

Foundations of Probabilistic Proofs

A course by **Alessandro Chiesa**

Lecture 06

Intro to PCPs



These slides are licensed under the [CC BY-SA 4.0 license](https://creativecommons.org/licenses/by-sa/4.0/).

Checking a proof without reading it?

GOAL in the next few lectures:

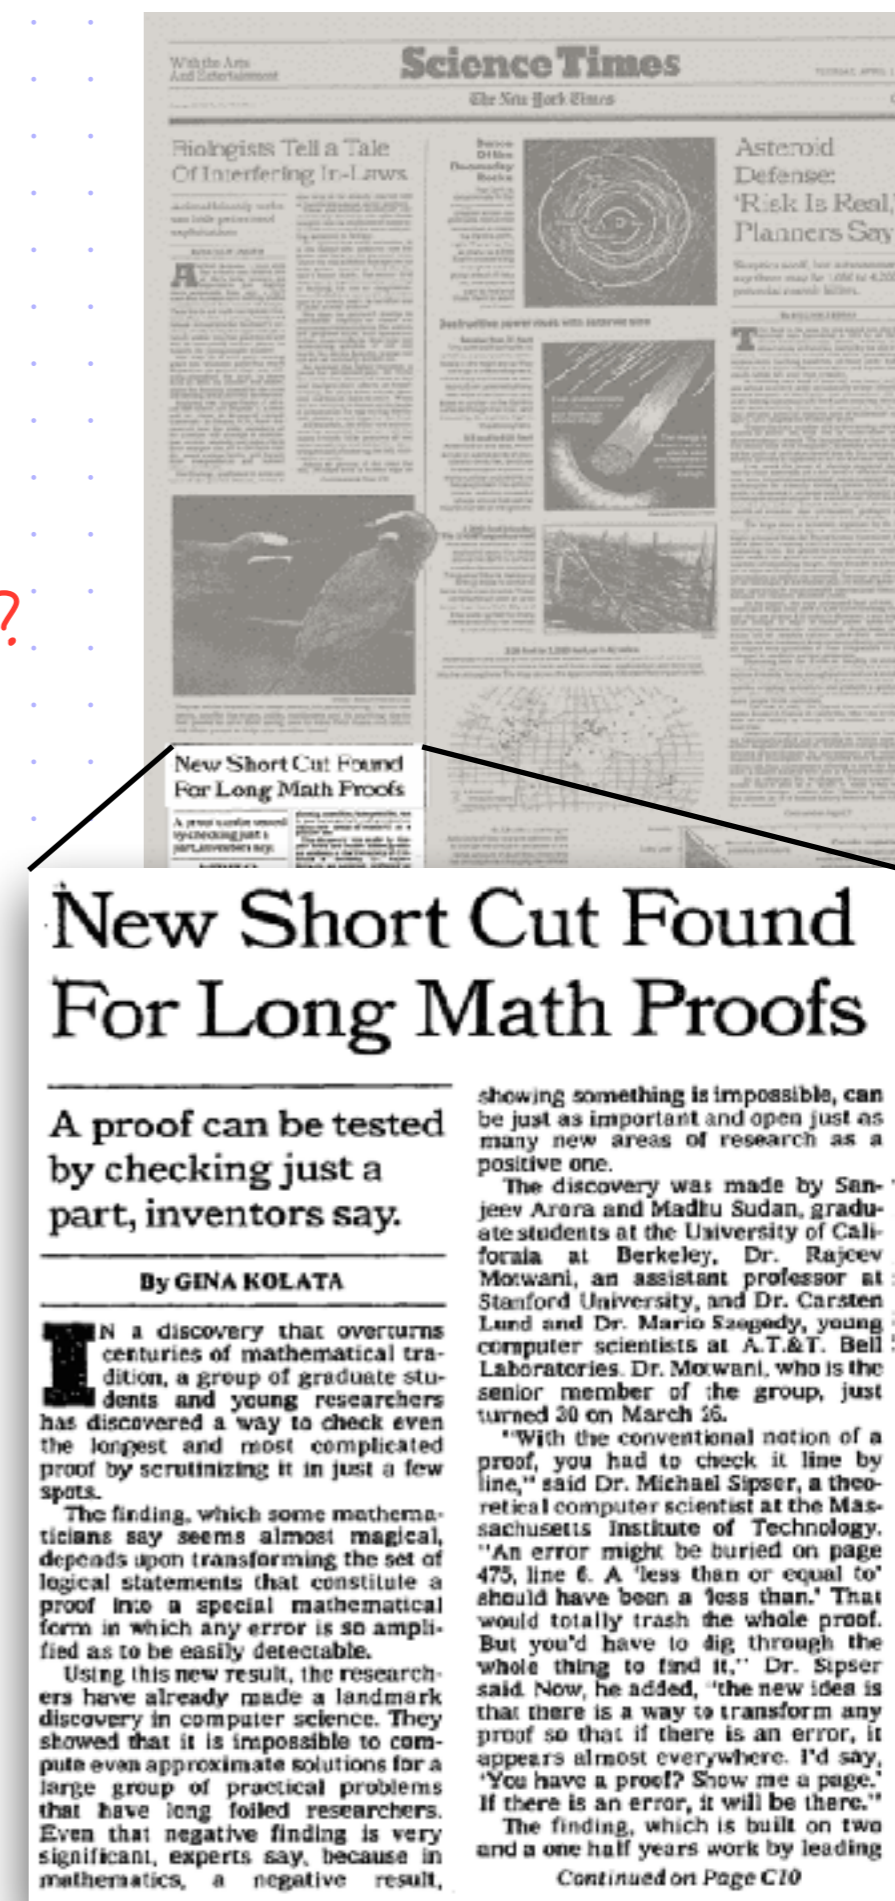
probabilistically check non-interactive proofs
at few locations with small error

CHALLENGE: what if the proof has only a single "mistake"?
(How to catch the mistake w.h.p.?)

EXAMPLE: given a graph $G=(V,E)$ and a coloring $\chi:V\rightarrow\{1,2,3\}$,
how to probabilistically check that χ is a 3-coloring of G ?
(If G is not 3-colorable, at worst all but one edges are valid.)


The theory of PROBABILISTICALLY CHECKABLE PROOFS (PCPs)
addresses this challenge!

Tools and ideas from interactive proofs, coding theory,
property testing, and more.



New Model: Probabilistically Checkable Proofs

NP represents proofs checkable via a deterministic polynomial-time verifier:




TODAY WE STUDY A NEW MODEL

INTERACTIVE PROOFS:

IPs are proofs checkable via a polynomial-time verifier that has these additional resources:


- ① randomness
- ② interaction



PROBABILISTICALLY-CHECKABLE PROOFS:

PCPs are proofs checkable via a polynomial-time verifier that has these additional resources:

- ① randomness
- ② oracle access to proof



Definition of PCP

[The definition for a language L is a special case.]

We say that (P, V) is a **PCP system** for a relation R

with completeness error ϵ_c and soundness error ϵ_s (with $1 - \epsilon_c > \epsilon_s$) if the following holds:

① **COMPLETENESS:** $\forall (x, w) \in R \quad \Pr[V^\pi(x) = 1 \mid \pi \leftarrow P(x, w)] \geq 1 - \epsilon_c.$

② **SOUNDNESS:** $\forall x \notin L(R) \quad \forall \tilde{\pi} \quad \Pr[V^{\tilde{\pi}}(x) = 1 \mid \tilde{\pi} \leftarrow \tilde{P}] \leq \epsilon_s.$ Equivalently: $\forall x \notin L(R) \quad \forall \tilde{\pi} \quad \Pr[V^{\tilde{\pi}}(x) = 1] \leq \epsilon_s.$


We use the notation $V^\pi(x; g)$ to make explicit that g is the randomness of V .

We call π a "PCP string". Intuitively, π is a "robust encoding" of a witness, which admits a probabilistic verification that reads a few symbols of it.

For IPs we cared about: round complexity, communication complexity, ...

For PCPs we have different parameters:

- Σ : proof alphabet
- l : proof length
- q : verifier query complexity
- r : verifier randomness complexity



(Queries to π are typically non-adaptive.)

Some Special Cases

We wish to understand $\text{PCP}[\epsilon_c, \epsilon_s, \Sigma, \ell, q, r, \dots]$ in different regimes.

Suppose that there is **no proof** ($q=0$):

- $\text{PCP}[q=0, r=0] = P$ (no proof and no randomness)
- $\text{PCP}[q=0, r=O(\log n)] = P$ (can try all random strings in polynomial time)
- $\text{PCP}[q=0, r=\text{poly}(n)] = \text{BPP}$ (any amount of randomness and no proof)

Suppose that there is **no randomness** ($r=0$):

- $\text{PCP}[q=\text{poly}(n), r=0] = \text{NP}$ (queries are fixed and polynomially many)

A trivial PCP:

- $3\text{SAT} \in \text{PCP}[\epsilon_c=0, \epsilon_s=1-\frac{1}{m}, \Sigma=\{0,1\}, \ell=n, q=3, r=\log m]$

proof: $\forall_{\varphi \in \{0,1\}^n} V_{\text{PCP}}(\varphi) :=$

1. Sample a random clause $j \in [m]$.
2. Check that φ satisfies the j -th clause of φ . ■

We denote by PCP the case with no restrictions (beyond V_{PCP} runs in polynomial time):

$$\text{PCP} := \text{PCP}[\epsilon_c=0, \epsilon_s=\frac{1}{2}, |\Sigma|=\exp(n), \ell=\exp(n), q=\text{poly}(n), r=\text{poly}(n)]$$

Upper Bound: NEXP

theorem: $PCP \subseteq NEXP$

- Σ : proof alphabet
- ℓ : proof length
- q : verifier query complexity
- r : verifier randomness complexity

REVIEW: $NTIME[T] = \left\{ L \mid \exists \text{ machine } M \text{ s.t. } \begin{cases} x \in L \rightarrow \exists w M(x,w)=1 \\ x \notin L \rightarrow \forall w M(x,w)=0 \end{cases} \text{ and } M(x,w) \text{ runs in } T(|x|) \text{ time} \right\}$

nondeterministic polynomial time: $NP = \bigcup_{c \in \mathbb{N}} NTIME[n^c]$

nondeterministic exponential time: $NEXP = \bigcup_{c \in \mathbb{N}} NTIME[2^{n^c}]$

lemma: $PCP[\epsilon_c, \epsilon_s, \Sigma, \ell, r] \subseteq NTIME[T = (2^r + \ell \cdot \log|\Sigma|) \cdot \text{poly}(n)]$

proof: Let (P, V) be a PCP system for L with the parameters above.

Define the decider as follows:

$D(x, \pi) := \begin{matrix} \prod_{\{0,1\}^r} \\ \prod_{\Sigma^\ell} \end{matrix} \begin{matrix} 1. \text{ For every } \rho \in \{0,1\}^r: \text{ compute } b_\rho := V^\pi(x; \rho) \in \{0,1\}. \\ 2. \text{ Output } 1 \text{ if and only if } (\sum_\rho b_\rho) / 2^r \geq 1 - \epsilon_c. \end{matrix}$

If $x \in L$ then $\exists \pi$ s.t. $D(x, \pi) = 1$. If $x \notin L$ then $\forall \pi D(x, \pi) = 0$. ■

lemma: (i) $\ell \leq 2^r \cdot q$ for non-adaptive verifiers

(ii) $\ell \leq 2^r \cdot |\Sigma|^q \cdot q$ for adaptive verifiers

(in PCP constructions ℓ is typically smaller than these upper bounds)

proof of (i): there are at most 2^r different query sets

proof of (ii): each query answer may lead to a different next query ■

Lower Bound: PSPACE

Theorem: $PSPACE \subseteq PCP$


We proved that $IP = PSPACE$. So the following lemma suffices.

Lemma: $IP[\epsilon_c, \epsilon_s, \underset{\text{rounds}}{K}] \subseteq PCP[\epsilon_c, \epsilon_s, q=K]$

proof: Let (P_{IP}, V_{IP}) be a k -round IP for L .

Consider PCP strings in this format:

$$\pi := (a_1, (a_2[b_1])_{b_1}, (a_3[b_1, b_2])_{b_1, b_2}, \dots, (a_k[b_1, \dots, b_{k-1}])_{b_1, \dots, b_{k-1}})$$



The PCP verifier (which adaptively makes k queries to π) works as follows:

$V^\pi(x) :=$ 1. Sample IP randomness g .

2. Simulate $V_{IP}(x; g)$ where, in round i , the IP prover message is $a_i[b_1, \dots, b_{i-1}]$.
prior $i-1$ messages by $V_{IP}(x; g)$

COMPLETENESS: the honest PCP string is $\pi := (P_{IP}(x), (P_{IP}(x, b_1))_{b_1}, \dots, (P_{IP}(x, b_1, \dots, b_{k-1}))_{b_1, \dots, b_{k-1}})$.

SOUNDNESS: any PCP string in the above format is an "unrolled" IP prover. ■

If V_{IP} is **PUBLIC-COIN** then each b_i is (independently) random.

Hence $g = (b_1, \dots, b_k)$ and the k queries $((b_1, \dots, b_{i-1}))_{i \in [k]}$ to π are **non-adaptively** determined.

Preview of PCP Capabilities

We can ask several questions about PCPs, inspired from our study of IPs.

- Which languages have PCPs? We learned that $PSPACE \subseteq PCP \subseteq NEXP$.

A: $PCP = NEXP$

- Do PCPs have benefits for NP languages? E.g. PCP query complexity \ll NP witness size

A: YES

- Do PCPs have benefits for P languages? E.g. PCP verifier time \ll P decision time

A: YES

- Are there ZK PCPs for NP?

A: YES

→ MANY GOOD NEWS!


We study most of these questions in the next few lectures.

Delegation of Computation via PCPs

We will prove (in a few lectures) the PCP Delegation Of Computation Theorem:

$$NTIME[T] = \left\{ L \mid \exists \text{ machine } M \text{ s.t. } \begin{cases} x \in L \rightarrow \exists w \ M(x,w)=1 \\ x \notin L \rightarrow \forall w \ M(x,w)=0 \end{cases} \text{ and } M(x,w) \text{ runs in } T(|x|) \text{ time} \right\}$$

theorem: $NTIME[T] \subseteq PCP \left[\begin{array}{ll} \text{proof length } \ell = \text{poly}(T) , & \text{query complexity } q = \text{poly}(\log T) \\ \text{prover time } p_t = \text{poly}(T) , & \text{verifier time } v_t = \text{poly}(n, \log T) \end{array} \right]$



Proving the theorem will involve:

- recycling techniques (arithmetization, sumcheck protocol, ...)
- new ideas (low-degree testing, succinct arithmetization, ...)

Checking Computations in Polylogarithmic Time

László Babai



Univ. of Chicago⁶ and
Eötvös Univ., Budapest

Lance Fortnow




Dept. Comp. Sci.
Univ. of Chicago⁶

Leonid A. Levin



Dept. Comp. Sci.
Boston University⁴

Mario Szegedy



Dept. Comp. Sci.
Univ. of Chicago⁶

The PCP Model Is Unrealistic

Checking Computations in Polylogarithmic Time

László Babai



Univ. of Chicago⁶ and
Eötvös Univ., Budapest

Lance Fortnow



Dept. Comp. Sci.
Univ. of Chicago⁶

Leonid A. Levin



Dept. Comp. Sci.
Boston University⁴

Mario Szegedy



Dept. Comp. Sci.
Univ. of Chicago⁶

In this setup, a single reliable PC can monitor the operation of a herd of supercomputers working with possibly extremely powerful but unreliable software and untested hardware.

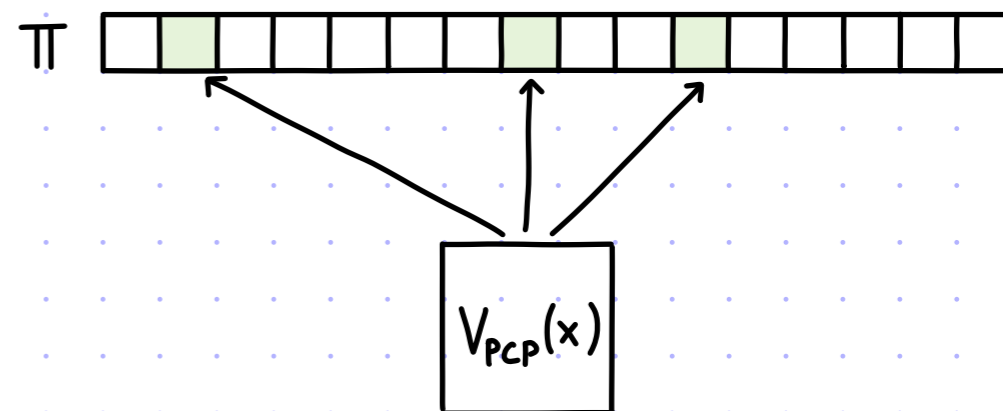
CHALLENGE: How can PCPs be useful?

The PCP model gives the verifier

an unrealistic "super" power:

ORACLE ACCESS to a long proof.

How can a verifier have such a capability?



Applications of PCPs


PCPs are useful as an **intermediate model**.

They are an invaluable **TOOL** in two major applications.

SUCCINCT ARGUMENTS



Cryptographic proofs with strong efficiency features (e.g. for delegation of computation).




TODAY

HARDNESS OF APPROXIMATION



Which problems remain hard even if we only seek an approximate solution?



LATER LECTURE

A Cryptography Jaunt:
Succinct Arguments
Based On PCPs

Two Attempts At A Succinct Argument

PROOF ATTEMPT #1:

$P_{IA}(x, w)$

- Produce PCP string: $\pi := P_{PCP}(x, w)$.
- Deduce PCP query set Q for $V_{PCP}^{\pi}(x; g)$.
- Set $a := \pi[Q] \in \Sigma^Q$.

\xleftarrow{g}

$V_{IA}(x)$

Sample PCP randomness $g \in \{0, 1\}^r$.

\xrightarrow{a}

$V_{PCP}^{[Q, a]}(x; g) \stackrel{?}{=} 1$

PROBLEM: a malicious prover can choose answers $a \in \Sigma^Q$ based on the randomness g .

(Separately, we cannot hope to succeed without cryptography.)

PROOF ATTEMPT #2:

$P_{IA}(x, w)$

- Produce PCP string: $\pi := P_{PCP}(x, w)$.
- Commit to PCP string: $cm := \text{Hash}(\pi)$.
- Deduce PCP query set Q for $V_{PCP}^{\pi}(x; g)$.
- Set $a := \pi[Q] \in \Sigma^Q$.

\xrightarrow{cm}

$V_{IA}(x)$

Sample PCP randomness $g \in \{0, 1\}^r$.

$\xrightarrow{a, \pi}$

$V_{PCP}^{[Q, a]}(x; g) \stackrel{?}{=} 1$ \wedge $\text{Hash}(\pi) \stackrel{?}{=} cm$
 \wedge $\pi[Q] \stackrel{?}{=} a$

PROBLEM: the honest prover sends the entire PCP string.

Merkle Commitment Schemes


We need a **short commitment with local openings**.

An example is a **MERKLE COMMITMENT**.

Let $h: \{0,1\}^{2^\lambda} \rightarrow \{0,1\}^\lambda$ be a hash function.

- **MT[h].Commit** ($m \in \{0,1\}^{l \cdot \lambda}$):

Pairwise hash the message until you obtain one block.



Output:

- Merkle root: $rt := m_{1,8} \in \{0,1\}^\lambda$.
- auxiliary info: $aux := (m, (m_{i,j})_{i,j})$.


- **MT[h].Open** ($aux, Q \subseteq [l]$):

For every $i \in Q$, set $pf_i := (m_v)_{v \in \text{copath}(i)}$. Output $pf := (pf_i)_{i \in Q}$.

- **MT[h].Check** ($rt, Q \subseteq [l], a \in (\{0,1\}^\lambda)^Q, pf = (pf_i)_{i \in Q}$):

For every $i \in Q$, check that pf_i authenticates $a[i]$ for position i relative to rt .

Copath(4) is green below:



lemma: $\mathcal{H} = \{H_\lambda = \{h_\lambda: \{0,1\}^{2^\lambda} \rightarrow \{0,1\}^\lambda\}_{\lambda \in \mathbb{N}}\}$ is collision resistant \rightarrow **MT[H]** is **position-binding**.

$$\forall \text{ efficient } A \quad \Pr_{h \leftarrow H_\lambda} [x \neq y \mid h(x) = h(y) \mid (x,y) \leftarrow A(h)] = \text{negl}(\lambda)$$

$$\forall \text{ efficient } A \quad \Pr_{h \leftarrow H_\lambda} \left[\begin{array}{l} \exists i \in Q_1, Q_2 : a_1[i] \neq a_2[i] \\ \text{MT}[h].\text{Check}(rt, Q_1, a_1, pf_1) = 1 \\ \text{MT}[h].\text{Check}(rt, Q_2, a_2, pf_2) = 1 \end{array} \mid (rt, Q_1, a_1, pf_1) \leftarrow A(h) \right] = \text{negl}(\lambda)$$

Kilian's Protocol

First commit to the PCP string and then locally open it.

$P_{IA}(1^\lambda, x, w)$

- Produce PCP string: $\Pi := P_{PCP}(x, w) \in \Sigma^\ell$.
- Commit to it: $(r, t, aux) := MT[h].Commit(\Pi)$.
[For simplicity here we assume $\log|\Sigma| \leq \lambda$.]
- Deduce PCP query set $Q \subseteq [\ell]$ for $V_{PCP}^\Pi(x; s)$.
- Set answers $a := \Pi[Q] \in \Sigma^Q$.
- Authenticate answers: $pf := MT[h].Open(aux, Q)$.

\xleftarrow{h}

$\xrightarrow{r, t}$

\xleftarrow{s}

$\xrightarrow{Q, a, pf}$

$V_{IA}(1^\lambda, x)$

Sample CRH: $h \leftarrow H_\lambda$.

Sample PCP randomness $s \in \{0, 1\}^r$.

$V_{PCP}^{[Q, a]}(x; s) \stackrel{?}{=} 1$

$MT[h].Check(r, t, Q, a, pf) \stackrel{?}{=} 1$

- round complexity: 2
- communication complexity: $\text{poly}(\lambda) + \lambda + r + q \cdot (\log \ell + \log |\Sigma| + \lambda \cdot \log \ell) \approx \text{poly}(\lambda) + r + q \cdot (\log |\Sigma| + \lambda \cdot \log \ell)$.
- prover time: $\text{time}(P_{PCP}) + \text{time}(V_{PCP}) + O_\lambda(\ell \cdot \log |\Sigma|) \approx pt + O_\lambda(\ell \cdot \log |\Sigma|)$.
- verifier time: $\text{poly}(\lambda) + r + \text{time}(V_{PCP}) + O_\lambda(\log \ell \cdot \log |\Sigma|) \approx vt + O_\lambda(\log \ell \cdot \log |\Sigma|)$.

SECURITY (beyond the scope of this class): a **REWINDING ARGUMENT**, assuming that h is collision resistant (more generally: a position-binding vector commitment)

Bibliography

PCPs

- [BFLS 1991]: [Checking computations in polylogarithmic time](#), by László Babai, Lance Fortnow, Leonid Levin, Mario Szegedy.
- [FGLSS 1996]: [Interactive proofs and the hardness of approximating cliques](#), by Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, Mario Szegedy.
- [AS 1998]: [Probabilistic checking of proofs: a new characterization of NP](#), by Sanjeev Arora, Shmuel Safra.
- [ALMSS 1998]: [Proof verification and the hardness of approximation problems](#), by Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, Mario Szegedy.
- [How computer scientists learned to reinvent the proof](#). Quanta 2022.
- ([▶Scientific revolutions, ToC and PCP](#)), by Avi Wigderson.
- [New short cut found for long math proofs](#), New York Times 1992.

Succinct arguments from PCPs

- [Kilian 1992]: [A note on efficient zero-knowledge proofs and arguments](#), by Joe Kilian.
- [Micali 2000]: [Computationally sound proofs](#), by Silvio Micali.
- [Merkle 1989]: [A certified digital signature](#), by Ralph Merkle.
- [CDGSY 2024]: [Untangling the security of Kilian's protocol: upper and lower bounds](#), by Alessandro Chiesa, Marcel Dall'Agnol, Ziyi Guan, Nicholas Spooner, Eylon Yogev.
- [CY 2024]: [Building cryptographic proofs from hash functions](#), by Alessandro Chiesa, Eylon Yogev.
([▶Video](#))