Lecture 25

Foundations of Probabilistic Proofs
Fall 2020
Alessandro Chiesa

Sublinear Verification for Any Computation

We have seen how to achieve sublinear verification via PCPs/IOPs for machine computations.
More generally, sublinear verification is achievable iff the description of a computation is shorter than the computation itself (informally, the computation is "structured").
(Indeed, the verifier must at minimum read the description of the computation!)

Q: How can we achieve sublinear verification for any computation?
(Including ones whose shortest description is the computation itself, like a random circuit.)

One approach: Holographic Proofs (a cool-sounding but not very descriptive historical term)

Consider an offline/online model where:

- in the offline phase the description of the computation is "encoded" into an oracle;
- in the online phase the PCP/IOP verifier has oracle access to this oracle, and may check multiple statements wrt different inputs to the computation.

Today we show how to formalize this idea and how to construct a protocol for it.

Indexed Relations and Languages

An indexed relation is a set $R = \{(\mathbb{i}, x, w) \mid \dots\}$ where $\mathbb{i}$ is the index, x the instance, and w the witness.

The corresponding indexed language is $L(R) = \{(\mathbb{i}, x) \mid \exists\, w \text{ s.t. } (\mathbb{i}, x, w) \in R\}$.

The valid witnesses for the index-instance pair $(\mathbb{i}, x)$ are $R[(\mathbb{i}, x)] = \{w \mid (\mathbb{i}, x, w) \in R\}$.

The index is to be interpreted as the "large" description of a computation.

Here are some examples:

- circuit satisfiability over $\mathbb{F}$

$$\mathrm{CSAT}(\mathbb{F}) = \{(\mathbb{i}, x, w) = (C, x, w) \mid C : \mathbb{F}^n \to \mathbb{F} \text{ is a circuit and } C(x, w) = 0\}$$

- quadratic equations over $\mathbb{F}$

$$\mathrm{QESAT}(\mathbb{F}) = \{(\mathbb{i}, x, w) = ((p_1, \dots, p_m), x, w) \mid p_1, \dots, p_m \in \mathbb{F}^{\leq 2}[X_1, \dots, X_n] \text{ and } p_1(x, w) = \dots = p_m(x, w) = 0\}$$

- rank-1 constraints over $\mathbb{F}$

$$\mathrm{R1CS}(\mathbb{F}) = \{(\mathbb{i}, x, w) = ((A, B, C), x, w) \mid A, B, C \in \mathbb{F}^{m \times n} \text{ and } A \cdot \binom{x}{w} \circ B \binom{x}{w} = C \cdot \binom{x}{w}\}$$

Holographic PCPs

A holographic PCP for an indexed language L is a tuple $(\mathcal{I}, \mathbf{P}, \mathbf{V})$ s.t.

① completeness: $\forall (\mathbb{i}, \mathbb{x}) \in L$, for $\Pi_0 := \mathcal{I}(\mathbb{i})$ and $\pi := \mathbf{P}(\mathbb{i}, \mathbb{x})$, $\Pr_{\rho}[\mathbf{V}^{\Pi_0, \pi}(\mathbb{x}; \rho) = 1] \geq 1 - \epsilon_c$.

② soundness: $\forall (\mathbb{i}, \mathbb{x}) \notin L$, for $\Pi_0 := \mathcal{I}(\mathbb{i})$, $\forall \tilde{\pi}$ $\Pr_{\rho}[\mathbf{V}^{\Pi_0, \tilde{\pi}}(\mathbb{x}; \rho) = 1] \leq \epsilon_s$.

Diagramatically:

- OFFLINE (once per index)
- ONLINE (any number of times)

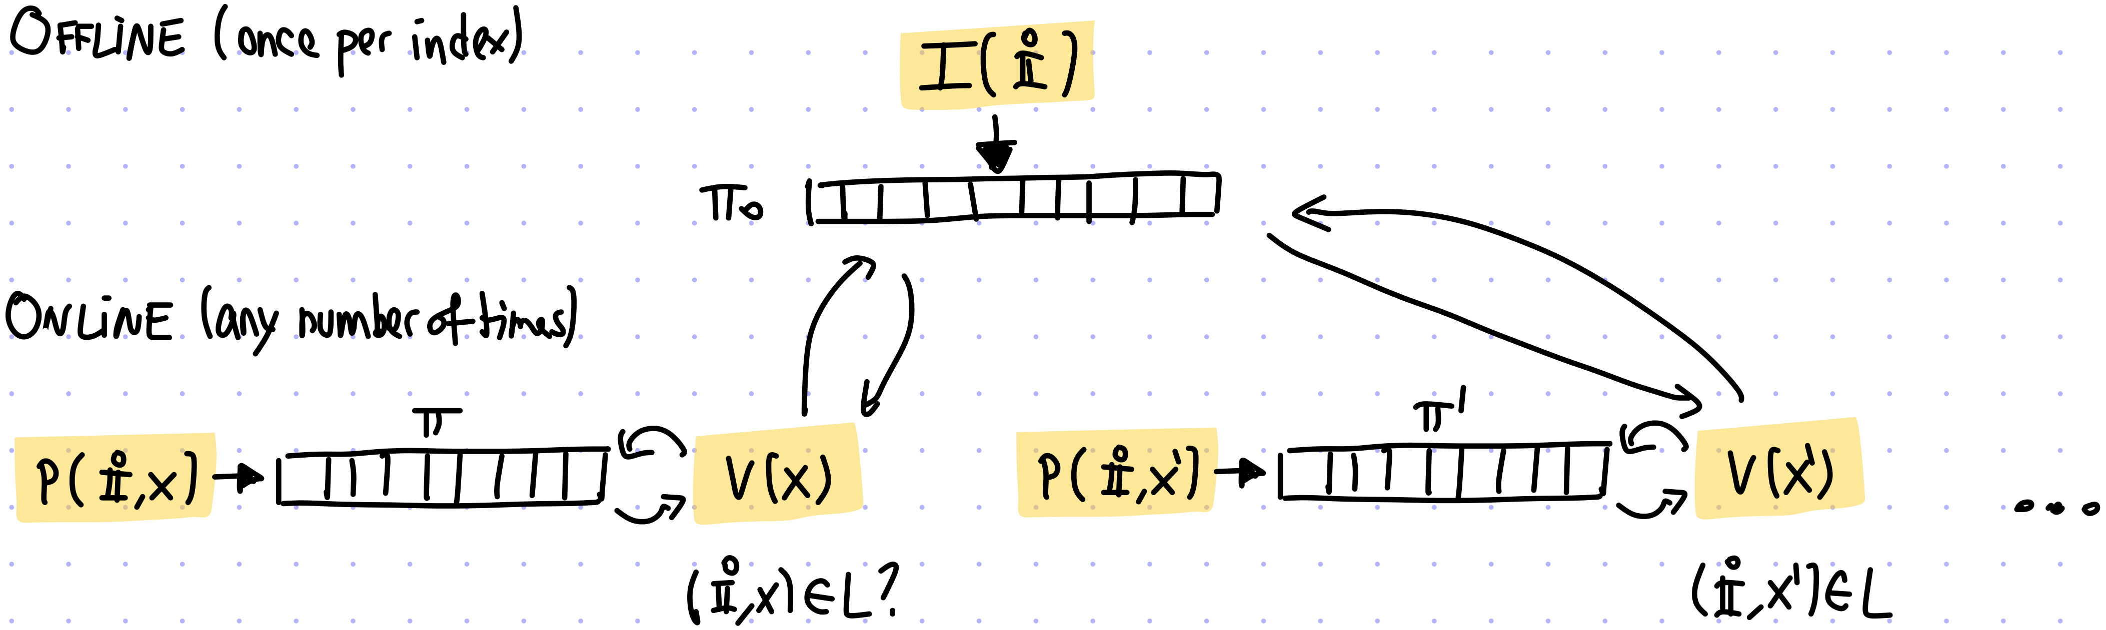

Efficiency: proof length is $|\Pi_0| + |\pi|$ and query complexity is $q_0 + q$.

We can similarly define holographic IPs and IOPs. (As well as variants for robustness, proximity, ...).

From Holography to Preprocessing [1/2]

One motivation to study holography is that it naturally leads to preprocessing arguments.
These enable sublinear verification for any computation, given a one-time (public) preprocessing step.

A preprocessing argument system for an indexed language L looks as follows:

OFFLINE (once)

ONLINE (multiple)

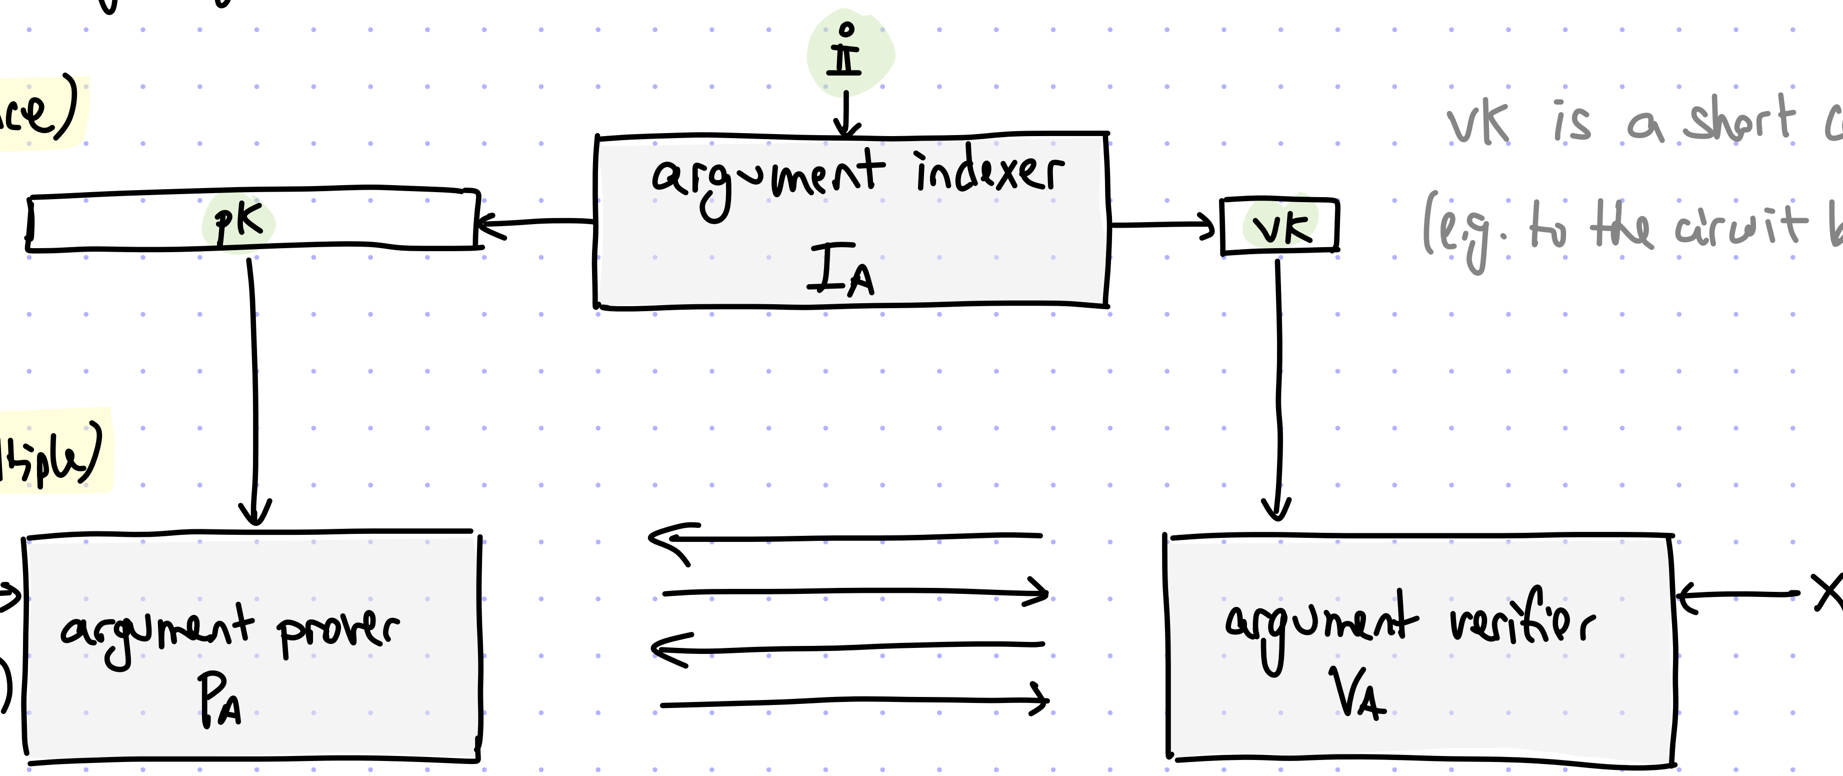

vk is a short commitment to $\mathbb{i}$
(e.g. to the circuit being proved)

In the past we have seen how: PCP (or IOP) + CRH $\longrightarrow$ succinct argument

Now we see how holographic PCP (or IOP) + CRH $\longrightarrow$ preprocessing succinct argument

From Holography to Preprocessing [2/2]

SETUP: Everyone has access to a collision-resistant function h (sampled from a family $\mathcal{H}_\lambda$).

OFFLINE: Anyone can compute the key pair for an index $\mathbb{i}$ (re-usable any number of times):

$I_A(\mathbb{i})$
1. Compute the encoded index: $\Pi_0 := I(\mathbb{i})$.
2. Commit to encoded index: $rt_0 := MT_h(\Pi_0)$.
3. Output key pair $(pk, vk) := ((\mathbb{i}, \Pi_0), rt_0)$.

ONLINE: Anyone can use the key pair to prove/verify statements of the form $(\mathbb{i}, x) \in L$:

$P_A(h, pk, x, w)$		$V_A(h, vk, x)$
1. Compute PCP string: $\Pi := P(\mathbb{i}, x, w)$		
2. Commit to PCP string: $rt := MT_h(\Pi)$	$\xrightarrow{rt}$	
	$\xleftarrow{\rho}$	Sample PCP randomness.
3. Deduce query set Q for $V^{\Pi_0, \Pi}(x; \rho)$.		
4. Produce auth path for each answer.	$\xrightarrow{(\Pi_0, \Pi)[Q], \text{auth}}$	$V^{(\Pi_0, \Pi)[Q]}(x; \rho)$ & check auth wrt (rt_0, rt).
$\text{time}(P_A) = \text{time}(P) + O_\lambda(\ell)$	$O_\lambda(q \cdot \log \ell)$	$\text{time}(V_A) = \text{time}(V) + O_\lambda(q \cdot \log \ell)$

Holographic PCP for NP

We have proved that NP has PCPs with polynomial proof length and polylogarithmic query complexity. The PCP verifier did not (and could not) run in sublinear time because it had to read the description of the NP statement being proved, in that case the list of quadratic equations.

We show how to achieve sublinear verification time with the help of an indexer:

<u>theorem</u>: $\mathrm{QESAT}(\mathbb{F}) \in \mathrm{HPCP}\left[\begin{matrix} \varepsilon_c = 0 \\ \varepsilon_s = 1/2 \end{matrix}, \begin{matrix} \Sigma = \mathbb{F} \\ l = |\mathbb{F}|^{O\left(\frac{\log n}{\log\log n}\right)} \end{matrix}, \begin{matrix} vt = \mathrm{poly}(|x|, \log n) \\ q = \mathrm{poly}(\log n) \end{matrix}\right]$

Here we mean the indexed language $\{(\mathbb{i}, x) = ((p_1, \ldots, p_m), x) \mid \exists w \text{ s.t. } p_1(x,w) = \cdots = p_m(x,w)\}$.

This implies, via the holography → preprocessing connection, a preprocessing succinct argument for QESAT($\mathbb{F}$) where $\mathrm{time}(\mathcal{I}_A) = \mathrm{poly}_\lambda(n)$, $\mathrm{time}(P_A) = \mathrm{poly}_\lambda(n)$, and $\mathrm{time}(V_A) = \mathrm{poly}_\lambda(|x|, \log n)$.

The ability to verify any (not necessarily structured) computation in sublinear time is convenient to "program" and is useful for cryptographic applications (e.g. recursive proofs).

We now prove the theorem by modifying the construction that we have already.

Holographic PCP for NP

Fix $H_v, H_e \subseteq \mathbb{F}$ of sizes $O(\log n), O(\log m)$ and set $s_v = \log n / \log|H_v|$ and $s_e = \log m / \log|H_e|$.

$\mathcal{I}(\mathbb{i} = (p_1, \dots, p_m))$

1. $\hat{p}(X_1, \dots, X_{s_e}) := \sum_{0 \le j_1, \dots, j_{s_e} < |H_e|} X_1^{j_1} \cdots X_{s_e}^{j_{s_e}} \cdot p_{j_1 \cdots j_{s_e}}$
2. $\hat{c}(X_1, \dots, X_{s_e}, Y_1, \dots, Y_{s_v}, Z_1, \dots, Z_{s_v}) := \sum_{a,b \in H_v^{s_v}} \hat{p}(X_1, \dots, X_{s_v})[a,b] \cdot I(a,Y) \cdot I(b,Z)$
3. Output $\Pi_0 : \mathbb{F}^{s_e + 2s_v} \to \mathbb{F}$ where $\Pi_0 := \hat{c}|_{\mathbb{F}^{s_e+2s_v}}$

$P(\mathbb{i} = (p_1, \dots, p_m), x, w)$

1. Output $\Pi_a : \mathbb{F}^{s_v} \to \mathbb{F}$ the $(\mathbb{F}, H_v, s_v)$-extension of $a : [n] \to \mathbb{F}$
2. For every $r_1, \dots, r_{s_e} \in \mathbb{F}$:
   - $p_r := \sum_{0 \le j_1, \dots, j_{s_e} < |H_e|} r_1^{j_1} \cdots r_{s_e}^{j_{s_e}} \cdot p_{j_1, \dots, j_{s_e}}$
   - output $\Pi_{SC}[r] :=$ eval table for sumcheck to show $p_r(a) = 0$
3. Output Π_{IC} that proves that Π_a is consistent with x.

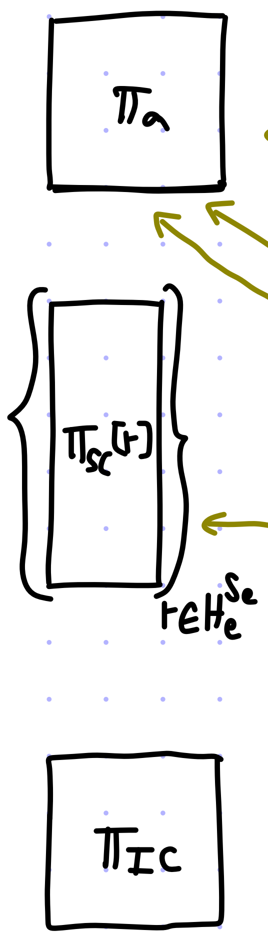

$V(~~\mathbb{i} = (p_1, \dots, p_m)~~, x)$

1. Run low-degree test on Π_a
   $V_{LDT}(\mathbb{F}, s_v, |H_v|)$
2. Sample $r_1, \dots, r_{s_e} \in \mathbb{F}$ and ~~compute $p_r := \sum_{0 \le j_1, \dots, j_{s_e} < |H_e|} r_1^{j_1} \cdots r_{s_e}^{j_{s_e}} \cdot p_{j_1, \dots, j_{s_e}}$~~
3. Run sumcheck protocol to check that
   $\sum_{\alpha, \beta \in H_v^{s_v}} \overset{\hat{c}(r,\alpha,\beta)}{\cancel{p_r(\alpha,\beta)}} \cdot \hat{a}(\alpha) \cdot \hat{a}(\beta) = 0$
   $\boxed{V_{SC}(\mathbb{F}, H_v, 2s_v, 0)}$
4. Check input consistency proof Π_{IC} (query Π_a and evaluate $\hat{x}$ at a point).

m equations to 1 equation

check 1 equation

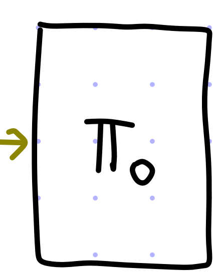

Holographic IOP for NP

We have obtained holographic PCPs for NP with polynomial size and polylogarithmic query complexity.

Can we improve the proof length by using IOPs instead of PCPs?

This requires some new ideas but can be done:

$S = \#$ of non-zero entries in A, B, C

Theorem: For "large smooth" $\mathbb{F}$, $\mathrm{R1CS}(\mathbb{F}) \in \mathrm{HIOP}\left[\begin{matrix}\varepsilon_c = 0 \\ \varepsilon_s = 1/2\end{matrix}, k = O(\log S), \begin{matrix}\Sigma = \mathbb{F} \\ \ell = O(S)\end{matrix}, \begin{matrix} vt = O(|x| + \log S) \\ q = O(\log S)\end{matrix}\right]$

Here we mean the indexed language $\{(\mathbb{i}, x) = ((A,B,C), x) \mid A,B,C \in \mathbb{F}^{m\times n} \ \& \ \exists\, w \text{ s.t. } A\cdot\binom{x}{w} \circ B\cdot\binom{x}{w} = C\cdot\binom{x}{w}\}$

This theorem builds on the non-holographic counterpart that we have already seen:

$$\mathrm{R1CS}(\mathbb{F}) \in \mathrm{IOP}\left[\begin{matrix}\varepsilon_c = 0 \\ \varepsilon_s = 1/2\end{matrix}, k = O(\log m), \begin{matrix}\Sigma = \mathbb{F} \\ \ell = O(m)\end{matrix}, \begin{matrix} vt = O(m) \\ q = O(\log m)\end{matrix}\right]$$

It requires one new idea: a holographic subroutine for linear equations.

$$\hat{g}|_H \equiv M \cdot \hat{f}|_H$$

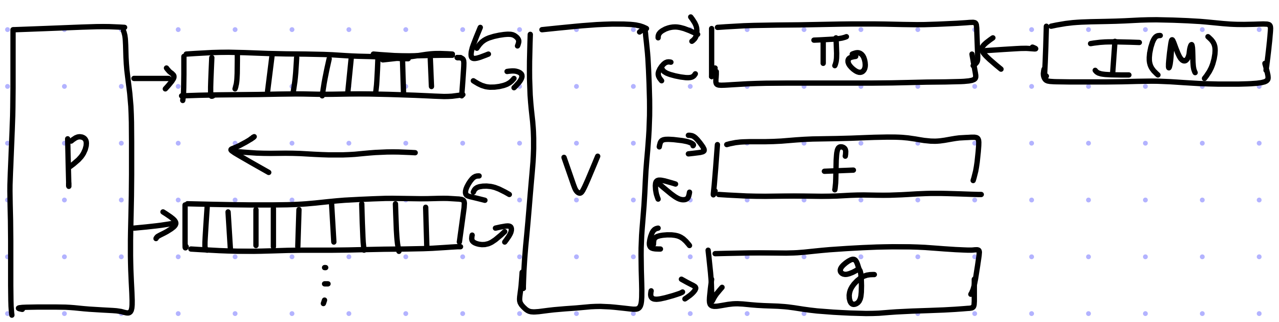

The goal is easy if we allow $\ell = O(m \cdot n)$.

To achieve $\ell = O(S)$ (and $vt = O(\log S)$)

we need to use algebraic tricks

related to Lagrange polynomials.

Interactive Proofs

arithmetization, sumcheck, low-degree extensions, GKR, IP=PSPACE, limitations, ZK

Probabilistically Checkable Proofs

exponential-size PCP, polynomial-size PCPs linearity testing, low-degree testing, zero testing

Interactive Oracle Proofs

linear-size proofs, univariate sumcheck, FRI protocol

Proof Composition

robust proofs, proximity proofs, composition, PCP Theorem

And more !

parallel repetition, sliding scale conjecture, PCP/IOP limitations, holography