JOURNAL OF COMBINATORIAL THEORY, Series A 54, 312–316 (1990)

Note

How to Communicate Efficiently

ALBRECHT BEUTELSPACHER

Mathematisches Institut, Justus-Liebig-Universität Giessen, Arndtstr. 2, D-6300 Giessen, West Germany

Communicated by A. Barlotti

Received February 3, 1989

We present an algorithm which makes it possible to exchange data in a network efficiently and quickly among all participants. This algorithm is based on finite geometries, in particular on affine spaces. © 1990 Academic Press, Inc.

1. INTRODUCTION

Consider a distributed system consisting of a set of N autonomous computers interconnected by a communication network. We suppose that any single transmitted message of one computer to another adds to the costs of the global communication. We shall simply speak of a network with N *knots*.

The problem we want to consider is known as "decentralized consensus problem." There is some information gathered from all knots, this information is synthesized, and the outcome must be made known to all knots. An example is that the knots are machines and status information of the machines must be known to every knot. We want to develop an algorithm which guarantees that after some time every knot of the net knows the critical value. This means that "in a way" any value must reach every participant. Such a procedure has been called [1] "broadcasting without broadcast."

Clearly, such algorithms exist. Here is one: Every knot sends his data to every other knot. Then every knot can evaluate the received data and knows the outcome. Since there are N knots, one needs a total of $N(N-1)/2$ transmit operations. Hence the complexity of this algorithm is $O(N^2)$.

The question is: What are the most efficient algorithms? Given N, we

0097-3165/90 $3.00

denote by t_N the number of transmits of the best algorithm which satisfies the rules. Hence the above algorithm implies $t_N \leqq c \cdot N^2$.

Using geometric methods (in particular projective planes) it was proved in [5] that $t_N \leqq c \cdot N\sqrt{N}$.

In a remarkable paper [1], Alon, Barak, and Manber proved (in a rather tricky way) that $t_N \leqq c \cdot N \cdot \log_2 N$.

The aim of this note is to prove the same result in a very easy way using finite geometry. Our main result reads as follows.

THEOREM. *Fix a prime power q. Let d be the smallest positive integer with $N \leqq q^d$. Then $t_N \leqq c \cdot N \cdot d$, where c is a constant depending on q. Hence $t_N \leqq c \cdot N \cdot \log_q N$.*

As the algorithms in [1] and [5], our algorithm is "completely symmetric," that is, every knot has the same duties.

2. GEOMETRIC LANGUAGE

Let $\mathbf{S} = (P, B, I)$ be an *incidence structure* consisting of *points*, *blocks*, and *incidences*. (Later on, the blocks will be subspaces of a classical geometry.) Let $\mathscr{F}_1, \mathscr{F}_2, \ldots, \mathscr{F}_n$ be families of blocks with the property that any point of $\mathbf{S}$ is on at least one block of $\mathscr{F}_i$ $(1 \leqq i \leqq n)$. (Later we shall use "parallel classes" as the $\mathscr{F}_i$'s.)

For an arbitrary set X of points we define $X^{\mathscr{F}_i}$ to be the set of those points which are connected to a point of X by an element of $\mathscr{F}_i$. Moreover,

$X^1 := X$,

$X^i :=$ the set of points P such that P is connected to a point of X^{i-1} by a line of $\mathscr{F}_{i-1}$ $(2 \leqq i \leqq n+1)$. In other words, $X^i = (X^{i-1})^{\mathscr{F}_{i-1}}$.

For a point P we set $P^i = \{P\}^i$.

EXAMPLE. If $\mathscr{F}_1$ is a parallel class, then P^2 is the set of points on the line of F_1 through P.

We are interested in the case where $P^1, P^2, P^3, \ldots$ eventually becomes the whole point set. We remark that this is not automatically the case. If, for example, $\mathscr{F}_1 = \mathscr{F}_2 = \cdots = \mathscr{F}_n$, then $X^3 = X^2$; so it is very unlikely that one reaches the whole point set.

We may now define the following *algorithm*. Identify the knots of the network with a subset of the points of $\mathbf{S}$. The idea is that in the ith step, knot P gives its information to all knots to which it is connected by an element of $\mathscr{F}_{i-1}$.

More formally: The algorithm runs in n steps, where the ith step is the following.

Any point P sends its data to all points of $P^{\mathscr{F}_{i-1}}$. It evaluates the received values and stores the outcome of the evaluation. (This is the datum to be used in Step $i+1$.)

THEOREM 1. *Suppose that there exists a positive integer n such that for any point P, P^{n+1} is the whole point set.*

(a) *After n steps every point knows the critical value. So, the algorithm does what is required.*

(b) *If any $F \in \mathscr{F}_1 \cup \cdots \cup \mathscr{F}_n$ has cardinality at most c_0, then there is a constant $c = c(c_0)$ such that*

$$t_N \leqq c \cdot n \cdot N.$$

Proof. (a) is obvious from the above discussion.

(b) Since the cardinalities of the F's are bounded, the complexity of transmitting in F and evaluating the data from F is bounded. ∎

3. GEOMETRIES

Let $\mathbf{A} = AG(d, q)$ be the affine space of dimension d of order q. We denote the hyperplane of infinity by H_∞. Let $P_1, \ldots, P_d$ be d points of H_∞ in general position. Then these points generate H_∞. Let $\Pi_1, \ldots, \Pi_d$ be the parallel classes (of lines) in $\mathbf{A}$ such that the lines of Π_i intersect in P_i $(i = 1, \ldots, d)$.

In our above terminology, $\mathbf{S}$ is the incidence structure of the points and lines of $\mathbf{A}$, and $\mathscr{F}_i := \Pi_i$.

THEOREM 2. *Given a prime power q. Let d be the smallest positive integer such that $N \leqq q^d$. Then*

$$t_N \leqq c \cdot d \cdot N.$$

Proof. In view of Theorem 1 we have only to show that for any point P, P^{n+1} is the whole point set. In order to show this, we prove

For any point P, P^i is a subspace of dimension $i-1$.

Namely: If $i = 1$, then $P^1 = P$ is a subspace of dimension 0.

Suppose now $i \geqq 1$ and assume that the assertion is true for i. Hence P^i is a subspace of dimension $i-1$ of $\mathbf{A}$. We consider this subspace also as a subspace of the projective space $\mathbf{P}$ associated with $\mathbf{A}$. Then P^i intersects

H_∞ in the subspace $\langle P_1, ..., P_{i-1}\rangle$ generated by $P_1, ..., P_{i-1}$. By the definition of P^{i+1} it follows that this can be described as a projective subspace,

$$P^{i+1} = \langle P^i, P_i\rangle.$$

(Note that $P_i \notin \langle P_1, ..., P_{i-1}\rangle$, since $P_1, ..., P_{i-1}, P_i$ are in general position.) Hence P^{i+1} is an i-dimensional subspace of **P** which intersects H_∞ in the $(i-1)$-dimensional subspace $\langle P_1, ..., P_{i-1}, P_i\rangle$. Hence P^{i+1} is also an i-dimensional subspace of **A**.

In particular, P^{d+1} is a d-dimensional subspace of **A**. Hence $P^{d+1} = \mathbf{A}$. ∎

Generalizing the above approach, one can do better. In order to formulate the corresponding statement, we need the following definitions and results.

Let **P** be a projective space of order q and dimension d. A t-spread of **P** is a set $\mathcal{S}$ of mutually skew t-dimensional subspaces of **P** with the property that any point of **P** is on a (unique) element of $\mathcal{S}$. It is well known (see for instance [3, 4]) that there is a t-spread in **P** if and only if $t+1$ divides $d+1$. A t-spread $\mathcal{S}$ is called *geometric* [2, 6] if for any two distinct elements $U, V \in \mathcal{S}$ the following assertion is true: Any element $X \in \mathcal{S}$ which has at least one point in common with $\langle U, V\rangle$ is totally contained in $\langle U, V\rangle$. The following assertions are (not so) well-known facts about geometric spreads.

1. A geometric t-spread exists in **P** if and only if $t+1$ divides $d+1$ [2, 6].

2. Let $\mathcal{S}$ be a geometric t-spread in **P**. Then the incidence structure whose points are the elements of $\mathcal{S}$ and whose lines are the subspaces of the form $\langle U, V\rangle$, where U and V are distinct elements of $\mathcal{S}$, is a (Desarguesian) projective space of order q^{t+1} and dimension $(d+1)/(t+1)-1$ [2, 6].

LEMMA. *Let H_∞ be a $(d-1)$-dimensional projective space of order q. Let t be a positive integer with $t < d-1$. Define $s = \lfloor (d+1)/t \rfloor$ to be the greatest integer not greater than $(d+1)/t$. Then there are $s+1$ subspaces of dimension $t-1$ of H_∞ which generate H_∞.*

Proof. By definition of s, there is an integer b with

$$d = s \cdot t + b \qquad \text{and} \qquad -1 \leqq b \leqslant t-2.$$

Consider a subspace K of H_∞ of dimension $st-1$. Then K has a geometric $(t-1)$-spread, whose elements are the points of a projective space of dimension $s-1$ and order q^t. Choosing a basis in this projective space we get s subspaces $U_1, U_2, ..., U_s$ of dimension $t-1$ generating K.

Since $\dim H_\infty - \dim K = b \leqq t-2$, we may choose U_{s+1} in such a way that

$$H = \langle K, U_{s+1} \rangle. \quad \blacksquare$$

THEOREM 3. *Fix a prime power q and a positive integer t. Let d be the smallest integer satisfying $N \leqq q^d$. Then*

$$t_N \leqq c \cdot N \cdot (\lfloor (d+1)/t \rfloor + 1).$$

Proof. Let $\mathbf{A} = AG(d, q)$, and denote by H_∞ the hyperplane at infinity. By the above lemma there exist $s+1$ subspaces $U_1, \ldots, U_{s+1}$ of dimension $t-1$ in H_∞ which span H_∞. Define $\mathscr{F}_i$ to be the set of t-dimensional subspaces of $\mathbf{A}$ which intersect H_∞ in U_i. As in the proof of Theorem 2 it follows that these sets satisfy the hypotheses of Theorem 1. Hence

$$t_N \leqq c \cdot N \cdot (s+1) = c \cdot N \cdot (\lfloor (d+1)/t \rfloor + 1). \quad \blacksquare$$

Remarks. 1. From the proof of the above lemma one sees immediately that also the following fact is true.

THEOREM 4. *If, in addition to the hypotheses of Theorem 4, one has that t divides $d+1$, then*

$$t_N \leqq c \cdot N \cdot (d+1)/t.$$

2. Note that the constant c counts the number of transmits which a knot in an $F \in \mathscr{F}_1 \cup \cdots \cup \mathscr{F}_n$ has to perform in order to send its data to any other knot in F. So, in the case of Theorem 2, we have $c \leqq q-1$, while in general (Theorems 3 and 4) it is $c \leqq q^t - 1$.

REFERENCES

1. N. ALON, A. BARAK, AND U. MANBER, On disseminating information reliably without broadcasting, *in* "IEEE, The 7th International Conference on Distributed Computing Systems," Berlin, September 21–25, 1987.
2. R. BAER, Partitionen abelscher Gruppen, *Arch. Math.* **14** (1963), 73–83.
3. A. BEUTELSPACHER, Einführung in die endliche Geometrie. II. Projektive Räume. B.I.-Wissenschaftsverlag, Mannheim/Wien/Zürich, 1983.
4. P. DEMBOWSKI, "Finite Geometries," Springer-Verlag, Berlin/Heidelberg/New York, 1968.
5. T. V. LAKSHMAN AND A. K. AGRAWALA, Efficient decentralized consensus protocols, *IEEE Trans. Software Engirg.* **12** (1985), 600–607.
6. B. SEGRE, Teoria di Galois, fibrazioni proiettive e geometrie non desarguesiane, *Ann. Mat. Pura Appl.* **64** (1964), 1–76.