Acta Informatica 28, 593–601 (1991)

Acta Informatica

Some order dimension bounds for communication complexity problems

U. Faigle and W. Kern

Faculty of Applied Mathematics, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

Received March 2, 1990 / March 6, 1991

Summary. We associate with a general (0, 1)-matrix M an ordered set $P(M)$ and derive lower and upper bounds for the deterministic communication complexity of M in terms of the order dimension of $P(M)$. We furthermore consider the special class of communication matrices M obtained as cliques vs. stable sets incidence matrices of comparability graphs G. We bound their complexity by $O((\log d)\cdot(\log n))$, where n is the number of nodes of G and d is the order dimension of an orientation of G. In this special case, our bound is shown to improve other well-known bounds obtained for the general cliques vs. stable set problem.

1. Introduction

Communication complexity in general deals with the following model in distributed computing. Player I governs some data x taken from a finite universe X and player II governs some data y from a finite universe Y. The pair (x, y) determines a unique value $M_{xy} \in \{0, 1\}$. The two players want to determine M_{xy} in a cooperative effort by exchanging as few bits of information as possible. The *communication complexity* of this "game" is the minimal number of bits needed in the worst case as (x, y) ranges over $X \times Y$. Thereby we assume in our model that the two players agree on a *protocol* which only depends on the *communication matrix* $M = (M_{xy})$ and is fixed before the actual choice of x and y.

To be more precise, we should talk about the *deterministic* communication complexity $cc(M)$ of the matrix M. There are also non-deterministic versions of the above problem, but we will not be concerned with these (see Lovász 1989; Yao 1979; Mehlhorn and Schmidt 1982; Aho et al. 1983, for more information and formal definitions).

There are two different approaches one can take in a combinatorial analysis of the communication complexity problem. M may be interpreted as the incidence matrix of a bipartite strict order relation $P(M)$. Thus the communication problem allows the following model: given a strict order relation P and two elements x and y in the underlying set, it is to be decided whether $x<y$ holds in P. This suggests to study $cc(M)$ in terms of order-theoretic parameters. We derive bounds for $cc(M)$ in terms of the order dimension in Sect. 3. In particular, we show that $cc(M)=O(\log r(M))$ if the order dimension of $P(M)$ is bounded. (For more order-theoretic aspects of the communication problem, see Faigle and Turán 1990).

The second approach is based on the following observation. If $M=UV$ is a factorization of M into (0, 1)-matrices U and V, U can be interpreted as the incidence matrix of some cliques vs. the nodes of a graph G while V is the incidence matrix of the nodes vs. some stable sets of G. M thus is the incidence matrix of the cliques vs. the stable sets of U and V. Note that $M=MI$ always yields a trivial such factorization. Thus we arrive at the general communication problem: given a graph G, a clique x and a stable set y of G, it is to be decided whether x and y have a node in common. Yannakakis (1989) has shown that the communication complexity of this problem is $O(\log^2 n)$, where n is the number of nodes of G. The fundamental open question is whether that bound can be improved to $O(\log n)$.

In Sect. 4, we investigate the communication problem in that setting for the special case where G is the comparability graph of the order relation P and derive a protocol of complexity $O((\log \dim P)\cdot \log n)$. We also discuss the relationship of our bound with the bound obtained by Lovász and Saks for the general communication problem.

2. The communication problem for ordered sets

Let M be a $(m\times n)$-matrix with entries in $\{0, 1\}$. Then M can be interpreted as the incidence matrix of a strict order relation $P=P(M)$ defined on the set of rows and columns of M with the only nontrivial relations:

$$x<y \quad \text{if and only if} \quad M_{xy}=1.$$

We will, therefore, consider an equivalent order-theoretic formulation of the communication problem: Given a strict (partial) order $P=(E, <)$ on the set E with $|E|=n$, players I and II independently choose an element x and an element y resp. in E and want to decide whether

$$x<y$$

holds by exchanging as few bits of information as possible. Again, we assume that, at the outset, both players have complete information about P and decide cooperatively on a strategy ("protocol") to settle the problem. The game stops as soon as at least one of the players knows the answer with certainty. The (*deterministic*) *communication complexity* $cc(P)$ is the minimum number of bits needed in the worst case.

Associating with P its $(n\times n)$ (0, 1)-incidence matrix $M=M(P)$, the logarithms (base 2) of several combinatorial parameters of M have been discovered to yield lower bounds on the communication complexity $cc(M)=cc(P)$: $k_1=k_1(M)$,

$k_0 = k_0(M)$, $\bar{k}_1 = \bar{k}_1(M)$, $\bar{k}_0 = \bar{k}_0(M)$ (Yao 1979) and $r = r(M)$ (Mehlhorn and Schmidt 1982), where

$$r(M) = \text{rank of } M$$
$$k_1(M) = \text{minimal number of submatrices of } M \text{ with only 1-entries covering all 1's of } M$$
$$\bar{k}_1(M) = \text{minimal number of disjoint submatrices of } M \text{ with only 1-entries covering all 1's of } M.$$

($k_0(M)$ and $\bar{k}_0(M)$ are similarly defined with respect to the 0-entries of M.)

An obvious upper bound on the communication complexity is achieved by the *trivial protocol:* Player II, say, communicates the "name" of (an isomorphic copy of) his chosen column. Denoting by n^* the number of distinct columns of the matrix M, we hence obtain

$$cc(M) \leq \lceil \log n^* \rceil.$$

Yao (1979) has shown that the trivial protocol is (up to 2 bits) optimal for random $(n \times n)$ communication matrices with probability at least

$$1 - 2^{-n^2/2}.$$

The best non-trivial general upper bound currently known is due to Lovász and Saks (see Lovász 1989). We will describe it in terms of the order $P = P(M)$. A *linear extension* of P is a linear arrangement

$$L = x_1 x_2 \dots x_n$$

of the elements x_i of the ground set underlying P such that $x_i < x_j$ in P implies $i < j$ in L. The *lineality* of L is the number

$$\mathrm{l}(L) = |\{i : x_i < x_{i+1} \quad \text{in} \quad P\}|$$

and the lineality of P is defined as

$$\mathrm{l}(P) = \max\{\mathrm{l}(L) : L \text{ linear extension of } P\}.$$

The Lovász-Saks bound now can be given as

$$cc(M) \leqq (1 + \lceil \log \mathrm{l}(P(M)) \rceil) \cdot (1 + \lceil \log k_0(M) \rceil).$$

Noting the relation

$$\mathrm{l}(P(M)) \leqq r(M) \leqq \bar{k}_1(M),$$

the Lovász-Saks bound implies, for example, other upper bounds due to Yannakakis (1989):

$$cc(M) = O(\log^2 \bar{k}_1(M))$$
$$cc(M) = O(\log^2 \bar{k}_0(M)),$$

where the second bound is obtained "by duality", i.e., by considering the complementary communication matrix $\bar{M}$ and noting $cc(M) = cc(\bar{M})$.

3. General order dimension bounds

We consider first an arbitrary communication matrix M and derive a lower bound for its communication complexity in terms of the dimension of the associated order $P(M)$.

Recall that any order P can be viewed as the intersection of all its linear extensions in the following sense: $x<y$ holds in P if and only if x occurs before y in every linear extension of P. Hence we define a *realizer* of P to be a collection $\mathscr{L}=\{L_1, \dots, L_k\}$ of linear extensions L_i of P such that for all $x, y \in E$,

$$x<y \text{ in } P \quad \text{if and only if} \quad x \text{ occurs before } y \text{ in every } L_i \in \mathscr{L}.$$

The *order dimension* $\dim P$ of P is the smallest number d such that P admits a realizer of size d.

The subset $A \subseteq E$ in an (*order*) *ideal* of P if for all $a \in A$, $x \in E$,

$$x<a \quad \text{in } P \quad \text{implies } x \in A.$$

Lemma 1. *Let A and B be two ideals of P with the induced order relation so that $A \cup B = E$. Then*

$$\dim A + \dim B \geqq \dim P.$$

Proof. We take minimal realizers $\mathscr{L}_A=(L_1, \dots, L_k\}$ for A and $\mathscr{L}_B=\{O_1, \dots, O_m\}$ for B and denote by $\bar{L}_1$ the restriction of L_1 to $A \setminus B$ and by $\bar{O}_1$ the restriction of O_1 to $B \setminus A$. Construct now linear extensions for P via the following concatenations:

$$L_i \oplus \bar{O}_1 \qquad (i=1, \dots, k)$$
$$O_j \oplus \bar{L}_1 \qquad (j=1, \dots, m).$$

Apparently, these $k+m$ linear extensions give rise to a realizer for P. □

Proposition 2. *Let M be a non-trivial communication matrix with associated order $P(M)$. Then*

$$\lceil \log \dim P(M) \rceil \leqq cc(M).$$

Proof. Consider an optimal protocol for the communication problem relative to M and assume w.l.o.g. that player II begins with sending a message to player I. This protocol partitions M into submatrices $\bar{M}_0$ and $\bar{M}_1$ as follows: Let the matrix M_i $(i=0, 1)$ consist of those columns y of M such that player II starts his message with the symbol "i" if he has chosen column y. Assume $\dim P(M_0) \geqq \dim P(M_1)$.

Note that no message is sent if player II picks a column consisting only of 1's or only of 0's. Add those columns to M_0 in order to obtain $\bar{M}_0$ and take $\bar{M}_1 = M_1$. Omitting the first symbol, it is clear that our optimal protocol for M turns into a feasible (though not necessarily optimal) protocol for both $\bar{M}_0$ and $\bar{M}_1$.

Now $P(\bar{M}_i)$ $(i=0, 1)$ is an order ideal of $P(M)$ and, because of $\dim P(\bar{M}_0) \geqq \dim P(\bar{M}_1)$, we conlcude from Lemma 1 that

$$\dim P(\bar{M}_0) \geqq (\dim P(M))/2.$$

By definition, we have $M_0 \neq \phi \neq M_1$ (otherwise the first bit in the communication protocol would be superfluous and hence the protocol not optimal). Moreover, $\bar{M}_0$ must be non-trivial, i.e., $cc(\bar{M}_0) \geqq 1$, because otherwise also M_0 would be trivial and our protocol for M not optimal. Induction on the size of M therefore yields the reduced protocol to require at least

$$\log(\dim P(M)/2 = \log \dim P(M) - 1$$

bits for $\bar{M}_0$, from which the Proposition follows. □

We next establish an upper bound.

Theorem 4. *Let $P = (E, <)$ be a strict order relation with incidence matrix M. Then*

$$\lceil \log r(M) \rceil \leqq cc(M) \leqq d \lceil \log r(M) \rceil,$$

where $d = \dim P$.

Proof. The lower bound for $cc(M)$ is the Mehlhorn-Schmidt bound (see Sect. 2). For the upper bound, it suffices to show that M has at most $r(M)^d$ distinct columns and then to invoke the trivial protocol.

Suppose to the contrary that the number of distinct columns of M is strictly larger than

$$r^d \geqq 1 + r + \binom{r}{2} + \ldots + \binom{r}{d} \quad (d \geqq 2).$$

Then, by Sauer's Lemma (see Lovász (1979), Problem 13.10c), M contains some $((d \times 2^d))$-submatrix M_d consisting of the incidence vectors of all subsets of a d-element set.

Let $A = \{a_1, a_2, \ldots, a_d\}$ be those elements of P corresponding to the rows of M_d. By $A' = \{a'_1, \ldots, a'_d\}$ we denote the elements of P that correspond to those columns of M_d having exactly $d-1$. Clearly, both A and A' consist of pairwise incomparable elements of P. It follows that P contains a suborder C_d isomorphic to

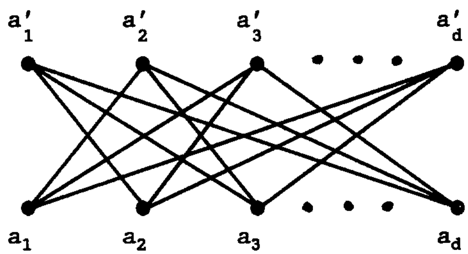

(C_d arises from the complete bipartite graph $K_{d,d}$ by removing a perfect matching.)

Already Dushnik and Miller (1941) discovered that $\dim C_d = d$. Hence we must also have $\dim P \geqq d$. □

We remark that Theorem 4 can be strengthened as follows. Say that the order $Q = (E, <)$ is an *interval order* if Q does not contain an induced suborder isomorphic to

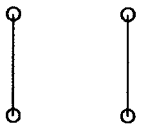

Now the *interval dimension* $i\_\dim P$ of P can be defined as the minimal number d such that P is the intersection of interval orders $Q_1, Q_2, \ldots, Q_d$. Each linear order is an interval order but there are interval orders with arbitrarily large order dimension. Hence $i\_\dim P$ might be substantially smaller than $\dim P$.

It is not hard to see that any interval order Q that contains C_d must include at least $d-1$ of the pairs (a_i, a_i'). Hence at least d interval orders are needed for a realizer of C_d by interval orders. Thus

$$i\_\dim C_d = \dim C_d = d.$$

Consequently, one may replace $\dim P$ by $i\_\dim P$ in the statement of Theorem 4 without changing the proof.

4. Chains vs. antichains

Consider again the order $P=(E, <)$ with $|E|=n$. Recall that a *chain* of P is a subset $C \subseteq E$ whose elements are pairwise comparable while an *antichain* of P is a subset $A \subseteq E$ whose elements are pairwise incomparable. An important parameter associated with P is the *width*

$$w(P)=\max\{|A|: A \text{ antichain}\}.$$

By Dilworth's Theorem (cf. Dilworth 1950), $w(P)$ is the minimal number of chains needed to cover E. Moreover, Hiraguchi's Theorem (cf. West 1985) exhibits $w(P)$ as an upper bound for the order dimension $\dim P$.

Let M be the incidence matrix of the chains vs. the antichains of P. That is, we investigate the communication problem where player I chooses a chain C and player II an antichain A. The objective is to decide whether $C \cap A \neq \phi$.

It is easy to see that $r(M)=n$. Indeed, $r(M) \geq n$ follows from the observation that M contains an $(|E| \times |E|)$-identity submatrix. $r(M) \leq n$ is implied by the factorization

$$M = U \cdot V,$$

where U is the chains vs. elements and V the elements vs. antichains incidence matrix.

The same argument also yields $\bar{k}_1(M)=n$. Yannakakis (1989) furthermore establishes $\bar{k}_0(M)=O(n^2)$. The Yannakakis upper bound (see Sect. 2), therefore, yields

$$cc(M)=O(\log^2 n)$$

for the communication complexity. We will now derive two upper bounds which, in the special case under consideration in this section, improve the Lovász-Saks bound.

Proposition 5. *Let $h = h(P)$ be the size of the longest chain in P. Then*

$$cc(M) \leqq (1 + \lceil \log h \rceil)(1 + \lceil \log n \rceil).$$

Proof. The bound is implied by the following recursive protocol.

Player I chooses an element a in his chain C such that the cardinality of either of the two chains

$$C_a = \{x \in C : x > a\}$$
$$C^a = \{x \in C : x < a\}$$

is at most $|C|/2$. He then sends the "name" of a to player II with $\lceil \log n \rceil$ bits.

If $a \in A$, the game stops since player II knows the answer with certainty.

If the antichain A of player II contains an element a' with $a' > a$, player II sends "1" and player I resumes the game with the chain C_a instead of C. Similarly, if there is an $a'' \in A$ with $a'' < a$, player II sends "0" and player I resumes the game with the chain C^a.

If no such a' nor a'' exists in A, then $A \cap C = \phi$ must hold and the game stops.

Note that this protocol is feasible because the antichain A cannot contain elements $a', a'' \in A$ such that $a' > a > a''$. □

Since our communication matrix M contains an $(n \times n)$-identity submatrix the order $P(M)$ has lineality $1(P(M)) \geqq n$ (in fact, $1(P(M)) = n$ because $1(P(M)) \leqq r(M)$). To obtain an estimate for $k_0(M)$, take a chain C with $|C| = h$ and let $\mathscr{C}$ be the system of the h subchains of C that can be constructed by removing an element from C. The incidence matrix M_c of $\mathscr{C}$ vs. the singletons of C is an $(h \times h)$-submatrix of M with entries "0" on the diagonal and entries "1" off the diagonal. Thus $k_0(M) \geqq k_0(M_c) \geqq h$. It follows that

$$(1 + \lceil \log h \rceil) \cdot (1 + \lceil \log n \rceil) \leqq (1 + \lceil \log 1(P(M)) \rceil) \cdot (1 + \lceil \log k_0(M) \rceil).$$

Theorem 6. *Let $d = \dim P$ be the order dimension of P. Then*

$$cc(M) \leqq (2 + \lceil \log d \rceil) \cdot \lceil \log n \rceil.$$

Proof. For simplicity of notation, let us assume that n is a power of 2. We will first describe a standard representation of P as a set of integral vectors in $\mathbb{R}^d$, endowed with the componentwise ordering.

Choose a realizer $\mathscr{L} = \{L_1, L_2, \ldots, L_d\}$ of P and associate with each element $e \in E$ the vector $(e_1, e_2, \ldots, e_d) \in \mathbb{R}^d$ in such a way that e_i is the position at which e occurs in the linear extension L_i $(i = 1, \ldots, d)$. By definition, we have for all $a, b \in E$,

$$a < b \text{ in } P \quad \text{if and only if} \quad a_i < b_i \quad (i = 1, \ldots, d).$$

We now identify P with its representation in $\mathbb{R}^d$ and describe a recursive protocol that yields the claimed upper bound on $cc(M)$.

For $i=1, \ldots, d$, consider the open half-spaces

$$H_i^+ = \left\{x \in \mathbb{R}^d : x_i > \frac{n+1}{2}\right\}$$

$$H_i^- = \left\{x \in \mathbb{R}^d : x_i < \frac{n+1}{2}\right\}$$

and the hyperplane

$$H_i = \left\{x \in \mathbb{R}^d : x_i = \frac{n+1}{2}\right\}.$$

Note that each such half-space contains exactly $n/2$ points of P. Moreover, an antichain A of P cannot have members in both of the following sets simultaneously:

$$H^+ = \left\{x \in \mathbb{R}^d : x_i > \frac{n+1}{2} \quad (i=1, \ldots, d)\right\}$$

$$H^- = \left\{x \in \mathbb{R}^d : x_i < \frac{n+1}{2} \quad (i=1, \ldots, d)\right\}.$$

Player I, who has selected the chain C, first checks whether his chain is completely contained in one of the $2d$ half-spaces defined above. If so, he sends the symbol "1" followed by the name of that half-space with $1+\lceil \log 2d \rceil$ bits. The game is thus reduced to an ordered set with $n/2$ elements and dimension at most d.

In case player I cannot find such a half-space, he connects the elements $c^1 < c^2 < \ldots < c^m$ of his chain by a polygonal line T. Obviously, T intersects each hyperplane H_i exactly once. Suppose that H_j is the first and H_k the last hyperplane to be intersected (ties may be broken arbitrarily). Player I now transmits the symbol "0".

When player II receives the latter message, the chain C is contained in

$$H^- \cup (H_j^+ \cap H_k^-) \cup H^+.$$

If the antichain A satisfies $A \cap H^- = \phi$, player II sends "1" – thus reducing the game to the $n/2$ points of P in H_j^+. Otherwise he sends "0" and the game may resume with respect to the suborder of P contained in H_k^-. Hence for a complete update of player II, player I now only needs to send the name of the index j or k respectively with $\lceil \log d \rceil$ bits and the next round can begin with $n/2$ points. □

To compare the bound of Theorem 6 with the general Lovász-Saks bound, recall that $n = l(P(M))$ and look at $k_0(M)$. If A is an antichain of size $w = w(P) = |A|$, we let $\mathscr{A}$ be the collection of all the w antichains obtained by removing a singleton from A. M contains as a submatrix the incidence matrix M_A of those singletons vs. $\mathscr{A}$. M_A is a $(w \times w)$-maxtrix with 0's on the diagonal and 1's off the diagonal. Thus, using Hiragushi's Theorem, we get

$$k_0(M) \geqq k_0(M_A) = w(P) \geqq \dim P.$$

The bounds of Proposition 5 and Theorem 6 cannot be compared because there is no monotone relationship between the parameters $h(P)$ and $\dim P$. Can both bounds be improved such that

$$cc(M) = O(\log n)?$$

Note that such an improvement is possible if one restricts one's attention to interval orders, for example.

Lemma 7. *If P is an interval order, then*

$$cc(M) \leqq 2 \cdot \lceil \log n \rceil.$$

Proof. It is well-known that an interval order P has at most n maximal antichains. This suggests the following protocol.

Player II augments his selected antichain A to a maximal antichain A' and sends the name of A' to player I with $\lceil \log n \rceil$ bits.

Player I then compares A' with his chain C. If $C \cap A' = \phi$, of course $C \cap A = \phi$. If $C \cap A' \neq \phi$, player I transmits the name of the unique element $a \in C \cap A'$ to player II. Player II can now verify whether $a \in A$. □

Lemma 7 raises the more modest question whether the order dimension can be replaced by the interval dimension in the bound of Theorem 6.

References

Aho A.V., Ullman J.D., Yannakakis, M.: On notions of information transfer in VLSI circuits. Proc. 15th ACM STOC, pp. 133–139 (1983)

Dilworth, R.P.: A decomposition theorem for partially ordered sets. Am. Ann. Math. **51**, 161–166 (1950)

Dushnik, B., Miller, E.: Partially ordered sets. Am. J. Math. **63**, 600–610 (1941)

Faigle, U., Turán, Gy.: Communication complexity. Computing (Suppl) **7**, 141–153 (1990)

Lovász, L.: Combinatorial problems and exercises. Amsterdam: North-Holland 1979

Lovász, L.: Communication complexity: A survey. Report No. 89555-OR, Institut für Operations Research, Universität Bonn (1989)

Mehlhorn, K., Schmidt, E.M.: Las Vegas is better than determinism in VLSI and distributed computing. Proc. 14th AMC STOC, pp. 330–337 (1982)

West, D.B.: Parameters of partial orders and graphs: packing, covering, and representation. In: Rival I (ed) Graphs and order, pp. 267–350. Dordrecht: D. Reidel 1985

Yannakakis, M.: Expressing combinatorial optimization problems by linear programs. Preprint, AT&T Bell Laboratories, Murray Hill (1989)

Yao, A.C.-C.: Some complexity questions related to distributive computing. Proc. 11th ACM STOC, pp. 209–213 (1979)