COMMUNICATION COMPLEXITY IN LATTICES

Rudolf Ahlswede, Ning Cai, and Ulrich Tamm

Department of Mathematics
University of Bielefeld
P.O. Box 100131
W-4800 Bielefeld
Germany

1. INTRODUCTION

Let $\mathcal{X}$ denote a finite lattice and let $\hat{f} : \mathcal{X} \to \mathcal{Z}$ be a function mapping $\mathcal{X}$ into some set $\mathcal{Z}$. In this note we determine the communication complexity of functions $f : \mathcal{X} \times \mathcal{X} \to \mathcal{Z}$ defined by

$$f(x, y) := \hat{f}(x \wedge y) \text{ for all } x, y \in \mathcal{X}. \tag{1}$$

The communication complexity of a function $f : \mathcal{X} \times \mathcal{Y} \to \mathcal{Z}$ (where $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{Z}$ are finite sets), denoted as $C(f)$, is the number of bits that two processors, P_1 and P_2 say, have to exchange in order to compute the function value $f(x, y)$, when initially P_1 only knows $x \in \mathcal{X}$ and P_2 only knows $y \in \mathcal{Y}$.

More specifically, let $\mathcal{Q}$ denote the set of protocols computing f such that finally both processors know the result and let $l_P(x, y)$ be the number of bits transmitted for the input (x, y), when the protocol $P \in \mathcal{Q}$ is used. Then the (worst-case) *communication complexity* is

$$C(f) := \min_{P \in \mathcal{Q}} \max_{(x,y) \in \mathcal{X} \times \mathcal{Y}} l_P(x, y). \tag{2}$$

A *protocol* P is a pair of mappings $\phi_1 : \mathcal{X} \times \{0,1\}^* \to \{0,1\}^*$, $\phi_2 : \mathcal{Y} \times \{0,1\}^* \to \{0,1\}^*$. So on input (x, y) the processors, starting with P_1, say, alternatively send binary messages N_1, N_2, N_3, etc., until they both know the result. Each message depends on the previous messages and on the current processor's input, hence $N_1 = \phi_1(x)$, $N_2 = \phi_2(y, \phi_1(x))$, $N_3 = \phi_1(x, \phi_1(x)\phi_2(y, \phi_1(x)))$, etc. . It is required that the set of messages a processor is allowed to send at an arbitrary moment in the course of the protocol is *prefix-free*, i. e., no possible message is the beginning (prefix) of another one. This property assures that the other processor immediately recognizes the end of the message and can hence start the transmission of its next message without delay.

An upper bound on $C(f)$ for any function $f : \mathcal{X} \times \mathcal{Y} \to \mathcal{Z}$ (w. l. o. g. $|\mathcal{X}| \leq |\mathcal{Y}|$) is always obtained from the following *trivial protocol*: P_1 transmits all the bits of its input $x \in \mathcal{X}$. P_2 now is able to compute the function value and returns the result $f(x, y) \in \mathcal{Z}$. Hence

$$C(f) \leq \lceil log|\mathcal{X}|\rceil + \lceil log|\mathcal{Z}|\rceil. \tag{3}$$

Throughout this paper the logarithm is always taken to the base 2.

The following lower bound is due to Mehlhorn and Schmidt [1]:

$$C(f) \geq \lceil log(\sum_{z \in \mathcal{Z}} \text{rank } M_z(f))\rceil, \tag{4}$$

where $M_z(f) := (m_{xy})_{x,y \in \mathcal{X}}$ is a Boolean matrix with $m_{xy} = 1$ exactly if $f(x, y) = z$.

2. THE MAIN RESULTS

In the following, we denote by $\preceq$ the underlying order of the lattice $\mathcal{X}$ and by μ the associated *Möbius function.* Further let

$$\mathcal{X}_z := \{x \in \mathcal{X} : \text{ there is some } \hat{x} \preceq x \text{ with } \hat{f}(\hat{x}) = z\}. \tag{5}$$

MAIN THEOREM: The communication complexity of the function f defined as in (1) is bounded from above and below as follows:

$$\lceil log(\sum_{z\in\mathcal{Z}} |\{x\in\mathcal{X}: \sum_{\hat{x}\preceq x, \hat{f}(\hat{x})=z} \mu(\hat{x},x)\neq 0\}|)\rceil \leq C(f) \leq \lceil log(\sum_{z\in\mathcal{Z}} |\mathcal{X}_z|)\rceil + 1. \tag{6}$$

If, additionally, $\sum\limits_{\hat{x}\preceq x, \hat{f}(\hat{x})=z} \mu(\hat{x},x)\neq 0$ for all possible $x\in\mathcal{X}$ and $z\in\mathcal{Z}$, i.e., for all x and z for which there exists some $\hat{x}\preceq x$ with $\hat{f}(\hat{x})=z$, then upper and lower bound differ by at most one bit, namely

$$\lceil log(\sum_{z\in\mathcal{Z}} |\mathcal{X}_z|)\rceil \leq C(f) \leq \lceil log(\sum_{z\in\mathcal{Z}} |\mathcal{X}_z|)\rceil + 1. \tag{7}$$

The lower bounds are based on the following theorem, which was discovered by Wilf [2] (see also Lindström [3]) and first used in the study of communication complexity by Lovasz [4]. We shall present Wilf's short proof from which the succeeding corollary is immediate, since the incidence matrix of a poset is nonsingular.

THEOREM (WILF): Let $\mathcal{X}$ be a finite lattice with order $\preceq$ and Möbius function μ. Further, let $\{a_x : x\in\mathcal{X}\}$ be a set of arbitrary numbers. Then

$$\det\ (a_{x\wedge y})_{x,y\in\mathcal{X}} = \det\ \mathrm{diag}\ (\sum_{\hat{x}\preceq x}\mu(\hat{x},x)\cdot a_{\hat{x}})_{x\in\mathcal{X}} = \prod_{x\in\mathcal{X}}(\sum_{\hat{x}\preceq x}\mu(\hat{x},x)\cdot a_{\hat{x}}).$$

PROOF: For arbitrary numbers $\{b_x : x\in\mathcal{X}\}$ consider the matrix $\zeta^T\cdot \mathrm{diag}(b_x)_{x\in\mathcal{X}}\cdot\zeta$, where $\zeta := (\zeta_{\hat{x}x})_{\hat{x},x\in\mathcal{X}}$, with $\zeta_{\hat{x}x}=1$ exactly if $\hat{x}\preceq x$ ($\zeta_{\hat{x}x}=0$ else) is the incidence matrix of $(\mathcal{X},\preceq)$. By the rules for matrix multiplication this is just the matrix $(\sum\limits_{\hat{x}\preceq x\wedge y} b_{\hat{x}})_{x,y\in\mathcal{X}}$.

Now let $a_x := \sum\limits_{\hat{x}\preceq x} b_{\hat{x}}$ for all $x\in\mathcal{X}$. By the Möbius inversion formula then $b_x = \sum\limits_{\hat{x}\preceq x}\mu(\hat{x},x)\cdot a_{\hat{x}}$ for all $x\in\mathcal{X}$ and the theorem follows.

COROLLARY: Let $\mathcal{X}$ and $\{a_x : x\in\mathcal{X}\}$ be as in the preceding theorem. Then

$$\mathrm{rank}\ (a_{x\wedge y})_{x,y\in\mathcal{X}} = \mathrm{rank}\ \mathrm{diag}\ (\sum_{\hat{x}\preceq x}\mu(\hat{x},x)\cdot a_{\hat{x}})_{x\in\mathcal{X}}. \tag{8}$$

PROOF OF THE LOWER BOUNDS IN THE MAIN THEOREM:
Observe that the function value matrices $M_z(f)$ are just of the form $(a_{x\wedge y})_{x,y\in\mathcal{X}}$ with

$$a_{x\wedge y} = 1 \text{ exactly if } \hat{f}(x\wedge y) = z.$$

With the above corollary for all $z\in\mathcal{Z}$ it is

$$\mathrm{rank} M_z(f) = |\{x\in\mathcal{X}: \sum_{\hat{x}\preceq x, \hat{f}(\hat{x})=z}\mu(\hat{x},x)\neq 0\}|\ .$$

The lower bound in (6) follows by application of the Mehlhorn - Schmidt lower bound (4).

PROOF OF THE UPPER BOUND IN THE MAIN THEOREM:
The upper bound in the Main Theorem is obtained via a natural and useful improvement of the trivial protocol, which was first introduced by Ahlswede and Cai [5]. As the trivial protocol, it consists of two rounds. In the first round the processor P_1 encodes its input $x\in\mathcal{X}$. The processor P_2 then knows both values x and y and hence is able to compute the result $f(x,y)$, which is returned to P_1. However, now the set of possible function values is reduced to

$$\hat{F}(x) := \{\hat{f}(\hat{x}) : \hat{x}\preceq x\}, \tag{9}$$

since the second processor already knows $x\in\mathcal{X}$.
Hence, only $\lceil log|\hat{F}(x)|\rceil$ bits have to be reserved for the transmission of the result $f(x,y)$ such that the first processor can assign longer messages (code words) to elements with few predecessors