Electronic Colloquium on Computational Complexity, Report No. 54 (2004)

Non-reducible descriptions for conditional Kolmogorov complexity

Andrej Muchnik* Alexander Shen†
Nikolai Vereshchagin‡ Michael Vyugin§

Abstract

Let a program p on input a output b. We are looking for a shorter program p' having the same property ($p'(a) = b$). In addition, we want p' to be simple conditional to p (this means that the conditional Kolmogorov complexity $K(p'|p)$ is negligible). In the present paper, we prove that sometimes there is no such program p', even in the case when the complexity of p is much bigger than $K(b|a)$. We give three different constructions that use the game approach, probabilistic arguments and algebraic (combinatorial) arguments, respectively.

1 Definitions and statements

Let a and b be binary strings. Consider programs p such that $p(a) = b$ (the program p on input a outputs b). What is the minimal length of such a program? If the programming language is chosen appropriately, this length is close to $K(b|a)$, the conditional Kolmogorov complexity of b given a. [We

*Institute of New Technologies; e-mail: muchnik@lpcs.math.msu.su. Work supported by RFBR grant 04-01-00427.

†The work was supported by CNRS (LIF, Marseille, France; the laboratory at Moscow Independent University), STINT foundation, Uppsala university (Sweden), Royal Holloway College (UK), RFBR (grants 02-01-22001, 03-01-00475) and Scientific schools supporting council (grant NSh-358.2003.1); e-mail: shen@mccme.ru.

‡Moscow State University, e-mail: ver@mccme.ru. The work was supported in part by the RFBR grants 02-01-22001, 03-01-00475, NSh-358.2003.1.

§Moscow State University, e-mail: misha@vyugin.mccme.ru

ISSN 1433-8092

will ignore additive terms of order $O(\log n)$ where n is the maximum length of the strings involved. With this precision all the versions of Kolmogorov complexity (the plain one, the prefix one etc.) coincide.]

To avoid references to a specific programming language we will consider "descriptions" instead of programs. A string p is called a conditional *description* of a string b given a if $K(b|a,p)$ is negligible. Here $K(b|a,p)$ stands for the conditional complexity of b given the pair $\langle a,p\rangle$. We will specify what is "negligible" in each case.

For given a and b consider all strings p such that $K(b|a,p) \approx 0$. One can easily verify that the length of any such p is at least $K(b|a)$. This bound is tight. (Both assertions are true with $O(\log n)$ precision; the same precision is required in the equality $K(b|a,p) \approx 0$.)

We say that a description p' is a *simplification* of a description p if $K(p'|p) \approx 0$ with logarithmic precision. The relation $K(p'|p) < \varepsilon$ is not transitive for a fixed ε: $K(p'|p) < \varepsilon$ and $K(p''|p') < \varepsilon$ imply only $K(p''|p) < 2\varepsilon + O(\log n)$. However, this relation resembles a preordering on strings and we are interested in the structure of the set of all conditional descriptions (for given a, b) with respect to this "pre-ordering".

The string b itself is a conditional description of b given a. Muchnik [1] has shown that (among all descriptions of b relative to a) there exists a description of minimal length ($\approx K(b|a)$) that is a simplification of b. We will prove that this is not true in the general case (for arbitrary description p instead of b): for some a, b there is a description p of complexity much larger than $K(b|a)$ that has no simplifications of length $\approx K(b|a)$.

The exact statement is as follows:

Theorem. There are constants $c_1 < c_2 < c_3 < c_4$, c and $\varepsilon > 0$ such that for all sufficiently large n there exist a, b, p of length at most $c_4 n$ having the following properties:

(a) $K(b|a,p) \leqslant c \log n$ ("the string p is a conditional description of b given a, with logarithmic precision");

(b) $K(b|a) \leqslant c_1 n$ ("the conditional complexity of b given a is small ..."");

(c) $K(p) \geqslant c_3 n$ ("... compared with the complexity of p");

(d) there is no string p' such that $K(p') \leqslant c_2 n$, $K(p'|p) \leqslant \varepsilon n$ and $K(b|a,p') \leqslant \varepsilon n$ ("... but p has no simplifications of complexity $c_2 n$").

Note that we are using linear upper bounds on $K(p'|p)$ and $K(b|a,p')$ instead of previously claimed bounds $O(\log n)$. This makes our statement stronger: there exists p having no simplifications p' even with linear upper bounds on conditional complexities. Note also that complexities $K(a)$, $K(b)$

of strings a, b provided by Theorem 1 are $\Theta(n)$ (and hence $|a|, |b| = \Theta(n)$). Indeed, if $K(a) < \delta n$ where δ is less than both ε and $c_2 - c_1$, then $p' = b$ is a counterexample to (d), since (a) and (b) imply $K(b|p) \leqslant \delta n + O(\log n)$ and $K(b) \leqslant (c_1 + \delta)n + O(\log n)$, respectively. And if $K(b) < \delta n$ (where $\delta < \varepsilon$), then the empty p' is a counterexample to (d), since (a) implies $K(b|a) \leqslant \delta n + O(1)$.

Let us mention also that for all our examples of strings a, b (except for the last example in Section 4 where random points and lines are used) the inequality (b) holds in a stronger form: $K(b) \leqslant c_1 n$.

In what follows we give three different proofs of the theorem, using three methods of constructing objects with given complexity properties (games, probabilistic arguments and combinatorial estimates).

In fact, our theorem is stated in a simple, but not the strongest, form. For example, our proof shows that for all $c_1 < c_2 < c_3 < c_4$ there exist c and ε satisfying the statement (we need only that ε is much less than differences $c_2 - c_1$ and $c_3 - c_2$).

Recently M. Ustinov has shown that *for all* a and b (except for trivial cases $K(a) \approx 0$ and $K(b|a) \approx 0$) there exists a program p that transforms a to b and cannot be simplified. This result was further improved by An. Muchnik (see [2]).

The authors are grateful to all participants of Kolmogorov seminar of the Department of Mathematics (Matematical Logic and Theory of Algorithms Division) at Moscow University.

2 Game approach

Consider the following game we play against an adversary.

Let P, P', A and B be finite sets (as we see later, they correspond to strings p, p', a, b respectively). On our moves we construct a partial function $\xi\colon P \times A \to B$. At the start of the game the function ξ is empty, and on each move we may define the value of ξ at one point (once defined values cannot be changed later). Or we may skip the move, that is, we may leave ξ unchanged.

The adversary on his moves constructs multi-valued functions $\varphi\colon P \to P'$ and $\psi\colon P' \times A \to B$. That is, the values of φ are subsets of P', and the values of ψ are subsets of B. Initially φ and ψ are empty (all their values are empty). At each move the adversary may add one new value to φ (adding a

new element to $\varphi(p)$ for some p) or ψ (adding a new element to $\psi(p', a)$ for some p', a). The existing elements cannot be removed. The adversary also may skip the move.

The adversary must obey the following rules: the function φ takes on every argument at most α values (i.e., $\#\varphi(p) \leqslant \alpha$ for any $p \in P$) and the function ψ takes on every argument at most β values ($\#\psi(p', a) \leqslant \beta$ for any p', a).

Players' moves alternate. Obviously, each player can make only finite number of non-trivial moves (moves that change the functions). Thus after a certain move all the three functions remain unchanged. The result of the game is defined as follows: we win if there exist $p \in P$, $a \in A$ and $b \in B$ such that $\xi(p, a) = b$ and p, a, b are not "covered" by the adversary: the latter means that there is no $p' \in \varphi(p)$ such that $b \in \psi(p', a)$.

So the game is determined by the sets A, B, P and P' (actually, only their cardinalities matter) and the parameters α and β. We represent the function ξ as a table with $\#P$ rows and $\#A$ columns. The cells of this table initially are empty; they are filled by elements of B (each cell may contain at most one element).

The adversary fills the table for function ψ. It has $\#P'$ rows of the same length $\#A$ as in our table. Each cell may contain up to β elements of B. The adversary also constructs the function φ. It is convenient to represent this function by arrows going from row p of our table to all rows of adversary's table that belong to $\varphi(p)$. The outdegree is bounded by α. We win if our table has a non-covered cell. A cell (p, a) is *covered* if its row is connected by an arrow to a row of adversary's table that has in the same column the same element of B (and, may be, some other elements). See Fig. 1.

The proof is based on the following simple observation:

Lemma. If $\alpha \cdot \beta < \#B$ and $\alpha \cdot \#P + \beta \cdot \#A \cdot \#P' < \#A \cdot \#P$ then we have a winning strategy in the game.

Proof of the lemma. The first inequality guarantees that if ξ is not yet defined on a pair p, a, then we can choose a value $b = \xi(p, a)$ so that the cell (p, a) is not covered (at the current step). Indeed, for each of at most α values $p' \in \varphi(a)$ there exist at most β values $b \in \psi(p', a)$, so there exists b that is different from all those values.

Choosing b in this way (assuming that there are empty slots in ξ-table), we guarantee that after each our move there exists a non-covered cell (p, a). Our move is non-trivial only when the previous adversary's move is non-trivial. The second inequality guarantees that the number of cells in ξ-table

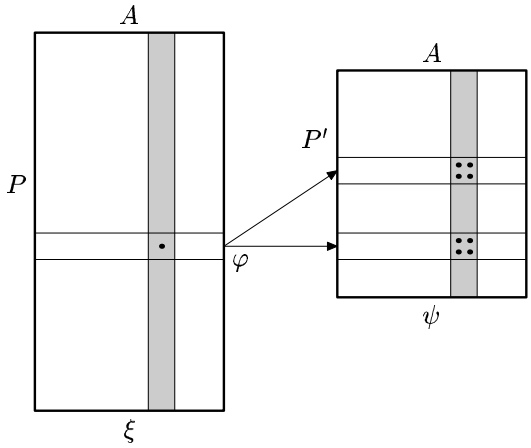

Figure 1: Cells of our table ξ and adversary's table ψ are filled with elements of B; each row of ξ has at most α outgoing edges, each cell of ψ contains at most β elements.

is greater than the number of adversary's non-trivial moves (so the empty slots do exist). Indeed, for each of $\#P$ arguments the value of φ may be changed at most α times and for each of $\#A \cdot \#P'$ pairs $\langle p', a\rangle$ the value of ψ may be changed at most β times.

Hence after every adversary's non-trivial move we can find an empty cell in ξ-table and enter a value in it so that the cell becomes non-covered. The lemma is proved.

Now we prove the theorem using Lemma. Fix some positive rational constants $c_1 < c_2 < c_3$ and $\varepsilon > 0$ such that ε is small compared with c_1, $c_2 - c_1$ and $c_3 - c_2$. Let B be the set of all strings of length at most $c_1 n$, let P' be the set of all strings of length at most $c_2 n$ and let P be the set of all strings of length at most $c_3 n$. The set A can be chosen in many ways, as we have almost no restrictions on a. For example, let A be equal to B.

Let us fix the adversary's strategy now. Assume that the adversary includes in $\varphi(p)$ (one by one) all $p' \in P'$ such that $K(p'|p) < \varepsilon n$, and includes in $\psi(p', a)$ all the strings $b \in B$ such that $K(b|a, p') < \varepsilon n$. One can do this effectively given n, as the function K is upper semi-computable (that is, the set $\{\langle x, y, l\rangle \mid K(x|y) < l\}$ is recursively enumerable). This strategy does not violate the rules provided $\alpha = \beta = 2^{\varepsilon n}$.

Let us verify that the conditions of the Lemma are satisfied:

$$\alpha \cdot \beta \leqslant 2^{2\varepsilon n+2} \ll 2^{c_1 n}$$

(we assume that ε is less than $c_1/2$), and both terms in the sum

$$\alpha \cdot \#P + \beta \cdot \#A \cdot \#P' \approx 2^{\varepsilon n + c_3 n} + 2^{\varepsilon n + c_1 n + c_2 n}$$

are much less than $\#A \cdot \#P = 2^{c_1 n + c_3 n}$ (we also assume that ε is less than $c_3 - c_2$). Therefore, by the Lemma, we have a winning strategy in the game.

The winning strategy is computable given n. Applying it against the adversary's strategy described above we obtain a function ξ that is computable given n (as the adversary's moves are computable, so are ours). To be precise we should write ξ_n indicating the dependence on n; complexity of algorithm that computes ξ_n is $O(\log n)$ since ξ_n is determined by n Since our strategy is a winning one, there exists a cell $\mathit{langlep}, a\rangle$ that is not covered after all non-trivial moves are performed. (It depends on n in a non-computable way, as we do not know which of the adversary's moves is the last non-trivial one.)

Let $b = \xi(p, a)$ be the element in the "winning" cell of our table. Then $K(b|a, p) = O(\log n)$. As the length of b is less than $c_1 n$ we have $K(b) \leqslant c_1 n + O(1)$. [This is $O(1)$ larger than the upper bound in the theorem but can be compensated by a small increase in c_1.] As the cell (p, a) is not covered, there is no string p' of *length* at most $c_2 n$ such that $K(p'|p) < \varepsilon n$ and $K(b|a, p') < \varepsilon n$. This is weaker than required: we want the statement to be true for all p' of *complexity* (not the length) less than $c_2 n$. However it is easy to fix this. Replacing p' by its shortest description we increase $K(b|a, p')$ and $K(p'|p)$ by $O(\log n)$ and this increase can be compensated by a small change in ε. Note also that lengths of all strings are at most $c_3 n$ so we may use any $c_4 > c_3$. It remains to fix only one problem: we want the complexity of p to be at least $c_3 n$ and the rules of the game do not provide any guarantee for this.

Let us change the game allowing the adversary at any step remove (="mark as unusable") any element of P; the total number of removed elements should not exceed $\#P/2$, so at least half of elements in P should remain intact. In the winning rule we require that element p has not been removed by the adversary. For the modified game the statement of the Lemma is changed as follows: in the right hand side of the inequality $\alpha \cdot \#P + \beta \cdot \#A \cdot \#P' < \#A \cdot \#P$ the term $\#A \cdot \#P$ is replaced by $\#A \cdot \#P/2$. The modified Lemma is still true: Indeed, if we cannot perform any move then all the non-removed p's

have been used with all a's, thus we have done $\#A \cdot \#P/2$ moves. And the conditions of the modified lemma are still fulfilled for large enough n.

Other changes are as follows: we let P be equal to the set of all strings of length at most $c_3 n + 2$, and the adversary removes all elements of P with complexity less than $c_3 n$. It is clear that at most half of elements could be removed, and all the other bounds remain true. After this modification we know that for the winning cell (p, a) the complexity of p is at least $c_3 n$, and the theorem is proved.

3 Probabilistic approach

Assume that finite sets A, B, P, P' are fixed. (They will play the same role as before.) Consider partial functions $\xi\colon P \times A \to B$ and multi-valued functions $\varphi\colon P \to P'$ and $\psi\colon P' \times A \to B$ having at most α and β values (respectively) for each argument.

Call a function ξ a *winning* function (cf. the game described above) if for all multi-valued φ and ψ (satisfying given bounds on the number of values) and for every set $\bar{P} \subset P$ of cardinality at most $\#P/2$ there exists a non-covered cell in a row outside $\bar{P}$, that is, there exist $p \in P \setminus \bar{P}$, $a \in A$ and $b \in B$ such that $\xi(p, a) = b$ but there is no $p' \in \varphi(p)$ such that $b \in \psi(p', a)$.

In other words, a function ξ is winning if we can put its values in the table ignoring the adversary's moves and be sure that we win. It is clear that without loss of generality we may assume that the functions φ and ψ always take maximum allowed number of values (if ξ wins in this case, it wins always). If a partial function ξ is a winning one, then any its total extension is also a winning function, so we consider only total winning functions in the sequel.

Thus if there is a winning function then there is a winning strategy. We will use probabilistic arguments to show that if the cardinalities of A, B, P satisfy certain requirements then a winning function exists. That is, we prove that with positive probability a randomly chosen function ξ is winning (assuming that all total functions ξ are equiprobable).

Let us estimate the probability that a random (total) function ξ does not win against given $\bar{P}$, φ and ψ; it is enough to show that this probability is so small that being multiplied by the number of different choices for $\bar{P}$, φ and ψ it is still less than 1.

Fix $\bar{P}$, φ and ψ. We need an upper bound for the probability that for all

$p \in P \setminus \bar{P}$ and all a the value $b = \xi(p, a)$ (that is chosen independently for all pairs $\langle p, a\rangle$) is covered by the functions φ and ψ. For a given pair $\langle p, a\rangle$ this probability is less than $\alpha\beta/\#B$, and the number of different pairs is at least $\#P \cdot \#A/2$. So we obtain the upper bound

$$(\alpha\beta/\#B)^{\#P\cdot\#A/2}.$$

Let us count now the number of different triples $\langle \bar{P}, \varphi, \psi\rangle$. We have at most $2^{\#P}$ choices for $\bar{P}$, at most $(\#P')^{\alpha\cdot\#P}$ choices for φ, and at most $(\#B)^{\beta\cdot\#A\cdot\#P'}$ choices for ψ. This gives a sufficient condition for the existence of a winning function:

$$(\alpha\beta/\#B)^{\#P\cdot\#A/2} \cdot 2^{\#P} \cdot (\#P')^{\alpha\cdot\#P} \cdot (\#B)^{\beta\cdot\#A\cdot\#P'} < 1.$$

What does this condition mean? Assume that $\alpha\beta < \#B/2$ (significantly larger $\alpha\beta$ do not satisfy the condition anyway). Let us focus on exponents in the inequality. The condition is true if all the exponents with bases greater than 1 are much less than the exponent with base less than 1:

$$\begin{aligned} \#P &\ll \#P \cdot \#A/2, \\ \alpha \cdot \#P &\ll \#P \cdot \#A/2, \\ \beta \cdot \#A \cdot \#P' &\ll \#P \cdot \#A/2. \end{aligned}$$

The first condition is true almost always, the second one means that $\alpha \ll \#A$, the third one means that $\beta \cdot \#P' \ll \#P$. We see that all these conditions (together with the inequality $\alpha\beta < \#B/2$) strengthen the conditions of the Lemma above (It could be expected since winning functions are special cases of winning strategies—those where all moves are fixed in advance and do not depend on the adversary's move).

In particular, a winning function exists if A, B, P, P', α, β are chosen as in the first proof of the theorem. Recall that we want $K(\xi(p, a)|a, p)$ to be $O(\log n)$. This can be achieved if the function ξ has Kolmogorov complexity $O(\log n)$, that is, the Kolmogorov complexity $K(\xi)$ of the graph of ξ is $O(\log n)$. To prove that there is a winning function ξ such that $K(\xi) = O(\log n)$ we can use the following (very general) argument: By a very long (but finite) exhaustive search we can check whether a given function is winning or not (checking all $\bar{P}$, φ and ψ). Thus we can probe all the functions ξ in some natural order until we find the first winning one. To run this algorithm

we need only to know n, hence the first winning function has Kolmogorov complexity $O(\log n)$.

The second proof of the theorem is completed.

What is the advantage of this (more complicated) proof? It shows that the theorem can be strengthened as follows: for every oracle X there exist p, a, b satisfying conditions (a)–(c) of the theorem (unchanged, without the oracle) such that there is no p' for which both $K^X(p'|p)$ and $K^X(b|a,p')$ are less than εn. Indeed, our winning function beats any adversary's strategy and its contruction (and the inequality $K(b|a,p) = O(\log n)$) does not depend on the enemy's strategy. [Instead of relativizing the Kolmogorov complexity by an oracle one can add any string as the extra condition in $K(p'|p)$ and $K(b|a,p')$.]

4 Algebraic construction

Although the proof in the previous section allows us to find the winning function by an exhaustive search, this search could be very long. We would like to have a more "explicit" example of the winning function. To this end we formulate certain conditions that guarantee that a function $\xi\colon P\times A\to B$ is a winning one. Then we will explicitly present a winning function satisfying those conditions.

Consider a function $\xi\colon P\times A\to B$. For every $p\in P$ consider the corresponding line in the table ξ, that is, the function $\xi_p\colon A\to B$ defined as $\xi_p(a)=\xi(p,a)$. We require that the functions ξ_p for different p (=different lines of the table ξ) are far away from each other. This requirement seems natural: if the number of different a's where $\xi_p(a)$ and $\xi_q(a)$ coincide is large, then the adversary may use the same p' for p and q.

Formally speaking, we give the following

Definition. A function ξ is γ-regular if for all $p\ne q$ the number of $a\in A$ such that $\xi_p(a)=\xi_q(a)$ is at most γ (=if the Hamming distance between corresponding lines is at least $\#A-\gamma$).

Lemma 1. If a function ξ is γ-regular,

$$8\alpha\beta^2<\#P/\#P' \quad\text{and}\quad 8\alpha\beta\sqrt{\gamma}<\sqrt{\#A},$$

then the function ξ is a winning one.

Proof. First we reduce the general case to the case $\beta=1$. To this end we replace every line in the table ψ by β lines (that contain the same elements

of B as the old line, one element per cell). The height of the table, $\#P'$, becomes β times bigger and the function φ has now β times more values (each arrow is replaced by β arrows). So α is replaced by $\tilde{\alpha} = \alpha\beta$. If a function ξ is winning in the modified game with $\tilde{P}' = \{1, \ldots, \beta\} \times P'$, $\tilde{\alpha} = \alpha\beta$ and $\tilde{\beta} = 1$ (all other parameters remain unchanged) then ξ is winning in the original game. Indeed, every $\bar{P}, \varphi, \psi$ for the original game can be transformed into $\bar{P}, \tilde{\varphi}, \tilde{\psi}$ for the modified game: let $\tilde{\varphi}(p)$ be the set $\{\langle i, p'\rangle \mid p' \in \varphi(p)\}$, and let $\tilde{\psi}(\langle j, p'\rangle, a)$ be equal to the jth value of $\psi(p', a)$), in some order. If ξ beats $\bar{P}, \tilde{\varphi}, \tilde{\psi}$ then it beats also $\bar{P}, \varphi, \psi$.

The conditions of the lemma translate into inequalities

$$8\tilde{\alpha} < \#P/\#\tilde{P}' \quad \text{and} \quad 8\tilde{\alpha}\sqrt{\gamma} < \sqrt{\#A}.$$

So we can assume that $\beta = 1$ from now on.

Let us split an α-valued function φ into α single-valued functions $\varphi_1, \ldots, \varphi_\alpha$. Each φ_i covers some cells of the table ξ. We will estimate the fraction of elements covered by φ_i and prove that it is less than $1/(2\alpha)$. This implies that less than half of all cells are covered.

Why any single-valued function φ covers few cells? The reason is that $\#P'$ is much less than $\#P$, thus the same line of the table ψ must correspond to many lines of the table ξ. By our assumption the lines of ξ have small intersection and hence cannot be easily covered by the same line. The formal argument use the following simple bound:

Lemma 2. Assume that a family of k subsets of an a-element set is given such that every two subsets in this family have at most γ common elements. Then the sum of cardinalities of all the subsets in the family is at most

$$2a + 2k\sqrt{a\gamma}.$$

Remark: For small k the first term of the sum $2a+2k\sqrt{a\gamma}$, not depending on k, is the main term; for large k the second term, linear in k, is the main term; two terms are equal for $k = \sqrt{a/\gamma}$.

Proof of Lemma 2. Let $a_1, \ldots, a_k$ be the cardinalities of the given subsets. The inclusions-exclusions formula implies that

$$a \geqslant a_1 + a_2 + \ldots + a_k - k^2\gamma$$

(there are at most k^2 pairwise intersections, each of cardinality at most γ). Therefore

$$a_1 + \ldots + a_k \leqslant a + k^2\gamma.$$

If $k \leqslant \sqrt{a/\gamma}$ then the second term ($k^2\gamma$) is bounded by a and the sum $a+k^2\gamma$ is at most $2a$. Hence the inequality of the lemma is true for all $k \leqslant \sqrt{a/\gamma}$. For $k = \sqrt{a/\gamma}$ we have also $a_1 + \ldots + a_k \leqslant 2k\sqrt{a\gamma}$, as in this case $2k\sqrt{a\gamma} = 2a$. Since the right hand side of the last inequality is linear in k, the inequality is true for all $k \geqslant \sqrt{a/\gamma}$. To demonstrate this let us delete from the sum $a_1 + \ldots + a_k$ all terms except for the $\sqrt{a/\gamma}$ largest ones. As the average of remaining terms is not smaller than the average of all terms, we are done.

Lemma 2 is proved.

In fact this proof works only if $\sqrt{a/\gamma}$ is an integer. This is not really important since one can easily adapt the arguments below and use Lemma 2 only for integer case, but we can still prove Lemma 2 in general case using more careful bounds. Namely, $a_1 + \ldots + a_k \leqslant a + (k(k-1)/2)\gamma$, since there are at most $k(k-1)/2$ pairwise intersections. Then for $k \leqslant \lceil\sqrt{a/\gamma}\rceil$ one has

$$a+(k(k-1)/2)\gamma \leqslant a+\sqrt{a/\gamma}(\sqrt{a/\gamma}+1)\gamma \leqslant a+\sqrt{a}(\sqrt{a}+\sqrt{gamma a}) \leqslant 2a \leqslant 2k\sqrt{a\gamma},$$

(since we may assume without loss of generality that $\gamma \leqslant a$), and the proof can be finished as before.

Let us continue the proof of Theorem 1. If k different lines of ξ are mapped by φ onto one line of ψ, then the sets of covered columns in any two of these lines have at most γ common elements. Hence the total number of covered cells in these k lines is at most

$$2\,\#A + 2k\sqrt{\#A\gamma}.$$

We have to sum this numbers for all $\#P'$ elements that can be values of the function φ, that is, over all lines of table ψ.

The first terms sum up to $2\,\#A \cdot \#P'$, the second ones sum up to $2 \cdot \#P\sqrt{\#A \cdot \gamma}$. So the total number of cells covered by each φ_i is at most

$$2\#A \cdot \#P' + 2 \cdot \#P\sqrt{\#A\gamma}.$$

Recalling that there are α functions φ_i we conclude that a function ξ is winning if

$$2\alpha\#A \cdot \#P' + 2\#P\alpha\sqrt{\#A\gamma} < \frac{1}{2}\#A \cdot \#P.$$

Lemma 1 is proved.

It is instructive to compare the requirements of Lemma 1 with those from the probabilistic argument. Note that the first requirement strengthens the

requirement $\beta \# P' \ll \# P$ and the second one strengthens the requirement $\alpha \ll \# A$.

It remains to construct a function ξ satisfying the conditions of Lemma 1. This can be done easily by the following algebraic construction. Let $A = B$ be the field of cardinality 2^n, and let P be the set of all linear functions $(x \mapsto a_1 x + a_2)$ from A to A. A linear function is determined by 2 coefficients, thus $\# P = 2^{2n}$. We can let $\gamma = 1$, as if two linear functions coincide in 2 points then they coincide everywhere. Let $P' = \{0,1\}^{1.5n}$. Let α and β be equal to $2^{\varepsilon n}$. For $\varepsilon < 1/6$ the conditions of Lemma 1 are fulfilled. We obtain a proof of the theorem with, say, $c_1 = 1.01$, $c_3 = 1.99$, $c_2 = 1.5$ and any $\varepsilon < 1/6$, $c > 2$ (small changes in c_1 and c_3 are needed to compensate for $O(\log n)$ terms). In place of linear functions we can take polynomials of small degree obtaining a proof with the same c_1, c_2 and larger c_3, ε.

Here is a more "geometric" example. Consider the two-dimensional vector space (the plane) over the finite field of cardinality 2^n. The set A consists of all points of this plane and the set B consists of all lines on it. The set P consists also of all points of this plane. The function ξ is defined as follows: $\xi(p, a)$ is the line passing through a and p. This time $\gamma = 2^n$, as the line ap_1 coincides with the line ap_2 only if a lies on the line $p_1 p_2$. Let $P' = \{0,1\}^{1.5n}$. If ε is small enough the conditions of Lemma 1 are satisfied. And the conditional complexity of $b = \xi(a, p)$ given a is at most $n + O(\log n)$, as there are about 2^n lines passing through any given point. Apply the winning strategy based on the function ξ against adversary's strategy from Section 2. The covered subset of $A \times P$ is small and can be enumerated given n. This implies that all the random pairs in $A \times P$ (those whose complexity is close to $4n$) are not covered. Therefore we can reformulate the result as follows (taking into account that the line passing through a pair of random independent points is random):

> any random line b on the plane over the field of cardinality 2^n has conditional complexity $\approx n$ given every its random point a; every other random point p on that line is a description of complexity $2n$ for b (given the point a) that cannot be reduced to a description of complexity $1.5n$.

(More precisely, we should require a and p be independent random points on b, i.e., $K(a, p|b) \approx 2n$.)

The constructions of this section have the following advantage compared with proofs from Sections 2 and 3: The complexity of $K(b|a)$ remains small

even if we consider time-bounded version of Kolmogorov complexity, i.e., require that the running time of the machine finding the object from its description is bounded by a polynomial in n. And the non-reducible program exists even for complexity relativized by any oracle, as in Section 3.

References

[1] Andrej A. Muchnik, Conditional complexity and codes, *Theoretical Computer Science*, **271** (2002), p. 97–109.

[2] An. Muchnik and M. Ustinov, *Constructing non-reducible programs for given pair of strings*, Preprint.

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'