COMBINATORICA
Bolyai Society – Springer-Verlag

COMBINATORICA **22** (4) (2002) 461–478

ERROR REDUCTION BY PARALLEL REPETITION – A NEGATIVE RESULT

URIEL FEIGE*, OLEG VERBITSKY†

Received July 24, 1996
Revised July 8, 1998

We show that no fixed number of parallel repetitions suffices in order to reduce the error in two-prover one-round proof systems from one constant to another. Our results imply that the recent bounds proven by Ran Raz, showing that the number of rounds that suffice is inversely proportional to the answer length, are nearly best possible.

1. Introduction

A two prover one round proof system [6], MIP(2,1), is a protocol by which two provers jointly try to convince a computationally limited probabilistic verifier that a common input belongs to a prespecified language. The verifier selects a pair of questions at random. Each prover sees only one of the two questions, and sends back an answer. The verifier evaluates a predicate on the common input and the two questions and answers, and accepts or rejects according to the output of the predicate. For inputs in the language, the provers have a strategy (where a strategy for a prover is a function from incoming messages to outgoing messages) that always causes the verifier to accept. For inputs not in the language, regardless of the strategy used by the provers, the verifier accepts with probability at most ϵ. The smaller the value of ϵ, known as the *error*, the more confidence the verifier has in using the proof system.

Mathematics Subject Classification (2000): 68Q10, 68R05, 91A20

* Incumbent of the Joseph and Celia Reskin Career Development Chair.

† Supported in part by grant No. MQT 300 from the International Science Foundation and by an AMS-FSU grant.

Reducing the error in MIP(2,1) proof systems is a subtle issue. A natural approach is to repeat an MIP(2,1) protocol n times, and accept only if all executions were accepting. Ideally, one would hope that this method reduces the error to ϵ^n. This is indeed true if each execution is performed with a fresh pair of provers, requiring n pairs of provers, or if the executions are performed sequentially (each prover must answer each question online, before seeing the question for the next execution), requiring n rounds of communication. However, parallel repetition, in which there are only two provers, and each prover sends out its answers only after receiving all its questions, is not guaranteed to reduce the error to ϵ^n [14]. Much work was invested in trying to analyse the rate at which parallel repetition reduces the error in MIP(2,1) proof systems.

There are many known results about error reduction by parallel repetition (see references in the end of this paper). Most of these results are surveyed in [11] (including the results of the current paper). In Section 3.1 we shall briefly recall some of these results. The most useful of these is a theorem of Raz [22], which informally says that the error in parallel repetition of MIP(2,1) proof systems goes down at a rate that is exponential in the number of repetitions, and the base of the exponent can be expressed as a function of the initial error and of the answer length. A natural question to ask is whether one can express the base of the exponent as a function of the initial error alone, without need of involving other parameters such as answer length. The main new result that we present is a negative answer to the above question.

In Section 2 we formally define the parallel repetition problem. In Section 3 we state our main results, and discuss related work. In Sections 4 and 5 we prove our results. Though these last two sections share similar techniques, it is possible to follow each of them without reading the other.

The results presented in this paper were previously announced in [24] and [11]. A preliminary version of the paper appeared in the proceedings of the Eleventh Annual IEEE Conference on Computational Complexity, 1996.

2. Definitions

Two-prover one-round proof systems are often modeled as a game between a verifier and two cooperating provers. It is not an adversarial game, as the strategy of the verifier is fixed in advance. Rather, it is a cooperative game with two players (provers) who coordinate a joint strategy that gives them the highest probability of winning. Modelling MIP(2,1) proof systems as games abstracts away the language recognition aspects of these proof

systems, but preserves the concept of error in a proof system, and the issue of how error can be reduced. One may best think of a game as an instantiation of an MIP(2,1) proof system on a single input that is not in the input language. We use the following notation.

$G = G(X, Y, Q, \pi, A, B, V)$ – a two-prover one-round game
X – set of questions to prover P_1.
Y – set of questions to prover P_2.
A – set of answers available to P_1.
B – set of answers available to P_2.
π – probability distribution on $X \times Y$.
Q – support of π. $Q \subseteq X \times Y$.
V – acceptance predicate on (X, Y, A, B).

Game G proceeds as follows. The verifier selects at random a question pair $(x, y) \in Q$, according to probability distribution π. Question x is sent to P_1 who replies with an answer $P_1(x) \in A$. Question y is sent to P_2 who replies with an answer $P_2(y) \in B$. (We identify between the name of a prover and the strategy that it employs.) The verifier then evaluates a predicate $V(x, y, P_1(x), P_2(y))$, and accepts if the predicate is satisfied. The goal of the provers is to select a strategy (namely, two functions, one for each prover, specifying an answer to each possible question) that maximizes the probability that the verifier accepts. The probability that the verifier accepts under the optimal strategy of the provers is denoted by $\omega(G)$. If $\omega(G) = 1$ the provers are said to have a *perfect strategy* for G, and the game G is *trivial*. For nontrivial games, $\omega(G)$ is also called the *error* of the game. We shall only be interested in nontrivial games.

An n-fold parallel repetition of game G is a new game, denoted by G^n, played on n coordinates. The verifier treats each coordinate of G^n as an independent copy of the original game G, and accepts in G^n only if it would have accepted all the n copies of G. The support set of G^n is Q^n, the answer sets are A^n and B^n. The verifier selects n question pairs $(x_i, y_i) \in Q$ independently, each according to the probability distribution π. We denote $\bar{x} = x_1 x_2 \ldots x_n$ and $\bar{y} = y_1 y_2 \ldots y_n$. (Hence $\bar{x} \in X^n$, and $\bar{y} \in Y^n$.) A strategy for the provers is $2n$ functions, $P_1^i : X^n \longrightarrow A$ and $P_2^i : Y^n \longrightarrow B$, where $1 \leq i \leq n$. The acceptance predicate is $V^n = \bigwedge_{i=1}^n V(x_i, y_i, P_1^i(\bar{x}), P_2^i(\bar{y}))$. That is, the verifier accepts if all n copies of the original game G are accepting.

Observe that even though the verifier treats each coordinate of G^n independently, the provers may not. In particular, the answer a prover gives in coordinate i may depend on the questions that the prover receives in other coordinates. For this reason, it is not true in general that $\omega(G^n) = (\omega(G))^n$.

We shall often refer to the following special classes of games. A game is *uniform* if π is uniform on Q. A game is *free* if it is uniform and has full support, i.e. $Q = X \times Y$. (The term *free* was introduced in [8].)

Remarks.

1. The game model of MIP(2,1) can be extended to allow for the acceptance predicate V to be a randomized predicate (that is, to have an additional input that is assigned randomly by the verifier). Known applications of MIP(2,1) proof systems do not make use of this extra option. Upper bounds on $\omega(G^n)$ were proved in the restricted model that we use, and the question of whether they hold in the extended model is most often ignored. Lower bounds that are proved in the restricted model hold also in the extended model.
2. Another extension of the MIP(2,1) model is to allow the provers to use randomized strategies. This is irrelevant for our purpose of analysing the error probability (since there always is a deterministic optimal strategy), but is important in the context of constructing *zero knowledge* MIP(2,1) proof systems [16,6]. We shall not deal with zero knowledge issues.

3. Our results, and related work

Our main result is the following:

Theorem 3.1. *There is a family $\{G_k\}$ of free games, in which $|X| = |Y| = k$, $|A| = |B| = 2^{\Theta(k)}$, $\omega(G_k) \le 3/4$, and $\omega((G_k)^n) \ge 1/8$, for all $n \le k/(4 \log k)$.*

Theorem 3.1 is proved in Section 5. It implies the following qualitative behavior of error reduction by parallel repetition.

Corollary 3.2. *For any constant $\alpha > 0$, there are games G, such that for large enough n, $\omega(G^n) > (\omega(G))^{\alpha n}$.*

Proof. Consider the game G_k given by Theorem 3.1, with k large enough so that $(3/4)^{\alpha k/(4 \log k)} < 1/64$. Assume for simplicity that $k/(4 \log k)$ is an integer. Theorem 3.1 implies that for every $n > k/(4 \log k)$, $\omega((G_k)^n) > (\omega(G_k))^{\alpha n}$. This can be seen as follows. Let $n = ck/(4 \log k)$ for some $c > 1$. Let c' be c rounded up to the nearest integer, and let $n' = c'k/(4 \log k)$. Then the fact that $\omega((G_k)^{k/(4 \log k)}) \ge 1/8$ implies that $\omega((G_k)^{n'}) \ge (1/8)^{c'}$ (the provers may repeat their optimal strategy for $(G_k)^{k/(4 \log k)}$ on each of its c' copies in $(G_k)^{n'}$). As $n \le n'$, this implies that $\omega((G_k)^n) \ge (1/8)^{c'}$ (the provers may just add $n' - n$ dummy coordinates and use the optimal strategy

of $(G_k)^{n'}$ as a strategy for $(G_k)^n$). However, for our choice of parameters, $(\omega(G_k))^{\alpha n} \leq (3/4)^{\alpha n} < (1/64)^{c'-1} \leq (1/8)^{c'}$. Hence $\omega((G_k)^n) > (\omega(G_k))^{\alpha n}$. ∎

Corollary 3.2 shows that if one wants to prove a relation of the form $\omega(G^n) \leq (\omega(G))^{\alpha n}$, then one cannot just take α as some universal constant, but must allow α to depend on some specific properties of G. Two specific properties that have been used are $|X|$ (the cardinality of the set of questions to prover P_1) and $|A|$ (the cardinality of the set of answers to prover P_1). W.l.o.g., we assume that $|X| \leq |Y|$, and that $|A| \leq |B|$, but also note that in the family of games referred to in Theorem 3.1, $|X| = |Y|$ and $|A| = |B|$. The proof of Corollary 3.2 together with the explicit values for $|X|$ and $|A|$ in Theorem 3.1 then shows:

Corollary 3.3. *Let α be a function of $|X|$ and assume that for every game G and every n, $\omega(G^n) \leq (\omega(G))^{\alpha n}$. Then $\alpha = O((\log |X|)/|X|)$. Let α be a function of $|A|$ and assume that for every game G and every n, $\omega(G^n) \leq (\omega(G))^{\alpha n}$. Then $\alpha = O(\log\log |A| / \log |A|)$. Furthermore, the above holds even for games in which $1/8 \leq \omega(G) \leq 3/4$.*

Corollary 3.3 is formulated in terms of games G. This corollary can be recast in terms of MIP(2,1) proof systems. As pointed out in Section 2, a game G with error bounded away from 1 can be thought of as an instantiation of an MIP(2,1) proof system on a single input that is not in the input language. A family of games $\{G_k\}$ (as in Theorem 3.1) can be thought of as instantiations of an MIP(2,1) proof system for the empty language. A technical condition that then arises is that for game G_k, which is an instantiation of the MIP(2,1) proof system with input I_k, the parameters of the game should be such that the acceptance predicate V can be evaluated in time polynomial in $|I_k|$ (as the verifier in MIP(2,1) proof systems is assumed to be polynomial time). In the proof of Theorem 3.1, we do not analyse the complexity of evaluating the predicate V in terms of the length of the inputs to V (which is roughly k). We do not know whether V can be evaluated in polynomial time, and in fact, this does not really matter. It suffices to have some upper bound on the complexity of V, and in our case, it can be seen that the complexity of V is at most doubly exponential in k. Now game G_k can correspond to an input I_k of size doubly exponential in k, and this makes the verifier in the corresponding MIP(2,1) proof system polynomial time. (One can also modify either questions or answers by padding with irrelevant bits so that their length becomes doubly exponential in k, and then the predicate V is certainly computable in time polynomial in the length of the inputs to the predicate. The dependence of the rate of error reduction on the length of the parameter that is not padded remains unchanged.)

As shown above, when parallel repetition is analysed, properties of G (such as question or answer length) must be taken into account if one wants to bound the error. One of the methods commonly used [8,17,21,1,10,23, 25] is the *forbidden subgraph method*, which attempts to bound $\omega(G^n)$ only as a function of X, Y, Q, π, $\omega(G)$ and n, and ignores A, B, V (for details, see Section 4). As the forbidden subgraph method abstracts away certain parameters of the game (such as answer length), it is not clear that the forbidden subgraph method is a *universal* proof technique, in the sense that any true bound on $\omega(G^n)$ that involves only the parameters considered by the method can indeed be proved via this method. The following theorem shows that in a special case (when G is just assumed to be nontrivial, but with no explicit upper bound on $\omega(G)$) the forbidden subgraph method is universal.

Theorem 3.4. *Let $FS(X,Y,Q,\pi,n)$ be the best possible upper bound on $\omega(G^n)$ that can be derived by the forbidden subgraph approach assuming that G is nontrivial. Then this upper bound is tight in the sense that there is some game G with the same parameters (X,Y,Q,π,n), augmented with suitable A,B, and V, such that G is nontrivial and $\omega(G^n)=FS(X,Y,Q,\pi,n)$.*

The value $FS(X,Y,Q,\pi,n)$ will be defined in combinatorial terms in Section 4, where Theorem 3.4 is proved after rephrasing it as Theorem 4.2.

An interesting implication of Theorem 3.4 is that the forbidden subgraph method can be used not only for proving upper bounds on $\omega(G^n)$, but also for proving lower bounds. This implication is used in Theorem 3.1, whose proof uses ideas from the proof of Theorem 3.4. Let us explain the connection between the two theorems in more detail. Consider a uniform game G that is nontrivial, implying $\omega(G)\le(1-1/|Q|)$. If the game is repeated n times *sequentailly*, the error becomes $(1-1/|Q|)^n$. We would like to prove that parallel repetition lowers the error at a much slower rate, that is, to show that $\omega(G^n)$ is substantially larger than $(1-1/|Q|)^n$. Theorem 3.4 shows that this would follow if we could show that $FS(X,Y,Q,\pi,n)$ is substantially larger than $(1-1/|Q|)^n$. Unfortunately, we do not know how to prove strong lower bounds on $FS(X,Y,Q,\pi,n)$. To make the task of proving strong enough lower bounds easier (or even possible), we use a theorem similar to 3.4, but for games in which $\omega(G)\simeq 1/2$ (rather than $\omega(G)=1-1/|Q|$). For this case, we only get an approximate correspondence between error reduction by parallel repetition and the forbidden subgraph problem. But the approximate correspondence that we get, together with a lower bound on the respective forbidden subgraph problem, suffice in order to prove Theorem 3.1.

3.1. Related work

Here we discuss some of the previous results on parallel repetition that are most related to our work. For more information on previous work on parallel repetition, the reader is referred to [11], or to other references in the end of this paper. Issues such as how MIP(2,1) proof systems are constructed, their applications to proving hardness of approximation results, zero knowledge versions and applications to cryptography, are beyond the scope of this paper. (See for example [6,14,17,9,3,10,18,13,2,19,4,12,5]).

MIP(2,1) proof systems were introduced by Ben-Or, Goldwasser, Kilian, and Wigderson [6]. The authors remarked that parallel repetition of these proof systems preserves zero knowledge properties, but the issue of its effect on the error was not touched upon. Refuting initial intuitions that $\omega(G^n) = (\omega(G))^n$, Fortnow, Rompel and Sipser [14] constructed an example showing $\omega(G^2) > (\omega(G))^2$. Lapidot and Shamir [17] showed the same effect in a natural zero knowledge MIP(2,1) proof system for NP. Feige [10] presented an example of a nontrivial game in which $\omega(G^2) = \omega(G)$. All the above examples are for free games. A nonfree (but uniform) game with $\omega(G^2) = \omega(G)$ is presented in [13]. Previous to our current work, there was no known example of a game G in which $\omega(G^n) > (\omega(G))^{n/2}$. For games with $k > 2$ provers, Raz shows an example where $\omega(G^k) = \omega(G)$ (details appear in [11]).

The most useful upper bound on $\omega(G^n)$ was proved by Raz [22]. He shows the existence of a function W that maps from the open interval $(0,1)$ to itself (in fact, to the open interval $(t,1)$, where t is some constant, $0 < t < 1$), such that $\omega(G^n) \leq (W(\omega(G)))^{n/\log|A|}$. (The bound stated in [22] is $(W(\omega(G)))^{n/\log(|A|\cdot|B|)}$, but the analysis itself does not use $|B|$. See also [20] for other improved parameterizations of the upper bound.) Raz does not explicitly describe W, but notes that for any fixed value of $\omega(G)$, the value of $W(\omega(G))$ is just another fixed constant, and hence the effect of W can be replaced by some constant in the exponent. The value of this constant depends on $\omega(G)$, and careful examination shows that it behaves like $(1-\omega(G))^c/\log(1+1/\omega(G))$, for some constant c.

Theorem 3.5. (Raz.) *For any game G, $\omega(G^n) \leq (\omega(G))^{\alpha n}$, for some constant α that depends on G. Explicitly,*

$$\alpha = \Omega\left(\frac{(1-\omega(G))^c}{\log|A| \cdot \log(1+1/\omega(G))}\right)$$

where c is some universal constant.

The most disturbing feature about this upper bound is that α is not a fixed constant (unlike the case of sequential repetition, where the constant

in the exponent is 1). Corollary 3.2 shows that Raz's bound is qualitatively optimal, in the sense that a dependency of α on G is unavoidable. Quantitatively, Corollary 3.3 shows a dependence of α on $|A|$, and this dependence nearly matches Raz's theorem – $O(\log\log|A|/\log|A|)$, compared to $\Omega(1/\log|A|)$.

Theorem 3.1 shows that the number of parallel repetitions needed in order to reduce the error from $\omega(G)$ to a desirable value ϵ does not depend on $\omega(G)$ and ϵ alone, giving a negative answer to a question that was asked in [8] and in [10]. It is interesting to note that for a slight variation on parallel repetition, in which each prover is requested to answer some of the questions that are originally intended to the other prover, the number of the questions that the verifier needs to send can be bounded in terms of $\omega(G)$ and ϵ alone. For details, see [12].

Theorem 3.1 also implies that for any fixed $\epsilon > 0$, there is no universal constant n such that for every nontrivial game G, $\omega(G^n) \leq (1-\epsilon)\omega(G)$. Otherwise, the number of parallel repetitions that suffice in order to reduce the error from 3/4 to 1/8 would be bounded by a constant. Interestingly, for nontrivial free games, Feige and Szegedy show that $\omega(G^3) < \omega(G)$ (the proof appears in [11]). It is an open question whether there is some universal constant n such that for every nontrivial game G, $\omega(G^n) < \omega(G)$.

The forbidden subgraph method, to be described in Section 4, is perhaps the approach most commonly used in order to upper bound $\omega(G^n)$. The above approach was used in [8,17,21,1,10] to analyse nontrivial free games. For the special case that $|X| = |Y| = 2$, Peleg [21] shows that $\omega(G^n) = O(2^{-(1-(\log 3/2))n}) \simeq O(2^{-0.21n})$, and that the approach cannot give anything better than $\omega(G^n) = 2^{-n/3}$. Combined with our Theorem 3.4, this last result implies that there are free games with $|X| = |Y| = 2$ for which indeed $\omega(G^n) \geq 2^{-n/3}$, for every n divisible by 3. For free games with arbitrary support size, Alon [1] and Feige [10] prove similar bounds, showing that $\omega(G^n) = 2^{-\Omega(n/(|X|\cdot|Y|\cdot\log|X|))}$. Upto the $O(\log|X|)$ factor in the exponent, this result matches the trivial lower bound. In addition, Feige [10] proves that $\omega(G^n) < e^{-k}$, where $k = n\log(1/\sqrt{\omega(G)})/2(|X| + \log(1/\sqrt{\omega(G)}) + 2)$. For constant $\omega(G)$, the lower bound of Corollary 3.3 matches this upper bound up to an $O(\log|X|)$ factor in the exponent.

The forbidden subgraph method was also used to analyse nonfree games in which π is uniform over the support. Verbitsky [23] uses Ramsey Theory [15] to give the first proof that as $n \to \infty$, $\omega(G^n) \to 0$, without giving a constructive bound. In [25], a constructive upper bound of $\omega(G^n) = 2^{-\Omega(n/|Q|^8)}$ is given for a family of nonfree games, in which Q induces a tree

structure on $X \bigcup Y$. It is an open question whether a bound of the form $\omega(G^n) \leq 2^{-n/f(|Q|)}$ is true in general, for some function f.

4. The forbidden subgraph method

The forbidden subgraph method is an approach for upper bounding $\omega(G^n)$ as a function of n, $|X|$, $|Y|$, and possibly also $|Q|$, π, and $\omega(G)$. It was first used in [8] for the special case of free games, and generalizes in a straightforward manner to any game (though the combinatorics becomes much harder). In this section we describe the method for the case that π is uniform over Q, and indicate at the end of the section the changes required when π is not uniform.

For a game G uniform on Q and for $n \geq 1$, we view Q^n as a bipartite graph with vertex classes X^n and Y^n by identifying an element $\langle (x_1, y_1), \ldots, (x_n, y_n) \rangle \in Q^n$ with an edge $\langle (x_1, \ldots, x_n), (y_1, \ldots, y_n) \rangle$. For an edge $\bar{e} \in Q^n$, we denote its vertex in X^n by $x(\bar{e})$ and its vertex in Y^n by $y(\bar{e})$. For an n-vector $\bar{v}$, we denote its i-th component by $\bar{v}|_i$.

Fix an integer k in the range $[1, |Q|]$. Suppose that $\bar{e}_1, \ldots, \bar{e}_k$ are distinct edges of Q^n. The l-projection of $\bar{e}_1, \ldots, \bar{e}_k$ is the set of edges $\{\bar{e}_1|_l, \ldots, \bar{e}_k|_l\}$ (belonging to Q^1 rather than Q^n). A subgraph F of Q^n with edges $\bar{e}_1, \ldots, \bar{e}_k$ is called *k-forbidden* if for some $l \leq n$, its l-projection satisfies the following requirements:

1. $|\{\bar{e}_1|_l, \ldots, \bar{e}_k|_l\}| = k$. That is, with respect to coordinate (copy of G) l, the edges of F correspond to k distinct question pairs.
2. Projecting the vertices of F on coordinate l above, all vertices are distinct. That is, for any $v \in X$, there is at most one vector $\bar{x} \in X^n$ such that $\bar{x}|_l = v$ and $\bar{x} = x(\bar{e})$ for some $\bar{e} \in \{\bar{e}_1, \ldots, \bar{e}_k\}$ (maybe more than one such $\bar{e}$). A similar condition holds for $v \in Y$.

For the special case of $k = |Q|$, the above requirements imply that the graphs F and Q are isomorphic. For $k < |Q|$, the above requirements imply that the graph F is isomorphic to a subgraph of Q that is induced on k edges.

By $E_{Q,k}(n)$ we denote the maximum density $|W|/|Q|^n$ of a graph $W \subset Q^n$ that does not contain any k-forbidden subgraph.

The following proposition, which is the essence of the forbidden subgraph method, was first used in [8] in the case of free nontrivial games.

Proposition 4.1. *For any uniform game G, and for any $k > \omega(G)|Q|$,*

$$\omega(G^n) \leq E_{Q,k}(n).$$

Proof. Let us choose and fix some strategies P_1 and P_2 optimal in G^n. We define $W \subset Q^n$ to be the set of successes of P_1 and P_2 in G^n. Formally, W is defined as

$$\left\{ \bar{e} \in Q^n : \bigwedge_{i=1}^{n} V(x(\bar{e}|_i), y(\bar{e}|_i), P_1^i(x(\bar{e})), P_2^i(y(\bar{e}))) = 1 \right\}.$$

Since the game G is uniform, $\omega(G^n) = |W|/|Q|^n$.

To prove the proposition, it suffices to show that W does not contain any k-forbidden subgraph. Suppose, to the contrary, that W contains a k-forbidden graph $F = \{\bar{e}_1, \ldots, \bar{e}_k\}$ whose l-projection satisfies requirements 1 and 2 above. Our goal is to prove that in this case $\omega(G)|Q| \geq k$.

Define the strategies p_1 and p_2 in the game G as follows. For any $v \in X$, if $v = \bar{x}|_l$ for some $\bar{x} = x(\bar{e})$ where $\bar{e} \in \{\bar{e}_1, \ldots, \bar{e}_k\}$, then $p_1(v) = P_1^l(\bar{x})$. Note that by requirement 2 above, the respective $\bar{x}$ is unique, and no conflict can arise. If no such $\bar{x}$ exists, assign an arbitrary value to $p_1(v)$. Similarly, for any $v \in Y$, if $v = \bar{y}|_l$ for some $\bar{y} = y(\bar{e})$ where $\bar{e} \in \{\bar{e}_1, \ldots, \bar{e}_k\}$, then $p_2(v) = P_2^l(\bar{y})$. The strategies p_1 and p_2 win for at least k possible question pairs from the support Q, as derived by projecting the k edges of F on their lth coordinate. Hence $\omega(G) \geq k/|Q|$. ∎

Proposition4.1 shows that $E_{Q,k}(n)$ provides an upper bound on $\omega(G^n)$ for any game G . The main theorem of this section (essentially Theorem 3.4) shows that in the special case that $k = |Q|$ (that is, all we know is that G is nontrivial), this bound cannot be improved. For notational convenience, when $k = |Q|$ we denote $E_{Q,k}(n)$ by $E_Q(n)$. When π is the uniform distribution on $Q \subseteq X \times Y$ we have that $E_Q(n) = FS(X, Y, Q, \pi, n)$, where $FS(X, Y, Q, \pi, n)$ is the notation used in Theorem 3.4. (For non-uniform π see Remark 3 below).

Theorem 4.2. *Let $Q \subseteq X \times Y$ be a connected bipartite graph. Then, for any n, there is a nontrivial uniform game G with support Q such that*

$$\omega(G^n) = E_Q(n).$$

Thus, given n, the maximum possible error $\omega(G^n)$ of a nontrivial uniform game G with support Q is equal to $E_Q(n)$.

Proof. Given a connected bipartite graph $Q \subseteq X \times Y$ and a natural n, we have to construct answer sets A and B, and acceptance predicate $V \subset X \times Y \times A \times B$. We set

$$A = \{1, \ldots, n\} \times X^n \quad \text{and} \quad B = \{1, \ldots, n\} \times Y^n.$$

In order to define V, we choose and fix a graph $W \subset Q^n$ without $|Q|$-forbidden subgraphs, of the maximum possible density $E_Q(n)$. For each $\bar{e} \in W$ and $i \le n$, we put into V the quadruple $\langle x(\bar{e})|_i, y(\bar{e})|_i, i \circ x(\bar{e}), i \circ y(\bar{e})\rangle$. In other words,

$$\langle x, y, i \circ \bar{x}, m \circ \bar{y}\rangle \in V \iff$$

$$(\bar{x}, \bar{y}) \in W,\ i = m,\ x = \bar{x}|_i,\ y = \bar{y}|_m.$$

Thus, the game G is completely specified.

Consider the strategies P_1 and P_2 in G^n defined by $P_1^i(\bar{x}) = i \circ \bar{x}$ and $P_2^i(\bar{y}) = i \circ \bar{y}$ for each $i \le n$. Obviously, P_1 and P_2 win on W. Hence, $\omega(G^n) \ge E_Q(n)$.

It remains to prove that G is nontrivial. By Proposition 4.1, this will imply the reverse inequality $\omega(G^n) \le E_Q(n)$. The proof is by reductio ad absurdum. Suppose, to the contrary, that some strategies $p_1 : X \longrightarrow A$ and $p_2 : Y \longrightarrow B$ always win.

Let $Q = \{e_1, \ldots, e_k\}$. Consider an arbitrary edge $e_j = (x, y)$ in Q. By our assumption, $\langle x, y, p_1(x), p_2(y)\rangle \in V$. Our construction of the acceptance predicate V implies that, for some $l \le n$, $p_1(x) = l \circ \bar{x}$ and $p_2(y) = l \circ \bar{y}$, where $\bar{x} \in X^n$, $\bar{y} \in Y^n$, $\bar{x}|_l = x$, and $\bar{y}|_l = y$. Moreover, the number l is the same for any two edges e_j and e_p adjacent in Q. As Q is connected, l is the same for all edges and, therefore, for all $x \in X$ and $y \in Y$. Another consequence of our definition of V is that $\bar{x}$ and $\bar{y}$ are joined by an edge in W. For $e_j = (x, y)$, denote this edge by $\bar{e}_j$.

Consider a graph F with edges $\bar{e}_1, \ldots, \bar{e}_k$. To obtain a contradiction, it suffices to show that F is a k-forbidden subgraph of W. Notice that $\bar{e}_j|_l = e_j$ for all $j \le k$ and, therefore, $\{\bar{e}_1|_l, \ldots, \bar{e}_k|_l\} = Q$. This implies the first condition in the definition of a forbidden subgraph. The second condition is met too, as each vertex v of Q is lifted to a unique vertex $\bar{v}$ of F with property $\bar{v}|_l = v$. ∎

Remarks.

1. We can construct a game G with $\omega(G^n) = E_Q(n)$ for *all* n, if we allow countable (rather than finite) answer sets A and B. The latter condition captures the situation when the answers of the provers are allowed to be arbitrarily long.
2. The condition of connectivity we imposed on Q in Theorem 4.2 is only for simplifying the notation in the proof. An analogous theorem holds if Q is not connected. In this case, the game G must be nontrivial on at least one of the connected components of Q, say Q'. The definition of a forbidden subgraph can be made with respect to this Q', and the proof applies virtually without change.

3. If π is not uniform on Q, then the forbidden subgraph approach still applies. A direct way of applying it is to consider Q^n as an edge weighted graph, where the weight of an edge is the probability of the respective question pair being selected under π^n. Now a forbidden subgraph is not measured in terms of k, the number of edges that it has on the projection to coordinate ℓ, but by the sum of the weights of these edges with respect to π. This direct way of using the forbidden subgraph approach with nonuniform π preserves Proposition 4.1 and Theorem 4.2, with the appropriate modifications to the relevant definitions. However, proving upper bounds on $\omega(G^n)$ with the weighted forbidden subgraph approach becomes very difficult.
   A way of avoiding the weighted case is by observing that any π can be decomposed into $\pi_1+\pi_2$, where π_1 is uniform over the support Q, and π_2 accounts for the nonuniformity of π. Playing n copies of G can be interpreted as playing (on average) $\mu(\pi_1)n$ copies of the uniform version of the nontrivial game G, and $\mu(\pi_2)n$ copies of G with π_2 (the latter copies may in fact be of a trivial game, because the support of π_2 is smaller than Q). The verifier can ignore the latter copies, and then we are left with $\Omega(n)$ copies of a uniform game, for which the uniform version of the forbidden subgraph approach applies.
4. Theorem 4.2 can be extended to characterize the maximum value of $\omega(G^n)$ for a game $G(X,Y,Q,\pi)$ with *any* specified error $\omega(G)$, though the statement of the result in this most general case becomes much more cumbersome.

We conclude this section with discussion of asymptotics of $E_Q(n)$. As shown in [23], $\omega(G^n)\to 0$ for $n\to\infty$ uniformly over all nontrivial games with support Q. By Theorem 4.2, it follows that $E_Q(n)\to 0$ as $n\to\infty$ for any connected Q. In fact, the bound suggested in [23] applies directly to $E_Q(n)$. For completeness, we describe this bound below (in Proposition 4.3).

We need some well-known notions from Ramsey theory. Let $A=\{a_1,\dots,a_k\}$ be a finite set, and z be a variable that can be replaced with any element of A. Let $u(z)$ be an n-vector from $(A\cup\{z\})^n$ with at least one component z, and let $u(a_i)$ denote the vector obtained from $u(z)$ by replacing each occurence of z by a_i. Then the set

$$L=\{u(a_1),\dots,u(a_k)\}$$

is called a *combinatorial line* in A^n. By $r_k(n)$ we denote the maximum possible density $|W|/k^n$ of a set $W\subseteq A^n$ without combinatorial lines. The Furstenberg-Katznelson theorem [15] says that $r_k(n)=o(1)$ for any k.

Proposition 4.3. *For any bipartite graph $Q \subseteq X \times Y$ with $|Q| = k$,*

$$E_Q(n) \leq r_k(n).$$

Thus, $E_Q(n) = o(1)$ for any Q.

Proof. As easily seen, any combinatorial line in Q^n is a $|Q|$-forbidden subgraph. ∎

The bound of $r_k(n)$ for $E_Q(n)$ is rather weak. Currently, there are no strong upper bounds known for $r_k(n)$. Clearly, $r_k(n)$ does not decrease at an exponential rate as n grows. For example, by choosing W as the set of n-vectors with exactly n/k coordinates that are a_1, it follows that $r_k(n) > \Omega(1/\sqrt{n})$.

Question 4.4. Is $E_Q(n)$ exponentially small in n for any Q? Equivalently, is $\omega(G^n)$ exponentially small for all games with infinite answer sets A and B (see Remark 1 above)?

The answer is affirmative for Q being a complete bipartite graph [8,1,10] or a tree [25]. This problem has an interesting weaker version.

Question 4.5. Let Q be a connected bipartite graph, and let $D_Q(n)$ denote the maximum possible density of $W \subseteq Q^n$ without subgraphs isomorphic to Q. Obviously, $D_Q(n) \leq E_Q(n)$. Is $D_Q(n)$ exponentially small in n for any Q?

If Q is a complete bipartite graph, an affirmative answer is given by the Zarankievicz theorem. If Q is a tree, a strong bound on $D_Q(n)$ follows from the well-known fact that every graph on v vertices with at least tv edges contains any tree of size t as a subgraph. If Q is a cycle C_{2l}, an exponentially small upper bound on $D_Q(n)$ follows from the Bondy-Simonovits theorem that any graph of order v without subgraphs C_{2l} has at most $O\left(v^{1+1/l}\right)$ edges. (See [7] for information on the bounds cited in this paragraph.)

5. The lower bound

In this section we prove Theorem 3.1. For any n, we exhibit a game for which n repetitions are required in order to reduce the error from one constant to another. First, we partially define a free game G, specifying only the size of $|Q|$, without defining the answer set or the acceptance predicate. Then we turn to the forbidden subgraph problem for G^n, in which we wish to select as many edges as possible without having a forbidden subgraph with $|Q|/2$ edges. A probabilistic argument shows that a constant fraction of the

edges can be selected (for an appropriate choice of $|Q|$ as a function of n). Now we return to the game G and define its answer set and acceptance predicate. This is done in a way similar to the proof of Theorem 4.2, and is consistent with the pattern of edges that are selected in the forbidden subgraph problem. Finally, we prove that for G defined in this way, $\omega(G) \leq 3/4$.

The proof builds upon the ideas of Section 4, but is presented in a self contained manner, without explicit referral to Section 4.

Proof. (**Theorem 3.1**.) We show the existence of an infinite family of games $\{G_k\}$ with the following properties.

1. The support of game G_k is $X_k \times Y_k$, where $X_k = \{x_1, \ldots, x_k\}$ and $Y_k = \{y_1, \ldots, y_k\}$. Hence the size of the support is k^2. For simplicity of notation, we shall omit the subscript k from X_k and Y_k, as the value of k will be clear from the context.
2. Each game G_k is free. That is, the verifier picks a question pair uniformly at random from $X \times Y$.
3. The answer length of game G_k is $\Theta(k)$. (Hence length of answer is exponential in length of question.)
4. The error of game G_k is constant. More specifically, $\omega(G_k) \leq 3/4$.
5. For all $n \leq k/(4 \log k)$, the error in the n-fold parallel repetition of G_k remains above some constant. More specifically, $\omega((G_k)^n) \geq 1/8$.

For arbitrary large enough k, we explain how to obtain G_k with the properties described above. The definition of the acceptance criteria in game G_k depends on a certain graph G_k^n, that will be described shortly. Recall that in G_k, the question set for prover P_1 is $X = \{x_1, \ldots, x_k\}$, and the question set for prover P_2 is $Y = \{y_1, \ldots, y_k\}$. To describe the answer set of G_k, let $n = k/(4 \log k)$ (rounded down to the nearest integer). The answer set for prover P_1 is $\{1, \ldots, n\} \times X^n$ (that is, an integer from 1 to n, followed by a list of n questions x_i), and the answer set for prover P_2 is $\{1, \ldots, n\} \times Y^n$. The verifier picks a question pair (x', y') uniformly at random from $X \times Y$. The verifier accepts the answer pair $(\ell \circ x_{i_1} x_{i_2} \ldots x_{i_n})$, $(m \circ y_{j_1} y_{j_2} \ldots y_{j_n})$, if and only if $x' = x_{i_\ell}$, $y' = y_{j_m}$, $\ell = m$, and the two vertices named $x_{i_1} \ldots x_{i_n}$ and $y_{j_1} \ldots y_{j_n}$ are connected by an edge in the bipartite graph G_k^n, to be described below.

Each side of the bipartite graph G_k^n has k^n vertices, where lefthand side vertices are labeled by the strings X^n, and righthand side vertices are labeled by the strings Y^n. The bipartite graph has $k^{2n}/8$ edges arranged in such a way that any k by k vertex induced subgraph of G_k^n has at most $k^2/2$ edges.

Proposition 5.1. *A graph G_k^n with the parameters above exists.*

Proof. We show that a random graph selected from an appropriate distribution has positive probability of satisfying the requirements of G_k^n (an existence proof by a probabilistic argument). The lefthand side vertex set is X^n and the righthand side vertex set is Y^n. Include each edge (connecting different sides) in the graph independently with probability 1/8. Then the expected total number of edges is $k^{2n}/8$, and with probability roughly 1/2 the number of edges will be at least its expectation. Consider now an arbitrary k by k vertex induced subgraph. The probability that it has $k^2/2$ edges (or more) is less than $8^{-k^2/2}\binom{k^2}{k^2/2} \leq 2^{-k^2/2}/4$ (for large enough k). There are less than k^{2kn} ways of choosing a k by k vertex induced subgraph. Hence, if $n \leq k/(4\log k)$, the probability that at least one of them has $k^2/2$ edges is at most 1/4. ∎

To make the description of G_k concrete, we fix G_k^n to be the first (under an arbitrary order) bipartite graph on k^n by k^n vertices that satisfies the restrictions specified for G_k^n.

Lemma 5.2. *For the game G_k and n as described above, $\omega((G_k)^n) \geq 1/8$.*

Proof. We describe a strategy of the provers for the game $(G_k)^n$. On receiving a question $x_{i_1} \dots x_{i_n}$, P_1 gives the n answers $\ell \circ x_{i_1} \dots x_{i_n}$, for $\ell = 1, \dots, n$. On receiving a question $y_{j_1} \dots y_{j_n}$, P_2 gives the n answers $m \circ y_{j_1} \dots y_{j_n}$, for $m = 1, \dots, n$. By the definition of G_k, the strategy of the provers succeeds on all n repetitions of G_k simultaneously, whenever $(x_{i_1} \dots x_{i_n}, y_{j_1} \dots y_{j_n})$ is an edge of G_k^n. By definition of G_k^n, exactly one eighth of the tuples $(x_{i_1} \dots x_{i_n}, y_{j_1} \dots y_{j_n})$ are edges of G_k^n. ∎

Lemma 5.3. *For the game G_k as described above, $\omega(G_k) \leq 3/4$.*

Proof. Consider an arbitrary strategy for the two provers for G_k. For each question x_i to P_1, this fixes an answer $\ell_i \circ x_{i_1} \dots x_{i_n}$. For each question y_j to P_2, this fixes an answer $m_j \circ y_{j_1} \dots y_{j_n}$. Construct the bipartite graph G with lefthand side vertex set X, righthand side vertex set Y, and an edge between x_i and y_j if the verifier accepts $(x_i, y_j, P_1(x_i), P_2(y_j))$.

Proposition 5.4. *If x_i and y_j are in the same connected component of G, then $\ell_i = m_j$.*

Proof. Follow the path connecting x_i and y_j in G. Since each edge (x', y') along the path signifies the fact that V accepts the provers' answers, for both endpoints of the edge the answers of the provers on x' and y' must respect $\ell' = m'$. The proof follows by transitivity. ∎

Proposition 5.5. *Any connected component of G has at most $k^2/2$ edges.*

Proof. Each answer of a prover is composed from an index (denoted by ℓ or m) and a name of a vertex in G_k^n. Within a connected component, all answers of the provers correspond to pairwise distinct vertices in G_k^n. This follows from Proposition 5.4, together with the fact that in order for the verifier to accept P_1's answer $m \circ x_{i_1} \dots x_{i_n}$ on x', it must hold that $x_{i_m} = x'$ (a similar condition holds for P_2). Each connected component can be viewed as (part of) a k by k vertex induced subgraph of G_k^n, and hence has at most $k^2/2$ edges (by definition of G_k^n). ∎

In Proposition 5.6 we prefer simplicity of proof over obtaining the best constant (which is $2-\sqrt{2}$ rather than $3/4$).

Proposition 5.6. *G has at most $3k^2/4$ edges.*

Proof. For connected component C_i in G, let a_i denote the number of vertices on the lefthand side and let b_i denote the number of vertices on the righthand side. Arrange the connected components C_i of G in order of decreasing b_i. C_1 contributes at most $k^2/2$ edges to G (by Proposition 5.5). For each C_i, $i>1$, we compare between the number of edges that it adds to G and the number of edges that it excludes from being in G. Component C_i contributes at most $a_i b_i$ edges to G. Now count the number of edges that are excluded from G, by the fact that they join a vertex of C_i with a vertex of C_{i-1}. This number is $a_i b_{i-1} + a_{i-1} b_i > a_i b_i$ (because $b_{i-1} \geq b_i$). Hence beyond the first connected component, for any edge that we count in G we can count another edge as not being in G. ∎

Hence any strategy of the provers succeeds on at most $3/4$ of the support of G_k, and this completes the proof of Lemma 5.3. ∎

This completes the proof of Theorem 3.1. ∎

Acknowledgements. We thank Ran Raz and Alex Russell for helpful discussions.

References

[1] N. Alon: Probabilistic methods in extremal finite set theory, in: *Extremal Problems for Finite Sets,* (P. Frankl, Z. Füredi, G. O. H. Katona and D. Miklós Eds.), Bolyai Society Mathematical Studies,3, Visegrád, Hungary, 1991, 39–57, 1991.

[2] S. Arora, C. Lund, R. Motwani, M. Sudan, M. Szegedy: Proof verification and the hardness of approximation problems, *Journal of the ACM,* **45(3)** (1998), 501–555.

[3] L. Babai, L. Fortnow, C. Lund: Non-deterministic exponential time has two-prover interactive protocols, *Computational Complexity,* **1** (1991), 3–40.

[4] M. Bellare, S. Goldwasser, C. Lund, A. Russell: Efficient probabilistic checkable proofs and applications to approximation, *Proc. of 25th Annual Symposium on the Theory of Computing,* 294–304, 1993.

[5] M. Bellare, O. Goldreich, M. Sudan: Free bits, PCPs and non-approximability – towards tight results, *SIAM J. Comput.,* **27(3)** (1998), 804–915.

[6] M. Ben-or, S. Goldwasser, J. Kilian, A. Wigderson: Multi prover interactive proofs: how to remove intractability, *Proc. of 20th Annual Symposium on the Theory of Computing,* 113–131, 1988.

[7] B. Bollobas: *Extremal graph theory,* Academic Press, 1978.

[8] J. Cai, A. Condon, R. Lipton: On games of incomplete information, *Theoretical Computer Science,* **103(1)** (1992), 25–38.

[9] J. Cai, A. Condon, R. Lipton: PSPACE is provable by two provers in one round, *Journal of Computer and System Sciences,* **48(1)** (1994), 183–193.

[10] U. Feige: On the success probability of the two provers in one round proof systems, *Proc. of 6th IEEE Symposium on Structure in Complexity Theory,* 116–123, 1991.

[11] U. Feige: Error reduction by parallel repetition – the state of the art, *Technical report CS95-32 of the Weizmann Institute,* 1995.

[12] U. Feige, J. Kilian: Two prover protocols – low error at affordable rates, *SIAM J. Comput.,* **30(1)** (2000), 324–346.

[13] U. Feige, L. Lovasz: Two-prover one-round proof systems, their power and their problems, *Proc. of 24th Annual Symposium on the Theory of Computing,* 733–744, 1992.

[14] L. Fortnow, J. Rompel, M. Sipser: On the power of multi-prover interactive protocols, *Theoretical Computer Science,* **134(2)** (1994), 545–557.

[15] A. Furstenberg, Y. Katznelson: A density version of the Hales–Jewett theorem, *Journal d'Analyse Mathematique,* **57** (1991), 64–119.

[16] S. Goldwasser, S. Micali, C. Rackoff: The knowledge complexity of interactive proof systems, *SIAM J. Comp.,* **18(1)** (1989), 186–208.

[17] D. Lapidot, A. Shamir: A one-round, two-prover, zero-knowledge protocol for NP, *Combinatorica,* **15(2)** (1995), 203–214.

[18] D. Lapidot, A. Shamir: Fully parallelized multi prover protocols for NEXP-time, *Journal of Computer and System Sciences,* **54(2)** (1997), 215–220.

[19] C. Lund, M. Yannakakis: On the hardness of approximating minimization problems, *Journal of the ACM,* **41(5)** (1994), 960–981.

[20] I. Parnafes, R. Raz, A. Wigderson: Direct product results and the GCD problem, in old and new communication models, *Proc. of 29th Annual ACM Symposium on Theory of Computing,* 363–372, 1997.

[21] D. Peleg: On the maximum density of 0–1 matrices with no forbidden rectangles, *Discrete Mathematics,* **140** (1995), 269–274.

[22] R. Raz: A parallel repetition theorem, *SIAM J. Comput.,* **27(3)** (1998), 763–803.

[23] O. Verbitsky: Towards the parallel repetition conjecture, *Theoretical Computer Science,* **157(2)** (1996), 277–282.

[24] O. Verbitsky: The parallel repetition conjecture for trees is true, *Electronic Colloquium on Computational Complexity,* TR95-013, 1995.

[25] O. Verbitsky: Remarks on a query-based variant of the parallel repetition theorem, *International Journal of Foundations of Computer Science,* **12(4)** (2001), 517–532.

Uriel Feige

Department of Computer Science
and Applied Mathematics
The Weizmann Institute
Rehovot 76100, Israel
feige@wisdom.weizmann.ac.il

Oleg Verbitsky

Department of Mechanics
and Mathematics
Lviv University
79000 Lviv, Ukraine
oleg@ov.litech.net