Electronic Colloquium on Computational Complexity
ECCC Technical Reports Series 1994 Report Nr: TR94–022

The Möbius Function, Variations Ranks, and $\Theta(n)$–Bounds on the Modular Communication Complexity of the Undirected Graph Connectivity Problem

Christoph Meinel
Lehrstuhl Theoretische Informatik
Fachbereich IV - Informatik
Universität Trier
D–54286 Trier

Stephan Waack
Institut für Numerische
und Angewandte Mathematik
Georg–August–Universität Göttingen
Lotzestr. 16–18
D–37083 Göttingen

Received *December 22, 1994*

Abstract. We prove that the modular communication complexity of the undirected graph connectivity problem UCONN equals $\Theta(n)$, in contrast to the well–known $\Theta(n \log n)$ bound in the deterministic case, and to the $\Omega(n \log\log n)$ lower bound in the nondeterministic case, recently proved by Raz and Spieker.
We obtain our result by combining Möbius function techniques due to Lovasz and Saxe with rank and projection reduction arguments.

Keywords: Computational Complexity, Communication Protocols, Modular Acception Modes, Undirected Graph Connectivity.

Introduction

During the last few years communication complexity theory gained popularity. In several papers many interesting questions of complexity theory were answered by reducing them to several kinds of communication games. Among others, this regards time–area tradeoffs for VLSI–circuits [1], [10], time–space tradeoffs for Turing machines, width–length tradeoffs for oblivious and usual Ω–branching programs ([2],[4]), branching programs of bounded alternation [14], and threshold circuits of depth 2 [11] and depth 3 [7].

The *graph connectivity problem for undirected graphs* $\mathrm{UCONN} = (\mathrm{UCONN}_{n(n-1)})_{n\in\mathbb{N}}$ in distributed form can be formulated as follows. Assume that we are given two not necessarily edge-disjoint undirected graphs $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ on a common n–set of vertices V, where both graphs are represented as Boolean vectors of length $\binom{n}{2}$. The question is whether or not the graph $G \stackrel{def}{=} G_1 \cup G_2 = (V, E_1 \cup E_2)$ is connected, i.e. each pair of vertices in G is connected. In [18] the major developments in understanding the complexity of the graph connectivity problem in several computational models are surveyed.

In the following we investigate the modular communication complexity of UCONN. Let two graphs $G_i = (V, E_i)$, for $i = 1, 2$, be given to two processors $\mathcal{P}_1$ and $\mathcal{P}_2$. In order to solve UCONN both processors have to communicate via a common communication tape. The computation of the whole structure, which is called a *communication protocol* or simply a *protocol*, is going on in *rounds*. Starting with $\mathcal{P}_1$ the processors write alternatingly bits on the communication tape. These bits depend on the input available to the processor which is to move and on the bits already written on the communication tape before. We assume without loss of generality, that in each round exactly one bit is written down on the communication tape. If the last bit written on the communication tape is "1" or "0" the computation is called *accepting* or *rejecting*, respectively. So co-operative computations can be thought of as to be Boolean strings. The length of the string is the *communication complexity* of the computation. Since we consider the worst–case–complexity in this paper, we assume without loss of generality, that all computations of a protocol are of equal length, say L. We shall assume the processors to be nondeterministic. That's why we have to define the *output* of a protocol via defining *acception modes*. As it is common use an acception mode is called a *counting mode* if the output of a protocol for a given input depends only on the numbers of accepting and rejecting computations performed by the protocol accessing this input. In this paper we discuss *the modular acception modes* in which the protocol accepts an input, if the number of accepting computations is not equal to 0 modulo m.

How to motivate the modular acception modes modulo m? In [20] it has been shown that all problems computable by constant depth, polynomial size circuits with MOD_m-gates for arbitrary integers m, are contained in certain counting communication complexity classes. In [5] these modes were formaly introduced and studied. Several papers (see e.g. [6]) deal with comparing the power of different counting acception modes. Roughly speaking, the computational power of the acception modes modulo m_i, $i = 1, 2$, is uncomparable, provided that $(m_1, m_2) = 1$ (see [8]).

We conclude this section by reviewing the results and methods which are strongly related to ours and by formulating the result of this paper. We use the notions and notations of Definition 1. Hajnal, Maass, and Turan proved in [9] the following theorem.

Theorem A $\mathrm{Comm}(\mathrm{UCONN}_{n(n-1)}) = \Theta(n \log n)$. □

Their method involves the use of the Möbius function μ for the lattice of partitions of an n–set. Lovasz and Saxe extended in [12] and [13] this ideas to a large class of problems, the so-called *meet problems*

for finite lattices, which can be formulated as follows. Let S be a finite lattice, and let both processor $\mathcal{P}_1$ and $\mathcal{P}_2$ be given an element x and y, respectively. Then they have to decide whether $x \wedge y = 0$.

Theorem B *Let* MEET_S *be the meet problem of a finite lattice. Let* S *have* a *atoms and* b *Möbius elements (i.e. elements* x *such that* $\mu(0,x) \neq 0$*). Then*

$$\log b \leq \text{Comm}(\text{MEET}_S) \leq (\log a)(\log b).$$

□

Recently, Raz and Spieker [15] proved

Theorem C *If processor* $\mathcal{P}_1$ *as well as processor* $\mathcal{P}_2$ *have an bipartite perfect matching on* $2n$ *vertices with two colors of size* n *as an input, and if their goal is to determine whether the union of the two matchings forms a Hamiltonian cycle, the nondeterministic communication complexity of the problem is* $\Omega(n \log\log n)$. □

Since the problem of Theorem C is a subproblem of UCONN (see Lemma 2), it follows

Corollary D $\text{N-Comm}(\text{UCONN}_{n(n-1)}) = \Omega(n \log\log n)$. □

It is the aim of this paper to show that modular acception modes help for detecting undirected graph connectivity.

Theorem *Let* m *be arbitrary. Then* $\text{MOD}_m\text{-Comm}(\text{UCONN}_{n(n-1)}) = \Theta(n)$.

Proof. The claim follows directly from Proposition 2 in Section 3 and from Proposition 4 in Section 4. □

We use the technique related to the Möbius function to prove the upper bound of Proposition 2. The lower bound of Proposition 4 follows from rank and reduction arguments.

1 The computational model

Let $f : S_1 \times S_2 \to \{0,1\}$ be given in distributed form. A *protocol of length* L consisting of two processors $\mathcal{P}_1$ and $\mathcal{P}_2$ that access inputs of S_1 and S_2, respectively, can be described by two functions

$$\Phi_i : S_i \times \{0,1\}^{*L} \to \{0,1\},$$

$i = 1,2$, and $\{0,1\}^{*L} = \{w \in \{0,1\}^* \mid 1 \leq |w| \leq L\}$. The interpretation is as follows. Let $\gamma = \gamma_1 \ldots \gamma_j$, $\gamma_k \in \{0,1\}$. If $\Phi_i(s_i,\gamma) = 1$, and if $|\gamma| - i$ is even, then the corresponding processor $\mathcal{P}_i$ is able to write γ_j on the communication tape provided that it has read $\gamma_1 \ldots \gamma_{j-1}$ on the communication tape and that it has s_i as input. If, however, $\Phi_i(s_i,\gamma) = 0$, then $\mathcal{P}_i$ is not able to write γ_j.

Now we define two $\#S_1 \times \#S_2$–matrices Acc^P and Rej^P associated with the protocol P of length L by

$$Acc^P_{s_1,s_2} \stackrel{def}{=} \sum_{\gamma_1 \ldots \gamma_L \in \{0,1\}^L,\ \gamma_L = 1} \prod_{j=1}^{L} \Phi_{1+((j+1) \bmod 2)}(s_{1+((j+1) \bmod 2)}, \gamma_1 \ldots \gamma_j) \tag{1}$$

$$Rej^P_{s_1,s_2} \stackrel{def}{=} \sum_{\gamma_1 \ldots \gamma_L \in \{0,1\}^L,\ \gamma_L = 0} \prod_{j=1}^{L} \Phi_{1+((j+1) \bmod 2)}(s_{1+((j+1) \bmod 2)}, \gamma_1 \ldots \gamma_j) \tag{2}$$

Clearly, $Acc^P_{s_1,s_2}$ gives the number of accepting computations of the protocol P on the input (s_1, s_2), whereas $Rej^P_{s_1,s_2}$ is the number of the rejecting computations. In order to make this approach unique, we agree that $\Phi_i(s_i, \gamma) = 1$, if $|\gamma| - i$ is odd, for $i = 1, 2$. We may give an equivalent definition of the above two matrices as follows. Let $\gamma \in \{0,1\}^L$ be a computation. Define

$$\chi^P_i(s_i, \gamma) \stackrel{def}{=} \prod_{\gamma' \in \{0,1\}^{*L},\ \gamma' \preceq \gamma} \Phi_i(s_i, \gamma'), \tag{3}$$

for $i = 1, 2$. Then we get directly from the equations (1) and (2)

$$Acc^P_{s_1,s_2} = \sum_{\gamma \in \{0,1\}^L,\ \gamma_L = 1} \chi^P_1(s_1, \gamma) \cdot \chi^P_2(s_2, \gamma) \tag{4}$$

$$Rej^P_{s_1,s_2} = \sum_{\gamma \in \{0,1\}^L,\ \gamma_L = 0} \chi^P_1(s_1, \gamma) \cdot \chi^P_2(s_2, \gamma) \tag{5}$$

Definition 1 *1. A* counting acception mode μ *for a protocol P is a function $\mu : \mathbb{N}^2 \to \{0,1\}$ such that P accepts an (s_1, s_2), if and only if, $\mu(Acc^P_{s_1,s_2}, Rej^P_{s_1,s_2}) = 1$. Otherwise P rejects the input. A protocol P equipped with an acception mode μ is called a μ-protocol. The function computed is sometimes denoted by* $\mathrm{Comp}(P, \mu)$. *If we are given a function $f : S_1 \times S_2 \to \{0,1\}$ then* μ-$\mathrm{Comm}(f) \stackrel{def}{=} \min\{L | \mathrm{Comp}(P, \mu) = f, L$ *is the length of* $P\}$.

2. We define the following acception modes.

Nondeterministic mode: $\mathrm{N}(n_1, n_2) = 1 \stackrel{def}{\iff} n_1 > 0,$

Modular modes: $\mathrm{MOD}_m(n_1, n_2) = 1 \stackrel{def}{\iff} n_1 \not\equiv 0 \pmod m,$

By the way, a deterministic communication protocol is not characterized by a special acception mode but by a property of the underlying protocol, namely $\Phi_i(s_i, \gamma 0) + \Phi_i(s_i, \gamma 1) \leq 1$, for $s_i \in S_i$, $i = 1, 2$, and $\gamma \in \{0,1\}^*$. For such protocols all reasonable counting modes coincide.

Lemma 1 *If $m_1 | m_2$, then* MOD_{m_2}-$\mathrm{Comm}(f) \leq \frac{m_2}{m_1} \cdot \mathrm{MOD}_{m_1}$-$\mathrm{Comm}(f)$, *for each function f.*

Proof. Clearly, $m_2 | \frac{m_3 \cdot m_2}{m_1}$, if and only if, $m_1 | m_3$.

Let P be the MOD_{m_1}–protocol for f. We describe the following protocol P'.

First, processor $\mathcal{P}_1$ chooses nondeterministically an index k, $1 \leq k \leq \frac{m_2}{m_1}$ and sheds k.

Second, $\mathcal{P}_2$ and $\mathcal{P}_1$ proceed in the same way as $\mathcal{P}_1$ and $\mathcal{P}_2$ do according to the protocol P. We get that $Acc^{P'}_{ij} = \frac{m_2}{m_1} \cdot Acc^P_{ij}$. Consequently, $Acc^{P'}_{ij} \equiv 0 \pmod{m_2} \iff Acc^{P'}_{ij} \equiv 0 \pmod{m_1}$. If L is the length of protocol P, then $\frac{m_2}{m_1} \cdot L$ is the length of protocol P'. □

Now we have to define what we mean by reductions. Fortunately, this is much easier here than in machine–based complexity theory.

Definition 2 *Let $F = (f_{2n} : \Sigma^n \times \Sigma^n \to \{0,1\})_{n \in \mathbb{N}}$ and $G = (g_{2n} : \Gamma^n \times \Gamma^n \to \{0,1\})_{n \in \mathbb{N}}$ be two decision problems. We say that F is rectangular reducible to G with respect to q, where $q : \mathbb{N} \to \mathbb{N}$ is a nondecreasing function, iff there are two transformations l_n, $r_n : \Sigma^n \to \Gamma^{q(n)}$ such that for all n and for all $\vec{x}, \vec{y} \in \Sigma^n$ we have $f_{2n}(\vec{x}, \vec{y}) = g_{2q(n)}(l_n(\vec{x}), r_n(\vec{y}))$. We write $F \leq^q_{rec} G$.*

We can utilize rectangular reductions for proving lower bounds. Let $q : \mathbb{N} \to \mathbb{N}$ be an unbounded nondecreasing function. Then we define $q^{(-1)}$ by $q^{(-1)}(i) = \max\{j | \ q(j) \leq i\}$. For example let

$\rho : \mathbb{R}^+ \to \mathbb{R}^+$ be an unbounded monotone increasing continous function and let $\rho^{-1} : \mathbb{R}^+ \to \mathbb{R}^+$ be a right–inverse to ρ, i. e. $\rho \circ \rho^{-1} = 1$. If we define $q : \mathbb{N} \to \mathbb{N}$ to be $q(i) = \lceil \rho(i) \rceil$, then $q^{(-1)}(i) = \lfloor \rho^{-1}(i) \rfloor$, for almost all natural numbers.

The proof of the following lower bound reduction argument is easy.

Lemma 2 *Assume that we are given two sequences of functions* $F = (f_{2n} : \Sigma^n \times \Sigma^n \to \{0,1\})_{n \in \mathbb{N}}$ *and* $G = (g_{2n} : \Gamma^n \times \Gamma^n \to \{0,1\})_{n \in \mathbb{N}}$. *If* μ-Comm$(F) \geq c(n)$ *and if* $F \leq^q_{rec} G$, *then* μ-Comm$(G) \geq c \circ q^{(-1)}(n)$. □

One efficient way to get rectangular reductions is to handle with projection reductions. The variables over $\{0,1\}^n$ are coordinate functions $x_i : \{0,1\}^n \to \{0,1\}$ such that $x_i(\sigma_1, \ldots, \sigma_n) = \sigma_i$. In accordance with Skyum and Valient (see [17]) we define.

Definition 3 *1. Let* $F_n : \{0,1\}^n \to \{0,1\}$ *and* $g_m : \{0,1\}^m \to \{0,1\}$. F_n *is called reducible to* g_m *via a projection* $\pi_n : \{y_1, \ldots, y_m\} \to \{x_1, \ldots, x_n, \neg x_1, \ldots, \neg x_n, 0, 1\}$ *and we write* $F_n \leq_{\pi_n} g_m$, *where the* y_j *and the* x_i *are the Boolean variables of* F_n *and* g_m, *resp., if*

$$F_n(x_1, \ldots, x_n) = g_m(\pi(y_1), \ldots, \pi(y_m)).$$

2. If F_n *and* g_m *are given in distributed form, i. e.* $F_n : \{0,1\}^{n/2} \times \{0,1\}^{n/2} \to \{0,1\}$ *and* $G_n : \{0,1\}^{m/2} \times \{0,1\}^{m/2} \to \{0,1\}$, *then we say that the reduction* π *respects the distribution of the variables, if*

$$\pi_n^{-1}\{x_1, \ldots, x_{n/2}, \neg x_1, \ldots, \neg x_{n/2}\} \subseteq \{y_1, \ldots, y_{m/2}\}$$

and

$$\pi_n^{-1}\{x_{n/2+1}, \ldots, x_n, \neg x_{n/2+1}, \ldots, \neg x_n\} \subseteq \{y_{m/2+1}, \ldots, y_m\}.$$

3. There is a transpose $\pi_n^t : \{0,1\}^n \to \{0,1\}^m$ *of the projection reduction* π. *It is defined by*

$$\pi_n^t(\vec{u}) = (\pi_n(y_1)(\vec{u}), \ldots, \pi_n(y_m)(\vec{u})),$$

where $\vec{u} = (x_1(\vec{u}), \ldots, x_n(\vec{u})) \in \{0,1\}^n$ *is any Boolean vector of length* n.

4. If $F = (F_n)_{n \in \mathbb{N}}$ *and* $G = (G_n)_{n \in \mathbb{N}}$ *are sequences of functions, if* $\Pi = (\pi_n)_{n \in \mathbb{N}}$ *is a sequence of projection reductions defined in the first item of this definition, i. e.* $F_n \leq_{\pi_n} g_m$, *and if* $m \leq p(n)$, *then we say that* Π *ist a* $p(n)$*–projection reduction and we write* $F \leq^p_\Pi G$. *If both* F *and* G *are given in distributed form, then the definition of the notion "* π *respects the distribution of the variables" can be done by analogy with the second item of this definition.*

If the elements of $\{0,1\}^n$ are representations of graphs, then we visualize the graph which is the transpose $\pi_n^t(\vec{\sigma})$ of a vector $\vec{\sigma} \in \Sigma^n$ in such a way that the edges which are not constant are labelled by the corresponding literal (see figures 1 and 2). The meaning is that such an edge belongs to the graph, if and only if, the labelling literal is true.

Due to Lemma 2 we get

Lemma 3 *Assume that we are given two sequences of functions* $F = (f_{2n} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\})_{n \in \mathbb{N}}$ *and* $G = (G_{2m} : \{0,1\}^m \times \{0,1\}^m \to \{0,1\})_{m \in \mathbb{N}}$ *such that* $F \leq^p_\Pi G$, *where* $p : \mathbb{N} \to \mathbb{N}$ *is increasing, and* $\Pi = (\pi_n)_{n \in \mathbb{N}}$ *is a sequence of projection reductions which respects the distribution of the variables. If* μ-Comm$(F) \geq c(n)$, *then* μ-Comm$(G) \geq c \circ q^{(-1)}(n)$. □

2 Rank arguments for upper and lower bounds

We shall derive rank arguments for proving upper and lower bounds on the length of protocols equipped with the modular acception modes from Definition 1. We adopt the concept of variation ranks of communication matrices developed in [11]. Throughout this section let f denote a function $f : S_1 \times S_2 \to \{0,1\}$, $N = \#S_1 = \#S_2$, and let M^f denote the communication matrix, where $M^f_{i,j} = f(i,j)$, for $i,j = 1,\ldots,N$.

Let the *sequence equality function* be defined by $\mathrm{SEQ}_{2n}(x_1,\ldots,x_n,y_1,\ldots,y_n) = \bigwedge_{i=1}^{n}(1 - ((x_i + y_i) \bmod 2))$. Here $S_1 = S_2 = \{0,1\}^n$.

Definition 4 *1. Two $N \times N$–matrices A and B over the ring of integers are defined to be mod_m–equivalent, where m is a natural number, if and only if, for all indices i, j,*

$$a_{ij} \equiv 0 \pmod{m} \quad \Longleftrightarrow \quad b_{ij} \equiv 0 \pmod{m}.$$

2. Let A be an integer matrix. We define $\text{var-rank}_{\mathbb{Z}/m\mathbb{Z}}(A \bmod m)$ *to mean the minimum of all numbers* $\mathrm{rank}_{\mathbb{Z}/m\mathbb{Z}}(B \bmod m)$*, where B is an integer matrix which is mod_m–equivalent to A.*

A 0-1 matrix is interpreted as an R–matrix, where R an arbitrary semiring, in the canonical way. As usual, *the R–rank of a $m \times n$–matrix A over R*, which we denote by $\mathrm{rank}_R A$, is defined to be the minimal number k such that $A = B \cdot C$, where B is a $m \times k$–matrix and C is a $k \times n$–matrix over R. A straightforward calculation yields the next lemma.

Lemma 4 *Let A be an integer matrix.*

1. $\mathrm{rank}_{\mathbb{Z}/m\mathbb{Z}}(A \bmod m) = \max\{\mathrm{rank}_{\mathbb{Z}/m_i\mathbb{Z}}(A \bmod m_i) \mid i = 1,\ldots,r\}$,
   provided that $m = m_1 \cdot \ldots \cdot m_2$, where $(m_i, m_j) = 1$, for all $i \neq j$.

2. $\mathrm{rank}_{\mathbb{Z}/m\mathbb{Z}}(A \bmod m) = \min\{\mathrm{rank}_{\mathbb{Z}} D \mid D$ *is mod_m–equivalent to $A\}$.* □

Lemma 5 *Let R be any semiring. Let P be a protocol of the length L on the input set $S_1 \times S_2$, $\#S_1 = \#S_2 = N$, and let Acc^P and be the $N \times N$–matrix defined in equation 1. Then* $\mathrm{rank}_R(Acc^P) \leq 2^{L-1}$.

Proof. The inequality follows directly from equation 4. □

Now we can fully charaterize the modular communication complexity im terms of variation ranks.

Proposition 1

$$\log_2\left(\text{var-rank}_{\mathbb{Z}/m\mathbb{Z}}(M^f)\right) \leq \mathrm{MOD}_m\text{-Comm}(f) \leq \log_2\left(\text{var-rank}_{\mathbb{Z}/m\mathbb{Z}}(M^f)\right) + 2\log_2 m + 1.$$

Proof. The left inequality follows directly from Lemma 5 and from Definition 4. Let us turn to the right one. We choose by Lemma 4 an integer matrix B which is mod_m–equivalent to M^f, such that $r = \mathrm{rank}_{\mathbb{Z}/m\mathbb{Z}}(B \bmod m) = \text{var-rank}_{\mathbb{Z}/m\mathbb{Z}}(M^f)$. Then $B = B^{(1)} + \ldots + B^{(r)}$, where the $B^{(k)}$ have $\mathbb{Z}/m\mathbb{Z}$–rank 1. This is equivalent to $B^{(k)}_{ij} \equiv U^{(k)}_i \cdot V^{(k)}_j \pmod{m}$, for $U^{(k)}_i, V^{(k)}_j \in \{1,\ldots,m\}$, and for $i,j = 1,\ldots N$.

Now we can describe the following protocol P. Assume that the input is $(i,j) \in S_1 \times S_2$.

First, processor $\mathcal{P}_1$ chooses nondeterministically some indices k, $1 \leq k \leq r$, and l_1, $1 \leq l_1 \leq U^{(k)}_i$, and sheds (k, l_1).

Second, processor $\mathcal{P}_2$ chooses nondeterministically some index l_2, $1 \leq l_2 \leq V_j^{(k)}$, and sheds $(l_2, 1)$.

Clearly, there are $\sum_{k=1}^{r} U_i^{(k)} \cdot V_j^{(k)} \equiv B_{ij} \pmod{m}$ many accepting computations assigned to the input (i,j). It follows that $\mathrm{Comp}(P, \mathrm{MOD}_m) = f$. Obviously, the length of the protocol is bounded above by $\log_2 r + 2\log_2 m + 1$. □

In the case of m being a prime number, we can even do better.

Corollary 1 *If $m = p$ is a prime number, we have*

$$\frac{1}{p-1} \cdot \log_2 \left(\mathrm{rank}_{\mathbb{Z}/p\mathbb{Z}}(M^f) \right) \leq \mathrm{MOD}_p\text{-Comm}(f) \leq \frac{1}{p-1} \cdot \left(\log_2 \left(\mathrm{rank}_{\mathbb{Z}/p\mathbb{Z}}(M^f) \right) + 2\log_2 p + 1 \right).$$

Proof. By means of Fermat's Little Theorem each protocol of length L can be transformed into a protocol P' of length $(p-1)L$ such that for all inputs (i,j)

$$Acc_{ij}^{P'} = \left(Acc_{ij}^{P} \right)^{p-1} \equiv \begin{cases} 0 \pmod{p} & \text{if } Acc_{ij}^{P} \equiv 0 \pmod{p}; \\ 1 \pmod{p} & \text{if } Acc_{ij}^{P} \not\equiv 0 \pmod{p}. \end{cases}$$

□

3 The Möbius function and upper bounds on the length of MOD_m–protocols for undirected graph connectivity

In this section we transform a method due to Lovasz and Saxe (see [12], [13]) for proving lower bounds on the length of deterministic protocols to the case of MOD_m–protocols in oder to prove upper bounds. We can only give a very brief treatment on Möbius functions. For more see [16].

Let S be a finite partially ordered set, R be a commutative ring with 1. The *R–valued incidence algebra* $\mathcal{A}(S,R)$ is defined as follows. Consider the set of functions of two variables $f(x,y)$, for x and y ranging over S having values in R, and with the property that $f(x,y) = 0$ whenever $x \not\leq y$. The sum and the multiplication by scalars are defined pointwise. The product of f and g is defined as follows.

$$(fg)(x,y) \stackrel{def}{=} \sum_z f(x,z)g(z,y)$$

Clearly, Kronecker's δ–function is the 1 of $\mathcal{A}(S,R)$. The *R–valued zeta function* $\zeta(x,y) \in \mathcal{A}(S,R)$ is defined by $\zeta(x,y) = 1$ if $x \leq y$ and $\zeta(x,y) = 0$ otherwise. The function $\iota(x,y) \stackrel{def}{=} \zeta(x,y) - \delta(x,y)$ is called the *incidence function*.

The following formula is the key to prove Lemma 6.

$$g(x,y) \quad = \quad -\frac{1}{f(x,x)} \sum_z f(x,z)g(z,y)\iota(z,y) \tag{6}$$

It allows a recursive definition of the inverse of f, provided that the $f(x,x)$ are units in R.

Lemma 6 *An element of $\mathcal{A}(S,R)$ is a unit, if and only if, $\prod_x f(x,x)$ is a unit in R.* □

Consequently, we can define *the R–valued Möbius function* to be the inverse of the zeta function. Let us denote this function for a moment by $\mu^{(R)}$.

Analogously to the standard real–valued case, we have the *Möbius inversion formula.* Let $f(x)$ be an R–valued function, for x ranging over the finite poset S, and let $g(x) = \sum_y f(y)\zeta(y,x)$. Then $f(x) = \sum_y g(y)\mu^{(R)}(y,x)$.

If μ denotes the real–valued Möbius function, then because of formula 6 μ takes values only in $\mathbb{Z}$. Consequently, if $R_0 \subseteq R$ is the prime ring of R, which equals either $\mathbb{Z}$ or $\mathbb{Z}/m\mathbb{Z}$, for some $m \in \mathbb{Z}$, then

$$\mu^{(R)}(x,y) = \begin{cases} \mu(x,y) & \text{if } R_0 = \mathbb{Z}; \\ \mu(x,y) \bmod m & \text{if } R_0 = \mathbb{Z}/m\mathbb{Z}. \end{cases}$$

Now, of course, we can drop the notation $\mu^{(R)}$.

Again from formula 6 it follows that $\mu(x,y)$ only depends on the the structure of the interval. Moreover, we know, that if μ^* is the Möbius function of the dual poset S^*, then $\mu^*(x,y) = \mu(y,x)$.

Let us assume from now on that the poset S is a lattice. In line with [12] we shall consider the meet problem $\mathrm{MEET}_S : S \times S \to \{0,1\}$ of the finite lattice S, defined by $\mathrm{MEET}_S(x,y) = \delta(0, x \wedge y)$. We proceed as follows. Let M be a 0–1 matrix. Check whether there are two equal rows or colomns in M and if this is the case, then delete one of them. Do that as long as possible. The resulting matrix $\tilde{M}$ is called the *core* of M. Clearly, the communication complexity of the underlying problems is the same. Now it is not difficult to see that the core of $M^{\mathrm{UCONN}_{n(n-1)}}$ equals the core of $M^{\mathrm{MEET}_{\mathcal{P}(n)^*}}$, where $\mathcal{P}(n)^*$ is the lattice dual to the lattice of partitions of an n–set.

Lemma 7 *Let M be the communication matrix of the meet problem assigned to the finite lattice S, and let p be a prime number. Then* $\mathrm{rank}_{\mathbb{Z}/p\mathbb{Z}}(M) = \#\{x \in S \mid \mu(0,x) \not\equiv 0 \pmod p\}$.

Proof. Let $\tilde{M}$ be the diagonal matrix $\mathrm{diag}(\mu(0,x))_{x\in S}$, and let $\zeta = (\zeta(x,y))_{x,y\in S}$ be the matrix associated with the zeta function. Wilf observed in [19], that $\zeta^T \cdot \tilde{M} \cdot \zeta = M$. The claim follows from the Möbius Inversion Formula. □

Now let us compute $\#\{x \in S \mid \mu(0,x) \not\equiv 0 \pmod p\}$ in a special case.

Lemma 8 *Let $\mathcal{P}(n)^*$ be the lattice dual to the lattice $\mathcal{P}(n)$ of partitions, let $p < n$ be a prime number, and let μ^* be the Möbius function of $\mathcal{P}(n)^*$. Then $\#\{x \in \mathcal{P}(n)^* \mid \mu^*(0,x) \not\equiv 0 \pmod p\} \leq p^n$.*

Proof: The following three facts are well-known.

Fact 1. If $x \in \mathcal{P}(n)$, and if $b(x)$ is the number of blocks of the partition x, then $[x,1] \cong \mathcal{P}(b(x))$.

Fact 2. If μ is the Möbius function of $\mathcal{P}(n)$, then $\mu^*(0,1) = \mu(0,1) = (-1)^{n-1}(n-1)!$.

Fact 3. Let $S(n,k)$ denote the number of partitions of an n–set into exactly k blocks (Stirling numbers of the second kind), then

$$\sum_{k=0}^{n} S(n,k)[X]_k = X^n,$$

where X is an indeterminant and $[X]_k = X \cdot (X-1) \cdot \ldots \cdot (X-k+1)$ is the falling factorial.

The next equality follows from Fact 1 and from Fact 2. The next but one from Fact 3.

$$\begin{aligned} \sum_{k=0}^{p} S(n,k) &= \#\{x \in \mathcal{P}(n)^* \mid \mu^*(0,x) \not\equiv 0 \pmod p\} \\ \sum_{k=0}^{p} S(n,k) &\leq \sum_{k=0}^{p} S(n,k)[p]_k = p^n \end{aligned}$$

□

Proposition 2 *Let m be arbitrary. Then* $\mathrm{MOD}_m\text{-Comm}(\mathrm{UCONN}_{n(n-1)}) = O(n)$.

Proof. Let p be a prime number such that $p|m$. By Lemma 1 we have

$$\mathrm{MOD}_m\text{-Comm}(\mathrm{UCONN}) \leq \frac{m}{p} \cdot \mathrm{MOD}_p\text{-Comm}(\mathrm{UCONN}).$$

The claim follows from Corollary 1, Lemma 7, and Lemma 8. □

4 Variation ranks and lower bounds on the length of MOD_m–protocols for undirected graph connectivity

The following lemma improves the corresponding one from [11].

Lemma 9 *Let I_N denote the identity $N \times N$–matrix. Let $m = p_1^{l_1} \cdot \ldots \cdot p_r^{l_r}$ be a natural number which is given by its primary decomposition. Then* $\text{var-rank}_{\mathbb{Z}/m\mathbb{Z}}(I_N) = \lceil N/r \rceil$.

Proof. First we prove that $\lceil N/r \rceil$ is a lower bound. Let A be an integer matrix such that A is mod_m–equivalent to I_N and $\text{var-rank}_{\mathbb{Z}/m\mathbb{Z}}(I_N) = \text{rank}_{\mathbb{Z}} A$, which exists by Lemma 4. By definition we have, for all i, $a_{ii} \not\equiv 0 \pmod{m}$, and $a_{ij} \equiv 0 \pmod{m}$, for all $j \neq i$. For all $i \in \{1, \ldots, N\}$ there is a $k \in \{1, \ldots, r\}$ such that $a_{ii} \not\equiv 0 \pmod{p^{l_k}}$. We conclude that there is a primary component $p_k^{l_k}$ of m, which we denote for simplicity by p^l, a set of indices

$$\mathcal{I} \subseteq \{1, 2, \ldots, N\}, \#\mathcal{I} \geq N' := \lceil N/r \rceil,$$

and, for all $i \in \mathcal{I}$, natural numbers $\nu_i \in \{1, \ldots, l_i\}$, such that

$$\begin{aligned} a_{ii} &\equiv 0 \pmod{p^{l-\nu_i}}, \\ a_{ii} &\not\equiv 0 \pmod{p^{l-\nu_i+1}}, \\ a_{ij} &\equiv 0 \pmod{p^l}, \end{aligned}$$

for all $j \in \mathcal{I}$, $j \neq i$. After deleting all rows and columns of A whose indices do not belong to $\mathcal{I}$, we get an integer $N' \times N'$–matrix B. It is sufficient to show that $\det B \neq 0$. It is easy to see that

$$\begin{aligned} b_{1,1} \cdot \ldots \cdot b_{N',N'} &\not\equiv 0 \pmod{p^{N'.l+1-\sum_{i=1}^{N'} \nu_i}}, \text{ but} \\ b_{1,\sigma(1)} \cdot \ldots \cdot b_{N',\sigma(N')} &\equiv 0 \pmod{p^{N'.l+1-\sum_{i=1}^{N'} \nu_i}}, \end{aligned}$$

for all permutations σ of the set $\{1, \ldots, N'\}$ different from the identity permutation. Consequently,

$$\det B \equiv b_{1,1} \cdot \ldots \cdot b_{N',N'} \not\equiv 0 \pmod{p^{N'.l+1-\sum_{i=1}^{N'} \nu_i}}.$$

Second let us prove that $\lceil N/r \rceil$ is an upper bound. Let $f_i = p_i^{-l_i} \prod_{\mu=1}^{r} p_\mu^{l_\mu}$, $F_j = (f_1, \ldots, f_j)$, and $A_j = F_j^T \cdot F_j$ for $i, j = 1 \in \{1, \ldots, r\}$. A_0 is defined to be the unique 0×0–matrix, which, of course, has rank 0. Clearly, $A_j \bmod m$ is a $j \times j$–diagonal matrix of $\mathbb{Z}/m\mathbb{Z}$–rank 1, for $j \in \{1, \ldots, r\}$. Define the matrix A to be the following direct sum of matrices.

$$A \stackrel{def}{=} A_r \oplus \overbrace{\ldots\ldots}^{\lfloor N/r \rfloor} \oplus A_r \oplus A_{r'},$$

where $r' \equiv N \pmod{r}$, and $r' \in \{0, \ldots, r-1\}$. It follows that $A \bmod m$ is a diagonal $N \times N$–matrix, and that $\text{rank}_{\mathbb{Z}/m\mathbb{Z}}(A \bmod m) \leq \lceil N/r \rceil$. □

Proposition 3 *For m arbitrary, we have that* $\mathrm{MOD}_m\text{-Comm}(\mathrm{SEQ}_{2n}) = \Theta(n)$.

Proof. The claim follows from Proposition 1 and from Lemma 9. □

Lemma 10 $\mathrm{SEQ} = (\mathrm{SEQ}_{2n})_{n\in\mathbb{N}}$ *is reducible to* $\mathrm{UCONN} = (\mathrm{UCONN}_{n(n-1)})_{n\in\mathbb{N}}$ *given in distributed form via a $O(n^2)$–projection reduction with respect to the partition of the variables.*

Proof. Consider an input $(t_1, \ldots, t_n, u_1, \ldots, u_n)$ of SEQ_{2n} The projection reduction

$$\pi_{n(n-1)} : \{x_{ij}, y_{ij} \mid i, j = 1, \ldots, n,\ i < j\} \to \{0, 1, t_\nu, u_\nu, \neg t_\nu, \neg u_\nu \mid \nu = 1, \ldots, n\},$$

where the values of the Boolean variables x_{ij} and y_{ij} define the graphs G_1 and G_2 accessible to the processors $\mathcal{P}_1$ and $\mathcal{P}_2$, is defined by the help of Figure 1 and Figure 2, in which the transpose

$$\pi^t_{n(n-1)}(t_1, \ldots, t_n, u_1, \ldots, u_n)$$

is shown. Clearly, this graph is connected, if and only if,

$$\mathrm{SEQ}_{2n}(t_1, \ldots, t_n, u_1, \ldots, u_n) = 1.$$

□

Now it is easy to prove the lower bound.

Proposition 4 *Let m be arbitrary. Then* $\mathrm{MOD}_m\text{-Comm}(\mathrm{UCONN}_{n(n-1)}) = \Omega(n)$.

Proof. The claim follows from Lemma 10, Lemma 3 and Proposition 3. □

References

[1] A. V. Aho, J. D. Ullman, M. Yannakakis, *On notions of information transfer in VLSI circuits*, in: Proc. 15th ACM STOC 1983, pp. 133–183.

[2] N. Alon, W. Maass, *Meanders, Ramsey theory and lower bounds for branching programs*, Journal of Computer and System Sciences **37**(1988), pp. 118–129.

[3] L. Babai, P. Frankl, J. Simon, *Complexity classes in communication complexity theory*, in: Proc. 27th IEEE FOCS, pp. 337–347, 1986.

[4] L. Babai, N. Nisan, M. Szegedy, *Multiparty protocols and logspace-hard pseudorandom sequences*, Journal of Computer and System Sciences **45**(1992), pp. 204–232.

[5] B. Halstenberg, R. Reischuk, *Relations between Communication Complexity Classes*, Journal of Computer and System Sciences **41**(1990), pp. 402–429.

[6] B. Halstenberg, *The Polynomial Communication Hierarchy and Protocols with Moderately Bounded Error,* Technical Report TI 1/90 of TH Darmstadt, 1990.

[7] J. Hastad, M. Goldmann, *On the power of small–depth threshold circuits*, in: Proc. 31st IEEE FOCS 1990, pp. 610–618.

[8] C. Damm, M. Krause, Ch. Meinel, St. Waack, *Separating counting communication complexity classes*, in: Proc. 9th STACS, Lecture Notes in Computer Science **577**, Springer Verlag 1992, pp. 281–293.

[9] A. Hajnal, W. Maass, G. Turan, *On the communication complexity of graph problems* in: Proc. 20th ACM STOC 1988, pp. 186–191.

[10] J. Hromkovic, M. Krause, Ch. Meinel, St. Waack, *Branching programs provide lower bounds on the area of multilective deterministic and nondeterministic VLSI circuits.*, Information and Computation **94**(2)(1992) pp. 168–178.

[11] M. Krause, St. Waack, *Variation ranks of communication matrices and lower bounds for depth two circuits having symmetric gates with unbounded fan–in*, in: Proc. 32th IEEE FOCS 1991, pp. 777–782.

[12] L. Lovasz, *Communication complexity: A survey*, in: *Paths, flows and VLSI–layouts*, Springer–Verlag 1990, pp. 235–266.

[13] L. Lovasz, M. Saks, *Communication complexity and combinatorial lattice theory*, Journal of Computer and System Sciences **47**(1993), pp. 322-349.

[14] Ch. Meinel, St. Waack, *Upper and lower bounds for certain graph-accessibility problems on bounded alternating ω-branching programs*,
in: *Complexity Theory – current research*, K. Ambros–Spies, St. Homer, U. Schöning (editors), Cambridge University Press 1993, 273–290.

[15] R. Raz, B. Spieker, *On the "log rank"–conjecture in communication complexity*, in: Proc. 34th IEEE FOCS 1993, pp. 168–176.

[16] G.-C. Rota, *On the foundation of combinatorial theory: I. Theory of Möbius functions*, Z. Wahrscheinlichkeitstheorie **2**(1964), pp. 340–368.

[17] L. Skyum, L. V. Valiant, *A complexity theory based on Boolean algebra*, in: Proc. 22th IEEE FOCS, pp. 244-253.

[18] A. Wigderson, *The complexity of graph connectivity*, TR 92-19, Leipniz Center for Research in Computer Science, Institute of Computer Science, Hebrew University, Jerusalem.

[19] S. Wilf, *Hadamard determinants, Möbius functions and the cromatic number of graphs*, Bull./ Amer. Math. Soc. **74**(1968), pp. 960–964.

[20] A. C.-C. Yao, *On ACC and threshold circuits*, in: Proc. 31st IEEE FOCS 1990, pp. 619–627.

Figure 1: The graph $\pi^t_{n(n-1)}(t_1, \ldots, t_n, u_1, \ldots, u_n)$
($K_{2,2}$ denotes full bipartite graph having 2×2 nodes, $\mathcal{G}(t_\mu, u_\mu)$ is defined in Figure 2.)

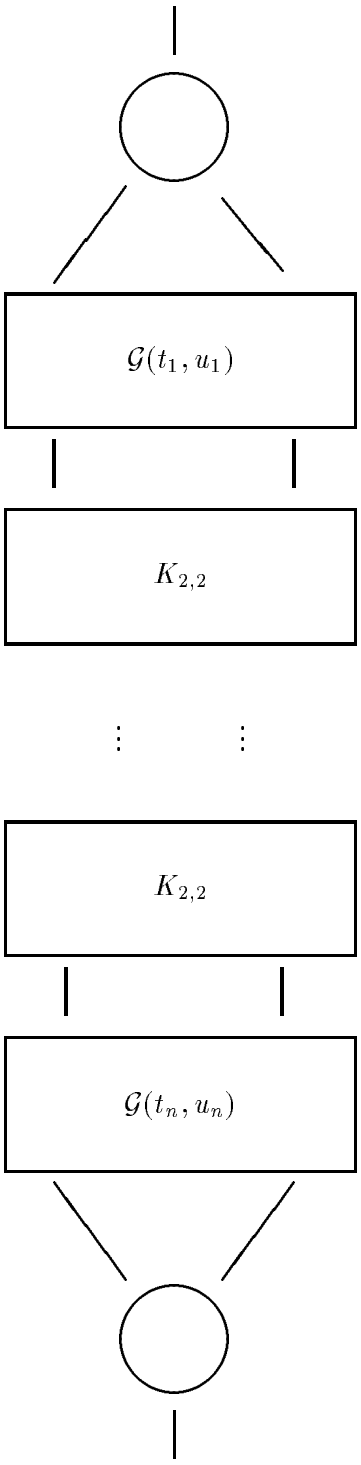

Figure 2: The graphs $\mathcal{G}(t_\mu, u_\mu)$ of Figure 1