ON BOUNDED ROUND MULTI-PROVER INTERACTIVE PROOF SYSTEMS

Jin-yi Cai
Anne Condon
Richard J. Lipton

CS-TR-237-89

December 1989

On Bounded Round Multi-Prover Interactive Proof Systems

Jin-yi Cai *
Department of Computer Science
Princeton University
Princeton, New Jersey 08544

Anne Condon †
Computer Science Department
University of Wisconsin-Madison

Richard J. Lipton ‡
Department of Computer Science
Princeton University
Princeton, New Jersey 08544

Abstract

We compare *bounded round* multi-prover interactive proof systems (MIP's) with *unbounded round* interactive proof systems (IPS's). We show that for any constant ϵ, any language accepted by an unbounded round IPS has a bounded round, 2-prover MIP that has error probability ϵ, resolving an open problem of Fortnow, Rompel and Sipser [7]. To obtain this result, we show that a certain 1-round MIP that simulates the computation of an unbounded round IPS can be executed many times in parallel to significantly reduce its probability of error.

1 Introduction

We study multi-prover interactive proof systems (MIP's), introduced by Ben-Or et al. [3]. For $k \geq 1$, a k-prover interactive proof system is a protocol between a computationally limited verifier V and k powerful provers $P_1, \ldots P_k$. The provers cannot communicate with each other during execution of the protocol. Informally, the provers and the verifier share an input w, and the provers wish to convince the verifier that the input is in a language L. The MIP accepts L with error probability ϵ if for all inputs $w \in L$, the provers can convince the verifier to accept w with probability $> 1 - \epsilon$, but for all $w \notin L$, no provers can convince the verifier to accept w with probability greater than ϵ. In this paper we assume that the verifier is probabilistic and is polynomially time bounded.

*Research supported by NSF grant CCR-8709818.
†Work supported by NSF grant number DCR-8402565
‡Research supported by DARPA and ONR contracts N00014-85-C-0456 and N00014-85-K-0465

The special case when $k = 1$ is the well-known interactive proof system (IPS) of Goldwasser, Micali and Rackoff [8] and Babai [2]. Clearly, MIP's generalize IPS's; however it is not known whether they are actually more powerful. To address this problem, we consider the number of rounds of an interactive proof as a resource. In a *round* of an interactive proof, the verifier sends a string to each prover in turn and receives a response from each prover. A protocol is a *bounded round protocol* if the number of rounds is a constant independent of the input size. The main result of this paper is that any *unbounded round* IPS can be simulated by a *bounded round 2-prover* MIP.

Theorem 1.1 *For any constant ϵ, any language accepted by an unbounded round IPS has a bounded round 2-prover MIP that has error probability ϵ.*

Limitations on the number of rounds of an interactive proof has already been studied in detail for IPS's. Results of Babai [2] and Goldwasser and Sipser [9] show that bounded round IPS's are equivalent to 2-round IPS's. However, it is not known whether unbounded round IPS's are equivalent to bounded round IPS's. In fact, Aiello et al. [1] showed that there exists an oracle A and a language L such that, relativized to A, L is accepted by an unbounded round IPS but not by a bounded round IPS. Thus our result shows that bounded round, 2-prover MIP's may be more powerful than bounded round IPS's.

The problem of simulating unbounded round interactive proof systems by bounded round multi-prover systems has previously been considered by Fortnow, Rompel and Sipser [7]. They proved the following result.

Theorem 1.2 (Fortnow, Rompel, Sipser, [7]) *If L is accepted by an unbounded round IPS, say (P_L, V_L), then there is a 1-round, 2-prover MIP (P_1, P_2, V) with the following property. For some polynomial $p(n)$ and some N, for all strings w of length $n \geq N$,*

1. *if $w \in L$, then the probability w is accepted by (P_1, P_2, V) is $\geq 1 - 1/2^n$; and*
2. *if $w \notin L$, then for all provers P_1^* and P_2^*, the probability w is accepted by (P_1^*, P_2^*, V) is $\leq 1 - 1/p(n)$.*

For completeness, we describe in Section 2 a protocol very similar to that of [7] and show that it satisfies these properties. Throughout, we call such a MIP an *IPS-simulation protocol.* However, for some pair of provers, this IPS-simulation protocol may accept strings not in L with probability as high as $1 - 1/p(n)$. To overcome this problem, Fortnow et al. proposed that this protocol be repeated many times in parallel to reduce the probability of error. More precisely, suppose we define the *d-product* of a MIP (P_1, P_2, V) to be the MIP $(P_{1,d}, P_{2,d}, V_d)$ obtained by executing (P_1, P_2, V) d times in parallel and accepting if and only if all executions of the MIP accept. Then Fortnow et al. claimed that if for all provers P_1^*, P_2^*, (P_1^*, P_2^*, V) accepts an input w with probability $\leq p$, then for all provers $P_{1,d}^*, P_{2,d}^*$, the probability that $(P_{1,d}^*, P_{2,d}^*, V_d)$ accepts w is $\leq p^d$. This claim, combined with Theorem 1.2, would be sufficient to prove Theorem 1.1.

However, Fortnow [6] later showed that surprisingly, this claim is false. To motivate our new results in this paper, we explain Fortow's argument in Section 2. Thus it remained an open problem whether Theorem 1.1 was true, and, more fundamentally, whether the probability of acceptance of a MIP can be decreased by taking the product of the MIP. The main technical contribution of this paper is to show that by taking the product of the IPS-simulation protocol, its probability of error on strings not in L can be lowered significantly.

Theorem 1.3 *For some polynomial $d(n)$, the $d(n)$-product of the IPS-simulation protocol of Theorem 1.2 for language L accepts any $w \notin L$ with probability at most $7/8$.*

This theorem is proved in Section 3. From this theorem, we obtain the main result, Theorem 1.1, easily. However, it is still an open problem whether for some polynomial $d(n)$, the error probability of the $d(n)$-product game is exponentially small in n. In Section 4, we discuss this and other interesting open problems.

2 Background

We begin this section with a definition of a multi-prover interactive proof system. A *k-prover interactive proof* is a tuple $(P_1, P_2, \ldots, P_k, V)$, where V is a probabilistic Turing machine, with a read-only input tape, a read-write worktape and a source of random bits (a coin). In addition, the verifier has k special communication tapes; informally, these allow the verifier to communicate with the provers.

The states of V are partitioned into two types, reading and communication states. A transition function describes the one-step transitions of the verifier. Whenever the verifier is in a reading state, the transition function of the verifier determines the next configuration of the verifier, based on the symbol under the tape heads, the state and the outcome of an unbiased coin toss. Whenever the verifier is in a communication state, the next configuration is determined as follows. For $1 \leq i \leq k$, the contents of the ith communication tape are replaced by a string written by the ith prover. Then, the next state of the verifier is determined just as when V is in a reading state.

Each prover P_i is specified by a prover transition function. This function determines what string is written by the prover, based on the input and the sequence of all past strings written by the verifier on the ith communication tape. Formally, if Σ is the tape alphabet of the verifier and provers, and $\mathcal{H}$ is the set of all finite sequences of strings over Σ^*, then P_i is a mapping $P_i : \Sigma^* \times \mathcal{H} \rightarrow \Sigma^*$.

We say $(P_1, \ldots P_k, V)$ is a k-prover interactive proof for language L with error probability $\epsilon < 1/2$ if there is some N such that for all strings w of length $\geq N$,

- if $w \in L$, the probability that $(P_1, \ldots P_k, V)$ accepts w is $> 1 - \epsilon$,
- if $w \notin L$, and all provers $P_1^*, \ldots, P_k^*$, the probability that $(P_1^*, \ldots, P_k^*, V)$ rejects w is $> 1 - \epsilon$.

The *number of rounds* of a protocol is the number of times V enters a communication state. If at the jth round, V enters a communication state with a string x written on ith communication tape, we say that *V sends the string x to P_i at the jth round.* Similarly, if P_i writes the string y on the ith communication tape, we say that *V receives the string y from P_i at the jth round.* If $s_1, \ldots, s_j$ are the strings sent by V to P_i in the first j rounds, then we denote the string sent by P_i to V at the jth round by $P_i(w, s_1, \ldots, s_j)$.

The following fact follows from results of Goldwasser and Sipser [9], and Babai [2], and provides a useful *normal form* for IPS's, the special case where there is only one prover. If L is accepted by an unbounded round IPS, then there is an interactive protocol (P_L, V_L) of the following form that accepts L with error probability $1/2^n$. On an input w of length n, the protocol has $m(n)$ rounds, where $m(n)$ is some polynomial in n. At the jth round, V_L sends a random bit b_j to P_L and receives a bit from P_L. Let x be the string $b_1 \ldots b_{m(n)}$ and let $\bar{P}_L(w, x)$ be the string $P_L(w, b_1)P_L(w, b_1, b_2) \ldots P_L(w, b_1, \ldots, b_{m(n)})$. When all the rounds are completed, V_L computes $\phi_L(w, x, \bar{P}_L(w, x))$, where ϕ_L is some polynomial time computable boolean function. If $\phi_L(w, x, \bar{P}_L(w, x)) = 1$, V_L accepts, else V_L rejects.

2.1 The IPS-Simulation Protocol

We now describe how an unbounded round IPS in the above normal form can be *simulated* by a 1-round, 2-prover MIP. We call this MIP the *IPS-simulation protocol.*

Theorem (1.2) (Fortnow, Rompel, Sipser, 1988) *If L is accepted by a $m(n)$-round IPS, then there is a 1-round, 2-prover MIP (P_1, P_2, V) that satisfies the following property. For some N and all strings w of length $n \geq N$,*

1. *if $w \in L$ then (P_1, P_2, V) accepts w with probability $\geq 1 - \frac{1}{2^n}$,*
2. *if $w \notin L$, then for any P_1^*, P_2^*, (P_1^*, P_2^*, V) accepts w with probability $\leq 1 - \frac{1}{2m(n)}$.*

Proof: We describe a 2-prover interactive proof that is a simple modification of the protocol of Fortnow et al. We first describe the protocol of the verifier V. Fix an input w of length n, and let $m = m(n)$.

(a) Generate a random binary string x of length m.
(b) Choose a number i randomly and uniformly from the set $\{1, \ldots, m\}$.
Let x' be the prefix of x of length i.
(c) Send x to the first prover and send x' to the second prover.
(d) Suppose V receives y from the first prover and y' from the second prover.
If $|y| = m$, $|y'| = |x'|$, y' is a prefix of y and $\phi_L(w, x, y) = 1$, then V accepts,
else V rejects.

The provers P_1, P_2 are defined to simulate the prover P_L. That is, $P_1(w, x) = \bar{P}_L(w, x)$ and $P_2(w, x')$ is the prefix of $\bar{P}_L(w, x)$ of length $|x'|$. This completes the description of the protocol (P_1, P_2, V).

We now show that this protocol satisfies properties 1 and 2 above. Fix an input w of length n and let $m = m(n)$. Suppose $w \in L$. Then the probability w is accepted by (P_1, P_2, V) is at least $1 - 1/2^n$. To see this, let X be the set of strings x of length m for which $\phi_L(w, x, \bar{P}_L(w, x)) = 1$. Then (P_1, P_2, V) also accepts w when V sends a string $x \in X$ to P_1, since $P_1(w, x) = \bar{P}_L(w, x)$ and $P_2(w, x)$ is a prefix of $P_1(w, x)$. Also, since the error probability of (P_L, V_L) is $1/2^n$, the probability that V chooses a string in X initially is at least $1 - 1/2^n$.

Next suppose that $w \notin L$. Fix provers P_1^* and P_2^*. We show that the probability (P_1^*, P_2^*, V) accepts w is at most $1 - 1/2m$. We need one definition: we say a string $x = b_1 \ldots b_m$ is *reasonable* if $\phi_L(w, x, P_2^*(w, x)) = 0$ or if for some $i, 1 \leq i < m$, $P_2^*(w, b_1 \ldots b_i)$ is not a prefix of $P_2^*(w, b_1 \ldots b_{i+1})$.

Next we show that if x is chosen randomly and uniformly from $\{0,1\}^m$, then x is reasonable with probability at least $1 - 1/2^n$. We do this by defining a prover P_L^* for the IPS. For any sequence of bits $b_1, \ldots, b_j$, let $P_L^*(w, b_1, \ldots, b_j)$ be the jth bit of the string $P_2^*(w, b_1 \ldots b_j)$. Note that if x is not reasonable then $\bar{P}_L^*(w, x) = P_2^*(w, x)$ and $\phi_L(w, x, \bar{P}_L^*(w, x)) = 1$, and thus (P_L^*, V_L) accepts w. Hence, if x is chosen randomly and uniformly, the probability that it is not reasonable must be at most $1/2^n$, since for $w \notin L$, the probability that (P_L^*, V_L) accepts w is at most $1/2^n$. So the probability that x is reasonable is at least $1 - 1/2^n$.

Finally, we show that if x is reasonable, the probability that V accepts w, when V sends x to P_1^*, is at most $1 - 1/m$. First suppose that $\phi_L(w, x, P_1^*(w, x)) \neq 1$. Then with probability 1 V rejects. Otherwise, for some i, $1 \leq i \leq m$, $P_2^*(w, b_1 \ldots b_i)$ is not a prefix of $P_1^*(w, x)$. With probability $1/m$, V sends $b_1 \ldots b_i$ to P_2^*, and V rejects.

Hence, (P_1^*, P_2^*, V) accepts w with probability at most 1 if V initially chooses a string that is not reasonable, and with probability at most $1 - 1/m$ otherwise; also the probability of choosing a string that is not reasonable is at most $1/2^n$. Therefore (P_1^*, P_2^*, V) accepts w with probability at most $1/2^n + (1 - 1/2^n)(1 - 1/m) \leq 1 - 1/2m$ for sufficiently large n, as required. □

2.2 The Product of a MIP

Let (P_1, P_2, V) be a 1-round, 2-prover MIP. We define the d-product of (P_1, P_2, V) to be the MIP $(P_{1,d}, P_{2,d}, V_d)$, where V_d simulates d independent copies of V and $P_{1,d}, P_{2,d}$ simulate P_1 and P_2 on each copy. For notational purposes, we assume that the provers and verifier use the # symbol to separate copies of the protocol on the communication tape. Thus, for example, if V_d sends $x_1\#\ldots\#x_d$ to a prover, then x_i is the string V_d sends to that prover for the ith copy of the game.

Suppose that for all provers P_1^* and P_2^*, (P_1^*, P_2^*, V) accepts a string w with probability $\leq p$. We might expect that for all provers $P_{1,d}^*, P_{2,d}^*$, $(P_{1,d}^*, P_{2,d}^*, V_d)$ accepts w with probability $\leq p^d$. Certainly this is true if for each i, the response of each prover on the ith copy of the protocol depends only on the string sent to that prover for the ith copy, and not on the strings for other copies. However Fortnow [6] constructed a simple MIP that accepts all inputs with probability 1/2, but there exist provers for the 2-product MIP which cause it to accept all inputs with probability 3/8.

Fortnow's construction is as follows. The protocol of the verifier does not depend on the input, so we ignore the input in this description. The verifier chooses two bits b, b' randomly and uniformly, sends b to P_1 and b' to P_2. Then V receives a bit y and y' from P_1 and P_2 respectively. The verifier computes the function $\phi(b, b', y, y') = ((b \vee y) \neq (b' \vee y'))$, and accepts if and only if it equals 1. The provers P_1 and P_2 are defined by $P_1(b) = b$ and $P_2(b') = b'$.

Clearly (P_1, P_2, V) accepts any input with probability 1/2, since V accepts when $b \neq b'$, which happens with probability 1/2. Next consider the 2-product of this MIP. In this game, the first prover receives a string $b_1\#b_2$ and the second prover receives a string $b_1'\#b_2'$; and the provers return strings $y_1\#y_2$ and $y_1'\#y_2'$, respectively. V accepts if

$$((b_1 \vee y_1) \neq (b_1' \vee y_1')) \wedge ((b_2 \vee y_2)) \neq (b_2' \vee y_2'))).$$

Suppose we define $P_{1,d}^*$ and $P_{2,d}^*$ as follows (where $d = 2$): $P_{1,d}^*(b_1\#b_2) = 0\#0$ if $b_1 = b_2 = 0$, otherwise $P_{1,d}^*(b_1\#b_2) = 1\#1$. $P_{2,d}^*$ is defined identically.

Then $(P_{1,d}^*, P_{2,d}^*, V)$ accepts all inputs with probability 3/8 because if $b_1 = b_2 = 0$, V accepts when b_1' and b_2' are not both 0; and by symmetry if $b_1' = b_2' = 0$, V accepts when b_1 and b_2 are not both 0. Hence V accepts on 6 of the 16 possible choices for the bits b_1, b_2, b_1', b_2'.

3 Main Result

In this section we show that the $d(n)$-product of the IPS-simulation protocol of Theorem 1.2 has error 7/8, for some polynomial $d(n)$.

Theorem (1.3) *For some polynomial $d(n)$, the $d(n)$-product of the IPS-simulation protocol of Theorem 1.2 for language L accepts any $w \notin L$ with probability at most 7/8. And for $w \in L$, it accepts w with probability $1 - O(\frac{d(n)}{2^n})$.*

Proof: We first describe the $d(n)$-product of the IPS-simulation protocol on a fixed input w. In this theorem we denote the product by (P_1, P_2, V), and as before we denote the $m(n)$-round IPS by (P_L, V_L). Let $m = m(n)$ and let $d = O(m)$; its exact value will be given later. The protocol of the verifier V is as follows.

(a) Generate d independent random binary strings $x_1, \ldots, x_d$, each of length m.
(b) For each string x_j, choose a number i_j randomly and independently from the range $\{1, \ldots, m\}$. For each j, let x'_j be the prefix of string x_j with length i_j.
(c) Send $x_1 \# \ldots \# x_d$ to the first prover and send $x'_1 \# \ldots \# x'_d$ to the second prover.
(d) Suppose V receives the string $y_1 \# \ldots \# y_d$ from the first prover and $y'_1 \# \ldots \# y'_d$ from the second prover. If for all j, $1 \leq j \leq m$, $|y_j| = m$, $|y'_j| = |x'_j|$, y'_j is a prefix of y_j and $\phi_L(w, x_j, y_j) = 1$, then V accepts, else V rejects.

The provers P_1 and P_2 are defined to simulate the prover P_L of the original IPS protocol. Thus,

$$P_1(w, x_1 \# \ldots \# x_d) = \bar{P}_L(w, x_1) \# \ldots \# \bar{P}_L(w, x_d) \text{ and } P_2(w, x'_1 \# \ldots \# x'_d) = \bar{P}_L(w, x'_1) \# \ldots \# \bar{P}_L(w, x'_d).$$

Fix an input w of length n and let $m = m(n)$. If $w \in L$, then with probability $(1-2^{-n})^d = 1-O(\frac{d(n)}{2^n})$, the verifier accepts.

Next suppose that $w \notin L$. Fix provers P_1^* and P_2^*. We show that the probability (P_1^*, P_2^*, V) accepts w is at most $7/8$. The proof is organized just as the proof of Theorem 1.2 and consists of three parts. We first define what it means for a sequence $x_1, \ldots, x_d$ to be reasonable. Then we show that most sequences are reasonable, and finally we show that if V sends a reasonable sequence to P_1^*, then V rejects w with high probability.

We first give the definition of a reasonable sequence of strings. Let T denote the set of all binary strings of length between 1 and m, $T = \{\ b_1 \ldots b_i \mid b_k \in \{0,1\} \text{ and } 1 \leq i \leq m\ \}$. We can express the function P_2^*, for the fixed input w as the concatenation of d functions $\ell_j : T^d \to T, 1 \leq j \leq d$, where

$$P_2^*(w, x'_1 \# \ldots \# x'_d) = \ell_1(x'_1, \ldots, x'_d) \# \ldots \# \ell_d(x'_1, \ldots, x'_d).$$

That is, $\ell_j(x'_1, \ldots, x'_d)$ is the response of P_2^* on the jth copy of the protocol. Fix any strings $x_2, \ldots, x_d$ of length m. We define a *majority function M_1, determined by* $x_2, \ldots, x_d$ as follows. $M_1 = M_1^{x_2, \ldots, x_d}$ is a function from T to T,

$$M_1(x) = \text{majority } \{\ \ell_1(x, x'_2, \ldots, x'_d) \mid x'_k \in T \text{ is a prefix of } x_k, 2 \leq k \leq d\ \}$$

(Ties, if any, are broken in favor of the lexicographically minimum string). One can similarly define majority functions M_j for any $d-1$ fixed strings $x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_d$ of length m: $M_j(x) =$ majority $\{\ \ell_j(x'_1, \ldots, x'_{j-1}, x, x'_{j+1}, \ldots, x'_d) \mid x'_k \in T$ is a prefix of $x_k, k \neq j\ \}$.

We say the sequence $x_1, \ldots, x_d$ is *j-reasonable* if the following conditions are satisfied: Let $x_j = b_1 \ldots b_m$, either $\phi_L(w, x_j, M_j(x_j)) = 0$ or for some $i, 1 \leq i < m$, $M_j(b_1 \ldots b_i)$ is not a prefix of $M_j(b_1 \ldots b_{i+1})$. We say that $x_1, \ldots, x_d$ is *reasonable* if it is j-reasonable for all $j, 1 \leq j \leq m$. The next lemma shows that most sequences are reasonable.

Lemma 3.1 Suppose $x_1, \ldots, x_d$ are binary strings of length m chosen independently and uniformly. Then

$$\text{Prob}[x_1, \ldots, x_d \textit{ is reasonable }] \geq 1 - d/2^n.$$

Proof: We first show that for any j and for any fixed $x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_d$, if x_j is chosen randomly then $x_1, \ldots, x_d$ is j-reasonable with probability at least $1 - 1/2^n$. Without loss of generality, suppose that $j = 1$, and fix any $x_2, \ldots, x_d$. From M_1, we define a prover P_L^* as follows. For any sequence of bits $b_1, \ldots, b_i$, let $P_L^*(w, b_1, \ldots, b_i)$ be the ith bit of the string $M_1(b_1 \ldots b_i)$.

Just as in Theorem 1.2, for any x, if $x, x_2, \ldots, x_d$ is not 1-reasonable, then $\bar{P}_L^*(w, x) = M_1(x)$, and $\phi_L(w, x, \bar{P}_L^*(w, x)) = 1$, and thus (P_L^*, V_L) accepts w. Hence, if x is chosen randomly and uniformly, the probability that $x, x_2, \ldots, x_d$ is not 1-reasonable must be at most $1/2^n$, since for $w \notin L$, the probability that (P_L^*, V_L) accepts w is at most $1/2^n$. So the probability that $x, x_2, \ldots, x_d$ is 1-reasonable is at least $1 - 1/2^n$.

To complete the proof, note that

$$\text{Prob}[x_1, \ldots, x_d \text{ is not reasonable }] \leq \sum_{j=1}^{d} \text{Prob}[x_1, \ldots, x_d \text{ is not } j\text{-reasonable }] \leq d/2^n.$$

□

Finally, we need to show that if $x_1, \ldots, x_d$ is reasonable then (P_1^*, P_2^*, V) rejects w with high probability when V sends $x_1, \ldots, x_d$ to P_1^*. Assume therefore that $x_1, \ldots, x_d$ is reasonable.

Let Q^d be the d-dimensional cube of lattice points up to m, $Q^d = \{ (i_1, \ldots, i_d) \mid 1 \leq i_j \leq m \}$. For each j, the function ℓ_j, when restricted on prefixes of $x_1, \ldots, x_d$ gives a natural labeling of the cube Q^d. That is, the label of $(i_1 \ldots i_d)$ is $\ell_j(x'_1, \ldots, x'_d)$, where x'_k is the prefix of x_k of length i_k. We denote this label by $\ell_j(i_1, \ldots i_d)$. Let $F_1, \ldots, F_m$ be the $d-1$ dimensional *flats* along the first dimension, defined by

$$F_k = \{ (k, i_2, \ldots, i_d) \mid 1 \leq i_2, \ldots, i_d \leq m \}, 1 \leq k \leq m.$$

We can similarly define flats along every other dimension.

To simplify notation in the rest of the proof, let $y_1 \# y_2 \ldots \# y_d = P_1^*(w, x_1 \# x_2 \# \ldots \# x_d)$.

Lemma 3.2 *For each dimension j, $1 \leq j \leq d$, either $\phi_L(w, x_j, y_j) = 0$, or there exists a flat F along dimension j, such that the set*

$$H_j = \{ p \in F \mid \ell_j(p) \text{ is not a prefix of } y_j \}$$

has cardinality $\geq m^{d-1}/2$.

Proof: Without loss of generality let $j = 1$ and let $x_1 = b_1 \ldots b_m$. Suppose $\phi_L(w, x_1, y_1) = 1$. All points p on each flat along dimension 1 are partitioned into equivalence classes according to their label $\ell_1(p)$. We ask the question: is there an i, such that the ith flat F_i along dimension 1 has at least $m^{d-1}/2$ points p with label $\ell_1(p)$ different from the majority function $M_1(b_1 \ldots b_i)$?

If so, then no string labels more than half the points on F_i. In particular, the prefix of y_1 of length i labels at most half the points on flat i and hence flat F_i satisfies the lemma.

Now suppose the answer to the question is no. Then, on each flat along dimension 1, more than half of the points are labeled $M_1(b_1 \ldots b_i)$. Since $x_1, \ldots, x_d$ is reasonable, in particular, 1-reasonable, and by assumption, $\phi_L(w, x_1, y_1) = 1$, it must be the case that either $M_1(x_1) \neq y_1$ or there exists an i, $1 \leq i < m$, such that $M_1(b_1 \ldots b_i)$ is not a prefix of $M_1(b_1 \ldots b_{i+1})$. To complete the proof, note that when $M_1(b_1 \ldots b_i)$ is not a prefix of $M_1(b_1 \ldots b_{i+1})$, then either $M_1(b_1 \ldots b_i)$ is not a prefix of y_1 or $M_1(b_1 \ldots b_{i+1})$ is not a prefix of y_1. Thus in this case, either F_i or F_{i+1} satisfies the lemma. □

We can now show that if $x_1, \ldots, x_d$ is reasonable, then when V sends $x_1\#\ldots\#x_d$ to P_1^*, the probability that V rejects is high. This is immediate if for some j, $\phi(w, x_j, y_j) = 0$, since then V rejects with probability 1. Hence suppose that for all j, $\phi_L(w, x_j, y_j) = 1$. Let $H = \cup_j H_j$, where H_j is defined as in Lemma 3.2. Note that for any point $(i_1, \ldots, i_d) \in H$, if V sends the string $x_1'\#\ldots\#x_d'$ to P_2^*, where $|x_k| = i_k$, then V rejects. This is because for some j, P_2^*'s response on the jth copy of the game, $\ell_j(x_1', \ldots, x_d')$ is not a prefix of y_j. Hence the probability that V rejects w is at least $(1/m^d)|H|$, since there are exactly m^d possible strings that V can send to P_2, with equal probability, given that it sends $x_1\#\ldots\#x_d$ to P_1. We next get a lower bound on $|H|$.

From the lemma, $|H_j| \geq m^{d-1}/2$. Also, $|H_j \cap H_k| \leq m^{d-2}$ for all $j \neq k$, since the flats are orthogonal. Hence

$$\begin{aligned} |\bigcup_{j=1}^{d} H_j| &= \sum_{j=1}^{d} |H_j - \bigcup_{k<j} H_k| \\ &\geq \sum_{j=1}^{d} [|H_j| - (j-1)m^{d-2}] \\ &\geq dm^{d-1}/2 - d(d-1)m^{d-2}/2 \\ &\geq \frac{m^d}{8}\left(1 + \frac{2}{m}\right), \text{ if } d = \lceil m/2 \rceil. \end{aligned}$$

Hence, (P_1^*, P_2^*, V) accepts w with probability at most 1 if V initially chooses a sequence of strings that is not reasonable, and with probability at most $7/8 - 1/4m$ otherwise; also the probability of choosing a sequence of strings that is not reasonable is at most $d(n)/2^n$. Therefore (P_1^*, P_2^*, V) accepts w with probability at most $d(n)/2^n + (1 - d(n)/2^n)(7/8 - 1/4m)) \leq 7/8$ for sufficiently large n, as required. This completes the proof of Theorem 1.3 □

From Theorem 1.3, our main result now follows easily.

Theorem (1.1) *For any constant ϵ, any language accepted by an unbounded round IPS has a bounded round 2-prover MIP that has error probability ϵ.*

Proof: Let k be an integer satisfying $(7/8)^k \leq \epsilon$. To accept L with probability at most ϵ, an MIP simply repeats the product protocol of Theorem 1.3 *sequentially* k times. □

4 Conclusions and Open Problems

Theorem 1.1 shows that for any constant ϵ, any language accepted by an unbounded round IPS has a bounded round 2-prover MIP that has error probability ϵ. However, the number of rounds of the protocol depends on ϵ. It is still open whether for any ϵ and any language accepted by an unbounded round IPS, there is a *one*-round 2-prover MIP that accepts L. A related question is whether the hierarchy of bounded-round, two-prover protocols collapses to 1-round protocols, as is the case for single provers.

More generally, if a MIP accepts an input w with probability p, what can we say about the probability that the d-product of the MIP accepts w? In [5], Cai, Condon and Lipton provide partial answers to this problem for restricted types of MIP's, showing for example that as $d \to \infty$, the probability that the d-product MIP accepts $w \to 0$. However, even for these restricted types of MIP's, it is not known if this probability is strictly decreasing as d increases!

References

[1] W. Aiello, S. Goldwasser and J. Hastad, On the power of interaction, Proceedings of 27th FOCS, 1986, pp 368-379.

[2] L. Babai, Trading Group Theory for Randomness, Proceedings of 17th STOC, 1985, pp 421-429.

[3] M. Ben-Or, S. Goldwasser, J. Killian and A. Wigderson, Multi-Prover Interactive Proofs: How to Remove Intractability, Proceedings of the 20th STOC, May, 1988.

[4] B. Bollobás, *Extremal Graph Theory*, Academic Press, 1978.

[5] Jin-Yi Cai, Anne Condon and Richard J. Lipton, Playing Games of Incomplete Information, To appear in STACS 90.

[6] L. Fortnow, Complexity-Theoretic Aspects of Interactive Proof Systems, Ph. D. Thesis, Tech Report #MIT/LCS/TR-447, MIT.

[7] L. Fortnow, J. Rompel and M. Sipser, On the Power of Multi-Prover Interactive Protocols, Proceedings of the conference on Structure in Complexity Theory, 1988, pp 156-161.

[8] S. Goldwasser, S. Micali and C. Rackoff, The knowledge complexity of interactive protocols, Proceedings of 17th STOC, 1985, pp 291-304.

[9] S. Goldwasser and M. Sipser, Private Coins versus Public Coins in Interactive Proof Systems, Proceedings of 18th STOC, 1986, pp 59-68.

570