COMBINATORICA
Akadémiai Kiadó – Springer-Verlag

COMBINATORICA **15** (2) (1995) 203–214

A ONE-ROUND, TWO-PROVER, ZERO-KNOWLEDGE PROTOCOL FOR NP

DROR LAPIDOT and ADI SHAMIR

Received July 29, 1992
Revised July 25, 1994

The model of zero-knowledge multi-prover interactive proofs was introduced by Ben-Or, Goldwasser, Kilian and Wigderson in [4]. A major open problem associated with this model is whether NP problems can be proven by one-round, two-prover, zero-knowledge protocols with exponentially small error probability (e.g. via parallel executions). A positive answer was claimed by Fortnow, Rompel and Sipser in [12], 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 much higher than the probability of cheating in n independent sequential rounds (with exponential ratio between them). In this paper we solve this problem: We show a new one-round two-prover interactive proof for Graph Hamiltonicity, we prove that it is complete, sound and perfect zero-knowledge, and thus every problem in NP has a one-round two-prover interactive proof which is perfectly zero knowledge under no cryptographic assumptions. The main difficulty is in proving the soundness of our parallel protocol namely, proving that the probability of cheating in this one-round protocol is upper bounded by some exponentially low threshold. We prove that this probability is at most $1/2^{n/9}$ (where n is the number of parallel rounds), by translating the soundness problem into some extremal combinatorial problem, and then solving this new problem.

1. Introduction

In [16] 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 [9] and [5] 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 [13] 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 [14] under the assumption that claw-free functions exist. The problem of perfect zero knowledge was solved for some special cases: Brassard, Crepeau and Yung [2]

Mathematics Subject Classification (1991): 94 A 60, 05 C 35

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 [6] exhibit perfect zero-knowledge proofs for Quadratic residuosity and graph isomorphism in 5 rounds.

Ben-Or, Goldwasser, Kilian and Wigderson [4] 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 without any complexity assumptions, 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 [12] 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 [8]. Moreover there are some examples of protocols (see [8] 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. Cai, Condon and Lipton showed in [7] several results regarding parallel repetition of multi-prover protocols and two-person games. In particular they pointed out the connection between the size of the messages (which are exchanged between the provers and the verifier) and the decreasing rate of the error probability.

In this paper we describe a simple one-round two-prover interactive proof for Graph Hamiltonicity, prove that it is complete sound and perfect zero knowledge, and thus every language in NP has a one-round two-prover perfect zero knowledge interactive proof under no intractability assumptions. Note however that we did not solve the more general problem whether every protocol, when executed in parallel, yields an exponentially decreasing error probability.

The main difficulty is in proving the soundness of our parallel protocol, i.e. proving that the probability of accepting a false statement is upper bounded by some exponentially low threshold. We prove that this probability is at most $1/2^{n/9}$ (where n is the number of parallel rounds), by translating the soundness problem into some extremal combinatorial problem, and then solving this particular problem.

We end up this paper by proving that our one-round two-prover protocol is also a perfect zero knowledge proof of knowledge, which can extract an actual witness from any sufficiently successful pair of provers.[1]

The paper is organized in the following way: In Section 2 we give several definitions. In Section 3 we present our 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.

[1] The formal notion of zero knowledge proof of knowledge was introduced by Feige, Fiat and Shamir in [11]. Other researchers later formalized slightly modified versions of this notion, but we follow the original definition.

2. Definitions

The concept of two-prover interactive proofs was introduced by Ben-Or, Goldwasser, Kilian and Wigderson in [4]. We follow their definitions.

Let P_1 and P_2 (the provers) be probabilistic Turing machines and V (the verifier) be a probabilistic polynomial-time Turing machine, all of which share the same read-only input tape. The provers' goal is to convince the verifier that some word x (which is the common input of P_1, P_2 and V) belongs to some language L. In order to do this, the verifier V shares communication tapes with each P_i but the provers P_1 and P_2 have no common tapes except the input tape. This means that 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 namely, they are not allowed to communicate with each other and each one of them can not see the messages which are exchanged between the verifier and the other prover. At the end of the conversation, V outputs either an *accept* (or *reject*) based on the input x, its random bits and the entire conversation it has had with both provers.

Definition 1 Let L be an NP-language. We say that L has a two-prover interactive proof system if:

1. Completeness: $\exists P_1,\ P_2 \forall x \in L$ V accepts x.

2. Soundness: $\forall P_1, P_2\ \ \forall x \notin L$, the probability that V accepts x is at most $2^{-|x|}$ (where $|x|$ is the size of the common input x).

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).

Definition 2 We say that L has a two-prover perfect zero-knowledge proof system if there exist two provers P_1, P_2 such that for all probabilistic polynomial time verifiers V, there exists a probabilistic Turing machine M (a simulator) such that for all $x \in L$, $M(x) = View_{P_1,P_2,V}(x)$ (where $M(x)$ denotes the output of M on the input x) and $M(x)$ terminates in expected polynomial time.

The protocols we consider in this paper can be divided into time steps: In each step either the verifier sends his queries to the provers (the verifier's step) or they send him their responses (the provers' step). Without loss of generality, we can assume that the first step (in any interactive proof) is that of V. A *round* is a pair of two consecutive steps: A verifier's step followed by a provers' step.

3. The One-Round Two-Prover Protocol

Definition 3 A directed graph (on t vertices) is Hamiltonian if it contains a directed closed path (cycle) of length t which passes through all the vertices.

Definition 4 Let H be a $t \times t$ matrix of zeros 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.

We now describe a single prover subprotocol of the full two prover protocol for Hamiltonicity. Let A and B be two $t \times t$ random matrices of zeros and ones whose pointwise XOR $A \bigoplus B = H$ 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 of the vertices 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 entries in $\pi(A)$ and $\pi(B)$ which do not correspond to edges in G. V accepts the proof iff all 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. A formal proof involves the notions of simulator and distinguisher. We refer the interested reader to [15] and [16].

However, a cheating prover can use a matrix H which is not exactly Hamiltonian in order to fool the verifier. We introduce a second prover into the protocol in order to check this possibility. Let (P_1, P_2, V) denote the two-prover protocol in which P_1, P_2 and V receive the graph G (which has t vertices) as a common input, and (P_1, P_2) try to prove its Hamiltonicity. Let P_1 and P_2 share two random $t \times t$ matrices A and B such that $A \bigoplus B = H$, where H is a random $t \times t$ exactly Hamiltonian matrix, i.e. before the protocol begins, P_1 randomly selects an exactly Hamiltonian matrix H, and a random matrix A, and sends A and $B = A \bigoplus H$ to P_2. Assume that P_1 has a witness for the Hamiltonicity of G (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 $A \bigoplus B$ 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. More precisely: At the initial stage, the provers share n independent pairs of random matrices, each pair has the same properties as those of the above (A,B). At the interactive stage the verifier randomly chooses $2n$ independent bits and sends the first n bits to the first prover and the last n bits to the second prover. Now, each prover has a sequence of n pairs of random matrices and a sequence of n random bits, and for each triple (a random bit and a pair of matrices) he independently replies according to the basic protocol BP.

In the next section we prove that FP_n is a perfect zero knowledge interactive proof for Hamiltonicity. Since this is an NP-Complete problem, we have the following theorem:

Theorem 5. *Every language in NP has a two prover perfect zero knowledge interactive proof of membership in one round without making any intractability assumptions.*

4. Correctness

Our 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 [4] 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 the query bits 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 $M(G)=View_{P_1,P_2,V}(G)$ namely, 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 any simultaneous success of (P_1,P_2) in answering all the four possible combinations of b_1 and b_2 queries of V implies the Hamiltonicity of G. Therefore, the cheating provers (when G is not Hamiltonian) can answer successfully at most 3 out of the 4 possible requests. The probability of cheating in each execution of BP is thus at most 3/4, and 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 [8] 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 $4^2 = 16$ possible requests of V.

Let (X,Y) and (Z,W) be two pairs of $t \times t$ 0/1 matrices such that $X \bigoplus Y$ and $Z \bigoplus W$ 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 $b_{i,j}$ $(1 \le i,j \le 2)$ 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:

requests		*responses*	
$b_{1,1}$	$b_{1,2}$	$A_{1,1}$	$A_{1,2}$
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:

requests		*responses*	
$b_{2,1}$	$b_{2,2}$	$A_{2,1}$	$A_{2,2}$
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 results of all the possible 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 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 $(\frac{3}{4})^2$). This strategy, applied 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 6. *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 $\sigma_k(\tau_k)$ be the k'th bit of $\sigma(\tau)$. 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:

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

Lemma 7. *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 $H = A \bigoplus B$ 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 $H = A \bigoplus B$. ∎

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 success matrix demonstrated at the end of Section 4.1 there are several quadruples σ', σ'', τ', τ'' which define four successful executions, but we cannot extract from them a witness of Hamiltonicity 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 form τ'' simultaneously.

Definition 8 We say that σ is good if

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

Lemma 9. *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. The result follows from an elementary counting argument. ∎

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 strings which differ from σ in at most $9n/40$ bits.

END

Lemma 10. *This algorithm outputs a set of* $4 \cdot 2^{n/9}$ *good* n*-bit strings, such that every two of them differ from each other in more than* $9n/40$ *bits.*

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

$$9n/40 \binom{n}{9n/40} < \frac{2^n}{2^{11n/48}}$$

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 at most $9n/40$ bits. The validity of the following inequality implies the success of the algorithm:

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

∎

Lemma 11. *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 can be simplified to:

$$2 \cdot 2^n - \frac{4}{5} \cdot 2^n \leq 2^n$$

which is obviously false. ∎

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

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

By construction, 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-bits 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. ∎

Remark: Recent improvements of the analysis (obtained independently by Peleg [18], Alon [1] and Feige [10]) tighten the constants in the upper bound on the probability of cheating and extend the analysis to other protocols based on constant-size queries.

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 [11] (a slightly modified definition appears in [3]).

Definition 12 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 (called "knowledge extractor") 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 13. *FP_n is a perfect zero knowledge interactive proof of knowledge for Hamiltonicity.*

Proof. According to the above definition the probability space consists of all the (equally likely) $2n$-bit strings which V may send to (P_1,P_2). By assumption, 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 14. *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 11, using the first two terms of the inclusion exclusion formula trivially gives 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 the condition of Lemma 14 (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 6 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:
(τ',σ'), (τ',σ''), (τ'',σ'), (τ'',σ'').
Lemma 7 describes how to extract the cycle from the four sets of answers.

Remark: There is a negligible probability that the main procedure fails to find such τ and σ's. To complete the formal proof, we can carry out an (exponentially long) exhaustive search, and stop when one of the two approaches finds a cycle. ∎

Since Hamiltonicity is an NP-Complete problem, Theorem 13 implies the following theorem:

Theorem 15. *Every language in NP has a two prover perfect zero knowledge interactive proof of knowledge in one round without any intractability assumptions.*

References

[1] N. ALON: Private communication, 1990.

[2] G. BRASSARD, C. CREPEAU, M. YUNG: *Everything in NP can be argued in perfect zero knowledge in a bounded number of rounds*, Proc. of 16th International Colloquium on Automata, Languages and Programming (ICALP) 1989.

[3] M. BELLARE, and O. GOLDREICH: *On Defining Proofs of Knowledge*, Proc. of Crypto, 390–420, 1992.

[4] M. BEN-OR, S. GOLDWASSER, J. KILIAN, and A. WIGDERSON: *Multi-Prover Interactive Proofs: How to Remove Intractability Assumptions*, Proc. 20th ACM Symposium on Theory of Computing, 113–131, 1988.

[5] R. BOPPANA, J. HASTAD and S. ZACHOS: Does co-NP Have Short Interactive Proofs?, *Inform. Process. Lett.*, **25** (1987), 127–132.

[6] M. BELLARE, S. MICALI and R. OSTROVSKY: *Perfect Zero-Knowledge in Constant Rounds*, Proc. of 22^{nd} ACM Symposium on Theory of Computing, 482–493, (1990).

[7] J. CAI, A. CONDON, and R. LIPTON: *Playing Games of Incomplete Information*, Proc. of 7th Symposium on Theoretical Aspects of Computer Science, 58–69, 1990.

[8] L. FORTNOW: *Ph. D. Thesis*, M.I.T./LCS/TR-447

[9] L. FORTNOW: *The Complexity of Perfect Zero-Knowledge*, Proc. of 19th ACM Symposium on Theory of Computing, 204–209, 1987.

[10] U. FEIGE: *On the Success Probability of the Two Provers in One Round Proof Systems*, Proc. of Structures in Complexity Theory Conf., 1991.

[11] U. FEIGE, A. FIAT, and A. SHAMIR: *Zero Knowledge Proofs of Identity*, Proc of 19th ACM Symposium on Theory of Computing, 210–217, 1987.

[12] L. FORTNOW, J. ROMPEL, and M. SIPSER: *On the power of Multi-Prover Interactive Protocols*, Proc. of Structures in Complexity Theory Conf., 156–161, 1988.

[13] U. FEIGE, and A. SHAMIR: *Witness Indistinguishable and Witness Hiding Protocols*, Proc. of 22^{nd} ACM Symposium on Theory of Computing, 416–426, 1990.

[14] O. GOLDREICH, and A. KAHAN: Private communication, 1989.

[15] S. GOLDWASSER, S. MICALI, and C. RACKOFF: The Knowledge Complexity of Interactive Proof Systems, *SIAM Journal of Computing*, **1** (1989), 186–208.

[16] O. GOLDREICH, S. MICALI, and A. WIGDERSON: *Proofs that Yield Nothing But Their Validity and a Methodology of Cryptographic Protocol Design*, Proc. of 27th Symposium on Foundations of Computer Science, 174–187, 1986.

[17] D. LAPIDOT, and A. SHAMIR: *Fully Parallelized Multi Prover Protocols for NEXP-time*, Proc. of 32'nd Symposium on Foundations of Computer Science, 13–18 1991.

[18] D. PELEG: Private communication, 1990.

Dror Lapidot

Department of Applied Math.
and Computer Science
The Weizmann Institute of Science
Rehovot 76100, Isreal
drorl@wisdom.weizmann.ac.il

Adi Shamir

Department of Applied Math.
and Computer Science
The Weizmann Institute of Science
Rehovot 76100, Isreal
shamir@wisdom.weizmann.ac.il