

Probabilistic Proof Systems – A Survey*

Oded Goldreich

Department of Computer Science and Applied Mathematics
Weizmann Institute of Science, Rehovot, ISRAEL.
E-mail: oded@wisdom.weizmann.ac.il

December 1996

Abstract

Various types of *probabilistic* proof systems have played a central role in the development of computer science in the last decade. In this exposition, we concentrate on three such proof systems — *interactive proofs*, *zero-knowledge proofs*, and *probabilistic checkable proofs* — stressing the essential role of randomness in each of them.

1 Introduction

The glory given to the creativity required to find proofs, makes us forget that it is the less glorified procedure of verification which gives proofs their value. Philosophically speaking, proofs are secondary to the verification procedure; whereas technically speaking, proof systems are defined in terms of their verification procedures.

The notion of a verification procedure assumes the notion of computation and furthermore the notion of efficient computation. This implicit assumption is made explicit in the definition of \mathcal{NP} , in which efficient computation is associated with (deterministic) polynomial-time algorithms.

Definition 1 (NP-proof systems): Let $S \subseteq \{0, 1\}^*$ and $\nu : \{0, 1\}^* \times \{0, 1\}^* \mapsto \{0, 1\}$ be a function so that $x \in S$ if and only if there exists a $w \in \{0, 1\}^*$ such that $\nu(x, w) = 1$. If ν is computable in time bounded by a polynomial in the length of its first argument then we say that S is an NP-set and that ν defines an NP-proof system.

Traditionally, NP is defined as the class of NP-sets. Yet, each such NP-set can be viewed as a proof system. For example, consider the set of satisfiable Boolean formulae. Clearly, a satisfying assignment π for a formula ϕ constitutes an NP-proof for the assertion “ ϕ is

*Parts of the material presented in this survey have appeared in the *Proceedings of the International Congress of Mathematicians 1994*, Birkhäuser Verlag, Basel, 1995, pages 1395–1406.

satisfiable” (the verification procedure consists of substituting the variables of ϕ by the values assigned by π and computing the value of the resulting Boolean expression).

The formulation of NP-proofs restricts the “effective” length of proofs to be polynomial in length of the corresponding assertions (since the running-time of the verification procedure is restricted to be polynomial in the length of the assertion). However, longer proofs may be allowed by padding the assertion with sufficiently many blank symbols. So it seems that NP gives a satisfactory formulation of proof systems (with efficient verification procedures). This is indeed the case if one associates efficient procedures with *deterministic* polynomial-time algorithms. However, we can gain a lot if we are willing to take a somewhat non-traditional step and allow *probabilistic* verification procedures. In particular,

- Randomized and interactive verification procedures, giving rise to *interactive proof systems*, seem much more powerful (i.e., “expressive”) than their deterministic counterparts.
- Such randomized procedures allow the introduction of *zero-knowledge proofs* which are of great theoretical and practical interest.
- NP-proofs can be efficiently transformed into a (redundant) form which offers a trade-off between the number of locations examined in the NP-proof and the confidence in its validity (see *probabilistically checkable proofs*).

In all the abovementioned types of probabilistic proof systems, explicit bounds are imposed on the computational complexity of the verification procedure, which in turn is personified by the notion of a verifier. Furthermore, in all these proof systems, the verifier is allowed to toss coins and rule by statistical evidence. Thus, all these proof systems carry a probability of error; yet, this probability is explicitly bounded and, furthermore, can be reduced by successive application of the proof system.

Notational Conventions

When presenting a proof system, we state all complexity bounds in terms of the length of the assertion to be proven (which is viewed as an input to the verifier). Namely, polynomial-time means time polynomial in the length of this assertion. Note that this convention is consistent with the definition of NP-proofs.

Denote by poly the set of all integer functions bounded by a polynomial and by \log the set of all integer functions bounded by a logarithmic function (i.e., $f \in \log$ iff $f(n) = O(\log n)$).

2 Interactive Proof Systems

In light of the growing acceptability of randomized and distributed computations, it is only natural to associate the notion of efficient computation with probabilistic and interactive polynomial-time computations. This leads naturally to the notion of an interactive proof

system in which the verification procedure is interactive and randomized, rather than being non-interactive and deterministic. Thus, a “proof” in this context is not a fixed and static object but rather a randomized (dynamic) process in which the verifier interacts with the prover. Intuitively, one may think of this interaction as consisting of “tricky” questions asked by the verifier, to which the prover has to reply “convincingly”. The above discussion, as well as the following definition, makes explicit reference to a prover, whereas a prover is only implicit in the traditional definitions of proof systems (e.g., NP-proofs).

2.1 Definition

Loosely speaking, an interactive proof is a game between a computationally bounded verifier and a computationally unbounded prover whose goal is to convince the verifier of the validity of some assertion. Specifically, the verifier is probabilistic polynomial-time. It is required that if the assertion holds then the verifier always accepts (i.e., when interacting with an appropriate prover strategy). On the other hand, if the assertion is false then the verifier must reject with probability at least $\frac{1}{2}$, no matter what strategy is being employed by the prover. A sketch of the formal definition is given in Item (1) below. Item (2) introduces additional complexity measures which can be ignored in first reading.

Definition 2 (Interactive Proofs – IP) [33]:

1. *An interactive proof system for a set S is a two-party game, between a verifier executing a probabilistic polynomial-time strategy (denoted V) and a prover which executes a computationally unbounded strategy (denoted P), satisfying*
 - *Completeness: For every $x \in S$ the verifier V always accepts after interacting with the prover P on common input x .*
 - *Soundness: For every $x \notin S$ and every potential strategy P^* , the verifier V rejects with probability at least $\frac{1}{2}$, after interacting with P^* on common input x .*
2. *For an integer function m , the complexity class $\mathcal{IP}(m(\cdot))$ consists of sets having an interactive proof system in which, on common input x , at most $m(|x|)$ messages are exchanged¹ between the parties. For a set of integer functions, M , we let $\mathcal{IP}(M)$ equal $\cup_{m \in M} \mathcal{IP}(m(\cdot))$. Finally, $\mathcal{IP} \stackrel{\text{def}}{=} \mathcal{IP}(\text{poly})$.*

In Item (1), we have followed the standard definition which specifies strategies for both the verifier and the prover. An alternative presentation only specifies the verifier’s strategy while rephrasing the completeness condition as follows:

there exists a prover strategy P so that, for every $x \in S$, the verifier V always accepts after interacting with P on common input x .

¹ We count the total number of messages exchanged regardless of the direction of communication. Thus, an interactive proof in which the verifier sends a single message answered by a single message of the prover corresponds to $\mathcal{IP}(2)$. Clearly, $\mathcal{NP} \subseteq \mathcal{IP}(1)$, yet the inclusion may be strict since the verifier may toss coins after receiving the prover’s single message.

Arthur-Merlin games² introduced in [5] are a special case of interactive proofs; yet, as shown in [34], this restricted case has essentially³ the same power as the general case previously introduced in [33]. Also, in some sources interactive proofs are defined so that two-sided error probability is allowed; yet, this does not increase their power [25].

2.2 The Role of Randomness

Randomness is essential to the formulation of interactive proofs; if randomness is not allowed (or if it is allowed but zero error is required in the soundness condition) then interactive proof systems collapse to NP-proof systems. The reason being that the prover can predict the verifier's part of the interaction and thus it suffices to let the prover send the full transcript of the interaction and let the verifier check that the interaction is indeed valid. (In case the verifier is not deterministic, the transcript sent by the prover may not match the outcome of the verifier coin tosses.) The moral is that there is no point to interact with predictable parties which are also computationally weaker⁴.

2.3 The Power of Interactive Proofs

A simple example demonstrating the power of interactive proofs follows. Specifically, we present an interactive proof for proving that two graphs are not isomorphic⁵. It is not known whether such a statement can be proven via an NP-proof system.

Construction 1 (Interactive proof system for Graph Non-Isomorphism) [28]:

- **Common Input:** A pair of two graphs, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$. Suppose, without loss of generality, that $V_1 = \{1, 2, \dots, |V_1|\}$, and similarly for V_2 .
- **Verifier's first step (V1):** The verifier selects at random one of the two input graphs, and sends to the prover a random isomorphic copy of this graph. Namely, the verifier selects uniformly $\sigma \in \{1, 2\}$, and a random permutation π from the set of permutations over the vertex set V_σ . The verifier constructs a graph with vertex set V_σ and edge set

$$E \stackrel{\text{def}}{=} \{ \{ \pi(u), \pi(v) \} : \{ u, v \} \in E_\sigma \}$$

and sends (V_σ, E) to the prover.

²In Arthur-Merlin games, the verifier must send the outcome of any coin it tosses (and thus need not send any other information).

³Here and in the next sentence, not only \mathcal{IP} remains invariant under the various definitions, but also $\mathcal{IP}(m(\cdot))$, for every integer function satisfying $m(n) \geq 2$ for every n .

⁴This moral represents the prover's point of view. Certainly, from the verifier's point of view it is beneficial to interact with the prover, since it is computationally stronger.

⁵Two graphs, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, are called *isomorphic* if there exists a 1-1 and onto mapping, ϕ , from the vertex set V_1 to the vertex set V_2 so that $\{u, v\} \in E_1$ if and only if $\{\phi(u), \phi(v)\} \in E_2$. The ("edge preserving") mapping ϕ , if existing, is called an *isomorphism* between the graphs.

- *Motivating Remark: If the input graphs are non-isomorphic, as the prover claims, then the prover should be able to distinguish (not necessarily by an efficient algorithm) isomorphic copies of one graph from isomorphic copies of the other graph. However, if the input graphs are isomorphic then a random isomorphic copy of one graph is distributed identically to a random isomorphic copy of the other graph.*
- *Prover's step: Upon receiving a graph, $G' = (V', E')$, from the verifier, the prover finds a $\tau \in \{1, 2\}$ so that the graph G' is isomorphic to the input graph G_τ . (If both $\tau = 1, 2$ satisfy the condition then τ is selected arbitrarily. In case no $\tau \in \{1, 2\}$ satisfies the condition, τ is set to 0). The prover sends τ to the verifier.*
- *Verifier's second step (V2): If the message, τ , received from the prover equals σ (chosen in Step V1) then the verifier outputs 1 (i.e., accepts the common input). Otherwise the verifier outputs 0 (i.e., rejects the common input).*

The verifier's strategy presented above is easily implemented in probabilistic polynomial-time. We do not know of a probabilistic polynomial-time implementation of the prover's strategy, but this is not required. The motivating remark justifies the claim that Construction 1 constitutes an interactive proof system for the set of pairs of non-isomorphic graphs. Recall that the latter is a coNP-set (not known to be in \mathcal{NP}).

Interactive proofs are powerful enough to prove *any* coNP assertion (e.g., that a graph is not 3-colorable) [41]. Furthermore, the class of sets having interactive proof systems coincides with the class of sets that can be decided using a polynomial amount of work-space [49].

Theorem 1 [41, 49]: $\mathcal{IP} = \mathcal{PSPACE}$.

Recall that it is widely believed that $\mathcal{NP} \subset \mathcal{PSPACE}$. Thus, under this conjecture, interactive proofs are more powerful than NP-proofs.

Concerning the finer structure of the IP hierarchy it is known that this hierarchy has a “linear speed-up” property [8]. Namely, for every integer function, f , so that $f(n) \geq 2$ for all n , the class $\mathcal{IP}(O(f(\cdot)))$ collapses to the class $\mathcal{IP}(f(\cdot))$. In particular, $\mathcal{IP}(O(1))$ collapses to $\mathcal{IP}(2)$. It is conjectured that coNP is *not* contained in $\mathcal{IP}(2)$, and consequently that interactive proofs with unbounded number of message exchanges are more powerful than interactive proofs in which only a bounded (i.e., constant) number of messages are exchanged. Still, the class $\mathcal{IP}(2)$ contains sets not known to be in \mathcal{NP} ; e.g., Graph Non-Isomorphism (as shown above).

2.4 How Powerful Should the Prover be?

Assume that a set S is in \mathcal{IP} . This means that there is a verifier V that can be convinced to accept any input in S but cannot be convinced to accept any input not in S (except with small probability). One may ask how powerful should a prover be so that it can convince the verifier V to accept any input in S . More interestingly, considering all possible verifiers which give rise to interactive proof systems for S , what is the minimum power required from

a prover which satisfies the completeness requirement with respect to one of these verifiers? We stress that, unlike the case of computationally-sound proof systems (see Sec. 5), we do not restrict the power of the prover in the soundness condition but rather consider the minimum complexity of provers meeting the completeness condition. Specifically, we are interested in *relatively efficient* provers which meet the completeness condition. The term ‘relatively efficient prover’ has been given three different interpretations.

1. A prover is considered *relatively efficient* if, when given an auxiliary input (in addition to the common input in S), it works in (probabilistic) polynomial-time. Specifically, in case $S \in \mathcal{NP}$, the auxiliary input maybe an NP-proof that the common input is in the set⁶. This interpretation is adequate and in fact crucial for applications in which such an auxiliary input is available to the otherwise-polynomial-time parties. Typically, such auxiliary input is available in cryptographic applications in which parties wish to prove in (zero-knowledge) that they have conducted some computation correctly resulting in some string x . In these cases the NP-proof is just the transcript of the procedure by which x has been computed and thus the auxiliary input is available to the proving party. See [28].
2. A prover is considered *relatively efficient* if it can be implemented by a probabilistic polynomial-time oracle machine with oracle access to the set S itself. (Note that the prover in Construction 1 has this property.) This interpretation generalizes the notion of self-reducibility of NP-sets. (By self-reducibility of an NP-set we mean that the search problem of finding an NP-witness is polynomial-time reducible to deciding membership in the set.) See [12].
3. A prover is considered *relatively efficient* if it can be implemented by a probabilistic machine which runs in time which is polynomial in the deterministic complexity of the set. This interpretation relates the difficulty of convincing a “lazy verifier” to the complexity of finding the truth alone. Hence, in contrast to the first interpretation which is adequate in settings where assertions are generated along with their NP-proofs, the current interpretation is adequate in settings in which the prover is given only the assertion and has to find a proof to it by itself (before trying to convince a lazy verifier of its validity). See [43].

3 Zero-Knowledge Proof Systems

Zero-knowledge proofs, introduced in [33], are central to cryptography. Furthermore, zero-knowledge proofs are very intriguing from a conceptual point of view, since they exhibit an extreme contrast between being convinced of the validity of a statement and learning anything in addition while receiving such a convincing proof. Namely, zero-knowledge proofs have the remarkable property of being both convincing while yielding nothing to the verifier, beyond the fact that the statement is valid. Formally, the fact that “nothing is

⁶Still, even in this case the interactive proof need not consist of the prover sending the auxiliary input to the verifier; e.g., an alternative procedure may allow the prover to be zero-knowledge (see Construction 2).

gained by the interaction” is captured by stating that whatever the verifier can efficiently compute after interacting with a zero-knowledge prover, can be efficiently computed from the assertion itself without interacting with anyone.

3.1 A Sample Definition

Zero-knowledge is a property of some interactive proof systems, or more accurately of some specified prover strategies. The formulation of the zero-knowledge condition considers two ensembles of probability distributions, each ensemble associates a probability distribution to each valid assertion. The first ensemble represents the output distribution of the verifier after interacting with the specified prover strategy P , where the verifier is not necessarily employing the specified strategy (i.e., V)—but rather any efficient strategy. The second ensemble represents the output distribution of some probabilistic polynomial-time algorithm (which does not interact with anyone). The basic paradigm of zero-knowledge asserts that for every ensemble of the first type there exist a “similar” ensemble of the second type. The specific variants differ by the interpretation given to ‘similarity’. The most strict interpretation, leading to *perfect zero-knowledge*, is that similarity means equality. Namely,

Definition 3 (perfect zero-knowledge) [33]: *A prover strategy, P , is said to be perfect zero-knowledge over a set S if for every probabilistic polynomial-time verifier strategy, V^* , there exists a probabilistic polynomial-time algorithm, M^* , such that*

$$(P, V^*)(x) = M^*(x), \quad \text{for every } x \in S$$

where $(P, V^*)(x)$ is a random variable representing the output of verifier V^* after interacting with the prover P on common input x , and $M^*(x)$ is a random variable representing the output of machine M^* on input x .

A somewhat more relaxed interpretation, leading to *almost-perfect zero-knowledge*, is that similarity means statistical closeness (i.e., negligible difference between the ensembles). The most liberal interpretation, leading to the standard usage of the term zero-knowledge (and sometimes referred to as *computational zero-knowledge*), is that similarity means computational indistinguishability (i.e., failure of any efficient procedure to tell the two ensembles apart). Since the notion of computational indistinguishability is a fundamental one, it is indeed in place to present a definition of it.

Definition 4 (computational indistinguishability) [32, 50]: *An integer function, f , is called negligible if for every positive polynomial p and all sufficiently large n , it holds that $f(n) < \frac{1}{p(n)}$. (Thus, multiplying a negligible function by any fixed polynomial yields a negligible function.)*

Two probability ensembles, $\{A_x\}_{x \in S}$ and $\{B_x\}_{x \in S}$, are indistinguishable by an algorithm D if

$$d(n) \stackrel{\text{def}}{=} \max_{x \in S \cap \{0,1\}^n} \{|\text{prob}(D(A_x)=1) - \text{Prob}(D(B_x)=1)|\}$$

is a negligible function. The ensembles $\{A_x\}_{x \in S}$ and $\{B_x\}_{x \in S}$ are computationally indistinguishable if they are indistinguishable by every probabilistic polynomial-time algorithm.

The definitions presented above are a simplified version of the actual definitions. For example, in order to guarantee that zero-knowledge is preserved under sequential composition it is necessary to slightly augment the definitions. For details see [30].

3.2 The Power of Zero-Knowledge

A simple example, demonstrating the power of zero-knowledge proofs, follows. Specifically, we will present a simple zero-knowledge proof for proving that a graph is 3-colorable⁷. The interactive proof will be described using “boxes” in which information can be hidden and later revealed. Such “boxes” can be implemented using one-way functions (see below).

Construction 2 (Zero-knowledge proof of 3-colorability) [28]:

- Common Input: A simple graph $G = (V, E)$.
- Prover’s first step: Let ψ be a 3-coloring of G . The prover selects a random permutation, π , over $\{1, 2, 3\}$, and sets $\phi(v) \stackrel{\text{def}}{=} \pi(\psi(v))$, for each $v \in V$. Hence, the prover forms a random relabelling of the 3-coloring ψ . The prover sends the verifier a sequence of $|V|$ locked and nontransparent boxes so that the v^{th} box contains the value $\phi(v)$;
- Verifier’s first step: The verifier uniformly selects an edge $\{u, v\} \in E$, and sends it to the prover;
- Motivating Remark: The verifier asks to inspect the colors of vertices u and v ;
- Prover’s second step: The prover sends to the verifier the keys to boxes u and v ;
- Verifier’s second step: The verifier opens boxes u and v , and accepts if and only if they contain two different elements in $\{1, 2, 3\}$;

The verifier strategy presented above is easily implemented in probabilistic polynomial-time. The same holds with respect to the prover’s strategy, provided it is given a 3-coloring of G as auxiliary input. Clearly, if the input graph is 3-colorable then the prover can cause the verifier to accept always. On the other hand, if the input graph is not 3-colorable then any contents put in the boxes must be invalid on at least one edge, and consequently the verifier will reject with probability at least $\frac{1}{|E|}$. Hence, the above game exhibits a non-negligible gap in the accepting probabilities between the case of 3-colorable graphs and the case of non-3-colorable graphs. To increase the gap, the game may be repeated sufficiently many times (of course, using independent coin tosses in each repetition). The zero-knowledge property follows easily, in this abstract setting, since one can simulate the real interaction by placing a random pair of different colors in the boxes indicated by the verifier. This indeed demonstrates that the verifier learns nothing from the interaction

⁷ A graph $G = (V, E)$ is said to be 3-colorable if there exists a function $\pi : V \mapsto \{1, 2, 3\}$ so that $\pi(v) \neq \pi(u)$ for every $\{u, v\} \in E$. Such a function, π , is called a 3-coloring of the graph.

(since it expects to see a random pair of different colors and indeed this is what it sees). We stress that this simple argument is not possible in the digital implementation since the boxes are not totally ineffectuated by their contents (but are rather effected, yet in an indistinguishable manner).

As stated above, the “boxes” need to be implemented digitally, and this is done using an adequately defined “commitment scheme”. Loosely speaking, such a scheme is a two phase game between a sender and a receiver so that after the first phase the sender is “committed” to a value and yet, at this stage, it is infeasible for the receiver to find out the committed value. The committed value will be revealed to the receiver in the second phase and it is guaranteed that the sender cannot reveal a value other than the one committed. Such commitment schemes can be implemented assuming the existence of one-way functions (i.e., loosely speaking, functions that are easy to compute but hard to invert, such as the multiplication of two large primes) [44, 37].

Using the fact that 3-colorability is NP-complete, one gets zero-knowledge proofs for any NP-set.

Theorem 2 [28]: *Assuming the existence of one-way functions, any NP-proof can be efficiently transformed into a (computational) zero-knowledge interactive proof.*

Theorem 2 has a dramatic effect on the design of cryptographic protocols (cf., [28, 29]). In a different vein and for the sake of elegance, we mention that, using further ideas and under the same assumption, any interactive proof can be efficiently transformed into a zero-knowledge one [38, 13].

The above results may be contrasted with the results regarding the complexity of *almost-perfect* zero-knowledge proof systems; namely, that almost-perfect zero-knowledge proof systems exist only for sets in $\mathcal{IP}(2) \cap \text{coIP}(2)$ [23, 2], and thus are unlikely to exist for all NP-sets. Also, a recent result seems to indicate that one-way functions are essential for the existence of zero-knowledge proofs for “hard” sets (i.e., sets which cannot be decided in average polynomial-time) [45].

3.3 The Role of Randomness

Again, randomness is essential to all the above mentioned (positive) results. Namely, if either verifier or prover is required to be deterministic then only BPP-sets can be proven in a zero-knowledge manner [30]. However, BPP-sets have trivial zero-knowledge proofs in which the prover sends nothing and the verifier just test the validity of the assertion by itself⁸. Thus, randomness is essential to the usefulness of zero-knowledge proofs.

⁸ Actually, this is slightly inaccurate since the resulting “interactive proof” may have two-sided error, whereas we have required interactive proofs to have only one-sided error. Yet, since the error can be made negligible by successive repetitions this issue is insignificant. Alternatively, one can use ideas in [25] to eliminate the error by letting the prover send some random-looking help.

4 Probabilistically Checkable Proof Systems

When viewed in terms of an interactive proof system, the probabilistically checkable proof setting consists of a prover which is memoryless. Namely, one can think of the prover as being an oracle and of the messages sent to it as being queries. A more appealing interpretation is to view the probabilistically checkable proof setting as an alternative way of generalizing \mathcal{NP} . Instead of receiving the entire proof and conducting a deterministic polynomial-time computation (as in the case of \mathcal{NP}), the verifier may toss coins and query the proof only at location of its choice. Potentially, this allows the verifier to utilize very long proofs (i.e., of super-polynomial length) or alternatively examine very few bits of an NP-proof.

4.1 Definition

Loosely speaking, a probabilistically checkable proof system consists of a probabilistic polynomial-time verifier having access to an oracle which represents a proof in redundant form. Typically, the verifier accesses only few of the oracle bits, and these bit positions are determined by the outcome of the verifier's coin tosses. Again, it is required that if the assertion holds then the verifier always accepts (i.e., when given access to an adequate oracle); whereas, if the assertion is false then the verifier must reject with probability at least $\frac{1}{2}$, no matter which oracle is used. The basic definition of the PCP setting is given in Item (1) below. Yet, the complexity measures introduced in Item (2) are of key importance for the subsequent discussions, and should not be ignored.

Definition 5 (Probabilistic Checkable Proofs – PCP):

1. A probabilistic checkable proof system (pcp) for a set S is a probabilistic polynomial-time oracle machine (called verifier), denoted V , satisfying
 - Completeness: For every $x \in S$ there exists an oracle set π_x so that V , on input x and access to oracle π_x , always accepts x .
 - Soundness: For every $x \notin S$ and every oracle set π , machine V , on input x and access to oracle π , rejects x with probability at least $\frac{1}{2}$.
2. Let r and q be integer functions. The complexity class $\mathcal{PCP}(r(\cdot), q(\cdot))$ consists of sets having a probabilistic checkable proof system in which the verifier, on any input of length n , makes at most $r(n)$ coin tosses and at most $q(n)$ oracle queries. We stress that here, as usual in complexity theory, the oracle answers are always binary (i.e., either 0 or 1). For sets of integer functions, R and Q , we let $\mathcal{PCP}(R, Q)$ equal $\bigcup_{r \in R, q \in Q} \mathcal{PCP}(r(\cdot), q(\cdot))$.

The above model was suggested in [24] and shown related to a multi-prover model introduced previously in [14]. The fine complexity measures were introduced and motivated in [20], and further advocated in [4]. A related model was presented in [7], stressing the applicability to program checking.

We stress that the oracle π_x in a pcP system constitutes a proof in the standard mathematical sense⁹. Yet, this oracle has the extra property of enabling a lazy verifier, to toss coins, take its chances and “assess” the validity of the proof without reading all of it (but rather by reading a tiny portion of it).

4.2 The Power of Probabilistically Checkable Proofs

Clearly, $\mathcal{PCP}(\text{poly}, 0)$ equals coRP , whereas $\mathcal{PCP}(0, \text{poly})$ equals \mathcal{NP} . It is easy to prove an upper bound on the non-deterministic time complexity of sets in the PCP hierarchy. In particular,

Proposition 1 : $\mathcal{PCP}(\log, \text{poly})$ is contained in \mathcal{NP} .

These upper bounds turn out to be tight, but proving this is much more difficult (to say the least). The following result is a culmination of a sequence of great works [6, 7, 20, 4, 3].¹⁰

Theorem 3 : \mathcal{NP} is contained in $\mathcal{PCP}(\log, O(1))$.

Thus, probabilistically checkable proofs in which the verifier tosses only logarithmically many coins and makes only a constant number of queries exist for every set in the complexity class \mathcal{NP} . It follows that NP-proofs can be transformed into NP-proofs which offer a trade-off between the portion of the proof being read and the confidence it offers. Specifically, if the verifier is willing to tolerate an error probability of ϵ then it suffices to let it examine $O(\log(1/\epsilon))$ bits of the (transformed) NP-proof. These bit locations need to be selected at random.

The characterization of \mathcal{NP} in terms of probabilistically checkable proofs plays a central role in recent developments concerning the difficulty of approximation problems (cf., [20, 3, 42, 11] and [35, 36]). To demonstrate this relationship, we first note that Theorem 3 can be rephrased without mentioning the class \mathcal{PCP} altogether. Instead, a new type of polynomial-time reductions, which we call *amplifying*, emerges.

Theorem 4 (Theorem 3 — Rephrased): *There exists a constant $\epsilon > 0$, and a polynomial-time computable function f , mapping the set of 3CNF formulae¹¹ to itself so that*

- *As usual, f maps satisfiable 3CNF formulae to satisfiable 3CNF formulae; and*

⁹Jumping ahead, the oracles in pcP systems characterizing \mathcal{NP} have the property of being NP proofs themselves.

¹⁰The sequence has started with the characterization of $\mathcal{PCP}(\text{poly}, \text{poly})$ as equal non-deterministic exponential-time [6], and continued with its scaled-down in [7, 20] which led to the $\mathcal{NP} \subseteq \mathcal{PCP}(\text{polylog}, \text{polylog})$ result of [20]. The first PCP-characterization of \mathcal{NP} , by which $\mathcal{NP} = \mathcal{PCP}(\log, \log)$, has appeared in [4] and the cited result was obtained in [3]. This sequence of works, directly related to the stated theorem, was built on and inspired by works from various settings such as interactive proofs [41, 49, 22], program-checking [16, 26, 48], and private computation with oracles [9]. The constant (number of queries) in Theorem 3 has been subsequently improved and is currently 9; cf., [36].

¹¹A 3CNF formula is a Boolean formula consisting of a conjunction of clauses, where each clause is a disjunction of upto 3 literals. (A literal is variable or its negation.).

- *f maps non-satisfiable 3CNF formulae to (non-satisfiable) 3CNF formulae for which every truth assignment satisfies at most a $1 - \epsilon$ fraction of the clauses.*

The function f is called an amplifying reduction.

proof sketch (Thm. 3 \Rightarrow Thm. 4): We start by considering a pcg system for 3SAT, and use the fact that the pcg system given by the proof of Theorem 3 is non-adaptive (i.e., the queries are determined as a function of the input and the random-tape – and do not depend on answers to previous queries).¹² Next, we associate the bits of the oracle with Boolean variables and introduce a (constant size) Boolean formula for each possible outcome of the sequence of $O(\log n)$ coin tosses, describing whether the verifier would have accepted given this outcome. Finally, using auxiliary variables, we convert each of these formulae into a 3CNF formula and obtain (as the output of the reduction) the conjunction of all these polynomially many clauses. \square

It is also easy to see that Theorem 4 implies Theorem 3: Given a reduction as in Thm. 4, we construct a pcg system for 3SAT by letting the verifier select a clause uniformly among the clauses of the reduced formula, and make three queries corresponding to the three variables in it. This yields a proof system with soundness error bounded by $1 - \epsilon$. Reducing the error by $O(1/\epsilon)$ successive applications of this proof system, we obtain Thm. 3.

As an immediate corollary to the formulation of Theorem 4 one concludes that it is NP-Hard to distinguish satisfiable 3CNF formulae from 3CNF formulae for which no truth assignment satisfies at least a $1 - \epsilon$ fraction of the clauses (as otherwise, using the reduction, one may decide membership in 3SAT). In general, probabilistic checkable proof systems for \mathcal{NP} yield strong non-approximability results for various classical optimization problems. In particular, quite *tight* non-approximability results have been shown for MaxClique (cf., [35]), Chromatic Number (cf., [21]), Set Cover (cf., [19]), and Max-Exact-3SAT (cf., [36]).

4.3 The Role of Randomness

No trade-off between the number of bits examined and the confidence is possible if one requires the verifier to be deterministic. In particular, $\mathcal{PCP}(0, q(\cdot))$ contains only sets that are decidable by a deterministic algorithm of running time $2^{q(n)} \cdot \text{poly}(n)$. It follows that $\mathcal{PCP}(0, \log) = \mathcal{P}$. Furthermore, since it is unlikely that all NP-sets can be decided by (deterministic) algorithms of running time, say, $2^n \cdot \text{poly}(n)$, it follows that $\mathcal{PCP}(0, n)$ is unlikely to contain \mathcal{NP} .

5 Other Probabilistic Proof Systems

In this section, we shortly review some variants on the basic model of interactive proofs. These variants include models in which the prover is restricted in its choice of strategy, a model in which the prover-verifier interaction is restricted, and a model in which one proves “knowledge” of facts rather than their *validity*.

¹² Actually, it is not essential to use this fact, since one can easily convert any adaptive system into a non-adaptive one while incurring an exponential blowup in the query complexity (which in our case is a constant).

5.1 Restricting the Prover's Strategy

We stress that the restrictions discussed here refer to the strategies employed by the prover both in case it tries to prove valid assertions (i.e., the completeness condition) and in case it tries to fool the verifier to believe false statements (i.e., the soundness condition). Thus, the validity of the verifier decision (concerning false statements) depends on whether this restriction (concerning “cheating” prover strategies) really holds. The reason to consider these restricted models is that they enable to achieve results which are not possible in the general model of interactive proofs (cf., [14, 17, 39, 43]). We consider restrictions of two types: computational or physical.

We start with a physical restriction. In the so-called *multi-prover interactive proof* model, denoted MIP (cf., [14]), the prover is split into several (say, two) entities and the restriction (or assumption) is that these entities cannot interact with each other. Actually, the formulation allows them to coordinate their strategies prior to interacting with the verifier¹³ but it is crucial that they don't exchange messages among themselves while interacting with the verifier. The multi-prover model is reminiscent of the common police procedure of isolating collaborating suspects and interrogating each of them separately. On the other hand, the multi-prover model is related to the PCP model [24]. Interestingly, the multi-prover model allows to present (perfect) zero-knowledge proofs for all NP-sets, without relying on any computational assumptions [14]. Furthermore, these proofs can be made very efficient in terms of communication complexity [18].

We now turn to computational restrictions. Since the effect of this restriction is more noticable in the soundness condition, we refer to these proof systems as being *computationally-sound*. Two variants have been suggested. In *argument* systems [17], the prover strategy is restricted to be probabilistic polynomial-time with auxiliary input (analogously to item (1) in Sec. 2.4). In *CS-proofs* [43], the prover strategy is restricted to be probabilistic and run in time polynomial in the time required to validate the assertion (analogously to item (3) in Sec. 2.4). Interestingly, computationally-sound interactive proofs can be much more communication-efficient than (regular) interactive proofs; cf. [39, 43, 27].

5.2 Non-Interactive Zero-Knowledge Proofs

Actually the term “non-interactive” is somewhat misleading. The model, introduced in [15], consists of three entities: a prover, a verifier and a uniformly selected sequence of bits (which can be thought of as being selected by a trusted third party). Both verifier and prover can read the random sequence, and each can toss additional coins. The interaction consists of a single message sent from the prover to the verifier, who then is left with the decision (whether to accept or not). Based on some reasonable complexity assumptions, one may construct non-interactive zero-knowledge proof systems for every NP-set (cf., [15, 22, 40]).

¹³This is implicit in the universal quantifier used in the soundness condition.

5.3 Proofs of Knowledge

The concept of a proof of knowledge, introduced in [33], is very appealing; yet, its precise formulation is much more complex than one may expect (cf. [10]). Loosely speaking, a knowledge-verifier for a relation R guarantees the existence of a “knowledge extractor” that on input x and access to any interactive machine P^* outputs a y , so that $(x, y) \in R$, within complexity related to the probability that the verifier accepts x when interacting with P^* . By convincing such a knowledge-verifier, on common input x , one proves that he knows a y so that $(x, y) \in R$. It can be shown that the protocol which results by successively applying Construction 2 sufficiently many time constitutes a “proof of knowledge” of a 3-coloring of the input graph.

5.4 Knowledge Complexity

Zero-knowledge is the lowest level of a knowledge-complexity hierarchy which quantifies the “knowledge revealed in an interaction” [33]. *Knowledge complexity* may be defined as the minimum number of oracle-queries required in order to (efficiently) simulate an interaction with the prover (cf. [31]). Results linking two different variants of this measure to other complexity measures are given in [1, 47], respectively.

Acknowledgement

I am grateful to Shafi Goldwasser for suggesting the essential role of randomness as the unifying theme for this exposition. Thanks also to Leonid Levin, Dana Ron, Madhu Sudan and Uri Zwick for commenting on earlier versions of this survey.

References

- [1] W. Aiello, M. Bellare and R. Venkatesan. Knowledge on the Average – Perfect, Statistical and Logarithmic. In *27th STOC*, pages 469–478, 1995.
- [2] W. Aiello and J. Hastad. Perfect Zero-Knowledge Languages can be Recognized in Two Rounds. In *28th FOCS*, pages 439–448, 1987.
- [3] S. Arora, C. Lund, R. Motwani, M. Sudan and M. Szegedy. Proof Verification and Intractability of Approximation Problems. In *33rd FOCS*, pages 14–23, 1992.
- [4] S. Arora and S. Safra. Probabilistic Checkable Proofs: A New Characterization of NP. In *33rd FOCS*, pages 1–13, 1992.
- [5] L. Babai. Trading Group Theory for Randomness. In *17th STOC*, pages 421–420, 1985.
- [6] L. Babai, L. Fortnow, and C. Lund. Non-Deterministic Exponential Time has Two-Prover Interactive Protocols. *Computational Complexity*, Vol. 1, No. 1, pages 3–40, 1991. Preliminary version in *31st FOCS*, 1990.

- [7] L. Babai, L. Fortnow, L. Levin, and M. Szegedy. Checking Computations in Polylogarithmic Time. In *23rd STOC*, pages 21–31, 1991.
- [8] L. Babai and S. Moran. Arthur-Merlin Games: A Randomized Proof System and a Hierarchy of Complexity Classes. *JCSS*, Vol. 36, pp. 254–276, 1988.
- [9] D. Beaver and J. Feigenbaum. Hiding Instances in Multioracle Queries. In *7th STACS*, Springer Verlag, LNCS Vol. 415, pages 37–48, 1990.
- [10] M. Bellare and O. Goldreich. On Defining Proofs of Knowledge. In *Crypto92*, Springer Verlag, LNCS Vol. 740, pages 390–420, 1992.
- [11] M. Bellare, O. Goldreich and M. Sudan. Free Bits, PCPs and Non-Approximability – Towards Tight Results. In *36th FOCS*, pages 422–431, 1995.
- [12] M. Bellare and S. Goldwasser. The Complexity of Decision versus Search. *SIAM Journal on Computing*, Vol. 23, pages 97–119, 1994.
- [13] M. Ben-Or, O. Goldreich, S. Goldwasser, J. Håstad, J. Kilian, S. Micali and P. Rogaway. Everything Provable is Probable in Zero-Knowledge. In *Crypto88*, Springer Verlag, LNCS Vol. 403, pages 37–56, 1990.
- [14] M. Ben-Or, S. Goldwasser, J. Kilian and A. Wigderson. Multi-Prover Interactive Proofs: How to Remove Intractability Assumptions. In *20th STOC*, pages 113–131, 1988.
- [15] M. Blum, P. Feldman and S. Micali. Non-Interactive Zero-Knowledge and its Applications. In *20th STOC*, pages 103–112, 1988.
- [16] M. Blum, M. Luby and R. Rubinfeld. Self-Testing/Correcting with Applications to Numerical Problems. *JCSS*, Vol. 47, No. 3, pages 549–595, 1993.
- [17] G. Brassard, D. Chaum and C. Crépeau. Minimum Disclosure Proofs of Knowledge. *JCSS*, pages 156–189, 1988. Extended abstract, by Brassard and Crépeau, in *27th FOCS*, 1986.
- [18] C. Dwork, U. Feige, J. Kilian, M. Naor and S. Safra, Low Communication Perfect Zero Knowledge Two Provers Proof Systems. In *Crypto92*, Springer Verlag, LNCS Vol. 740, pages 215–227, 1992.
- [19] U. Feige. A Threshold of $\ln n$ for Approximating Set Cover. In *28th STOC*, pages 314–318, 1996.
- [20] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Approximating Clique is almost NP-complete. In *32nd FOCS*, pages 2–12, 1991.
- [21] U. Feige and J. Kilian. Zero knowledge and the chromatic number. In *11th IEEE Conference on Computational Complexity*, pages 278–287, 1996.

- [22] U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero knowledge proofs based on a single random string. In *31st FOCS*, pages 308–317, 1990.
- [23] L. Fortnow, The Complexity of Perfect Zero-Knowledge. In *19th STOC*, pages 204–209, 1987.
- [24] L. Fortnow, J. Rompel and M. Sipser. On the Power of Multi-Prover Interactive Protocols. In *Proc. 3rd IEEE Symp. on Structure in Complexity Theory*, pages 156–161, 1988.
- [25] M. Furer, O. Goldreich, Y. Mansour, M. Sipser, and S. Zachos, “On Completeness and Soundness in Interactive Proof Systems”, *Advances in Computing Research: a research annual*, Vol. 5 (Randomness and Computation, S. Micali, ed.), pp. 429–442, 1989.
- [26] P. Gemmell, R. Lipton, R. Rubinfeld, M. Sudan, and A. Wigderson. Self-Testing/Correcting for Polynomials and for Approximate Functions. In *23th STOC*, pages 32–42, 1991.
- [27] O. Goldreich and J. Håstad. On the Message Complexity of Interactive Proof Systems. Available as TR96-018 of *ECCC*, <http://www.eccc.uni-trier.de/eccc/>, 1996.
- [28] O. Goldreich, S. Micali and A. Wigderson. Proofs that Yield Nothing but their Validity or All Languages in NP Have Zero-Knowledge Proof Systems. *JACM*, Vol. 38, No. 1, pages 691–729, 1991. Extended abstract in *27th FOCS*, 1986.
- [29] O. Goldreich, S. Micali and A. Wigderson. How to Play any Mental Game or a Completeness Theorem for Protocols with Honest Majority. In *19th STOC*, pages 218–229, 1987.
- [30] O. Goldreich and Y. Oren. Definitions and Properties of Zero-Knowledge Proof Systems. *Journal of Cryptology*, Vol. 7, No. 1, pages 1–32, 1994.
- [31] O. Goldreich and E. Petrank. Quantifying Knowledge Complexity. In *32nd FOCS*, pp. 59–68, 1991.
- [32] S. Goldwasser and S. Micali. Probabilistic Encryption. *JCSS*, Vol. 28, No. 2, pages 270–299, 1984. Extended abstract in *14th STOC*, 1982.
- [33] S. Goldwasser, S. Micali and C. Rackoff. The Knowledge Complexity of Interactive Proof Systems. *SIAM Journal on Computing*, Vol. 18, pages 186–208, 1989. Extended abstract in *17th STOC*, 1985.
- [34] S. Goldwasser and M. Sipser. Private Coins versus Public Coins in Interactive Proof Systems. In *18th STOC*, pages 59–68, 1986.
- [35] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *37th FOCS*, pages 627–636, 1996.

- [36] J. Håstad. Getting optimal in-approximability results. Unpublish manuscript, June 1996. (Revised October 1996.)
- [37] J. Håstad, R. Impagliazzo, L.A. Levin and M. Luby. Construction of Pseudorandom Generator from any One-Way Function. Manuscript, 1993. See preliminary versions by Impagliazzo et. al. in *21st STOC* and Håstad in *22nd STOC*.
- [38] R. Impagliazzo and M. Yung. Direct Zero-Knowledge Computations. In *Crypto87*, Springer Verlag, LNCS Vol. 293, pages 40–51, 1987.
- [39] J. Kilian. A Note on Efficient Zero-Knowledge Proofs and Arguments. In *24th STOC*, pages 723–732, 1992.
- [40] J. Kilian and E. Petrank. An Efficient Non-Interactive Zero-Knowledge Proof System for NP with General Assumptions. To appear in the *Journal of Cryptography*. Available as TR95-038 of *ECCC*, <http://www.eccc.uni-trier.de/eccc/>, 1995.
- [41] C. Lund, L. Fortnow, H. Karloff, and N. Nisan. Algebraic Methods for Interactive Proof Systems. *JACM*, Vol. 39, No. 4, pages 859–868, 1992. Preliminary version in *31st FOCS*, 1990.
- [42] C. Lund and M. Yannakakis. On the Hardness of Approximating Minimization Problems, In *25th STOC*, pages 286–293, 1993.
- [43] S. Micali. CS Proofs. In *35th FOCS*, pages 436–453, 1994.
- [44] M. Naor. Bit Commitment using Pseudorandom Generators. *Journal of Cryptology*, Vol. 4, pages 151–158, 1991.
- [45] R. Ostrovsky and A. Wigderson. One-Way Functions are essential for Non-Trivial Zero-Knowledge, In *Proc. 2nd Israel Symp. on Theory of Computing and Systems (ISTCS93)*, IEEE Computer Society Press, pages 3–17, 1993.
- [46] C. H. Papadimitriou and M. Yannakakis. Optimization, Approximation, and Complexity Classes. In *20th STOC*, pages 229–234, 1988.
- [47] E. Petrank and G. Tardos. On the Knowledge Complexity of NP. In *37th FOCS*, pages 494–503, 1996.
- [48] R. Rubinfeld and M. Sudan. Robust Characterizations of Polynomials with Applications to Program Checking. *SIAM J. of Computing*, Vol. 25, No. 2, pages 252–271, 1996. Preliminary version in *3rd SODA*, 1992.
- [49] A. Shamir. $IP=PSPACE$. *JACM*, Vol. 39, No. 4, pages 869–877, 1992. Preliminary version in *31st FOCS*, 1990.
- [50] A.C. Yao. Theory and Application of Trapdoor Functions. In *23rd FOCS*, pages 80–91, 1982.