

An $O(n^{0.4732})$ upper bound on the complexity of the GKS communication game

MARIO SZEGEDY

Department of Computer Science
Rutgers, The State University of New Jersey
szegedy@cs.rutgers.edu

June 21, 2015

Abstract

We give an $5 \cdot n^{\log_{30} 5}$ upper bound on the complexity of the communication game introduced by G. Gilmer, M. Koucký and M. Saks [3] to study the Sensitivity Conjecture [4], improving on their $\sqrt{\frac{999}{1000}}\sqrt{n}$ bound. We also determine the exact complexity of the game up to $n \leq 9$.

1 The $O(n^{0.4732})$ upper bound

The GKS communication game, defined by G. Gilmer, M. Koucký and M. Saks [3] is played by two cooperating players, Alice and Bob, against an all-powerful adversary, Merlin. The game has a single parameter n . Merlin has a permutation $\pi = \pi_1\pi_2 \dots \pi_n$ of $[n]$ and a bit b . Alice has a strategy $S : \{\text{partial permutations on } [n]\} \rightarrow \{0, 1\}$ and Bob has a strategy $T : \{0, 1\}^n \rightarrow 2^{[n]}$.

In Phase 1. Alice assigns zeroes and ones to all but one entries of an array $A[1..n]$ and Merlin sets the remaining entry to 0 or 1 according to the *Alice-Merlin protocol* described below. In Phase 2 Bob has to guess which entry was set by Merlin by merely looking at A_{final} , where A_{final} is the setting of A when Phase 1 is finished. Bob's guess, $T(A_{\text{final}}) \subseteq [n]$, is a subset of entries of A that must include the entry Merlin has set. This has to hold for every strategy π, b of Merlin.

Alice-Merlin protocol: For $1 \leq i \leq n - 1$ Alice sets

$$A[\pi_i] = S(\pi_1 \dots \pi_i)$$

In the end Merlin sets $A[\pi_n] = b$.

Definition 1. A (k, n) strategy for the GKS game with parameter n is a pair S, T as above such that in addition $|T(\sigma)| \leq k$ for every $\sigma \in \{0, 1\}^n$.

$$k(n) = \min_k \text{ There is a } (k, n) \text{ strategy for the GKS game with parameter } n$$

The relevance of the GKS communication game is that $k(n)$ gives a lower bound on the sensitivity, $s(f) = \max_x |\{i \mid f(x \oplus e_i) \neq f(x)\}|$ for any Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ with Fourier degree $\deg(f) = n$. See [3]. See [5, 1, 2] for more background. In turn, any lower bound $k(n) \in \Omega(n^\alpha)$ for some $\alpha > 0$ would positively resolve the long-standing Sensitivity Conjecture [4] which says that the sensitivity and the Fourier degree are polynomially related.

Example 1. There exists a $(2, 4)$ strategy for the GKS game as follows: Alice (mentally) decomposes $\{1, 2, 3, 4\}$ into blocks, $\{1, 2\}$ and $\{3, 4\}$. When Merlin gives a position in a yet untouched block, Alice answers with 0, and the second time a block is touched she answers with 1 (unless it is the last entry overall and so Merlin's turn). Assume now that what Bob sees is $A = [1, 0, 0, 1]$. Then he knows that the last bit had to be at position 1 or at 4 (at 2 or 3 it could not be, since in that case the first position touched in that block is 1, contrary to Alice's strategy). If $A = [0, 0, 1, 0]$ or $A = [1, 1, 1, 0]$, the last bit had to be at position either 1 or 2, since Bob can deduce that Merlin had to interfere in the first block (in both final results) for the situation to arise. These are the only possible cases up to a permutation that respects the blocks.

Lemma 1. *If (k, n) and (k', n') strategies exist, then there is also a (kk', nn') strategy.*

Proof. Let S, T be a (k, n) strategy and S', T' be a (k', n') strategy. We design a (kk', nn') strategy as follows: Decompose the nn' elements into n' blocks of size n , e.g. in the fashion

$$\{1, 2, 3\}, \{4, 5, 6\}, \{7, 8, 9\}, \{10, 11, 12\} \quad (n = 3, n' = 4)$$

Alice's strategy:

1. Until the last element of any given block is reached, Alice follows strategy S restricted to that block. She does this independently for all blocks.
2. When the last element of a block i is reached, Alice decides at the value of the corresponding bit in such a way that the *sum of the bits modulo two* in the block agrees with the bit that strategy S' would give to the single entry, i , in a corresponding situation. More precisely: Alice mentally runs strategy S' on an array $A'[1..n']$ with indices corresponding to the blocks of the compound game. Every time when a block i in the compound game is about to be completed, she computes the entry i of A' , by strategy S' . At the same time she gives an assignment to the last entry of block i in the compound game in such a way that with the new bit the mod 2 entry-sum of block i equals to $A'[i]$. However, if block i is the very last one to be completed, it is Merlin's turn to assign the last bit. The mod 2 entry-sum of block i now gives some arbitrary evaluation of $A'[i]$, which is fine, since the last entry of A' is Merlin's move in the S', T' game as well.

Assume for instance that blocks 3, 4 are completed (in this order) and now Alice is to evaluate the single remaining bit of block 2 to Merlin's order. Then she finds out the bit that strategy S' gives to $A'[2]$ in the situation when $A'[3]$ and $A'[4]$ were set in this order. If at this point the assignment (in the composed game) is $(**0)(*10)(001)(101)$, and strategy S' says that with the 342 permutation-start Alice's evaluation of position 2 is $0 = S'(342)$, Alice must evaluate the last entry of block 2 as 1, because this makes the mod 2 sum of the second block $0 = S'(342)$. Thus the new assignment is $(**0)(110)(001)(101)$.

Bob's strategy:

Bob computes the mod 2 sum for each block of the array of the compound game to get an array A'_{final} of length n' . Then he computes $T'(A'_{\text{final}})$. This gives him at most k' indices. The crucial

last entry of the compound game must come from a block indexed from $T'(A'_{\text{final}})$. By applying T on each of these blocks separately, Bob gets at most $k'k$ entries total, and it is easy to see that they are the only candidates for the last entry. \square

G. Gilmer, M. Koucký and M. Saks have proven the existence of a $\left(\left\lceil\sqrt{\frac{999}{1000}}\sqrt{n}\right\rceil, n\right)$ strategy [3]. We describe a modification of the construction in [3], which together with our first lemma will give the $5 \cdot n^{\log_{30} 5}$ upper bound.

Lemma 2. *There is a $(5, 30)$ strategy for the GKS game.*

Proof. We first describe Alice's strategy. Before the game she mentally decomposes the entries of A into 5 blocks, each of length 6, in the fashion

$$\{1, 2, \dots, 6\}, \{7, 8, \dots, 12\}, \{13, 14, \dots, 18\}, \{19, 20, \dots, 24\}, \{25, 26, \dots, 30\}$$

For each block j ($0 \leq j \leq 4$) Alice performs the following (identical) protocol, independently: When Merlin tells Alice to access the block for the first time, and the first requested element from the block has index $6j + i$, Alice looks up w_i from

$$\begin{aligned} w_1 &= 0 & 0 & 0 & 0 & 0 & 0 \\ w_2 &= 1 & 0 & 0 & 1 & 1 & 0 \\ w_3 &= 0 & 1 & 0 & 1 & 0 & 1 \\ w_4 &= 0 & 0 & 1 & 0 & 1 & 1 \\ w_5 &= 1 & 1 & 1 & 0 & 0 & 0 \\ w_6 &= 1 & 1 & 1 & 1 & 1 & 1 \end{aligned}$$

When she evaluates any element with index $6j + i'$, except the last one from the block, she sets $A[6j + i']$ to $w_i[i']$. She sets the last entry from the block to $1 - w_i[i']$.

Notice that the Hamming distance between any $w_{i'}$, $w_{i''}$ ($i' \neq i''$) is at least three. This gives Bob the following recovery strategy: Regardless whether Merlin or Alice has set the last bit of block j , Bob can decode $i = i(j)$ for that block, because the Hamming distance of the block from $w_{i(j)}$ is at most one in both cases. If the last remaining bit of block j was controlled by Alice, it must be the only i' with

$$A[6j + i'] = 1 - w_i[i'] \tag{1}$$

If this equation fails to hold for all $1 \leq i' \leq 6$ for some block j_0 , Bob knows that the last bit of that block was set by Merlin. In this case Bob's output is the set $\{6j_0 + i' \mid 1 \leq i' \leq 6 \wedge i' \neq i(j_0)\}$. He could exclude $i' = i(j_0)$, since he knows that $6j_0 + i(j_0)$ was the first (and so not the last) requested element of that block to evaluate. If there is no such block j_0 (if there is, it must be unique), Bob outputs the unique i' for each $0 \leq j \leq 4$ for which Equation (1) holds. Either ways the set he outputs has size five. \square

Putting Lemmas 1 and 2 together we get:

Lemma 3. *There is a $(5^\ell, 30^\ell)$ strategy for the GKS game for $\ell = 1, 2, \dots$*

Corollary 1. $k(n) < 5 \cdot n^{\log_{30} 5}$

2 Further directions

The strategy in Lemma 2 generalizes as follows.

Definition 2. A (k, k_A, n) strategy is a usual (k, n) strategy equipped with an additional “Alice-mode.” In this Alice evaluates the last position as well (but otherwise she makes the exact same steps as in the usual mode). Bob does not know if the game is in Alice-mode or not. We denote the set of outcomes that may arise when the game runs in Alice mode by \mathcal{O}_A . The size of the set that Bob sends as an answer to any evaluation of A that is in \mathcal{O}_A has to be at most k_A .

Example 2. Consider the trivial strategy, where Alice evaluates all (non-final) entries of A to 0. In Alice-mode she also evaluates the final entry, to which she always gives the value 1. Bob’s strategy is the following: his output set includes all n entries of A when he sees only zeroes, but when he sees a 1, his set contains only that entry (he cannot see two ones regardless whether the game has run in Alice-mode or not). It is easy to see that this is an $(n, 1, n)$ strategy.

The above refined parametrization of strategies are useful for the following reason:

Lemma 4. *Assume there is a (b, a, n) strategy. Then there is also a $(b, \lfloor \frac{b}{a} \rfloor \times n)$ strategy.*

Proof. Assume that \mathcal{O}_A is set of possible final settings of the array in Alice-mode. Alice plays the (b, a, n) strategy on $\lfloor \frac{b}{a} \rfloor$ disjoint blocks, each of size n , independently, in such a way that in all blocks except in the final one she plays in Alice-mode. Bob looks at the blocks individually. If there is a block where the final setting is not in \mathcal{O}_A , Bob outputs only the output set associated with that block. Otherwise he outputs the union of the output sets for all blocks. In the former case the output size is at most b , in the latter it is at most $\lfloor \frac{b}{a} \rfloor \times a \leq b$. \square

The above lemma and a $(5, 1, 6)$ strategy (described in the previous page) yields the proof of Lemma 2. In general, it is worthwhile to design $(k, 1, n)$ strategies even when k is close to n . We suspect that there is a $(11, 1, 15)$ strategy. In this Alice exploits a 1-1 correspondence τ between the set of all four-element subsets of $\{1, \dots, 15\}$ and an appropriate $\binom{15}{4} = 1365$ size subset of the 1-error correcting Hamming code with length 15 and dimension 11. What is missing (but is very likely true) is that τ can be designed so that for every $H \subseteq [15]$, $|H| = 4$, and for $x = \tau(H) \in \{0, 1\}^{15}$ it holds that $x_H = 1111$, i.e. x restricted to any coordinate in H is 1. This type of construction goes back to [3]. The strategy: Alice’s first four answers are always 1, and her other answers make the final vector (when in Alice-mode) to be $\tau(\{i_1, i_2, i_3, i_4\})$, where i_1, \dots, i_4 are the first four entries that Merlin requests Alice to evaluate. If a $(11, 1, 15)$ strategy exists, by Lemma 4 it should give a $(11, 165)$ strategy and by Lemma 1 an $O(n^{0.47})$ upper bound on the GKS game.

Even better parameters could result from the 1-error correcting, non-linear code of length 9 that has 40 code words (see [6]). Here a map, similar to τ would yield a $(7, 1, 9)$ strategy. We remark that the problem of finding τ is a bipartite matching problem, thus it can potentially be done with a computer, but in the case of the Hamming code there are also more direct ways to try.

Problem 1. For given $n > 0$ compute or estimate the minimal k such that a $(k, 1, n)$ strategy exists.

Problem 2. Is there a $(O(\sqrt{n}), n)$ strategy in which Alice always answers with zero in the first $n - O(\sqrt{n})$ steps?




Figure 1: The minimum (2, 2) subgraph.

3 Exact bounds for dimensions up to 9

The GKS game also has a more combinatorial form [3]:

Lemma 5. *There is a (k, n) strategy for the GKS game if and only if there is a subgraph G with maximum degree at most k , of the n dimensional hypercube, $K_2^{\square n}$, such that in the Game below, the Chooser has a winning strategy.*

Game [3]: *The game is played by the Divider and the Chooser on $\{0, 1\}^n$. At each step the Divider picks a coordinate $i \in [n]$ not picked before and the Chooser decides whether to delete nodes x with $x_i = 0$, or with $x_i = 1$ from the current set of nodes. The game lasts $n - 1$ rounds and the Chooser wins if the remaining two vertices form an edge of G .*

Let us call a subgraph of $K_2^{\square n}$ as above with maximum degree k a (k, n) subgraph.

Lemma 6. *The structure of the best strategies up to $n = 4$ is characterized by:*

1. Any (1, 1) subgraph is an edge.
2. Any (2, 2) subgraph must contain a subgraph that is by an automorphism of the square equivalent to $\{*0, 0*\}$ (Figure 1).
3. Any (2, 3) subgraph must contain a subgraph that is by an automorphism of the cube equivalent to $\{10^*, *00, 0^*0, 01^*\}$. This is a (particular) path with four edges (Figure 2).
4. Any (2, 4) subgraph must contain a subgraph that is by an automorphism of the 4-hypercube equivalent to

$$\{ *000, 0^*00, 01^*0, 011^*, 0^*11, *011, 101^*, 10^*0 \}$$

This is a (particular) cycle of length 8 (Figure 2).

Proof. We give a sketch of the proof. Items 1.-3. are easy. Proving 4. seems to require a case separation. Consider a minimum (2, 4) subgraph G (which must be a vertex-disjoint union of cycles and paths). After the Chooser leaves out half of the vertices, the graph reduces to a (2, 3) subgraph, so by item 3, without loss of generality we can assume that G contains $\{10^*0, *000, 0^*00, 01^*0\}$. We show that the only extension of this edge set to any minimal (2, 4) subgraph is the one drawn in Figure 2.

Towards this goal we further exploit the structure that item 3. gives. For $1 \leq i \leq 4$, if we delete all edges in direction i from G , the remaining graph must still contain a path P_i of the shape




Figure 2: The minimum (2,3) and (2,4) subgraphs.




Figure 3: A six-edge configuration that is ruled out.


described in 3. By the minimality of G we have $G = P_1 \cup P_2 \cup P_3 \cup P_4$. These P_i s may share nodes and/or edges. We first rule out:

Case. Node 0010 occurs in the middle of one of P_i ($1 \leq i \leq 3$).

In this case G must either contain the subgraph in Figure 3 or its mirror-symmetric version where the 0^*10 edge is replaced with the $*010$ edge. Without loss of generality let us assume the first. Then the path that contains 0010 in the middle can be either


- (A) $P_1 = \{01^*0, 0^*10, 001^*, 00^*1\}$ or
- (B) $P_1 = \{0^*10, 001^*, 00^*1, 0^*01\}$ or
- (C) $P_3 = \{01^*0, 0^*10, 001^*, 00^*1\}$

We first show the impossibility of case (A), i.e. if G contains the pattern:




The undoing of (A) is that we cannot accommodate P_3 . To see this, imagine that we take out all of the vertical edges and look for P_3 either in the upper half ($x_3 = 1$) or in the lower half ($x_3 = 0$). In the upper half we have four free nodes and the node 1010, which has degree one. The latter should be a starting node of P_3 (it is one of the only five nodes we can use, so P_3 must use it). For adjacent edge only 101* or 1*10 come into question. The other final edge of P_3 then would be 011* or 0*11, respectively, each resulting in nodes of degree three. In the lower half we face a similar problem. Since node 1100 must be used in P_3 as an end point with adjacent edge 110*, the other final edge of P_3 must be 000*, creating not even only one, but two nodes with degree three (0000 and 0001). Case (B) is subsumed by (A), since it gives a strictly larger $P_4 \cup P_1$.


Similar argument rules out case (C):



We try to accommodate P_1 . Accommodating P_1 in the half $x_1 = 0$ is entirely out of question, since there are only three vertices available. The picture below shows the existing edges and degrees in the $x_1 = 1$ half (the x_4 direction is changed to horizontal) and the only way to accommodate P_1 (dotted line) in that half:



In the large picture:



But then P_2 cannot be accommodated.


This does not only finish the impossibility of the **Case**, but it gives the general statement that the pattern



should not occur in G . The reason is that the above pattern *in a minimal G* occurs if and only if the shape in Figure 3 occurs in G . The argument supporting this goes that in a minimal G the above pattern must be covered by a union of two P_i s. The one that contains two of the highlighted edges (at least one of the P_i s must be such) without loss of generality can be identified P_4 where the identification also has the property that the two edges in question are 0^*00 and 01^*0 . The other $P_{i'}$ must then contain 0010 as a middle point, referring us to **Case**.

We now develop a new representation for the shape of a path: we go through the path from one end-edge to the other and list the directions in which the edges go. In the case of a cycle we pick an arbitrary starting edge. When the cycle is a connected component of G (this is the case we are interested in), we can always start at the start of a P_i . By further exploiting symmetries we set the start (whether a path or a cycle) to 3123 (of course, only when we are searching for shapes of *single* connected components of G). In our new representation for instance the


path and cycle in Figure 2 have shapes 3123 and 31234214, respectively. We now build a tree representing the shapes of all potential path- or cycle- components of G (up to isomorphism) taking into consideration that a.) the path must be simple or a cycle b.) the forbidden pattern "aba" should not occur, and c.) the path (or cycle) must be a union of some P_i s:



We get that any path or cycle that *starts with* 3123 and which can occur in G must be one of:

Type 1:	3123	contains P_4 ;	5 nodes path
Type 2:	31231421	contains P_4 and P_3 ;	9 nodes path
Type 3:	3123143	contains P_4 and P_2 ;	8 nodes path
Type 4:	31234124	contains P_4 and P_3 ;	8 nodes cycle
Type 5:	3123413	contains P_4 and P_2 ;	8 nodes path
Type 6:	312342	contains P_4 and P_1 ;	7 nodes path
Type 7:	31234214	contains P_1, P_2, P_3, P_4 ;	8 nodes cycle

The shapes of all components of G must be *equivalent* to one of these types, where on equivalence we mean that we allow to permute $\{1, 2, 3, 4\}$. From now on, on a shape we mean the entire equivalence class. A cycle of Type 7 satisfies the criteria for G , and our goal is to show that there is no other solution. To look for a further (minimal) solution we can omit Type 7 from our list. By counting nodes we obtain that the connected components of G cannot be all of Type 1. A single Type 1- component and a single other one from types 2-6 contain only three of the P_i s. On the other hand, a three or more component G with any other than Type 1 components would have too many nodes. So we can eliminate Type 1 components altogether from our consideration. Perhaps the simplest way to proceed from here is just to look at each type from 2 to 6 and check if we can accommodate all P_i s. As an example we show how to exclude Type 6, which may seem the most economic of all the types among 2-6:



The solid line is now a whole *connected component* (of Type 6), so P_2 and P_3 , the remaining missing P_i s cannot share any vertex with this path. We argue that it is impossible to accommodate P_3 . The $x_3 = 1$ half has only four free nodes, thus we have to accommodate P_3 in the $x_3 = 0$ half. The latter has 5 free nodes, but the starting point of P_3 has to be 1100 with attached edge 110*. But then P_3 would also have to contain the edge 000*, which collides with our component. In fact, the above argument almost without any change works for types 2-5 as well: we cannot accommodate P_3 on the $x_3 = 1$ and $x_3 = 0$ halves for the very same (or in some cases even simpler) reasons. We are done, since we have excluded all types of components except Type 7. We can also observe (in order to prove the exact statement promised in the beginning), that there is a unique Type 7 cycle that contains our initial P_4 , the one drawn in Figure 2. \square

Lemma 7. *There is no (2, 5) strategy.*

Proof. Assume on the contrary that there is a (2, 5) subgraph G (corresponding to a (2, 5) strategy). When the Divider picks x_i , and the Chooser picks a side (either $x_i = 0$ or $x_i = 1$), the remaining subgraph must be a (2, 4) subgraph, hence it must contain a cycle C_i of length 8 of the shape described in Lemma 6. Moreover, the cycle C_i contains edges in all directions except in direction i . Because the maximum degree of G is at most two, any two cycles in G must either coincide or must be disjoint. For the above two reasons for $1 \leq i \neq j \leq 5$ the cycles C_i and C_j have to be disjoint. But $5 \times 8 = 40 > 32$, a contradiction. \square

Since it is known that there is a (3, 9) strategy (in general, (k, k^2) strategy), we have:

n	1	2	3	4	5	6	7	8	9	10
best k	1	2	2	2	3	3	3	3	3	?

References

- [1] H. Buhrman and R. de Wolf, *Complexity measures and decision tree complexity: a survey*. Theor. Comput. Sci., 288(1):21-43, 2002.
- [2] P. Hatami, R. Kulkarni, D. Pankratov, *Variations on the Sensitivity Conjecture*. Theory of Computing, Graduate Surveys 4: 1-27 (2011)
- [3] Justin Gilmer, Michal Koucký and Michael E. Saks, *A New Approach to the Sensitivity Conjecture*. ITCS 2015: 247-254
- [4] Nathan Linial, Noam Nisan, personal communication
- [5] N. Nisan and M. Szegedy, *On the degree of boolean functions as real polynomials*. Computational Complexity, 4:301-313, 1994.
- [6] <http://www.win.tue.nl/~aeb/codes/binary-1.html>