A One-Round, Two-Prover, Zero-Knowledge Protocol for NP

Dror Lapidot
Department of Applied Math.
The Weizmann Institute of Science
Rehovot, Israel
drorl@wisdom.bitnet

Adi Shamir
Department of Applied Math.
The Weizmann Institute of Science
Rehovot, Israel
shamir@wisdom.bitnet

Abstract

The model of zero knowledge multi prover interactive proofs was introduced by Ben-Or, Goldwasser, Kilian and Wigderson. A major open problem associated with these protocols is whether they can be executed in parallel. A positive answer was claimed by Fortnow, Rompel and Sipser, but its proof was later shown to be flawed by Fortnow who demonstrated that the probability of cheating in n independent parallel rounds can be exponentially higher than the probability of cheating in n independent sequential rounds. In this paper we use refined combinatorial arguments to settle this problem by proving that the probability of cheating in a parallelized BGKW protocol is at most $1/2^{n/9}$, and thus every problem in NP has a one-round two prover protocol which is perfectly zero knowledge under no cryptographic assumptions.

1 Introduction

In [GMW] Goldreich, Micali and Wigderson show that under the assumption that one way functions exist, every NP language has a computational zero knowledge interactive proof system. They prove it by giving a sequential zero knowledge protocol for an NP-complete statement. Results in [F2] and [BHZ] imply that if perfect zero-knowledge interactive proof-systems for NP exist (i.e. which do not rely on the fact that the verifier is polynomial time bounded), then the polynomial time hierarchy would collapse to its second level. This provides strong evidence that it will be very hard to show that NP has perfect zero-knowledge interactive proofs. As a result, considerable effort was devoted in the last few years to the design of alternative models in which it would be possible to solve the problems of perfect zero-knowledge proofs for NP, zero-knowledge proofs for NP without intractability assumptions, and zero-knowledge proofs for NP in a constant number of rounds.

Feige and Shamir [FS] solved the problem of zero-knowledge argument (namely, when the prover is polynomially bounded) for NP in a constant number of rounds, under the assumption that one-way functions exist. The counterpart problem with respect to an unbounded prover has been solved by Goldreich and Kahan [GK] under the assumption that claw-free functions exist. The problem of perfect zero knowledge was solved for

J. Feigenbaum (Ed.): Advances in Cryptology - CRYPTO '91, LNCS 576, pp. 213-224, 1992.

some special cases: Brassard, Crepeau and Yung [BCY] show the existence of parallel perfect zero knowledge arguments for NP under the Certified Discrete Log Assumption (or alternatively, under a generalization of this assumption), and Bellare, Micali and Ostrovsky [BMO] exhibit perfect zero-knowledge proofs for Quadratic residuosity and graph isomorphism in 5 rounds.

Ben-Or, Goldwasser, Kilian and Wigderson [BGKW] suggested the novel concept of multi-prover zero-knowledge interactive proof system for NP, solved the perfect zero-knowledge problem by exhibiting a *sequential* two-prover protocol which achieves this aim, and remarked that the *parallel* execution of their protocol is also a perfect zero-knowledge proof system with a single round under a weak definition which requires only a constant probability of cheating. Fortnow, Rompel and Sipser [FRS] claimed a similar result under the stronger definition which requires a negligible probability of cheating, but their proof of soundness was later shown to be faulty by Fortnow [F1], and no alternative parallel protocol is currently known to be sound in this strong sense. Moreover, there are some examples of protocols (see [F1] and (4.1) here) for which the probability of cheating in their parallel version is known to be exponentially better than in their sequential version.

In this paper we solve this open problem: we prove the soundness of the *parallel* two prover zero knowledge interactive proof for NP suggested by Ben-Or, Goldwasser, Kilian and Wigderson in [BGKW]. As a first step we describe a simpler one-round two-prover interactive proof for Hamiltonicity, and prove that it is sound, complete and perfect zero knowledge under no intractability assumptions. We then show that the same techniques can be applied to the original [BGKW] protocol.

In section 2 we give some definitions. In section 3 we present our simplified parallel protocol for Hamiltonicity, and prove its correctness in section 4. In section 5 we prove that our protocol is also a perfect zero knowledge proof of knowledge, which can extract an actual witness from any sufficiently successful pair of provers.

2 Definitions

Definition 1:

Let L be an NP-language. We say that L has a two-prover interactive proof system if there exists an interactive BPP machine V (the verifier) capable of interacting with two other machines P_1 and P_2 (the provers). The provers can cooperate and choose a common strategy before the interaction with the verifier starts, but are isolated from each other during the execution of the protocol. The protocol has to satisfy the following conditions:

1. $\exists P_1, P_2 \quad \forall x \in L$ the probability that V accepts x is overwhelming.
2. $\forall P_1, P_2 \quad \forall x \notin L$, the probability that V accepts x is negligible.

Definition 2:
Let (P_1, P_2, V) be a two-prover interactive proof system for L. Let $View_{P_1,P_2,V}(x)$ denote the verifier's view during the protocol (namely the sequence of messages exchanged between the verifier and the provers along with the private random bits of V). This is a probability space taken over the coin tosses of V and the random tapes of (P_1, P_2). We say that two-prover interactive protocol (P_1, P_2, V) is perfect zero knowledge for V if there exists a BPP machine M such that $M(x) = View_{P_1,P_2,V}(x)$. We say that L has a two-prover perfect zero-knowledge proof system if there exist independent provers P_1, P_2 such that for all BPP verifiers $\hat{V}$, there exists a probabilistic Turing machine M such that for all $x \in L$, $M(x) = View_{P_1,P_2,\hat{V}}(x)$ and $M(x)$ terminates in expected polynomial time.
Definition 3: Let H be a $t \times t$ matrix of zeroes and ones (which can be thought of as an adjacency matrix of a directed graph). We say that H is exactly Hamiltonian if there is exactly a single 1 in every row and in every column, and these t ones define a permutation with a single cycle.
The Basic Step of Proofs of Hamiltonicity (with a single prover):
Let A and B be two $t \times t$ random matrices of zeroes and ones whose pointwise XOR $A \oplus B = H$ is a random exactly Hamiltonian matrix. Denote by S the Hamiltonian cycle on t nodes whose adjacency matrix is H. Assume now that an honest prover wants to use H in order to prove to V the Hamiltonicity of some graph G with t nodes, and assume that only the prover knows A, B and H but V is convinced that H is exactly Hamiltonian. Let π be a permutation that maps S onto the Hamiltonian cycle of G (i.e. $\pi(S) \subseteq G$). P sends V the permutation π and the values of all the entries in $\pi(A)$ and $\pi(B)$ which do not correspond to edges in G. V accepts the proof iff all the revealed pairs $((\pi(A)_{i,j}, \pi(B)_{i,j})$ such that (i,j) is non-edge in $G)$ are $(0,0)$ or $(1,1)$. P's proof implies that the t ones that remain unrevealed in $\pi(H)$ correspond to edges of G, and thus G contains a Hamiltonian cycle.
Informally, this protocol is zero knowledge since all the verifier gets is a collection of (pairs of equal) random bits and a random permutation, and both things can be simulated in random polynomial time.

3 The Two Prover Protocol

Let (P_1, P_2, V) denote the two-prover protocol which receives as input the graph $G = (V, E)$, $|V| = t$ and tries to prove its Hamiltonicity. Let P_1 and P_2 share two random $t \times t$

matrices A and B such that $A \oplus B = H$, where H is a random $t \times t$ exactly Hamiltonian matrix, and assume that P_1 has a witness for this statement on his auxiliary tape.
The basic two prover protocol (BP) of Hamiltonicity is:

- V randomly and independently chooses two bits b_1 and b_2. He sends b_1 to P_1 and b_2 to P_2.
- If $b_1 = 0$ then P_1 sends A and B to V, otherwise he executes the basic step of the previous section.
- If $b_2 = 0$ then P_2 sends A to V, otherwise he sends B to V.
- According to b_1 V either checks that $A \oplus B$ is exactly Hamiltonian or checks that the basic step was done correctly, and in both cases he verifies the consistency of the revealed entries with P_2's response. V accepts iff these checks are successful.

The full protocol FP_n is a one-round protocol which consists of n *parallel* independent executions of BP, where n is a security parameter.
In the next two sections we prove that FP_n is a perfect zero knowledge interactive proof for Hamiltonicity, and that it is also a perfect zero knowledge proof of knowledge, which directly gives the following theorems:
Theorem 1: *Every language in NP has a two prover perfect zero knowledge interactive proof of membership in one round without making any intractability assumptions.*
Theorem 2: *Every language in NP has a two prover perfect zero knowledge interactive proof of knowledge in one round without making any intractability assumptions.*

4 Correctness

Our first goal is to prove that the parallel protocol FP_n is a perfect zero knowledge proof for Hamiltonicity.
Completeness: P_1 (which is either infinitely powerful or polynomial time bounded with knowledge of a Hamiltonian cycle in G) can determine the permutation π of the basic step and perform the protocol. Notice that unlike the [BGKW] protocol, only P_1 has to know the actual input graph, while P_2 should only know its size t.
Zero-Knowledge: We construct a probabilistic polynomial time simulator M which without knowledge of a cycle in G can give a response to every $2n$-bit query of V which is perfectly indistinguishable from the answers of the real provers. This simulation can be easily carried out because b_1 and b_2 are chosen by V before it gets any messages from the provers, and thus M can use them in choosing A and B. If $b_1 = 0$ then M sends V two random $t \times t$ 0/1 matrices whose XOR is an exactly Hamiltonian matrix

and according to b_2 he sends V one of these matrices. If $b_1 = 1$, M randomly chooses a 0/1 matrix A and a permutation π, simulates P_1's basic step (with the pair (A, A) and $\pi(G)$) and sends A as a simulation of P_2. It is easy to verify that this simulation is perfectly indistinguishable from a real execution, which means that our protocol is perfect zero-knowledge.
The main difficulty (and therefore the motivation of this paper) is how to prove the soundness.

4.1 Where is The Problem?

Consider first the basic protocol BP. It is easy to see that simultaneous success of (P_1, P_2) in answering the four possible requests of V implies the Hamiltonicity of G. Moreover, one can verify that the probability of cheating (when G is not Hamiltonian) is at most 3/4, and thus the probability of cheating in n sequential independent executions of BP is at most $(\frac{3}{4})^n$.
We would like to get the same result with respect to parallel executions but its falsehood is the motivation of this paper. Fortnow [F1] constructed a (somewhat artificial) two prover protocol that accepts all inputs with probability 1/2 such that there exists a strategy for the parallel execution of two rounds which causes the verifier to accept all inputs with probability 3/8. We now show that this problem can in fact arise in our protocol by showing that if G is not Hamiltonian then the probability of cheating in FP_2 is greater than $(\frac{3}{4})^2$. We demonstrate this fact by specifying a strategy for cheating (P_1, P_2) which succeeds in 10 out of the 16 possible requests of V.
Let (X, Y) and (Z, W) be two pairs of $t \times t$ 0/1 matrices such that $X \oplus Y$ and $Z \oplus W$ are exactly Hamiltonian matrices. Let $\hat{W}$, $\hat{X}$ and $\hat{Y}$ be sets of $t^2 - |E(G)|$ entries of W, X, Y, respectively, which correspond to non-edges in $\psi(G)$ for some arbitrary permutation ψ. Let $b_{i,j}$ $(1 \leq i, j \leq 2)$ be the bit sent by V to P_i in the j'th round, and $A_{i,j}$ be the corresponding answer of P_i.
The strategy is:
Instructions for P_1:

$b_{1,1}$	$b_{1,2}$	$A_{1,1}$	$A_{1,2}$
0	0	(X, Y)	(Z, W)
0	1	(X, Y)	$(\hat{W}, \hat{W})$ *and* ψ
1	0	$(\hat{Y}, \hat{Y})$ *and* ψ	(Z, W)
1	1	$(\hat{X}, \hat{X})$ *and* ψ	$(\hat{W}, \hat{W})$ *and* ψ

Instructions for P_2:

$b_{2,1}$	$b_{2,2}$	$A_{2,1}$	$A_{2,2}$
0	0	X	W
0	1	X	W
1	0	Y	Z
1	1	Y	W

It is easy to check that the following matrix represents the successful executions of FP_2 whenever the provers follow the above strategy:

	00	01	10	11
00	0	1	0	1
01	1	1	0	1
10	1	0	1	0
11	1	1	1	0

In this matrix the pairs at the top are the $(b_{1,1}, b_{1,2})$ requests, those on the left are the $(b_{2,1}, b_{2,2})$ requests, and the ten 1-entries represent the successful executions in which V accepts the provers' messages. Since all the choices of $b_{i,j}$ quadruples are equally likely, the cheating (P_1, P_2) succeed with probability $\frac{10}{16}$ (which is greater than $(\frac{3}{4})^2$). A simple extension of this strategy to n parallel rounds (which succeeds with probability $(\frac{10}{16})^{n/2} > (\frac{3}{4})^n$) demonstrates the difficulty of proving the soundness of parallel executions by using standard techniques. In the next subsection we show how to overcome this problem.

4.2 The Proof of Soundness.

Our main theorem uses novel techniques to show that the parallel protocol is sound, by proving that the probability of cheating decreases exponentially fast:

Theorem 3: *If G is not Hamiltonian then*

$$\forall(\hat{P}_1, \hat{P}_2) \;\; Pr\{FP_n \; succeeds\} < \frac{1}{2^{n/9}}$$

where the probability is taken over the coin tosses of V.

Proof: Without loss of generality we can assume that $\hat{P}_1$ and $\hat{P}_2$ are deterministic, and use their best strategy against the particular verifier V. Denote by σ a random n-bit string sent by V to $\hat{P}_2$, and by τ a random n-bit string sent by V to $\hat{P}_1$. Let σ_k (τ_k) be the k'th bit of σ (τ). For each σ denote by A_σ the set of all those τ's for which FP_n succeeds on (σ, τ). We prove the theorem by proving that if:

$$Pr\{FP_n \; succeeds\} \geq \frac{1}{2^{n/9}}$$

then there exists a successful quadruple, i.e. $(\sigma', \sigma'', \tau', \tau'')$ such that FP_n succeeds on each one of the following pairs: (σ', τ') , (σ', τ'') , (σ'', τ') , (σ'', τ'') , and there exists $1 \leq k \leq n$ such that:

$$\sigma'_k \neq \sigma''_k \quad \text{and} \quad \tau'_k \neq \tau''_k \ .$$

Lemma 4: The existence of a successful quadruple implies the Hamiltonicity of G.
Proof: Assume that $(\sigma', \sigma'', \tau', \tau'')$ is a successful quadruple and without loss of generality assume that for some $1 \leq k \leq n$

$$\sigma'_k = 0,\ \sigma''_k = 1,\ \tau'_k = 0,\ \tau''_k = 1.$$

We concentrate now on the answers of the provers at the k'th stage of the parallel protocol: As a response for σ', P_1 sends V the $0/1$ $t \times t$ matrices A and B, where $H = A \oplus B$ is an exactly Hamiltonian matrix. The success of the executions implies that P_2 sends A as a response to τ', and B as a response to τ''. It also implies that while executing the basic step in response to σ'', P_1 sends a permutation π and pairs of equal bits which are identical to their counterparts in P_2's matrices, and thus identical also to their counterparts in P_1's answer on σ' (i.e. A and B). Therefore, by executing this protocol just against P_1 on σ' and on σ'' we can extract the Hamiltonian cycle (HC) in G by concentrating on his answers at the k'th stage, and comparing the adjacency matrix of $\pi(G)$ to $H = A \oplus B$. ■

The existence of a successful quadruple was shown to contradict the assumption that G is not Hamiltonian. Note that the condition on k is essential, since in the concrete matrix demonstrated at the end of section 4.1 there are several quadruples $\sigma', \sigma'', \tau', \tau''$ which define four successful executions, but we cannot extract the witness since none of them satisfies the condition on k. For example: $\sigma' = (01)$, $\sigma'' = (11)$, $\tau' = (00)$, $\tau'' = (01)$ define four successes but there is no index $1 \leq k \leq 2$ on which σ' differs from σ'' and τ' differs from τ'' simultaneously.
Definition: We say that σ is good if

$$|A_\sigma| \geq \frac{2^n}{2 \cdot 2^{n/9}}.$$

Lemma 5: If $Pr\{FP_n succeeds\} \geq \frac{1}{2^{n/9}}$ then there exist at least $\frac{2^n}{2 \cdot 2^{n/9}}$ good σ's.
Proof: The provers are deterministic, therefore there are at least $(2^{2n}/2^{n/9})$ $2n$-bit strings for which (P_1, P_2) succeed. Therefore trivially by applying an elementary counting argument we get the result. ■

Denote by T the set of all the good σ's ($|T| \geq \frac{2^n}{2 \cdot 2^{n/9}}$). Our goal now is to show that it is possible to choose a set $S \subseteq T$ of $4 \cdot 2^{n/9}$ good n-bit strings σ's such that every two strings in S differ from each other in more than $9n/40$ bits.
An algorithm for choosing S:
BEGIN

- $S \longleftarrow \phi$.
- Repeat $4 \cdot 2^{n/9}$ times: Choose an arbitrary good n-bit string σ in T, add it to S, and remove from T all the strings which differ from σ in less than $9n/40$ bits.

END

Lemma 6: *This algorithm cannot stop prematurely.*

Proof: First we have to notice that the total number of n-bit strings $x = (x_1, x_2, \ldots, x_n)$ for which $\sum_{i=1}^{n} x_i < \frac{9n}{40}$ is less than:

$$9n/40 \left(\begin{array}{c} n \\ 9n/40 \end{array} \right) < \frac{2^n}{2^{11n/48}}$$

for all sufficiently large n.

Therefore for each n-bit string there are at most $\frac{2^n}{2^{11n/48}}$ strings which differ from it in less than $9n/40$ bits. Therefore the validity of the following inequality implies the success of the algorithm:

$$4 \cdot 2^{n/9}(1 + \frac{2^n}{2^{11n/48}}) \leq \frac{2^n}{2 \cdot 2^{n/9}} \leq |T|. \qquad \blacksquare$$

Lemma 7: *There exist $\sigma', \sigma'' \in S$, such that:*

$$|A_{\sigma'} \bigcap A_{\sigma''}| \geq \frac{2^n}{10 \cdot 2^{2n/9}}.$$

Proof: According to the inclusion exclusion formula we have:

$$\sum_{\sigma \in S} |A_\sigma| - \sum_{\sigma', \sigma'' \in S} |A_{\sigma'} \bigcap A_{\sigma''}| \leq 2^n.$$

If the Lemma is not true then we get:

$$4 \cdot 2^{n/9} \cdot \frac{2^n}{2 \cdot 2^{n/9}} - \frac{4^2 \cdot 2^{2n/9}}{2} \cdot \frac{2^n}{10 \cdot 2^{2n/9}} \leq 2^n.$$

which is false. $\blacksquare$

Denote by σ' and σ'' two strings in S for which:

$$|A_{\sigma'} \bigcap A_{\sigma''}| \geq \frac{2^n}{10 \cdot 2^{2n/9}}.$$

We showed that every two strings in S differ from each other in at least $9n/40$ bits, and in particular these two σ', σ'' have this property. Denote by I the set of $9n/40$ indices in which σ' differs from σ''. Choose an arbitrary $\tau' \in A_{\sigma'} \bigcap A_{\sigma''}$. There are exactly $\frac{2^n}{2^{9n/40}}$

n-bit strings which are identical to τ' on each of the indices of I. Therefore the total number of strings in the intersection which are identical to τ' on each of the indices of I is bounded by:

$$\frac{2^n}{2^{9n/40}} < \frac{2^n}{10 \cdot 2^{2n/9}} \leq |A_{\sigma'} \bigcap A_{\sigma''}|.$$

Therefore there exists $\tau'' \in A_{\sigma'} \bigcap A_{\sigma''}$ which differs from τ' in at least one of the indices of I, and we have thus found a successful quadruple. ■

Remarks:

1. Recent improvements of the analysis (obtained independently by Peleg, Alon and Feige [Fe]) reduced the upper bound on the probability of cheating and extended the analysis to other protocols based on constant-size queries.
2. The same technique of successful quadruples can be used to prove the soundness of the original [BGKW] protocol. The analysis is slightly more complicated due to their use of three valued "bits" as messages, and will be given in the full version of this extended abstract.

5 The Protocol is a Proof of Knowledge.

In this section we prove that our protocol for Hamiltonicity is also a perfect zero knowledge proof of knowledge. We follow the definition suggested by Feige, Fiat and Shamir in [FFS].

Definition:

Let (P_1, P_2, V) be a two-prover perfect zero knowledge interactive proof system for an NP-language L such that P_1 and P_2 are probabilistic polynomial time bounded. We say that (P_1, P_2, V) is an interactive proof of knowledge if there exists an interactive probabilistic machine T (knowledge extractor with complete control over $(\hat{P}_1, \hat{P}_2)$) such that for all $(\hat{P}_1, \hat{P}_2)$ and for all input x, if V accepts the proof (that $x \in L$) with non negligible probability, then the output produced by T at the end of polynomially many executions of $(\hat{P}_1, \hat{P}_2, T)$ on input x is a witness for $x \in L$ and T terminates in expected polynomial time. More formally:

$$\exists T \; \forall (\hat{P}_1, \hat{P}_2) \; \forall x \;\; \forall a \; \exists b \; \exists N \; \forall n > N$$

$$Pr\{(\hat{P}_1, \hat{P}_2, V) \; succeeds \; on \; x\} > 1/n^a \Longrightarrow$$

$$Pr\{output \; of \; (\hat{P}_1, \hat{P}_2, T) \; on \; x \; is \; a \; witness \; for \; x\} = 1$$

and the expected running time of T is $O(n^b)$, where the probability is taken over the coin tosses of V.

Theorem 8: FP_n is a perfect zero knowledge interactive proof of knowledge for Hamiltonicity.

Proof: Without loss of generality we can assume that P_1 and P_2 are deterministic. Therefore the probability space consists of all the (equally likely) $2n$-bit strings which V may send to (P_1, P_2), and at least $2^{2n}/n^a$ of them result in successful executions.

Due to the same argument (and using the same notation) of the previous section we conclude that there exist at least $2^n/2n^a$ σ's whose $|A_\sigma| \geq 2^n/2n^a$, and call them good σ's.

Lemma 9: *For every set $\hat{S}$ of $4n^{2a}$ good σ's there exist $\sigma', \sigma'' \in \hat{S}$ such that:*

$$|A_{\sigma'} \bigcap A_{\sigma''}| \geq 2^n/(2n^a)^3.$$

Proof: As in the proof of Lemma 7, the first two terms of the inclusion exclusion formula trivially give the result. ■

We now specify the knowledge extractor T: Choose a random set $\hat{S}$ of $4n^{2a}$ good σ's. This step can be performed by an expected polynomial number of statistical experiments of the following type: randomly choose an n-bit string σ; for this string choose independently polynomially many random n-bit strings (τ's), execute the protocol for each such pair (τ, σ) and estimate the probability of success with respect to this σ.

Choose an arbitrary pair $\sigma', \sigma'' \in \hat{S}$ which satisfies Lemma 9 (there are only $O(n^{4a})$ pairs for which we have to execute statistical experiments). Choose an arbitrary n-bit string in $A_{\sigma'} \bigcap A_{\sigma''}$, and call it τ'.

Notice the following crucial point: In order to choose $\hat{S}$, we randomly choose *polynomially* many n-bit strings, each one by n unbiased and independent coin tosses, and thus every two chosen strings differ from each other in at least $n/3$ bits with overwhelming probability. In particular, the σ', σ'' chosen from $\hat{S}$ satisfy this property with overwhelming probability. Denote by J the set of indices on which σ' differs from σ'' $(|J| \geq n/3)$. The same argument used in the proof of *Theorem* 1 yields that almost all the strings in $A_{\sigma'} \bigcap A_{\sigma''}$ differ from τ' in at least one of the indices of J, therefore we can easily choose a string in this intersection which has this property, and call it τ''.

Now all T has to do in order to extract the Hamiltonian cycle in G is to execute the protocol FP_n against (P_1, P_2) on the following four pairs:
$(\tau', \sigma'), (\tau', \sigma''), (\tau'', \sigma'), (\tau'', \sigma'')$, and use Lemma 4. To complete the proof, we can execute in parallel an exhaustive search for a witness to handle the negligible probability that the main procedure fails to find a witness. ■

6 Acknowledgments

We are grateful to Uri Feige for pointing out the parallelization problem, and for interesting discussions and helpful comments on this work.

References

[BCY] G. Brassard, C. Crepeau, M. Yung. *Everything in NP can be argued in perfect zero knowledge in a bounded number of rounds,* Proceedings of 16th ICALP (1989).

[BMO] M. Bellare, S. Micali and R. Ostrovsky. *Perfect Zero-Knowledge in Constant Rounds,* Proceedings of the 22nd annual ACM Symposium on Theory of Computing (1990), 482–493.

[BGKW] M. Ben-Or, S. Goldwasser, J. Kilian, A. Wigderson. *Multi-Prover Interactive Proofs: How to Remove Intractability Assumptions.* Proceedings of the 20th annual ACM Symposium on Theory of Computing (1988), 113–131.

[BHZ] R. Boppana, J. Hastad and S. Zachos. *Does CoNP Have Short Interactive Proofs?,* Information Processing Letters **25 2** (1987), 127–132.

[F1] L. Fortnow. *Complexity Theoretic Aspects of Interactive Proof Systems, Ph.D. Thesis,* MIT/LCS/TR-447, (1989).

[F2] L. Fortnow. *The Complexity of Perfect Zero-Knowledge,* Proceedings of the 19th annual ACM Symposium on Theory of Computing (1987), 204–209.

[Fe] U. Feige. *On the Success Probability of the Two Provers in One-Round Proof Systems,* Proceedings of the 6th Structure in Complexity Theory Conference (1991), IEEE.

[FFS] U. Feige, A. Fiat and A. Shamir. *Zero Knowledge Proofs of Identity,* Proceedings of the 19th annual ACM Symposium on Theory of Computing (1987), 210–217.

[FRS] L. Fortnow, J. Rompel and M. Sipser. *On the Power of Multi-Prover Interactive Protocols,* Proceedings of the 3rd Structure in Complexity Theory Conference (1988), 156–161.

[FS] U. Feige and A. Shamir. *Witness Indistinguishable and Witness Hiding Protocols,* Proceedings of the 22th annual ACM Symposium on Theory of Computing (1990), 416–426.

[GK] O. Goldreich and A. Kahan, private communication (1990).

[GMW] O. Goldreich, S. Micali and A. Wigderson. *Proofs that Yield Nothing But their Validity and a Methodology of Cryptographic Protocol Design*, Proceedings of the 27th Symposium on Foundations of Computer Science (1986), IEEE, 174–187.

[GMR] S. Goldwasser, S. Micali, and C. Rackoff. *The Knowledge Complexity of Interactive Proofs*, Proceedings of the 17th annual ACM Symposium on Theory of Computing (1985), 291–304.