle \tilde{P}, V^R(x)\rangle = 1]$

$X_1,\dots,X_m$ iid in $[0,1]$

$\Pr[|\bar{X} - \mathbb{E}[X_1]| > \varepsilon] \leq 2\cdot e^{-2\cdot\varepsilon^2\cdot m}$

We can conclude the proof by a Chernoff bound:

$$\Pr_R\left[\,|\Delta(\tilde{P},R)| > \frac{1}{10}\right] = \Pr_R\left[\,\left|\Pr[\langle \tilde{P}, V^R(x)\rangle = 1] - \Pr[\langle \tilde{P}, V(x)\rangle = 1]\right| > \frac{1}{10}\right]$$

$$= \Pr_R\left[\,\left|\frac{z_1+\cdots+z_m}{m} - \mathbb{E}[z_1]\right| > \frac{1}{10}\right] \leq 2\cdot e^{-2\cdot\left(\frac{1}{10}\right)^2\cdot m} \qquad \blacksquare$$

<u>claim:</u> $\forall \tilde{P},\ \Pr_{R}\left[|\Delta(\tilde{P},R)| > \frac{1}{10}\right] \leq 2\cdot e^{-2\cdot(\frac{1}{10})^2\cdot m}$.

Any prover $\tilde{P}$ is a function from transcript so far to next message.
So there are at most $(2^c)^{2^c} = 2^{c\cdot 2^c}$ provers (as input and output sizes are $\leq 2^c$).

By a union bound on all such provers, and taking $m = \Theta(2^c\cdot c)$ large enough,

$$\Pr_{R}\left[\exists \tilde{P} : |\Delta(\tilde{P},R)| > \frac{1}{10}\right] \leq \sum_{\tilde{P}} \Pr_{R}\left[|\Delta(\tilde{P},R)| > \frac{1}{10}\right] \leq 2^{c2^c}\cdot 2\cdot e^{-2\cdot(\frac{1}{10})^2\cdot m} < \frac{1}{100}.$$

We conclude the proof by noting that:

$$\Pr_{R}\left[|\mathrm{val}(T[R]) - \mathrm{val}(T)| > \frac{1}{10}\right] \leq \Pr_{R}\left[\exists \tilde{P} : |\Delta(\tilde{P},R)| > \frac{1}{10}\right] \quad \left(< \frac{1}{100}\right).$$

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

- $\mathrm{val}(T[R]) > \mathrm{val}(T) + \frac{1}{10} \rightarrow \Pr[\langle P^*_R, V^R(x)\rangle = 1] > \Pr[\langle P^*, V(x)\rangle = 1] + \frac{1}{10} \geq \Pr[\langle P^*_R, V(x)\rangle = 1] + \frac{1}{10}$

- $\mathrm{val}(T) > \mathrm{val}(T[R]) + \frac{1}{10} \rightarrow \Pr[\langle P^*, V(x)\rangle = 1] > \Pr[\langle P^*_R, V^R(x)\rangle = 1] + \frac{1}{10} \geq \Pr[\langle P^*, V^R(x)\rangle = 1] + \frac{1}{10}$