Two Prover Protocols – Low Error at Affordable Rates

(Preliminary version)

Uriel Feige * Joe Kilian †

Abstract

Known algebraic methods for reducing the error in two-prover one-round proof systems prove unsatisfactory for many applications, as they have a large overhead in communication and in computation, and they do not preserve zero knowledge.

We show that for a general class of proof systems, *confuse or compare* proof systems, parallel repetition reduces the error at a polynomial rate. Using this result we show that NP has two-prover one-round proof systems with logarithmic communication and arbitrarily small error (bounded away from 0). The same result holds for zero knowledge proof systems for NP. Using these results we obtain improved NP-hardness results for a wide variety of approximation problems.

We also show how the known algebraic error reduction techniques can be modified so as not to disrupt the zero knowledge property. In fact, they can be made to enhance zero knowledge, transforming any protocol which is zero knowledge with respect to honest verifiers, to a protocol which is zero knowledge with respect to cheating verifiers as well. This implies that NEXPTIME has two-prover one-round perfect zero knowledge proof systems with exponentially small error, and that NP has such proof systems with polylogarithmic communication, and superpolynomially small error.

*Dept. of Applied Math., The Weizmann Institute, Rehovot, Israel. feige@wisdom.weizmann.ac.il. Supported by a Koret Foundation fellowship. Work was done while visiting NEC Research Institute.

†NEC Research Institute, 4 Independence Way, Princeton, New Jersey. joe@nec.research.com

STOC 94- 5/94 Montreal, Quebec, Canada
© 1994 ACM 0-89791-663-8/94/0005..$3.50

1 Introduction

We consider single-round two prover proof systems, as introduced by Ben-Or *et al.* [9]. In these proof systems, V generates a pair of questions (q_1, q_2) based on the input x and his random coin tosses R, and sends q_i to P_i, who makes a response a_i. V accepts or rejects based on x, q_1, q_2, a_1, a_2 and R. More formally, we define $MIP(2,1)$ and $mip(2,1)$ ("small $MIP(2,1)$") as follows.

Definition 1 *$[V, P_1, P_2]$ is an $MIP(2,1)$ proof system with error ϵ for language L if the following two conditions hold:*

- Completeness: *If $x \in L$, $V(x, P_1, P_2)$* accepts.
- Soundness: *If $x \notin L$, then for all $\tilde{P}_1, \tilde{P}_2$, $Pr[V(x, \tilde{P}_1, \tilde{P}_2)$* accepts*$] < \epsilon$, where the probability is taken over the random coin tosses of V.*

$[V, P_1, P_2]$ is an $mip(2,1)$ proof system if in addition V's questions and random coins are of size $O(\log |x|)$, and the length of the provers' replies is a constant, independent of $|x|$. We say that $L \in MIP(2,1)$ $\{mip(2,1)\}$ if it has an $MIP(2,1)$ $\{mip(2,1)\}$ proof system with error $1/3$.

An important property possessed by some two prover proof systems is of *zero knowledge*. This property was first introduced by Goldwasser *et al.* [23] in the context of single prover protocols, and later adapted to the multiple-prover scenario [9].

Definition 2 *An $MIP(2,1)$ proof system is* zero knowledge *if there exists a random polynomial time simulator M, such that for any input $x \in L$, and for any possibly cheating $\tilde{V}$, the probability distributions induced by the protocols $[\tilde{V}, P_1, P_2](x)$ and $[\tilde{V}, M](x)$ are identical. The probabilities are taken over the random coin tosses of $\tilde{V}$, P_1, P_2, and M.*

It follows from [4, 28, 19] that $MIP(2,1) = NEXPTIME$. Lapidot and Shamir [27] showed that any language L in NP has a zero knowledge $MIP(2,1)$ proof system with exponentially small error. Arora *et al.* [2] showed that any NP language has an

$mip(2,1)$ proof system with a constant error $\epsilon < 1$. Dwork *et al.* [16] constructs zero-knowledge $MIP(2,1)$ proof systems for NEXPTIME and zero knowledge $mip(2,1)$ proof systems for NP, with some constant error $\epsilon < 1$. However, it was previously open whether $NP \in mip(2,1)$ or whether $NEXPTIME$ had zero-knowledge $MIP(2,1)$ proofs with exponentially small error.

As well as their cryptographic applications (See [9, 10, 27, 16] for a discussion of these issues), two-prover one-round proof systems have been used as a starting point to prove that certain optimization problems are hard to approximate. This was demonstrated in [5, 7, 19], and in a more spectacular way, by Lund and Yannakakis [30] who showed that SETCOVER is hard to approximate within a ratio of $\Omega(\log n)$. Also, the protocols in [2, 6] use constructions of two prover one round proof systems as building blocks. Using the machinery that has been developed, our results imply a number of improved hardness results.

1.1 Error reduction

A major problem in the study of two-prover one-round proof systems is how to reduce the probability that V accepts when $x \notin L$. Given an $MIP(2,1)$ proof system with error ϵ, how do we decrease the error to $\delta < \epsilon$? Sequential repetition can decrease the error exponentially, but requires multiple rounds. Parallel repetition preserves the number of rounds, but its efficacy at error reduction is an outstanding open problem. If the original error probability is less than 1, Verbitsky has recently shown that it can be made arbitrarily small by sufficiently many parallel repetitions [34]. The number of repetitions required in Verbitsky's proof (for achieving a desired error probability) is expressed as a Ramsey-Theory type function of the number of random bits used by the verifier, and no constructive bound is known for this function. Very weak convergence bounds have been obtained for the case where the questions to the two provers are chosen independently at random [13], and have been used in the context of zero knowledge proof systems for NP [27] (see also [1, 17, 31] for improved bounds.). However, due to the weakness of these bounds, the communication and computational efficiency one obtains are comparable to having the provers enumerate and send their entire strategy.

Techniques somewhat similar to parallel repetition can more efficiently achieve error close to 1/2 [17, 24]. Like parallel repetition, these methods leave intact the set of questions the provers are required to answer. However, they require that the provers behave deterministically, and are thus inapplicable to zero-knowledge proof systems.

More recently, algebraic techniques have been developed that transform the original protocol in a nontrivial manner. These techniques can make the error exponentially small [28, 19], but have a number of drawbacks. The bit-complexity of the question size, answer size and randomness used by the verifier are increased by a superconstant factor in order to achieve even a modest error reduction. For this reason, these techniques cannot be used to reduce error within the class $mip(2,1)$. Sophisticated recursive constructions improve the bit complexity of algebraic error reduction techniques, but increase the number of provers to 4 [6][1]. Furthermore, these techniques require the provers to perform complicated computations, such as computing multilinear extensions, which may be much more difficult than those required of the original provers. Finally, zero-knowledge is not maintained by the algebraic techniques, due to the additional information revealed by the provers.

1.2 Our results

1.2.1 A new error reduction technique

Our first result is an error reduction technique based on parallel repetition. We analyze parallel repetition on a special class of proof systems, which we call *confuse or compare* proof systems. All two prover proof systems may be converted into a proof system from this class. We show that parallel repetition reduces the error at a polynomial rate.

Many known $MIP(2,1)$ proof systems are of the following form, which arises when one simulates an "oracle proof system" by a two-prover proof system (see e.g., [21, 16]). The verifier picks a set of questions $q_1, \ldots, q_k$, chooses $r \in \{1, \ldots, k\}$ at random and sends $(q_1, \ldots, q_k)$ to P_1 and q_r to P_2. The verifier checks that P_1's replies, $a_1, \ldots, a_k$ satisfy some predicate (depending on the actual proof system), and that P_2's is equal to a_r. We call this a *compare* round, since all V does with P_2's reply is to compare it for exact equality with one of P_1's replies. We will also be interested in a related protocol in which instead of sending q_r to P_2, V randomly picks an additional set of question $q'_1, \ldots, q'_k$, and sends q'_r instead. In this case, V ignores P_1 and P_2's answers, and accepts. We call this a *confuse* round. A proof system is in *confuse or compare* (C) form, if with probability 1/2 the verifier executes a *confuse* round, and with probability 1/2 the verifier executes a *compare* round.

Proposition 1 *Any $MIP(2,1)$ proof system with c communication bits, r random bits, and error ϵ, can be transformed into C form, with $2c$ communication bits, $2r+2$ random bits, and error at most $(3+\epsilon)/4$.*

[1] See a recent improvement in [33]

Theorem 2 *Let G be an $MIP(2,1)$ proof system in C form with error p. Then parallel repetition reduces the error of G at a polynomial rate. That is, to obtain error δ, it suffices to repeat the protocol in parallel $O(poly(1/(1-p), 1/\delta))$ times.*

Note that the above bound is independent of the size of the questions. It follows that the class of languages having $mip(2,1)$ proof systems is insensitive to the error probability ϵ (for constant ϵ bounded away from 0 and 1). Thus, using [2, 21, 16] we obtain:

Corollary 3 *$NP \subseteq mip(2,1)$. Furthermore, any language in NP has an $mip(2,1)$ zero knowledge proof system with arbitrarily small constant error.*

1.2.2 Applications to NP-hardness results

The results stated above have corollaries that relate to approximation of optimization problems. A direct corollary (using the reductions described in [7, 19]) is the following:

Corollary 4 *It is NP hard to approximate quadratic programming within any constant factor.*

Applying these results to the general framework of [6], we obtain more efficient probabilistically checkable proofs (PCP's). Let $PCP'(f,t)$ denote the class of languages with PCP proofs that achieve error $\frac{1}{2}$ using $O(f(|x|))$ randomness and expected t queries by the verifier. We obtain the following improved bound for t:

Corollary 5 *$NP \in PCP'(\lg, 24)$.*

Most recent results in this area ultimately depend on reductions to probabilistically checkable proofs or two-prover proof systems for NP, and the strength of these results depends on the efficiency of these proof systems. Hence, our results give improved bounds for most of these theorems. For example, we have the following hardness result for the approximatability of MAX 3SAT.

Corollary 6 *Assuming $P \neq NP$, there is no polynomial-time algorithm for satisfying 93/94 of the clauses in a 3SAT expression (i.e., approximating MAX 3SAT to within 93/94).*

Remark: The technique in [6] implies the above bound, under the somewhat stronger assumption that $NEXPTIME \neq EXPTIME$. We can alternatively view our results as weakening the assumptions required by these approximation results rather than improving the bounds.

As well as the automatic improvements we obtain from our results, we obtain additional improvements on the hardness of finding cliques and coloring. These results are obtained by the following four-step procedure: reduce PCP' to clique as in [18]; shrink the size of the resulting graph by deleting vertices that do not correspond to accepting computations; amplify the approximation gap by taking $\Theta(\log n)$ products of the graph with itself; reduce the size of the resulting graph by random sampling (as in [11]).

Theorem 7 *Assuming that $co-RP \neq NP$, there is no probabilistic polynomial time algorithms for*

- *approximating the maximum sized clique in a graph to within $n^{1/15}$, or*
- *approximating the chromatic number of a graph to within $n^{1/71}$.*

Under the $co-RP \neq NP$ assumption, it was previously known to be hard to approximate the maximum sized clique of a graph to within $n^{1/29}$ or to approximate the chromatic number of a graph to within $n^{1/146}$ [6].

The above theorem can be further improved, by optimizing the parameters of the PCP' proof system of Corollary 5 with this specific objective in mind, rather than the objective of minimizing the number of bits read from the proof[2].

1.2.3 Zero-knowledge proofs with low error

We modify the [19] construction so that it does not disrupt zero knowledge properties, obtaining the following theorem.

Theorem 8 *NEXPTIME has $MIP(2,1)$ zero knowledge proof systems with exponentially small error probability. NP has $MIP(2,1)$ zero knowledge proof systems with superpolynomially small error probability, and polylogarithmic communication.*

We note that our results on $MIP(2,1)$ zero knowledge proof systems can be used to show the hardness of generating probability spaces that approximately satisfy given pairwise statistics [25].

1.2.4 Questions that remain open

The rate at which parallel repetition decreases the error of arbitrary $MIP(2,1)$ proof systems is still poorly understood. Any new result on this question would be interesting.

[2] Indeed, Bellare and Sudan [8] have skillfully done such an optimization of parameters, and report impressive improvements to Theorem 7.

Saving random bits: for applications relating $MIP(2,1)$ proof systems to the hardness of approximation, it is desirable to reduce the error by a polynomial factor, while increasing the length of the provers' replies by a logarithmic factor, and the number of random bits used by the verifier only by a constant factor. Can this be done? A somewhat complementary question arises in the context of zero knowledge protocols, where it is desirable to minimize the number of random bits shared by the provers (see remark, end of Section 3).

2 "Confuse or Compare" protocols

Consider an arbitrary $MIP(2,1)$ proof system. For error reduction purposes, we are interested only in the case that $x \notin L$. Fix an arbitrary $x \notin L$, and let G denote the protocol that results. $\omega(G)$ denotes the success probability of the optimal strategy for the two provers in protocol G. In our case, $\omega(G) = p < 1$. Given a desired error probability $\delta < p$, we transform G into a new protocol $\hat{G}$ with $\omega(\hat{G}) \leq \delta$.

If G is not in C form, then the following procedure is used to transform G into G_C of this form. For protocol G, let Q be the set of possible question pairs, and let π be the probability distribution with which V selects questions from Q. For G_C , the verifiers picks at random $(q_1, q_2) \in_\pi Q$, and sends it to P_1. V also picks another pair $(q_1', q_2') \in_\pi Q$. All together, V picked four random questions. V sends one of the four (chosen uniformly at random) to P_2.

In order to accept, V checks that the replies of P_1 satisfy G. Moreover, if the single question sent to P_2 happens to be identical to one of the two questions sent to P_1 (observe that this happens with probability greater than $1/2$), than V checks that the corresponding replies are also identical.

Note that $\omega(G_C) \leq \frac{3+p}{4}$, and that the overhead of this transformation in terms of computational resources and bit complexity is only a constant multiplicative factor.

To obtain $\hat{G}$, the verifier repeats G_C $3m = poly(1/(1-p), 1/\delta)$ times in parallel, selecting his questions in each round independently of his choices of other rounds. Likewise, the verifier checks each round independently, not making any cross checks between different rounds. (Without this last condition, parallel repetition is not applicable to zero knowledge proof systems.)

We distinguish between two types of rounds. Recall that the question sent to P_2 is chosen from one of four possible questions (not necessarily all distinct). If V chose to send one of the first two questions (those sent to P_1), then the round is called a *compare round.* If V chose to send one of the last two questions (those not sent to P_1, though they might get sent to P_1 by chance, if they are identical to one of the first two questions), then the round is called a *confuse round.*

For protocol G, let G^k denote its k-fold parallel repetition.

Proposition 9 *Let G be an $MIP(2,1)$ protocol in C form, with $\omega(G) < 1$. Then*

- $\omega(G^3) < \omega(G)$.
- *There exists some constant q_G that depends on G, $q_G < 1$, such that $\omega(G^k) < (q_G)^k$.*

Proof: The first part of the proposition is a direct consequence of [20], where it is shown that for any protocol with full support, $\omega(G^3) < \omega(G)$. Confuse and compare protocols have full support.

The second part of the proposition follows from the effectiveness of parallel repetition if the questions to the two provers are not correlated [1, 13, 17]. This situation holds for confuse rounds. Thus the second part of the proposition holds even if V supplies the provers with extra information, pointing out the compare rounds. □

Despite the exponential reduction in error, Proposition 9 is not strong enough for decreasing the error of $MIP(2,1)$ proof systems. The value of q_G in our proof turns out to be too close to 1. In order to decrease the error below $1/2$, the required number of repetitions becomes exponential in the length of the questions.

2.1 Proof of Theorem 2

The protocol $\hat{G}$ has $3m$ rounds, about half of which are compare rounds, and the other confuse rounds. To simplify the presentation of the proof, we supply the provers with information. In some rounds, we reveal to each prover the question that each other received. We leave m random rounds untouched, n of which are compare rounds, and $M = m - n$ confuse rounds. Hence it suffices to analyze a protocol which has M confuse rounds and n compare rounds, in random order (unknown to the provers). We require that $n << M$.

Fix the strategies of P_1 and P_2 to ones that are both optimal and deterministic (such strategies always exist). We start by first analyzing the possible behaviors of P_2, completely ignoring P_1, and the fact that the strategies of P_1 and P_2 need to be coordinated. q^i denotes the question that P_2 receives in round i, and a^j denotes the answer of P_2 in round j. P_2 is a collection of m functions, where f_i is a function from m-tuples (the m questions that P_2 receives) to the answer a^i. Given an assignment to the m question, P_2 replies with an *answer sequence.* In our analysis, we shall consider the sequence of answers of P_2 on selected rounds, which

we call an *answer subsequence*. We will also consider answer subsequences when we have only partial information about the m-tuple of questions (only some of the rounds are revealed to us). This induces a probability distribution on the possible answer subsequences, and we shall use the term PASS (*probabilistic answer subsequence*) to describe it.

Definition 3 *Let $I \subset \{1,m\}$ and $J \subset I$ be index sets. Given an assignment $Q = q_1,...,q_I$ of questions to the rounds specified by I, the probability of the sequence of answers $A = a_1,...,a_J$ to the rounds specified by J is denoted by $PASS[(J,A)|(I,Q)]$.*

$$PASS[(J,A)|(I,Q)] = Pr[a^{j_1} = a_1,...,a^{j_J} = a_J \mid q^{i_1} = q_1,...,q^{i_I} = q_I]$$

where the probability is taken over the choice of questions in rounds $\{1,m\} \setminus I$.

Opening all rounds determines $PASS[(J,A)]$ completely (it becomes either 0 or 1). A central point in our proof is the following lemma, which shows that opening of just one randomly chosen round has only limited influence on PASS. It is best understood as a lemma on the influence of a single random variable on the value of a multi-variate function. The variables of the function are the rounds $\bar{I}$. The single variable is round $i \in \bar{I}$, which is chosen at random, and is given a random value q. The function is the indicator function that is 1 if P_2 replies to rounds J with answer set A, and 0 otherwise.

Lemma 10 *Let $\bar{I} = \{1,m\} \setminus I$, and let $M = |\bar{I}|$. Let $(I,Q)(i,q)$ denote the extension of (I,Q) that results from opening question q at round $i \in \bar{I}$. Then for any PASS (I,Q,J,A),*

$$Pr_{(i,q)}\left[\begin{array}{l} PASS[(J,A)|(I,Q)(i,q)] \\ -PASS[(J,A)|(I,Q)] \end{array} > \epsilon\right] < \epsilon$$

where $\epsilon = \max[\frac{128 \ln M}{M}, (\frac{8 \ln M}{M})^{1/3}]$.

Proof: A pair (i,q) for which $[PASS[(J,A)|(I,Q)(i,q)] - PASS[(J,A)|(I,Q)] > \epsilon]$ is called *biased*. Let P be the probability of choosing a biased pair. Assume that $P > \epsilon$, and derive a contradiction. Let x_i be the random variable denoting the event of choosing a biased pair, conditioned on round i being chosen, and let $p_i = Pr[x_i]$. Then $\sum_{i \in \bar{I}} p_i = P \cdot M$. Consider $S = \sum_{i \in \bar{I}} x_i$. Its expectation is $E[S] = P \cdot M$. The x_i are independent, and we use the following Chernoff bound (see e.g., [29]):

$$Pr[S \geq \beta E[S]] \leq e^{(1-1/\beta-\ln\beta)\beta E[S]}$$

Expressing $\beta = 1 + x$, we can use:

$$\frac{1}{1+x} > 1 - x + x^2 - x^3 \quad \text{and}$$

$$\ln(1+x) > x - \frac{x^2}{2}$$

One obtains:

$$Pr[S \geq (1+x)E[S]] \leq e^{(-x^2/2+x^3)E[S]}$$

For $E[S] > 128 \ln M$ (corresponding to $\epsilon > \frac{128 \ln M}{M}$), and $x = \sqrt{\frac{8 \ln M}{E[S]}} < 1/4$, we get

$$Pr[S \geq \beta E[S]] \leq M^{-2}$$

$E[S]$ expresses the expectation of the number of biased questions over all question sequences that on coordinates I have question set Q, regardless of whether the answer set A was given on coordinates J. We now bound the expectation of S, conditioned on the fact that answer set A was indeed given:

$$\begin{aligned} E[S|(J,A)] &\leq M \cdot Pr[S \geq \beta E[S]] + \beta E[S] \\ &\leq \frac{1}{M} + PM(1 + \sqrt{\frac{8 \ln M}{PM}}) \\ &\simeq PM(1 + \sqrt{\frac{8 \ln M}{PM}}) \end{aligned}$$

Hence, when an average biased pair (i,q) is picked:

$$\begin{aligned} &PASS[(J,A)|(I,Q)(i,q)] \leq \\ &\quad \frac{PASS[(J,A)|(I,Q)]E[S|(J,A)]}{E[S]} \\ &\leq PASS[(J,A)|(I,Q)](1 + \sqrt{\frac{8 \ln M}{PM}}) \\ &\leq PASS[(J,A)|(I,Q)] + \sqrt{\frac{8 \ln M}{PM}} \end{aligned}$$

Contradicting the choice of $P > \epsilon > (\frac{8 \ln M}{M})^{1/3}$. □

Lemma 10 shows that knowing an additional question (i,q) does not provide much help in predicting the answer set A on a set of questions indexed by J. Nevertheless, the next lemma shows a strong correlation between the answers to the set of questions indexed by J and the answer to the additional question q.

Definition 4 *The PASS (I,Q,J,A) is* alive *if $PASS[(J,A)|(I,Q)] \geq \varepsilon$. The* challenge *$(I,Q,J)$ is* alive *if one of its PASS is alive. Otherwise, it is* dead.

Definition 5 *The challenge (I, Q, J) and the question (i, q) $(i \in \bar{I})$ are* ε-correlated *if for every live PASS (I, Q, J, A), there exists an answer a such that:*

$$PASS[(J, A)(i, a)|(I, Q)(i, q)] \geq (1 - \varepsilon)PASS[(J, A)|(I, Q)(i, q)]$$

That is, if answer sequence A is given, this practically forces a particular answer a to q.

Two different error parameters were introduced, ϵ in Lemma 10, and ε in Definitions 5 and 4. We require that $\epsilon << \varepsilon$, or more specifically, $8\epsilon < \varepsilon^5$. We also require $n > 4\varepsilon^{-5}$.

Lemma 11 *There exists a* good *j, $1 \leq j \leq n/2$, such that one of the following two conditions holds:*

1. *A fraction of $(1-\varepsilon)$ of the challenges (I, Q, I), for $|I| = j$, are dead.*
2. *A fraction of $(1-\varepsilon)$ of the live challenges (I, Q, I), for $|I| = j$, are ε-correlated with a $(1 - \varepsilon)$ fraction of the respective (i, q), $i \in \bar{I}$.*

Proof: For each j, let E_j denote the following expectation:

$$E_j = E_{(I,Q|\,|I|=j)}[\sum_A (PASS[(I, A)|(I, Q)])^2]$$

For every j, $0 < E_j \leq 1$. We will show that if there is no good j then $E_{n/2} < 0$, contradiction.

Assume that j is not good, and compute $E_j - E_{j+1}$. At least an ε fraction of the challenges (I, Q, I) are live. At least an ε fraction of the live challenges are not ε-correlated with at least an ε fraction of their respective (i, q). Then for live strategy A (having probability at least ε), for any a:

$$\begin{aligned} & PASS[(I, A)(i, a)|(I, Q)(i, q)] \\ \leq \quad & (1 - \varepsilon)PASS[(I, A)|(I, Q)(i, q)] \end{aligned}$$

This amounts to a loss of an ε-fraction in $\sum_a (PASS[(I, A)(i, a)|(I, Q)(i, q)])^2$ relative to

$$\max[(PASS[(I, A)|(I, Q)(i, q)])^2, (PASS[(I, A)|(I, Q)])^2]$$

Altogether, there is an expected loss of at least ε^5.

This is offset to some extent by some gain. The extent of this gain is bounded by Lemma 10. Fix (I, Q). Then:

$$\begin{aligned} & E_{j+1}[(I, Q)] \\ & = \sum_{A,i,q,a} Pr[(i, q)](PASS[(I, A)(i, a)|(I, Q)(i, q)])^2 \\ & \leq \sum_{A,i,q} Pr[(i, q)](PASS[(I, A)|(I, Q)(i, q)])^2 \end{aligned}$$

We extract away all the cases that $PASS[(I, A)|(I, Q)(i, q)] \geq PASS[(I, A)|(I, Q)] + \epsilon$. For these cases we assume the worst, that $PASS[(I, A)|(I, Q)(i, q)] = 1$. By Lemma 10, the probability of these cases (over the choice of (i, q)) is at most ϵ, and hence their contribution to $E_{j+1}[(I, Q)]$ is at most ϵ. We evaluate the rest of the sum, which we denote by $E^-_{j+1}[(I, Q)]$, under the condition that for every (A, i, q) that participate in the sum, $PASS[(I, A)|(I, Q)(i, q)] \leq PASS[(I, A)|(I, Q)] + \epsilon$. We obtain:

$$E^-_{j+1}[(I, Q)] \leq \sum_{i,q} Pr[(i, q)] \sum_A (PASS[(I, A)|(I, Q)(i, q)])^2$$

We use the notation $\Delta_{A,i,q}$ to denote $PASS[(I, A)|(I, Q)(i, q)] - PASS[(I, A)|(I, Q)]$, if this difference is between 0 and ϵ, and $\Delta_{A,i,q} = 0$ otherwise. By our condition that for every (A, i, q) that participate in the sum, $PASS[(I, A)|(I, Q)(i, q)] \leq PASS[(I, A)|(I, Q)] + \epsilon$. We obtain:

$$\begin{aligned} & \textstyle\sum_A (PASS[(I, A)|(I, Q)(i, q)])^2 \\ & \leq \textstyle\sum_A (PASS[(I, A)|(I, Q)] + \Delta_{A,i,q})^2 \\ & \leq \textstyle\sum_A (PASS[(I, A)|(I, Q)])^2 + 2\epsilon + \sum_A \Delta^2_{A,i,q} \end{aligned}$$

Note that $\Delta_{A,i,q} \leq \epsilon$, and that for any (i, q), $\sum_A \Delta_{A,i,q} \leq 1$ (this follows from definition of $\Delta_{A,i,q}$ and the fact that $\sum_A PASS[(I, A)|(I, Q)(i, q)] = 1$). Hence $\sum_A \Delta^2_{A,i,q} \leq \epsilon$, and

$$\begin{aligned} & E^-_{j+1}[(I, Q)] \\ & \leq \textstyle\sum_{i,q} Pr[(i, q)](\sum_A (PASS[(I, A)|(I, Q)])^2 + 3\epsilon) \\ & \leq E_j[(I, Q)] + 3\epsilon \end{aligned}$$

Using $E_{j+1}[(I, Q)] \leq E^-_{j+1}[(I, Q)] + \epsilon$, the proof for $8\epsilon < \varepsilon^5$, and $4\varepsilon^{-5} < n$. □

Now we are ready to present the proof of our main theorem. We select the questions to the provers in a different (but equivalent) order. Let $j < n/2$ be a *good* j, as guaranteed to exist by Lemma 11. We first select at random a set I of j compare rounds. For these rounds we select at random a corresponding question set Q to be received by P_2. Now by the fact that j is good, with probability $1 - \varepsilon$ one of the following two events hold:

1. The challenge (I, Q, I) is dead.
2. The challenge (I, Q, I) is ε-correlated with a $(1-\varepsilon)$ fraction of the respective (i, q), $i \in \bar{I}$.

We first address event (1) above. Cluster the answer sets for (I, Q) in at most $2/\varepsilon$ clusters, such that the sum

of probabilities of answer sets in each cluster is between $\varepsilon/2$ and ε. We select $n-j$ additional compare rounds at random (N denotes the set of compare rounds) together with their respective questions Q^N to P_2. By Lemma 10, with probability at least $1-2n\epsilon/\varepsilon$, for any cluster $\mathcal{A}$ on I, it holds that $PASS[(I,\mathcal{A})|(N,Q^N)] < \varepsilon + n\epsilon$. Now complete all the questions to P_1. No matter what P_1's answers are on Q, the probability (over the possible completions of questions to P_2) that P_2 answers similarly is at most $2n\epsilon/\varepsilon + \varepsilon + n\epsilon \simeq \varepsilon$, assuming that $\epsilon << \varepsilon^2/n$.

We now address event (2) above. There are at most $1/\varepsilon$ live answer sets for (I,Q,I). An argument similar to event (1) above shows that we can ignore answer sets that are not live, paying at most ε in the error probability. For each live answer set A, and for each (i,q), $i \in \bar{I}$, call P_2's most likely answer a the *majority answer*. Then for $(1-\varepsilon)$ of the choices of (i,q), the majority answer has probability at least $(1-\varepsilon)$. Select the remaining $n-j$ corresponding compare rounds, together with their respective questions. By Markov's inequality, with probability at least $1-\delta$, P_2 answers a fraction of at least $1-2\varepsilon/\delta$ of the questions according to the majority strategy.

Now we turn to consider P_1. We complete all the questions for P_1, and tell him for free which are the first j compare rounds, and which questions P_2 received in these rounds. We may assume that P_1's answers on these rounds correspond to a live answer set for P_2. The fact that there are $1/\varepsilon$ such live answer sets does not significantly influence our analysis. Now we tell P_1 which are the other $n-j$ compare rounds, without showing P_1 the questions that P_2 was asked on these rounds. Prover P_1 must answer a fraction of nearly $1-2\varepsilon/\delta$ of the questions according to the majority strategy. Otherwise, there is overwhelming probability that P_1's answers would somewhere contradict P_2's answers. But with overwhelming probability, the majority strategy fails on a fraction of at least $\frac{1-p}{2}$ of the $n-j$ rounds. Hence P_1 cannot succeed if $1-2\varepsilon/\delta > \frac{1+p}{2}$.

The parameters for the proof system:

$n >> \frac{1}{1-p}\log\delta^{-1}$, so that any fixed set of strategies will have probability at most δ of succeeding on more that $\frac{1+p}{2}$ of the rounds.

$\varepsilon < \frac{1-p}{4}\delta$, so that P_1 does not satisfy the $n-j$ remaining compare rounds.

$\epsilon << \varepsilon^2/n$, since this is assumed in event (1) above.

$n > 4\varepsilon^{-5}$, by Lemma 11.

$8\epsilon < \varepsilon^5$, by Lemma 11.

$\epsilon = \max[\frac{128\ln M}{M}, (\frac{8\ln M}{M})^{1/3}]$, by Lemma 10. Observe that if $n << m$, then $m/M \simeq 1$.

If we assume that $p = 1/2$, one would ideally hope to obtain error probability δ by $O(\log\delta^{-1})$ repetitions. However, the proof for our construction may require $m > \delta^{-21}$.

3 2-provers, 1-round, 0-knowledge.

We modify the error reduction method of [28, 19] so as to preserve zero knowledge. An interesting property of the new transformation is that it only requires the original proof system to be zero knowledge with respect to an especially honest class of verifiers that we denote as *verifreiers*. This term is a play on the Israeli slang term "freier," roughly meaning one who knowingly allows himself to be manipulated by others due to his anachronistic ideals. For more details on the term freier, see [32]. After our transformation, the resulting interactive proof system is zero-knowledge against *all* verifiers.

A verifreier is like an honest verifier, except that instead of having a *query* operation it has *query* and *read* operation. When a verifreier executes a query operation on a set of bits of the oracle proof, it conceptually does not look at these bits and its view from this operation is null. To obtain the value of a bit that has been queried, it must explicitly execute a read operation on a that bit. While in a one-round proof system the query is performed nonadaptively, the strategy for executing the reads may be adaptive and designed to limit the information obtained by the verifreier. In the appendix, we give a simple interactive proof system for 3-SAT that is zero-knowledge against its verifreier.

3.1 Review of the Feige-Lovasz protocol.

We now review the Feige-Lovasz transformation from multi-prover one-round proof systems, to two-prover one-round proof systems. For simplicity of notation, we describe the transformation in terms of oracle protocols. We assume that the oracle proof is as follows. On input x, V uses his random bits R to generate queries $q_1,\ldots,q_m$, where $|q_i| = \ell$, and sends them to the oracle O, viewed as a function $O : \{0,1\}^\ell \to \{0,1\}$. He then receives the value of $a_i = O(q_i)$ for $1 \le i \le m$, computes $\mathbf{accept}(x,R,q_1,\ldots,q_m,a_1,\ldots,a_m)$ and accepts or rejects accordingly. If $x \in L$ then (O,V) always accepts (perfect completeness) and if $x \notin L$ then for all $\hat{O}$, $(\hat{O},V)$ accepts with probability at most p^2. Note that in the oracle model, it is straightforward to obtain very low error probabilities.

The Feige-Lovasz transformation results in a protocol with perfect completeness and an error probability of at most p. Let $k = max(\ell, -\log p, \log(4m))$, let $N > 2^{9k}$ be a prime, and let $F = GF(N)$. Let O_F denote the unique multilinear extension of O onto F^ℓ. Given a point

$$\alpha = (\alpha_1,\alpha_2,\ldots,\alpha_\ell) \in \{0,1\}^\ell$$

and a set of slopes $\beta = (\beta_1, \ldots, \beta_\ell)$, let

$$\begin{aligned} \mathcal{L}_{\alpha,\beta}(t) &= (\alpha_1 + \beta_1 t, \ldots, \alpha_\ell + \beta_\ell t) \text{ and} \\ \mathcal{P}_{\alpha,\beta}(t) &= O_F(\mathcal{L}_{\alpha,\beta}(t)). \end{aligned}$$

$\mathcal{P}_{\alpha,\beta}$ is a polynomial of degree at most ℓ. The new protocol (P_1, P_2, V') is as follows:

1. V' simulates V, generating queries $q_1, \ldots, q_m$ from his random bits R. V' then generates $\beta_{i,j}, t_i \in_R F$ for $1 \leq i \leq m$ and $1 \leq j \leq \ell$. V' sends R (implicitly $q_1, \ldots, q_m$) and $\beta_{i,j}$ to P_1, for $1 \leq i \leq m$ and $1 \leq j \leq \ell$. Writing $q_i = (q_{i,1}, \ldots, q_{i,\ell})$, V' sends $\gamma_i = \mathcal{L}_{q_i,\beta_i}(t_i)$ to P_2, for $1 \leq i \leq m$.

2a. P_1 sends $\mathcal{P}_{q_i,\beta_i}$ to V', for $1 \leq i \leq m$, where $\beta_i = (\beta_{i,1}, \ldots, \beta_{i,\ell})$. Note that $\mathcal{P}_{q_i,\beta_i}$ can be represented compactly as $\ell + 1$ elements of F and that $a_i = \mathcal{P}_{q_i,\beta_i}(0)$.

2b. For $1 \leq i \leq m$, P_2 sends $z_i = O_F(\gamma_i)$ to V'.

3. V' accepts if $\mathbf{accept}(x, R, q_1, \ldots, q_m, a_1, \ldots, a_m)$ is true (V would have accepted), and $z_i = \mathcal{P}_{q_i,\beta_i}(t_i)$ for $1 \leq i \leq m$. Otherwise, V' rejects.

3.2 A zero-knowledge protocol

We achieve zero-knowledge by augmenting O and restricting V''s questions. The modification to O will not affect the original proof of soundness and the restriction on V''s questions will affect the soundness in an easily quantifiable manner. First, we assume that the original protocol has perfect completeness and is zero-knowledge: Given any set of random bits R, one can simulate the values of $a_i = O(q_i)$ that V would have read (e.g., as in [16]). For each i, $1 \leq i \leq m$, the provers randomly construct a new oracle $O^{[i]} : \{0,1\}^{\ell+2} \to \{0,1\}$ such that $O^{[i]}(x00) = O(x)$ and $O^{[i]}(xb_1b_2)$ is uniformly distributed over F for $(b_1, b_2) \neq (0,0)$. We define $O_F^{[i]}$ as the unique multilinear extension of $O^{[i]}$, and we define

$$P_{\alpha,\beta}^{[i]}(t) = O_F^{[i]}(\alpha_1 + \beta_1 t, \ldots, \alpha_{\ell+2} + \beta_{\ell+2} t).$$

V^* converts a question q_i into q_i00. P_1 requires that $\beta_1, \ldots, \beta_{\ell+2}$ are nonzero and P_2 requires that no test slopes be nonzero. If either prover sees a violation of these conditions, they abort. For ease of analysis, we assume that V^* proceeds according to the Feige-Lovasz protocol, and accepts if he happens to break the constraints on the slopes or the test points (this will happen only with low probability if $|F|$ is large). The resulting protocol is given in Figure 1.

zero-knowledge(O, V^*, x)
Our convention is that i ranges over $\{1, \ldots, m\}$ and j ranges over $\{1, \ldots, \ell+2\}$ unless stated otherwise.

0. Using their shared randomness, P_1 and P_2 randomly generate $O^{[i]}$ as described above.

1. V^* generates $\beta_i \in_R F^{\ell+2}$, $t_i \in_R F$ and queries q_i as V would have on input x, using random bits R. We write $\gamma_i = \mathcal{L}_{q_i,\beta_i}(t)$. If for any i, some component of γ_i or β_i is 0, then V^* halts and accepts. Otherwise, V^* sends R (and hence $q_1, \ldots, q_m$) and $\beta_{i,j}$ to P_1, and sends γ_i to P_2.

2a. If for any i, some component of β_i is 0, then P_1 aborts. Otherwise, P_1 computes the set of i for which the verifreier would have read a_i, and sends $\mathcal{P}_{q_i,\beta_i}$ to V^*. Note that $\mathcal{P}_{q_i,\beta_i}$ can be represented compactly as $\ell + 2$ elements of F and that $a_i = \mathcal{P}_{q_i,\beta_i}(0)$. For those answers which V would not have looked at, P_1 sends a null message.

2b. If for any i, some component of γ_i is 0, then P_2 aborts. Otherwise, P_2 sends $z_i = O_F(\gamma_i)$ to V^*.

3. V^* accepts if $\mathbf{accept}(x, R, q_1, \ldots, q_m, a_1, \ldots, a_m)$ is true (i.e., V would have accepted) and for every i such that q_i would have been read by V, $z_i = \mathcal{P}_{q_i,\beta_i}(t_i)$. Otherwise, V^* rejects.

Figure 1: A low-error one round zero-knowledge proof system.

3.3 Analysis

The analysis of soundness and our completeness for our protocol is straightforward given the analysis in [19] (which in turn, is based on the analysis in [28]).

Theorem 12 *If $x \in L$, then (P_1, P_2, V^*) accepts with probability 1. If $x \notin L$ then for any $(\hat{P}_1, \hat{P}_2)$, $(V^*, \hat{P}_1, \hat{P}_2)$ accepts with probability at most $p + 2lm/|F|$.*

Proof: The perfect completeness property inherits straightforwardly from the original oracle-proof. Note that if V^* follows the protocol, he will never give P_1 or P_2 an input that would cause them to abort. Furthermore, the consistency checks made by V^* on $\mathcal{P}_{q_i,\beta_i}$ and

z_i will always be met, by the properties of multilinear functions.

To bound the probability of V^* accepting $x \notin L$, modify the protocol so that V^* does not automatically accept if a component of β_i or γ_i is 0. In this case, then by the same analysis as was done by Feige-Lovasz, V^* will accept with probability at most p.[3] By a simple analysis each component of β_i or γ_i is 0 with probability at most $1/|F|$. The probability of any of these events occurring is thus bounded by $2lm/|F|$, and the probability of V^* accepting is thus bounded above by $p + 2lm/|F|$. □

Next, we show that the resulting protocol is zero-knowledge if the original protocol is.

Theorem 13 *Suppose that a probabilistic oracle proof $(O, \hat{V})$ is zero knowledge with respect to an honest verifier. That is, there exists a simulator M_O that when given any R can perfectly simulate the values of $a_i = O(q_i)$ that V would have read given R. Then there exists a simulator M that perfectly simulates the view obtained by any verifier $\hat{V}$ from talking to P_1 and P_2.*

Proof: Before giving the simulator for our protocol, we first prove a useful lemma concerning multilinear extensions of random functions over F.

Lemma 14 *Let Q be a random function from $\{0,1\}^k \to F$, where $|F| > k+1$ and let Q_F be the multilinear extension of Q. Let $\alpha = (\alpha_1, \ldots \alpha_k) \in F^k$, $\beta = (\beta_1, \ldots, \beta_k) \in (F-0)^k$ and $\gamma = (\gamma_1, \ldots, \gamma_k) \in F^k$. Let the parametric line $\mathcal{L}_{\alpha,\beta}$ be defined by*

$$\mathcal{L}_{\alpha,\beta}(t) = (\alpha_1 + \beta_1 t, \ldots, \alpha_k + \beta_k t)$$

for $t \in F$ and let the polynomial $\mathcal{P}_{\alpha,\beta}$ be defined by $\mathcal{P}_{\alpha,\beta}(t) = Q_F(\mathcal{L}_{\alpha,\beta}(t))$. If γ does not lie on $\mathcal{L}_{\alpha,\beta}$, then the induced distribution on $(\mathcal{P}_{\alpha,\beta}, Q_F(\gamma))$ is uniform over all pairs (P, z), where P is a polynomial of degree at most k over F and $z \in F$.

Proof: Our proof is by induction on k. For $k = 1$, the lemma is vacuous. For $k = 2$, we note that choosing Q uniformly is equivalent to uniformly choosing the multilinear polynomial computing Q_F. We write $Q_F(x, y) = axy + bx + cy + d$, where $a, b, c, d \in F$ are uniformly distributed. First, we consider the special case where $\alpha = \vec{0}$. Substituting $x = \beta_1 t$ and $y = \beta_2 t$ we have

$$\begin{aligned} \mathcal{P}_{\alpha,\beta}(t) &= a\beta_1\beta_2 t^2 + (b\beta_1 + c\beta_2)t + d \text{ and} \\ Q_F(\gamma_1, \gamma_2) &= a\gamma_1\gamma_2 + b\gamma_1 + c\gamma_2 + d \end{aligned}$$

Since a is uniformly distributed and $\beta_1\beta_2$ is nonzero, the t^2 term of $\mathcal{P}_{\alpha,\beta}(t)$ is uniformly distributed. Fixing a, the free term of $\mathcal{P}_{\alpha,\beta}(t)$ is uniformly distributed since d is. It remains to show that if we fix a and d, the distribution on $(b\beta_1 + c\beta_2)$ and $a\gamma_1\gamma_2 + b\gamma_1 + c\gamma_2 + d$ will be uniformly and independently distributed. It suffices to show that $(b\beta_1 + c\beta_2)$ and $(b\gamma_1 + c\gamma_2)$ will be uniformly distributed. Since $\alpha = (0,0)$ and (γ_1, γ_2) is not on $\mathcal{L}_{\alpha,\beta}(t)$, (β_1, β_2) and (γ_1, γ_2) are linearly independent. Therefore, for each distinct value of $(b, c) \in F \times F$, $(b\beta_1 + c\beta_2, b\gamma_1 + c\gamma_2)$ has a distinct value in $F \times F$, and thus is uniformly distributed since (b, c) is.

For $\alpha \neq \vec{0}$ we note that multilinear functions are closed under translations. That is, if $Q_F(x, y)$ is multilinear, then $Q'_F(x, y) = Q_F(x_1 + \alpha_1, x_2 + \alpha_2)$ is multilinear, and there is a bijection between Q and Q' defined by $Q'(x, y) = Q_F(x_1 + \alpha_1, x_2 + \alpha_2)$ for $x_1, x_2 \in \{0, 1\}$. In particular, is Q is uniformly distributed, then so is Q'. Defining $\mathcal{P}'_{\alpha,\beta}(t) = Q'_F(\mathcal{L}_{\alpha,\beta}(t))$ and $\gamma'_i = \gamma_i - \alpha_i$, we have the identities

$$\begin{aligned} Q_F(\gamma_1, \gamma_2) &= Q'_F(\gamma'_1, \gamma'_2) \text{ and} \\ \mathcal{P}_{\alpha,\beta}(t) &= \mathcal{P}'_{\vec{0},\beta}(t). \end{aligned}$$

Note that (γ'_1, γ'_2) is on $\mathcal{L}_{\vec{0},\beta}$ iff (γ_1, γ_2) is on $\mathcal{L}_{\alpha,\beta}$. Since $Q'_F(\gamma'_1, \gamma'_2)$ and $\mathcal{P}'_{\vec{0},\beta}$ will be uniformly distributed by the previous argument, then so will $Q_F(\gamma_1, \gamma_2)$ and $\mathcal{P}_{\alpha,\beta}$.

For $k > 2$, we assume without loss of generality that the projection of γ onto the first $k-1$ dimensions is not contained in the projection of $\mathcal{L}_{\alpha,\beta}(t)$ onto the first $k-1$ dimensions. Otherwise, we can make the same argument, just for a different dimension. Let $\mathcal{L}^{(b)}_{\alpha,\beta}(t)$ denote the projection of $\mathcal{L}_{\alpha,\beta}(t)$ onto the hyperplane $x_k = b$ (i.e., what we obtain by setting $\alpha_k = b$ and $\beta_k = 0$). Similarly, we define

$$\mathcal{P}^{(b)}_{\alpha,\beta}(t) = Q_F(\mathcal{L}^{(b)}_{\alpha,\beta}(t)).$$

By the properties of multilinear functions, we have

$$\begin{aligned} Q_F(\gamma_1, \ldots, \gamma_{k-1}, \gamma_k) &= \gamma_k Q_F(\gamma_1, \ldots, \gamma_{k-1}, 1) \\ &\quad + (1 - \gamma_k) Q_F(\gamma_1, \ldots, \gamma_{k-1}, 0), \quad \text{and} \end{aligned}$$

$$\mathcal{P}_{\alpha,\beta}(t) = (\alpha_k + \beta_k t)\mathcal{P}^{(1)}_{\alpha,\beta}(t) + (1 - (\alpha_k + \beta_k t))\mathcal{P}^{(0)}_{\alpha,\beta}(t)$$

Also, for $i \in \{0, 1\}$ the extension Q_F restricted to the hyperplane $x_k = i$ is simply the multilinear extension of Q restricted to the hyperplane $x_k = i$. Now the values of Q on the hyperplane $x_k = 1$ are independent of the values of Q restricted to the hyperplane $x_k = 0$. By this and our inductive hypothesis, we have $Q_F(\gamma_1, \ldots, \gamma_{k-1}, 0)$, $Q_F(\gamma_1, \ldots, \gamma_{k-1}, 1)$, $\mathcal{P}^{(0)}_{\alpha,\beta}(t)$ and $\mathcal{P}^{(1)}_{\alpha,\beta}(t)$ are distributed uniformly and independently. The lemma follows. □

We now state our main lemma concerning the view of a possibly malicious verifier $\hat{V}$.

[3] The only deviation from the Feige-Lovasz protocol is that parts of the oracle O is randomized for each query, and this has no effect on the analysis.

Lemma 15 *Let $q = q'00$, where $q' \in \{0,1\}^\ell$, let $\beta, \gamma \in (F-0)^{\ell+2}$ and let $O^{[i]}$ be generated as in the above protocol. If for some t, $\gamma = \mathcal{L}_{q,\beta}(t)$, then $\left(\mathcal{P}^{[i]}_{q,\beta}, O^{[i]}_F(\gamma)\right)$ will be distributed uniformly over all pairs $(\mathcal{P}, z)$, where P is a polynomial over $F[t]$ of degree at most $\ell+2$ with free term $O(q')$ and $z = P(t)$. If γ is not on $\mathcal{L}_{q,\beta}(t)$, then $\left(\mathcal{P}^{[i]}_{q,\beta}, O^{[i]}_F(\gamma)\right)$ will be distributed uniformly over all pairs $(\mathcal{P}, z)$, where $\mathcal{P}$ is as before and $z \in F$.*

Proof: Let $q^{(b)}$ denote q projected onto the hyperplane $x_{\ell+2} = b$ and similarly for $\mathcal{L}^{(b)}_{\alpha,\beta}$, $\mathcal{P}^{(b)}_{\alpha,\beta}$ and $\gamma^{(b)}$. Writing $\beta_i = (\beta_{i,1}, \ldots, \beta_{i,\ell+2})$ we have the identity

$$\mathcal{P}_{q,\beta} = (\beta_{i,\ell+2}t)\mathcal{P}^{(1)}_{q,\beta} + (1 - \beta_{i,\ell+2}t)\mathcal{P}^{(0)}_{q,\beta}.$$

Here we use multilinearity and the fact that the last two components of q are 0. Note that for $b \in \{0,1\}$, $\mathcal{P}^{(b)}_{q,\beta}$ depends only on the values of $O^{[i]}$ on the hyperplane $x_{\ell+2} = b$, and that the values of $O^{[i]}$ on the hyperplane $x_{\ell+2} = 1$ are uniformly and independently distributed. Thus, by Lemma 14, $\mathcal{P}^{(1)}_{q,\beta}$ will be uniformly distributed, independently of $\mathcal{P}^{(0)}_{\alpha,\beta}$. Since $\beta_{i,\ell+2} \neq 0$, it follows that all but the free term of $\mathcal{P}_{q,\beta}$ will be uniformly distributed. The free term of $\mathcal{P}_{q,\beta}$ is equal to $O^{[i]}(q)$, which is equal to $O(q)$ since the last two components of q are 0.

If γ is equal to $\mathcal{L}_{q,\beta}(t)$ for some $t \in F$ then $O^{[i]}_F(\gamma) = \mathcal{P}_{q,\beta}(t)$ by definition. If γ is not on $\mathcal{L}_{q,\beta}(t)$ then either one can project them both on all but the last component without a collision, or one can project them both on all but the second to last component without a collision (this follows from the fact that $\beta_1 \neq 0$). We assume the former case without loss of generality; the proof for the latter case is identical. Writing $\gamma = (\gamma_1, \ldots, \gamma_{\ell+2})$, we have

$$O^{[i]}_F(\gamma) = \gamma_{\ell+2}O^{[i]}_F(\gamma^{(1)}) - (\gamma_{\ell+2} - 1)O^{[i]}_F(\gamma^{(0)}).$$

As in the analysis of $\mathcal{P}_{q,\beta}$, the value of $O^{[i]}_F(\gamma^{(1)})$ will be independent of $O^{[i]}_F(\gamma^{(0)})$ and $\mathcal{P}^{(0)}_{\alpha,\beta}$. Furthermore, by Lemma 14, $O^{[i]}_F(\gamma^{(1)})$ will be uniformly distributed and independent of $\mathcal{P}^{(1)}_{\alpha,\beta}$ as well. Since $\gamma_{\ell+2} \neq 0$, it follows that $O^{[i]}_F(\gamma)$ is distributed uniformly and independently of $\mathcal{P}_{q,\beta}$. The lemma follows. □

The simulatability of our protocol follows immediately from Lemma 15 and the simulatability of the original oracle proof. We include an explicit construction of our simulator in Figure 2. □

Remark: In two-prover zero knowledge proof systems, it is often desirable to to have the provers share as few random bits as possible, so as to limit the communication requirements between the two provers before the protocol begins. It follows from [26] that the expected number of shared random bits used by the provers need not exceed the sum of the number of random bits used by the verifier and length of answers of the provers. In our construction, the number of random bits shared by the provers is much larger than necessary. We do not know of a computationally efficient way of reducing the number of shared random bits.

zero-knowledge(M_O, V, x) /* simulator */
Simulation works even if cheating verifier waits for the reply of one prover before deciding on question to the other prover.
M_O denotes the simulator for the original oracle proof.

A. On input R, M generates $q_1, \ldots, q_m$ as truthful V would have (given R), and invokes M_O to generate simulated replies $a_1, \ldots, a_m$. Note that some of these answers will be null (unread by V).

B. On input β_i, M checks that all of the components of β_i are nonzero and simulates P_1 aborting the protocol if not. Otherwise, if a_i is not null, M generates a polynomial $\mathcal{P}_{q_i,\beta_i}$ as follows: If γ_i and z_i have been determined and $\gamma_i = \mathcal{L}_{q_i,\beta_i}(t_i)$ for some t_i, then M generates $\mathcal{P}_{q_i,\beta_i}$ uniformly from all polynomials over F of degree at most $\ell+2$ such that $\mathcal{P}_{q_i,\beta_i}(t_i) = z_i$ and $\mathcal{P}_{q_i,\beta_i}(0) = a_i$. Otherwise, M generates $\mathcal{P}_{q_i,\beta_i}$ uniformly from all polynomials over F of degree at most $\ell+2$ with $\mathcal{P}_{q_i,\beta_i}(0) = a_i$. M sends $\mathcal{P}_{q_i,\beta_i}$ to $\hat{V}$. If a_i is null, then M sends a null message to $\hat{V}$.

C. On input γ_i, M checks that all of the components of γ_i are nonzero, and simulates P_2 aborting the protocol if not. Otherwise, M generates z_i as follows: If q_i, β_i and $\mathcal{P}_{q_i,\beta_i}$ have been determined and $\gamma_i = \mathcal{L}_{q_i,\beta_i}(t_i)$ for some t_i, then M computes $z_i = \mathcal{P}_{q_i,\beta_i}(t_i)$. Otherwise, M generates $z_i \in_R F$.

Figure 2: Simulator for our zero-knowledge proof system.

4 Acknowledgments:

We thank Mihir Bellare, Shafi Goldwasser, Marcos Kiwi, Carsten Lund, Mario Szegedy, for useful discussions, and for comments on earlier versions of our

manuscript. We thank Dina Kravets for help in the proof of Lemma 15.

References

[1] N. Alon, "Probabilistic Methods in Extremal Finite Set Theory", *Proc. of the Conference on Extremal Problems for Finite Sets, Hungary, 1991.*

[2] S. Arora, C. Lund, R. Motwani, M. Sudan, M. Szegedy, "Proof verification and intractability of approximation problems", *FOCS 92, 14–23.*

[3] L. Babai, L. Fortnow, L. Levin, M. Szegedy, "Checking computations in polylogarithmic time", *STOC 91, 21–31.*

[4] L. Babai, L. Fortnow, C. Lund, "Non-Deterministic Exponential Time has Two-Prover Interactive Protocols", *FOCS 90, 16–25.*

[5] M. Bellare, "Interactive proofs and approximation", *ISTCS 93, 266-274.*

[6] M. Bellare, S. Goldwasser, C. Lund, A. Russell, "Efficient probabilistic checkable proofs and applications to approximation", *STOC 93, 294–304.*

[7] M. Bellare, P. Rogaway, "The complexity of approximating a nonlinear program", In: *Complexity in numerical optimization*, P. Pardalos, ed., World Scientific, 1993.

[8] M. Bellare, M. Sudan, "Improved non-approximability results", *STOC 94.*

[9] M. Ben-or, S. Goldwasser, J. Kilian, A. Wigderson, "Multi Prover Interactive Proofs: How to Remove Intractability", *STOC 88, 113–131.*

[10] M. Ben-or, S. Goldwasser, J. Kilian, A. Wigderson, "Efficient identification schemes using two prover interactive proofs", *Crypto 89, 498–506.*

[11] P. Berman, G. Schnitger, "On the Complexity of Approximating the Independent Set Problem", *Information and Computation, Vol. 96, 77–94, 1992.*

[12] J. Cai, A. Condon, R. Lipton, "On Bounded Round Multi-Prover Interactive Proof Systems", *Structures 90, 45–54.*

[13] J. Cai, A. Condon, R. Lipton, "Playing Games of Incomplete Information", *STACS 90.*

[14] J. Cai, A. Condon, R. Lipton, "PSPACE is Provable by Two Provers in One Round", *Structures 91, 110–115.*

[15] H. Chernoff, "A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the Sum of Observations", *Annals of Math. Stat., 23:493–509, 1952.*

[16] C. Dwork, U. Feige, J. Kilian, M. Naor, S. Safra, "Low Communication, 2-Prover Zero-Knowledge Proofs for NP" *Crypto 92, 217–229.*

[17] U. Feige, "On the Success Probability of the Two Provers in One Round Proof Systems", *Structures 91, 116–123.*

[18] U. Feige, S. Goldwasser, L. Lovasz, M. Safra, M. Szegedy, "Approximating Clique is Almost NP-Complete", *FOCS 91, 2–12.*

[19] U. Feige, L. Lovasz, "Two-prover one-round proof systems, their power and their problems", *STOC 1992, 733–744.*

[20] U. Feige, M. Szegedy, *unwritten manuscript.*

[21] L. Fortnow, J. Rompel, M. Sipser, "On the Power of Multi-Prover Interactive Protocols", *Structures 88, 156–161.*

[22] L. Fortnow, J. Rompel, M. Sipser, "Errata for On the Power of Multi-Prover Interactive Protocols", *Structures 90, 318–319.*

[23] S. Goldwasser, S. Micali, C. Rackoff, "The Knowledge Complexity of Interactive Proof Systems", *SIAM Journal on Computing, 18 1, 1989, 186–208.*

[24] J. Kilian, "Strong Separation Models of Multi Prover Interactive Proofs" *DIMACS Workshop on Cryptography, October 1990.*

[25] J. Kilian, M. Naor, "On the complexity of statistical reasoning", *manuscript.*

[26] D. Koller, N. Megiddo, "Constructing small sample spaces satisfying given constraints", *STOC 93, 268–277.*

[27] D. Lapidot, A. Shamir, "A One-Round, Two-Prover, Zero-Knowledge Protocol for NP", *Crypto, 1991.*

[28] D. Lapidot, A. Shamir, "Fully Parallelized Multi Prover Protocols for NEXP-time" *FOCS 91, 13–18.*

[29] F. T. Leighton, "Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes" *Morgan Kaufmann, 1992, pp. 167-168.*

[30] C. Lund, M. Yannakakis, "On the hardness of approximating minimization problems", *STOC 93, 286–293.*

[31] D. Peleg, "On the Maximal Number of Ones in Zero-One Matrices with No Forbidden Rectangles", *manuscript, 1990.*

[32] L. Roniger, M. Feige, "From Pioneer to Freier: the Changing Models of Generalized Exchange in Israel", *European Journal of Sociology, 1992, 33(2), 280-307.*

[33] G. Tardos. "Multi-prover encoding schemes, and 3-prover interactive proofs." *Structures, 1994.*

[34] O. Verbitsky, "Towards the Parallel Repetition Conjecture", *Structures, 1994.*

A Proving satisfiability to a verifreier

Consider a 3-CNF formula

$$\phi = (t_{10}v_{10}, t_{11}v_{11}, t_{12}v_{12}), \ldots, (t_{m0}v_{m0}, t_{m1}v_{m1}, t_{m2}v_{m2})$$

where $v_{ij} \in \{1..n\}$ specifies the jth variable in the ith clause for $1 \leq i \leq m$ and $0 \leq j \leq 2$, and $t_{ij} = 1$ when

this variable should be negated and $t_{ij} = 0$ otherwise. Let $a_1, \ldots, a_n$ be a valid assignment to the variables and let $p_1, \ldots, p_m$ specify for each clause a literal that is satisfied by the assignment. Thus, if $p_5 = 2$, $v_{5,2} = 7$ and $t_{5,2} = 1$ specifies negation, then $a_7 = 0$.

The oracle decides on one random bit b and one random trit s. The proof consists of (b, s) and bits

$$\begin{aligned} x_{ij} &= b \oplus a_{v_{i,j+s}} \oplus t_{i,j+s} \\ y_i &= a_i \oplus b \\ z_i &= p_i + s \end{aligned}$$

where $1 \leq i \leq m$ and $0 \leq j \leq 2$, and $j + s$ and $p_i + s$ are computed mod 3.

The verifreier V picks a random clause $i \in_R \{1..m\}$ and queries

$$b, s, (x_{i0}, \ldots, x_{i2}), (y_{v_{i,0}}, \ldots, y_{v_{i,2}}), z_i.$$

With probability $\frac{1}{2}$, he

(checks that clause was properly constructed)

- reads s,
- chooses $j \in_R \{0, \ldots, 2\}$,
- reads x_{ij} and $y_{v_{i,j+s}}$, and
- accepts if and only if $x_{ij} = y_{v_{i,j+s}} \oplus t_{i,j+s}$, where $j + s$ is computed mod 3.

With probability $\frac{1}{2}$, he

(checks that clause was satisfied)

- reads b, z_i,
- reads x_{z_i}, and
- accepts iff $x_{z_i} \oplus b = 1$.

It is not hard to see that if ϕ is satisfiable, than V accepts, and if at most $m - \alpha m$ clauses of ϕ can be simultaneously satisfied, then V rejects with probability at least $\alpha/6$. Using the MAX-SNP hardness of 3-SAT and [2], it follows that the error probability can be made constant. V uses a logarithmic number of random bits, and the sum of the lengths of the replies of the oracle is constant. The oracle proof system is zero knowledge with respect to the verifreier, though a cheating verifier can obtain the value of a variable by reading just two bits and an honest verifier would also obtain the value of a variable if it read all of the bits that it queried. The verifreier proof system for NP can be scaled up to a proof system for NEXPTIME with polynomial communication.

Any oracle proof system with a verifreier can be converted into a "traditional" oracle proof system [21] (the verifier simply chooses not to base its actions on a bit until it has been officially read), and thereafter to a one round multiple-prover proof system (by asking each question to a different prover, adding one prover to check for consistency with a random previous prover, and repeating this construction many times with independent provers so as to reduce the error). This is transformed into a two-prover proof system, using our modification of [19]. If the original oracle proof system is zero knowledge with respect to the verifreier, then the resulting two-prover proof system is zero knowledge with respect to all verifiers.