Complete Classifications for the Communication Complexity of Regular Languages

Pascal Tesson[1] and Denis Thérien[1*]

School of Computer Science, McGill University
{ptesso,denis}@cs.mcgill.ca

Abstract. We show that every regular language L has either constant, logarithmic or linear two-party communication complexity (in a worst-case partition sense). We prove a similar trichotomy for simultaneous communication complexity and a "quadrichotomy" for probabilistic communication complexity.

1 Introduction

Communication complexity has grown to be one of the most ubiquitous tools of theoretical computer science [5] and the work pioneered by Babai, Frankl and Simon in [1] (see also [4]) has provided a rich structure of communication complexity classes to understand the relationships between various models of communication complexity.

The focus of this paper is the communication complexity of regular languages which, in some sense, are the simplest languages in the usual Time/Space complexity framework. Perhaps surprisingly, some regular languages have very high communication complexity even in powerful extensions of the deterministic model: There exist regular languages that are complete (in the sense of [1]) for every level of PH^{cc}, PSPACE^{cc} (using a result of [2]) and for $\mathrm{MOD_qP}^{cc}$ (see [4]). Moreover, a majority of the most well-studied examples in the field such as Equality, Disjointness or Inner Product modulo m, are equivalent, from a communication complexity point of view, to regular languages (in these examples, equivalent to the regular languages $\{00 \cup 11\}^*$, $\{01 \cup 10 \cup 00\}^*$ and $\{\{\{01 \cup 10 \cup 00\}^* 11 \{01 \cup 10 \cup 00\}^*\}^m\}^*$ respectively).

The algebraic theory of finite monoids and semigroups provides a powerful framework and tools to analyze and classify regular languages [6] and has also been used to uncover connections between algebraic properties of finite monoids and classes of Boolean circuits using the program over monoid formalism of Barrington and Thérien [2, 3]. Szegedy found a surprising connection between monoid programs and communication complexity by showing that a language has constant two-party deterministic communication complexity in the worst-case partition of the input if and only if it can be recognized by a program over some finite commutative monoid [9].

* Research supported in part by NSERC and FCAR grants.

It was established in [8] that the class of regular languages having $O(f)$ communication complexity forms a variety and so the question of the communication complexity of regular languages has an algebraic answer. Amazingly, this point of view allows us to determine the complexity of every regular language in the deterministic, probabilistic and simultaneous models of communication. This complete classification unexpectedly features complexity gaps in all three cases: a regular language has either $O(1)$, $\Theta(\log n)$ or $\Theta(n)$ communication complexity in the deterministic two-party model, $O(1)$, $\Theta(\log n)$ or $\Theta(n)$ complexity in the simultaneous model and $O(1)$, $\Theta(\log\log n)$, $\Theta(\log n)$ or $\Theta(n)$ in the bounded-error probabilistic model. The classes of languages thus defined have fairly simple algebraic and combinatorial descriptions and are thus decidable. Moreover, some of these classes coincide: a regular language has $O(\log n)$ probabilistic complexity only if it has $O(\log n)$ deterministic complexity and further has $\Theta(\log\log n)$ probabilistic complexity if and only if it has $\Theta(\log n)$ simultaneous complexity. In proving these theorems we use rectangular reductions to and from only four different problems: Disjointness, Inner Product modulo m, Greater Than and Index which both highlights and, in retrospect, explains their importance as fundamental examples in communication complexity. We should note that, in contrast with the regular case, it is easy to artificially define, for any function $1 \leq f \leq n$, some non-regular language of communication complexity $\Theta(f)$ (in any of these models).

The main technical content of our paper is algebraic in nature and is presented partly in Section 2 along with some background on finite monoids and partly in the Appendix due to space limitations. Next, Section 3 introduces the definitions and results in communication complexity needed to prove the bounds of our main theorems in Section 4.

2 Algebraic Background and Results

We refer the reader to e.g. [6] for further background on algebraic automata theory and semigroup theory.

A monoid M is a set together with a binary associative operation and a distinguished identity element 1_M for this operation. Throughout this paper, M will denote a finite monoid and A some finite alphabet.

The following equivalence relations on M, called Green's relations, describe whether two elements generate the same ideals: $x\mathcal{J}y$ if $MxM = MyM$, $x\mathcal{L}y$ if $Mx = My$, $x\mathcal{R}y$ if $xM = yM$ and $x\mathcal{H}y$ if both $x\mathcal{R}y$ and $x\mathcal{L}y$. It is known that $\mathcal{L}$ and $\mathcal{R}$ commute and that for finite monoids $\mathcal{L} \circ \mathcal{R} = \mathcal{J}$. We will denote by $\mathcal{J}_x$ the $\mathcal{J}$-class of x and similarly for the other relations.

If we instead ask whether element's ideals are contained in one another, we can define natural pre-orders $\leq_\mathcal{J}, \leq_\mathcal{R}, \leq_\mathcal{L}$ on M with e.g. $x \leq_\mathcal{J} y$ if and only if $MxM \subseteq MyM$. We will say that "x is (strictly) $\mathcal{J}$-above y" if $x \geq_\mathcal{J} y$ (resp. $x >_\mathcal{J} y$), and so on. Note that $x \leq_\mathcal{J} y$ if and only if there exists u, v such that $x = uyv$ and similarly $x \leq_\mathcal{R} y$ if and only if there is u with $x = yu$.

An element $m \in M$ is *idempotent* if $m^2 = m$ and since M is finite there must exist a smallest integer ω such that m^ω is idempotent for all $m \in M$. It can be shown that a subset T of M is a maximal subgroup in M if and only if T is an $\mathcal{H}$-class containing an idempotent. We call a $\mathcal{J}$-class regular if it contains at least one idempotent.

For a group G, the exponent of G is the smallest q such that $g^q = 1_G$ for all $g \in G$. More generally, we will call the *exponent of* M the least common multiple of the exponents of all subgroups of M.

A monoid N *divides* M if it is the morphic image of a submonoid of M and we denote this as $N \prec M$. A class $\mathbf{V}$ of monoids is said to form a *(pseudo-)variety* if it is closed under direct product and division. If $\mathbf{H}$ is a variety of finite groups, we denote by $\overline{\mathbf{H}}$ the variety of monoids whose subgroups all belong to $\mathbf{H}$. We will in particular mention the variety $\overline{\mathbf{Ab}}$ of monoids in which all subgroups are Abelian.

We say that $L \subseteq A^*$ can be *recognized* by M if there exists a morphism $\phi : A^* \to M$ and a subset $F \subseteq M$ such that $L = \phi^{-1}(F)$. A trivial variant of Kleene's Theorem states that a language L can be recognized by some finite monoid if and only if it is regular. Every such L admits some minimal monoid recognizer which we call L's *syntactic monoid* and denote $M(L)$. In particular, $M(L)$ recognizes L and for any N also recognizing L, we have $M(L) \prec N$. Moreover, for any class of regular languages $\mathbf{L}$ closed under Boolean operations, left and right quotients and inverse morphisms from A^* to B^*, there exists a variety of monoids $\mathbf{V}$ such that $L \in \mathbf{L}$ if and only if $M(L) \in \mathbf{V}$.

For a finite M, we say that $u, v \in A^*$ are *M-equivalent* if for all morphisms $\phi : A^* \to M$ we have $\phi(u) = \phi(v)$. For example, ab and ba are M-equivalent for any commutative M.

2.1 Combinatorial Descriptions of Subvarieties of DO

We denote $\mathbf{DS}$ the variety of monoids where every regular $\mathcal{J}$-class J is closed under multiplication. This is provably equivalent to the requirement that every $\mathcal{H}$-class within a regular J contain an idempotent. We denote by $\mathbf{DO}$ the subvariety of $\mathbf{DS}$ where every monoid satisfies $(xy)^\omega (yx)^\omega (xy)^\omega = (xy)^\omega$ for all x, y. Alternatively, one can characterize $\mathbf{DO}$ as the class of monoids such that the product of any two $\mathcal{J}$-related idempotents e, f is itself idempotent and $\mathcal{J}$-related to e and f. The following two lemmas, whose proofs are included in the appendix, describe some of the useful properties of monoids in $\mathbf{DS}$ and $\mathbf{DO}$.

Lemma 1. *Let $M \in \mathbf{DS}$ then*

- *If e, f are $\mathcal{J}$-related idempotents then $(ef)^\omega e = e$ (but $efe = e$ if and only if ef is idempotent).*
- *for any $a, b \in M$ such that $a \leq_{\mathcal{J}} b$ with a regular then $a \,\mathcal{J}\, ab \,\mathcal{J}\, ba$.*
- *for any $x, y \in M$ such that $xy \,\mathcal{R}\, x$ we have in fact $R_x y \subseteq R_x$.*

Lemma 2. *Suppose $M \in \mathbf{DO}$ and let $w_1, w_2 \in A^*$ be G-equivalent for any subgroup G of M. For any morphism $\phi : A^* \to M$ and any $x \in M$ such that $x\phi(w_1) \,\mathcal{H}\, x\phi(w_2) \,\mathcal{R}\, x$ we have in fact $x\phi(w_1) = x\phi(w_2)$.*

For a word $u \in A^*$, we denote $\alpha(u)$ the set of letters occurring in u. For $a \in \alpha(u)$, the *a-left* (resp. a-right) *decomposition* of u is the unique factorization $u = u_0 a u_1$ with $a \notin \alpha(u_0)$ (resp. $a \notin \alpha(u_1)$). For a finite group G, we define $\sim^G_{n,k}$ on A^* where $n = |A|$ by induction on $n + k$. First, we have $x \sim^G_{n,0} y$ if and only if x, y are G-equivalent. Next, we let $x \sim^G_{n,k} y$ when and only when:

1. $x \sim^G_{n,k-1} y$;
2. $\alpha(x) = \alpha(y)$;
3. For any $a \in \alpha(x) = \alpha(y)$, if $x = x_0 a x_1$ and $y = y_0 a y_1$ are the a-left decompositions of x and y then $x_0 \sim^G_{n-1,k} y_0$ and $x_1 \sim^G_{n,k-1} y_1$;
4. For any $a \in \alpha(x) = \alpha(y)$, if $x = x_0 a x_1$ and $y = y_0 a y_1$ are the a-right decompositions of x and y then $x_0 \sim^G_{n,k-1} y_0$ and $x_1 \sim^G_{n-1,k} y_1$.

This equivalence relation is well-defined since $|\alpha(x_0)| < |\alpha(x)|$ in (3) and $|\alpha(x_1)| < |\alpha(x)|$ in (4). It is easy to check that $\sim^G_{n,k}$ is in fact a congruence of finite index.

For $a_1, \ldots, a_s \in A$ and $K_0, \ldots, K_s \subseteq A^*$, we say that the concatenation $K = K_0 a_1 K_1 \ldots a_s K_s$ is *unambiguous* if for all $w \in K$ there exists a unique factorization $w = w_0 a_1 \ldots a_s w_s$ with $w_i \in K_i$. The following theorem due to Schützenberger gives a combinatorial description of languages recognized by monoids in $\mathbf{DO} \cap \overline{\mathbf{H}}$:

Theorem 1. *Let $L \subseteq A^*$, then $M(L) \in \mathbf{DO} \cap \overline{\mathbf{H}}$ if and only if L is the disjoint union of languages of unambiguous concatenations $K = K_0 a_1 K_1 \ldots a_s K_s$ where each K_i can be recognized by the direct product of a group of $\mathbf{H}$ and an idempotent and commutative monoid.*

In the appendix, we prove a slight refinement of this theorem by relating these languages with the congruences $\sim_G$; specifically:

Theorem 2. *Let $M = A^*/\gamma$, with $|A| = n$. Then $M \in \mathbf{DO} \cap \overline{\mathbf{H}}$ iff $\sim^G_{n,k} \subseteq \gamma$ for some $k \in \mathbb{N}$ and $G \in \mathbf{H}$.*

Thus, a language is recognized by a monoid in $\mathbf{DO} \cap \overline{\mathbf{H}}$ if and only if it is the disjoint union of $\sim^G_{n,k}$-classes for some $G \in \mathbf{H}$ and some $n, k \in \mathbb{N}$. It is not hard to see that each $\sim^G_{n,k}$ class can be described by an unambiguous concatenation as in the statement of Schützenberger's Theorem.

We will also be interested in the following small subclass of $\mathbf{DO}$.

Definition 1. *We call $\mathbf{W}$ the variety of monoids M satisfying:*

1. *If p is the exponent of M then $exw^p yf = exyf$ for all $w, e, f, x, y \in M$ such that e, f are idempotents lying $\mathcal{J}$-below w. (i.e. M satisfies the identity $(swt)^\omega x w^p y (uwv)^\omega = (swt)^\omega xy (uwv)^\omega$)*
2. *$ewxf = exwf$ for all $w, e, f, x, y \in M$ such that e, f are idempotents lying $\mathcal{J}$-below w. (i.e. M satisfies $(swt)^\omega wx (uwv)^\omega = (swt)^\omega xw (uwv)^\omega$)*

In particular, condition 1 implies $\mathbf{W} \subseteq \mathbf{DO}$ since for all $x, y \in M$:

$$(xy)^\omega (yx)^\omega (xy)^\omega = (xy)^\omega (yx)^{p\omega} (xy)^\omega = (xy)^\omega (xy)^\omega = (xy)^\omega .$$

On the other hand condition 2 implies that any two elements w, x of the maximal subgroup with identity e commute since $xw = exwe = ewxe = wx$ and so $\mathbf{W} \subseteq \overline{\mathbf{Ab}}$. Also from 2 we get $exwyf = ewxyf$ for e, f idempotent $\mathcal{J}$-below w.

For a word $u \in A^*$ and $a \in A$ we denote by $|u|_a$ the number of occurrences of a in u and let $RED_t(u)$ be the unique word of A^* obtained by keeping in u only the first and last t occurrences of each letter a with $|u|_a \geq 2t$ and all occurences of letters a with $|u|_a < 2t$. For example, $RED_2(abcbabbababba) = abcbaabba$. We will show that languages recognized by monoids in $\mathbf{W}$ have a useful combinatorial characterization: we set $u \approx_{t,p} v$ if and only if:

1. $RED_t(u) = RED_t(v)$;
2. For all $a \in A$ we have $|u|_a \equiv |v|_a \pmod{p}$.

Alternatively, we could define $RED_{t,p}(u)$ as the word obtained from u by the following process: For every $a \in \alpha(u)$ with $|u|_a \geq 2t$ mark the first and last t occurrences of a then move all *other* occurrences of a, if any, next to the t^{th} one and then reduce that block of a's modulo p. If $|u|_a < 2t$, all occurrences of a are left untouched. Note that we clearly have $RED_{t,1}(u) = RED_t(u)$ and $u \approx_{t,p} v$ if and only if $RED_{t,p}(u) = RED_{t,p}(v)$.

Theorem 3. *Let $M = A^*/\gamma$, then $M \in \mathbf{W}$ if and only if $\approx_{t,p} \subseteq \gamma$ for some t, p.*

Proof. For one direction, we need to show that $M = A^*/\approx_{t,p}$ satisfies both conditions in the definition of $\mathbf{W}$. We start with the second: Consider the words $q = (uwv)^{tp} wx (ywz)^{tp}$ and $r = (uwv)^{tp} xw (ywz)^{tp}$. For any $a \in A$, $|q|_a \equiv |x|_a + |w|_a \equiv |r|_a \pmod{p}$ and

$$\begin{aligned} RED_t((uwv)^{tp} wx (ywz)^{tp}) &= RED_t((uwv)^{tp} x (ywz)^{tp}) \\ &= RED_t((uwv)^{tp} xw (ywz)^{tp}) \end{aligned}$$

since for any letter a occurring in w, the first t occurrences of a lie in $(uwv)^{tp}$ and its last t occurrences lie in $(ywz)^{tp}$. Thus, $q \approx_{t,p} r$ so M satisfies condition 3. Similarly, $(uwv)^{tp} xw^p y (swz)^{tp} \approx_{t,p} (uwv)^{tp} xy (swz)^{tp}$ and M thus satisfies the first condition.

Conversely, suppose M is in $\mathbf{W}$. We need to show that there exist t, p such that for any morphism $\phi : A^* \to M$ we have $\phi(q) = \phi(r)$ for any $q \approx_{t,p} r$ and it is in fact sufficient to establish $\phi(q) = \phi(RED_{t,p}(q))$ since $RED_{t,p}(q) = RED_{t,p}(r)$. In particular, we choose p as the exponent of M and t as $|M| + 1$.

Recall that to obtain $RED_{t,p}(q)$, one successively considers all $a \in A$ with $|q|_a \geq 2t$, "groups" together the "middle" a's and reduces their number modulo p. We will show that the image under ϕ is preserved by this process. Consider a word $u = u_0 a u_1 a \ldots a u_t$ with at least t occurences of a. Since $t = |M| + 1$, there must exist $1 \leq i < j \leq t$ such that $s_i = \phi(u_0 a u_1 a \ldots u_i) = \phi(u_0 a u_1 a \ldots u_j) =$

s_j. This means that $s_i = s_i\phi(au_{i+1}a \ldots au_j) = s_i\phi(au_{i+1}a \ldots au_j)^\omega$. Therefore there exist $g, h \in M$ such that $\phi(u)$ can be written as geh where $e = \phi(au_{i+1}a \ldots au_j)^\omega$ is an idempotent lying $\mathcal{J}$-below a.

Suppose now that q contains at least $2t+1$ occurrences of a. We can thus factor q as $q = u_0a \ldots u_{t-1}axayav_{t-1}a \ldots av_0$ where the u_i's and v_i's do not contain a. From the remarks of the preceding paragraph, we can now use condition 2 in Definition 1 to obtain $\phi(q) = \phi(u_0a \ldots u_{t-1}a\,a\,xy\,av_{t-1}a \ldots av_0)$. Repeating the same process for all occurrences of a in x or y we can get $\phi(q) = \phi(u_0a \ldots u_{t-1}a\,a^{kp+d}z\,av_{t-1}a \ldots av_0)$ where $a \notin \alpha(z)$ and from condition 1: $\phi(q) = \phi(u_0a \ldots u_{t-1}a\,a^dz\,av_{t-1}a \ldots av_0)$ where $0 \leq d < p$ is such that $|q|_a - 2t \equiv d \pmod{p}$. If the same manipulation is made for every $a \in A$, we obtain $\phi(q) = \phi(RED_{t,p}(q))$ as we needed. □

3 Communication Complexity

We only provide here a short introduction to three communication complexity models but refer the reader to the excellent book of Kushilevitz and Nisan [5] for further details and formal definitions.

In the deterministic model, two players, Alice and Bob, wish to compute a function $f : S^{n_A} \times S^{n_B} \to T$ where S and T are finite sets. Alice is given $x \in S^{n_A}$ and Bob $y \in S^{n_B}$ and they collaborate in order to obtain $f(x, y)$ by exchanging bits using a common blackboard according to some predetermined *communication protocol* $\mathcal{P}$. This protocol determines whose turn it is to write, furthermore what a player writes is a function of that player's input and the information exchanged thus far. When the protocol ends, its output $\mathcal{P}(x, y) \in T$ is a function of the blackboard's content. We say that $\mathcal{P}$ computes f is $\mathcal{P}(x, y) = f(x, y)$ for all x, y and define the *cost* of $\mathcal{P}$ as the maximum number of bits exchanged for any input. The *deterministic communication complexity* of f, denoted $D(f)$ is the cost of the cheapest protocol computing f. We will be interested in the complexity of functions $f : S^* \times S^* \to T$ and will thus consider $D(f)$ as a function from $\mathbb{N} \times \mathbb{N}$ to $\mathbb{N}$ and study its asymptotic behavior.

In a *simultaneous protocol* $\mathcal{P}$, we disallow any interaction between Alice and Bob: Each of them *simultaneously* sends a message to a trusted referee which has access to none of the input and the referee produces the output $\mathcal{P}(x, y) \in T$. We denote $D^{\|}(f)$ the simultaneous communication complexity of f, i.e. the cost of the cheapest simultaneous protocol computing f.

In a *probabilistic communication protocol* $\mathcal{P}$, Alice and Bob have access to private random bits which determine their behavior. The protocol is said to compute f if for all x, y, the probability over the choices of these random bits that $\mathcal{P}(x, y) = f(x, y)$ is at least $3/4$. We denote $R(f)$ the randomized communication complexity of f.

All three of these models have been extensively studied. We will use classical communication complexity lower bounds for the following functions [5]:

- For $x, y \in \{0,1\}^n$, we define Disjointness as: $DISJ(x, y) = 1$ if and only if $\bigvee_{1 \leq i \leq n} x_i y_i = 0$;

– For $x, y \in \{0,1\}^n$, and any $q \in \mathbb{N}$ we define Inner Product $\pmod q$ as: $IP_q(x,y) = 1$ if and only if $\sum_{1 \le i \le n} x_i y_i \equiv 0 \pmod q$;
– For two n-bit numbers $x, y \in [2^n]$ we define Greater Than as: $GT(x,y) = 1$ if and only if $x \ge y$.
– For $x \in \{0,1\}^n$ and a $\log n$-bit number $p \in [n]$ we define $INDEX(x,p) = x_p$;

The known communication complexity bounds for these problems can be summed up in the following table:

Theorem 4.

	D	R	$D^{\parallel}$
$DISJ$	$\Theta(n)$	$\Theta(n)$	$\Theta(n)$
IP_q	$\Theta(n)$	$\Theta(n)$	$\Theta(n)$
GT	$\Theta(n)$	$\Theta(\log n)$	$\Theta(n)$
$INDEX$	$\Theta(\log n)$	$\Theta(\log n)$	$\Theta(n)$

It should be noted that non-trivial work is necessary to prove the GT randomized upper bound and the $DISJ$ and IP randomized lower bounds.

In general, we want to study the communication complexity of functions which do not explicitly have two inputs. In the case of regular languages and monoids we will use a form of *worst-case partition* definition. Formally, the deterministic (resp. randomized, simultaneous) *communication complexity of a finite monoid* M is the communication complexity of evaluating in M the product $m_1 \cdot m_2 \cdot \ldots \cdot m_{2n}$ where the odd-indexed $m_i \in M$ are known to Alice and the even-indexed m_i are known to Bob.

Similarly, the *communication complexity of a regular language* $L \subseteq A^*$ is the communication complexity of the following problem: Alice and Bob respectively receive $a_1, a_3, \ldots a_{2n-1}$ and $a_2, a_4, \ldots, a_{2n}$ where each a_i is either an element of A or the empty word[1] ϵ and they want to determine whether $a_1 a_2 \ldots a_{2n}$ belongs to L.

The following basic facts from [8] are proved for completeness in the appendix:

Lemma 3. *Let $L \subseteq A^*$ be regular with $M(L) = M$. We have $D(M) = \Theta(D(L))$ and similarly for $D^{\parallel}$ and R.*

In particular the deterministic (resp. simultaneous, randomized) complexity of a monoid M is, up to a constant, the maximal communication complexity of any regular language that it can recognize.

Lemma 4. *For any increasing $f : \mathbb{N} \to \mathbb{N}$ the class of monoids such that $D(M)$ (resp. $D^{\parallel}(M)$, $R(M)$) is $O(f)$ forms a variety.*

In order to compare the communication complexity of two languages K, L in different models, Babai, Frankl and Simon [1] defined *rectangular reductions* from K to L which are, intuitively, reductions which can be computed privately by Alice and Bob without any communication cost. We give here a form of this definition which specifically suits our needs:

[1] See comment in appendix

Definition 2. *Let $f : \{0,1\}^{n_A} \times \{0,1\}^{n_B} \to \{0,1\}$ and M some finite monoid. A rectangular reduction of length t from f to M is a sequence of $2t$ functions $a_1, b_2, a_3, \ldots, a_{2t-1}, b_{2t}$, with $a_i : \{0,1\}^{n_A} \to M$ and $b_i : \{0,1\}^{n_B} \to M$ and such that for every $x \in \{0,1\}^{n_A}$ and $y \in \{0,1\}^{n_B}$ we have $f(x,y) = 1$ if and only if $eval_M(a_1(x)b_2(y) \ldots b_{2t}(y)) \in T$ for some target subset T of M.*

Such a reduction transforms an input (x,y) of the function f into a sequence of $2t$ monoid elements $m_1, \ldots, m_{2t}$ where the odd-indexed m_i are obtained as a function of x only and the even-indexed m_i are a function of y.

It should be clear that if f has communication complexity $\Omega(g(n_A, n_B))$ and has for any n_A, n_B a reduction of length $t(n_A, n_B)$ to M then M has complexity $\Omega(g(t^{-1}(n)))$.

We will write $f \leq_r^t M$ to indicate that f has a rectangular reduction of length t to M and will drop the t superscript whenever $t = O(n)$.

4 Communication Complexity Bounds for Regular Languages and Monoids

Our goal is to prove tight bounds on the communication complexity of any regular language or, equivalently using Lemma 3, of any finite monoid in the deterministic, randomized and simultaneous models. We start with a number of useful upper bounds, the first of which is an easy observation due to [8].

Lemma 5. *If M is commutative then $D^{||}(M) = O(1)$.*

Lemma 6. *Let $L \subseteq A^*$ be such that $M(L) \in \mathbf{DO} \cap \overline{\mathbf{Ab}}$. Then $D(L) = O(\log n)$.*

Proof. By Lemma 2, L is a union of $\sim^G_{|A|,k}$-classes for some Abelian group G. We claim that any such class has logarithmic communication complexity and argue by induction on $t = |A| + k$. For $t = 1$ there is nothing to prove. For $t > 1$, let $u \in A^*$ be some predetermined representative of the class. Given the input $x = x_1 x_2 \ldots x_{2n} \in A^*$, Alice and Bob can check whether $\alpha(x) = \alpha(u)$ by exchanging $|A| + 1$ bits. Next, they need to verify that u and x are G-equivalent. It is well known that for any Abelian group G of exponent p, we get u and x G-equivalent if and only if $|u|_a \equiv |x|_a \pmod{p}$ for all $a \in A$. The latter condition can easily be verified with communication cost about $|A|\lceil \log p \rceil$, a constant. Let $u = vaw$ be the a-left decomposition of u and i, j denote the locations of the leftmost occurrence of a in x that is seen respectively by Alice and Bob. These indices can be exchanged at logarithmic communication cost so that if e.g. i is smaller than j Alice and Bob can conclude that $x = x_1 \ldots x_{i-1} a x_{i+1} \ldots x_n$ is the a-left decomposition of x and, by induction, can verify that $x_1 \ldots x_{i-1} \sim^G_{|A|-1,k} v$ and $x_{i+1} \ldots x_n \sim^G_{|A|,k-1} w$ using only $O(\log n)$ communication. Left-right symmetry completes the proof. ◻

Lemma 7. *Let $L \subseteq A^*$ be such that $M(L) \in \mathbf{W}$. Then $D^{||}(L) = O(\log n)$ and $R(L) = O(\log \log n)$.*

Proof. As in the previous proof, we obtain these upper bounds for the $\approx_{t,p}$ classes. Let u be some representative of the target class and x the common input of Alice and Bob. Checking whether $|u|_a \equiv |x|_a \pmod p$ is easily done at constant cost so we only need to show that verifying $RED_t(x) = RED_t(u)$ can be done efficiently. For the simultaneous case, the players send to the referee the locations of the first and last t occurrences that they see of each letter $a \in A$. Given this information, the referee can reconstruct $RED_t(x)$ and compare it to $RED_t(u)$.

For the probabilistic case we use a subprotocol of cost $O(\log\log n)$ to determine for any $k \leq t$ which of Alice or Bob holds the k^{th} (or symmetrically the k^{th} to last) occurrence of some letter a in x, provided of course that $|x|_a \geq k$. We argue by induction on k: For $k = 1$, let i, j be the positions of the first occurrence of a seen by Alice and Bob respectively. Of course Alice holds the first occurrence of a if and only if $i < j$ and by Theorem 4 this can be tested by a randomized protocol at cost $O(\log\log n)$ since i, j are only $\log n$-bits long. For $k > 1$ we can assume from induction that Alice and Bob have marked, in their respective inputs, the occurrences of a which are among the first $k-1$ of a in x. The k^{th} occurrence must be either the first unmarked a that Alice sees or the first unmarked a that Bob sees, whichever comes first in x. Once again, Alice and Bob are left with comparing two $\log n$-bit numbers and apply the $O(\log\log n)$ cost protocol.

For $i, j \leq t$, the i^{th} occurrence of a in x comes before the j^{th} occurrence of b in x if and only if the i^{th} occurrence of a in $RED_t(x)$ comes before the j^{th} occurrence of b in $RED_t(x)$. This means that Alice and Bob can check $RED_t(x) = RED_t(u)$ by verifying that for all $i, j \leq t$ and all $a, b \in A$ the i^{th} occurrence of a precedes the j^{th} occurrence of b in $RED_t(u)$ if and only if the i^{th} occurrence of a precedes the j^{th} occurrence of b in x. Since they can determine which of them holds these occurrences, they can check precedence either privately (when one player holds both occurrences) or by using once more the $O(\log\log n)$ randomized protocol to compare two $\log n$ bit numbers.

It should be noted that in any event, the GT protocol is used only a constant number of times (depending on t and $|A|$) so we need not worry about the dwindling of the overall probability of correctness in the protocol. □

Lemma 8. *1. If M is non-commutative then $GT \leq_r^{2^n} M$ (Notice that the reduction has exponential length);*
2. If M is not in **DS** *then $DISJ \leq_r M$.*
3. If M lies in **DS** *but is not in* **DO** *then $IP_q \leq_r M$ for some integer q.*
4. If G is a non-commutative group then $IP_q \leq_r G$ for some integer q.
5. If M is in **DO** *but not in* **W** *then $INDEX \leq_r M$.*

Proof. Since 2 and 4 are already proved explicitly in [8], we sketch their proofs in the Appendix for completeness.
1- Let $a, b \in M$ be such that $ab \neq ba$. We obtain an exponential length reduction from $GT(x, y)$ by building $m_1 m_2 \ldots m_{2^{n+1}}$ where $m_i = a$ for $i = m_{2x}$; $m_i = b$ for $i = m_{2y-1}$ and $m_i = 1_M$ otherwise. The product of the m_i is then ba if and

only if $x \geq y$ and is ab otherwise.
3- Since M is in **DS** but not **DO**, there must exist two $\mathcal{J}$-related idempotents e, f such that ef is not idempotent. Recall from Lemma 1 that we thus have $efe \neq e = (ef)^\omega e$. Let q be minimal such that $(ef)^q e = e$: The reduction produces elements $m_1 \dots m_{2n}$ where $m_{2i-1} = e$ if $x_i = 1$ and $m_{2i-1} = e(ef)^\omega$ otherwise and $m_{2i} = fe$ if $y_i = 1$ and $m_{2i} = (ef)^\omega e$ otherwise. In particular the product $m_{2i-1}m_{2i}$ is efe if and only if $x_i = y_i = 1$ and is e otherwise and so the product $m_1 \dots m_{2n}$ equals $(ef)^{\sum_{1 \leq i \leq n} x_i y_i \equiv 0 \pmod q} e$ which equals e if and only if $IP_q(x, y) = 1$.
5- Suppose first that there exist e, f, u, v, w with e, f idempotent and $\mathcal{J}$-below w such that $euw^p vf \neq euvf$. Since we assume membership in **DO**, we can use Lemma 2 to show that $ew^p e = e$ and $fw^p f = f$. We obtain a rectangular reduction from $INDEX(x, s)$ by creating $m = m_1 m_2 \dots m_{2n+1}$ as follows:

$$m_i = \begin{cases} e & \text{for } i = 1, 3, \dots 2s-3 \text{ (the first } s-1 \text{ odd-indexed } m_i\text{'s);} \\ (eu) & \text{for } i = 2s-1; \\ (vf) & \text{for } i = 2s+1; \\ f & \text{for } i = 2s+3 \dots 2n+1 \text{ (all other odd indexed } m_i\text{'s);} \\ 1_M & \text{for } i = 2j \text{ and } x_j = 0; \\ w^p & \text{for } i = 2j \text{ and } x_j = 1. \end{cases}$$

The values of the odd indexed and even indexed m_i depend respectively on s and x as required and one can see that since $ew^p e$ is e and $fw^p f$ is f, the product of the m_i's is equal to $euvf$ when x_s is 0 and $euw^p vf$ when x_s is 1 which shows the correctness of the reduction.

Suppose now that there exist $e, f, u, w \in M$ with e, f idempotent $\mathcal{J}$-below w but $ewuf \neq euwf$. On the other hand, we can assume that M satisfies the first condition and consequently that $ewuf = ew^{p+1}uf$ and $euf = ew^p uf = euw^p f$. We obtain a rectangular reduction from $INDEX(x, s)$ (assuming w.l.o.g. that $x_n = 1$) as follows:

$$m_i = \begin{cases} e & \text{for } i = 1 \\ f & \text{for } i = 2n+1 \\ u & \text{for } i = 2s+1; \\ 1_M & \text{for all other odd-indexed } i; \\ w^p & \text{for even } i = 2j \text{ such that } x_j = x_{j-1}; \\ w & \text{for even } i = 2j \text{ such that } x_j = 1 \text{ and } x_{j-1} = 0; \\ w^{p-1} & \text{for even } i = 2j \text{ such that } x_j = 0 \text{ and } x_{j-1} = 1. \end{cases}$$

Again, the values of the odd indexed and even indexed m_i's depend respectively on s and x. The value of the even indexed m_i's are such that the product $m_2 m_4 \dots m_{2i}$ is w^{kp+1}, for some k, if and only if x_i is 1 and w^{kp} if x_i is 0. Similarly, using the fact that $x_n = 1$, the product $m_{2i+1} \dots m_{2n}$ is $w^{k'p}$ if and only if $x_i = 1$ and $w^{k'p+1}$ otherwise. Using the values assigned by the reduction to the odd-indexed m_i's we have

$$m_1 m_2 \dots m_{2n+1} = em_2 m_4 \dots m_{2s} u m_{2s+1} \dots m_{2n} f$$

and by our previous remark this is $ew^{kp+1}uw^{k'p}f = ewuf$ if $x_s = 1$ and $ew^{kp}uw^{k'p+1}f = euwf$ if $x_s = 0$ so our reduction is correct. $\square$

In particular, if M does not lie in $\mathbf{DO} \cap \overline{\mathbf{Ab}}$ then it either admits a linear length reduction from $DISJ$ (if it is outside $\mathbf{DS}$) or from IP_q for some q (if it is either in $\mathbf{DS}$ but not in $\mathbf{DO}$ or if it is outside $\overline{\mathbf{Ab}}$). Combining our last lemma with the upper bounds above we can obtain the three following theorems.

Theorem 5. *Let $L \subseteq A^*$ be a regular language with $M = M(L)$. Then*

$$D(L) = \begin{cases} O(1) & \text{if and only if } M \text{ is commutative;} \\ \Theta(\log n) & \text{if and only if } M \text{ is in } \mathbf{DO} \cap \overline{\mathbf{Ab}} \text{ but not commutative;} \\ \Theta(n) & \text{otherwise.} \end{cases}$$

Proof. We know $D(L) = O(1)$ if M is commutative and $D(L) = \Omega(\log n)$ otherwise since in that case $GT \leq_r^{2^n} M$. Lemma 6 gives the upper bound when $M \in \mathbf{DO} \cap \overline{\mathbf{Ab}}$. Finally, when M is not in $\mathbf{DO} \cap \overline{\mathbf{Ab}}$, then it admits a linear length reduction from $DISJ$ or IP_q which yields the last $\Omega(n)$ lower bound. $\square$

Theorem 6. *Let $L \subseteq A^*$ be a regular language with $M = M(L)$. Then*

$$R(L) = \begin{cases} O(1) & \text{if and only if } M \text{ is commutative;} \\ \Theta(\log\log n) & \text{if and only if } M \text{ is in } \mathbf{W} \text{ but not commutative;} \\ \Theta(\log n) & \text{if and only if } M \text{ is in } \mathbf{DO} \cap \overline{\mathbf{Ab}} \text{ but not in } \mathbf{W}; \\ \Theta(n) & \text{otherwise.} \end{cases}$$

Proof. When M is in $\mathbf{W}$ but not commutative we put together Lemma 7 and part 1 of Lemma 8 to get the tight $\log\log n$ bound. Similarly, if M in $\mathbf{DO} \cap \overline{\mathbf{Ab}}$ but not in $\mathbf{W}$ then it admits a reduction from $INDEX$ which proves the $\Omega(\log n)$ lower bound matching Lemma 6 and when M is not in $\mathbf{DO} \cap \overline{\mathbf{Ab}}$ we again use the linear lower bounds on the probabilistic complexity of $DISJ$ and IP_q. $\square$

Theorem 7. *Let $L \subseteq A^*$ be a regular language with $M = M(L)$. Then*

$$D^{\|}(L) = \begin{cases} O(1) & \text{if and only if } M \text{ is commutative;} \\ \Theta(\log n) & \text{if and only if } M \text{ is in } \mathbf{W} \text{ but not commutative;} \\ \Theta(n) & \text{otherwise.} \end{cases}$$

Proof. When M is in $\mathbf{W}$ but not commutative we combine the upper bound of Lemma 7 with the lower bound obtained from part 1 of Lemma 8. When M is not in $\mathbf{W}$ then it admits a linear length reduction from $INDEX$, $DISJ$ or IP_q which all have $\Omega(n)$ simultaneous complexity. $\square$

5 Concluding Remarks

We have proved that the communication complexity of a regular language in the deterministic, simultaneous and probabilistic two-party models can be decided by simple and decidable algebraic properties of its syntactic monoid and that, in each of the three cases, this classification only defines a small number of classes. Because regular languages are so well-behaved in this context, they are a natural starting point for further investigations in communication complexity. For instance, it is natural to ask whether such gaps in the complexity of regular languages also occur in other communication models. Recently, we have considered the MOD_p-counting communication complexity model (see e.g. [4] for definitions), where a similar constant to logarithmic to linear trichotomy seems to hold, and the non-deterministic case. The latter requires some extra algebraic machinery since we need to deal with ordered monoids and positive varieties to reflect the fact that a language might have a non-deterministic complexity exponentially larger than that of its complement. There is also ongoing work concerning the multiparty (number in the forehead) model where one might still hope to prove a generalization of the theorem of Szegedy mentioned in our introduction. Some partial results along those lines appear in [8] and [7].

References

1. L. Babai, P. Frankl, and J. Simon. Complexity classes in communication complexity theory. In *Proc. 27^{th} IEEE FOCS*, pages 337–347, 1986.
2. D. A. Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1. *J. Comput. Syst. Sci.*, 38(1):150–164, 1989. Preliminary version appeared in STOC'86.
3. D. A. M. Barrington and D. Thérien. Finite monoids and the fine structure of NC^1. *Journal of the ACM*, 35(4):941–952, Oct. 1988.
4. C. Damm, , M. Krause, C. Meinel, and S. Waack. On relations between counting communication complexity classes. In *9th Annual Symp. on Theor. Asp. of Comp. Sci. (STACS'92)*, 1992. Full version to appear in Journal of Computer and Systems Sciences. Currently available e.g. from www.num.math.uni-goettingen.de/damm/.
5. E. Kushilevitz and N. Nisan. *Communication Complexity.* Cambridge University Press, 1997.
6. J.-E. Pin. *Varieties of formal languages.* North Oxford Academic Publishers Ltd, London, 1986.
7. P. Pudlák. An application of Hindman's theorem to a problem on communication complexity. Draft, 2002.
8. J.-F. Raymond, P. Tesson, and D. Thérien. An algebraic approach to communication complexity. *Lecture Notes in Computer Science (ICALP'98)*, 1443:29–40, 1998.
9. M. Szegedy. Functions with bounded symmetric communication complexity, programs over commutative monoids, and ACC. *J. Comput. Syst. Sci.*, 47(3):405–423, 1993.

Appendix

Algebra

We prove two lemmas about the algebraic properties of monoid in **DS** and **DO** whose proofs could be simplified by introducing some more background on monoids which we prefer to avoid.

Lemma 2. *Let $M \in \mathbf{DS}$ then*

- *If e, f are $\mathcal{J}$-related idempotents then $(ef)^\omega e = e$ (but $efe = e$ if and only if ef is idempotent).*
- *for any $a, b \in M$ such that $a \leq_\mathcal{J} b$ with a regular then $a \mathrel{\mathcal{J}} ab \mathrel{\mathcal{J}} ba$.*
- *for any $x, y \in M$ such that $xy \mathrel{\mathcal{R}} x$ we have in fact $R_x y \subseteq R_x$.*

Proof. First it is clear that e is both $\mathcal{L}$ and $\mathcal{R}$ above $(ef)^\omega e$ and $\mathcal{J}$-related to it since the $\mathcal{J}$-class of e and f is closed under product. Thus, $(ef)^\omega e \mathrel{\mathcal{H}} e$. Noting that $(ef)^\omega e$ is itself idempotent, we conclude equality since $\mathcal{H}_e$ has a unique idempotent. However, if $efe = e$ then $efef = ef$ and ef is idempotent and conversely if ef is idempotent then so is efe and since $efe\mathcal{H}e$ then $efe = e$.

Suppose $a = sbt$. Then we have

$$a \mathrel{\mathcal{J}} a^\omega = (sbt)^\omega \mathrel{\mathcal{J}} sb(sbt)^\omega \leq_\mathcal{J} ba^\omega \leq_\mathcal{J} ba \leq_\mathcal{J} a$$

so $a \mathrel{\mathcal{J}} ba$ and by symmetry $a \mathrel{\mathcal{J}} ab$.

Suppose now $xy \mathrel{\mathcal{R}} x$. We let u be such that $xyu = x$ and call e the idempotent $(yu)^\omega$. Suppose $z \mathrel{\mathcal{R}} x$, that is there are s, t such that $xs = z$ and $zt = x$. Then $z(set) = z$ and $z(set)^\omega = z$. By part 1, $(set)^\omega \mathrel{\mathcal{R}} (set)^\omega y$ so $z = z(set)^\omega \mathrel{\mathcal{R}} z(set)^\omega y = zy$. Hence, $R_x y \subseteq R_x$. □

Lemma 3. *Suppose $M \in \mathbf{DO}$ and let $w_1, w_2 \in A^*$ be G-equivalent for any subgroup G of M. For any morphism $\phi : A^* \to M$ and any $x \in M$ such that $x\phi(w_1) \mathrel{\mathcal{H}} x\phi(w_2) \mathrel{\mathcal{R}} x$ we have in fact $x\phi(w_1) = x\phi(w_2)$.*

Proof. We first observe that for any idempotent $e \in M$ and any $u, v \in M$ lying $\mathcal{J}$-above e we have $(eu)^\omega e(ve)^\omega = (eu)^\omega (ve)^\omega$ since $e, (eu)^\omega$ and $(ve)^\omega$ are $\mathcal{J}$-related idempotents. Similarly, if f is another idempotent $\mathcal{J}$-related to e we have $exyf = exfeyf$.

Let $\phi(w_1) = y_1$ and $\phi(w_2) = y_2$. Since $xy_1 \mathrel{\mathcal{H}} xy_2 \mathrel{\mathcal{R}} x$ there must exist $s, t \in M$ lying $\mathcal{J}$-below y_1 and y_2 with $xst = x = x(st)^\omega$ and $xy_i ts = xy_i = xy_i(ts)^\omega$. Let T be the submonoid of elements of M lying $\mathcal{J}$-above $(st)^\omega$ and $\psi : T \to H_{(st)^\omega (ts)^\omega}$ given by $\psi(x) = (st)^\omega x (ts)^\omega$. Using our earlier observation we conclude that ψ is a morphism since

$$\psi(xy) = (st)^\omega xy(ts)^\omega = (st)^\omega x(ts)^\omega (st)^\omega y(ts)^\omega = \psi(x)\psi(y).$$

Since w_1 and w_2 are equivalent with respect to the group $H_{(st)^\omega(ts)^\omega}$, we must have $(st)^\omega y_1(ts)^\omega = \psi(y_1) = \psi(y_2) = (st)^\omega y_2(ts)^\omega$. Thus,

$$xy_1 = x(st)^\omega y_1(ts)^\omega = x(st)^\omega y_2(ts)^\omega = xy_2.$$

$\square$

Theorem 8. *2 Let $M = A^*/\gamma$, with $|A| = n$. Then $M \in \mathbf{DO} \cap \overline{\mathbf{H}}$ iff $\sim^G_{n,k} \subseteq \gamma$ for some $k \in \mathbb{N}$ and $G \in \mathbf{H}$.*

Proof. For one direction, we need to show that $A^*/\sim^G_{n,k}$ is in $\mathbf{DO} \cap \overline{\mathbf{H}}$ for any integer k and any group $G \in \mathbf{H}$. Showing membership in $\overline{\mathbf{H}}$ however requires the introduction of more algebraic background so we will appeal to the theorem of Schützenberger which we previously cited as Theorem 1: It is an easy exercise to verify that each $\sim^G_{n,k}$-class can be described by an unambiguous concatenation $K_0a_1 \dots a_sK_s$ where the K_i can be recognized by the direct product of an idempotent and commutative monoid (to verify whether $\alpha(x) = \alpha(y)$) and a group in $\mathbf{H}$ (to verify that x, y are G-equivalent). Schützenberger's Theorem thus insures that any disjoint union of these classes forms a language whose syntactic monoid lies in $\mathbf{DO} \cap \overline{\mathbf{H}}$.

For the second part of our proof, let us denote as $[u]$, for any $u \in A^*$, the γ equivalence class of u. We define the $\mathcal{R}$-decomposition (with respect to the congruence γ) of a string $u \in A^*$ as the unique factorization $u = u_0a_1u_1 \dots a_tu_t$, with $a_i \in A$ and $u_i \in A^*$ such that:

1. $[u_0a_1 \dots a_su_s] \not\mathcal{R} [u_0a_1 \dots a_su_sa_{s+1}]$ for any $s < t$.
2. $[u_0a_1 \dots a_s] \mathcal{R} [u_0a_1 \dots a_su_s]$ for any $s \leq t$.
3. $[u_0] \mathcal{H} 1_M$.

Because M lies in $\mathbf{DS}$ we know by Lemma 1 that $a_i \notin \alpha(u_{i-1})$. In particular, $u = u_0\, a_1\, u_1a_2 \dots a_tu_t$ is the a_1-left-decomposition of u. Symmetrically, we can define $\mathcal{L}$-decompositions of strings, which will relate to right-decompositions.

Let now k be the maximum of the number of $\mathcal{R}$-classes and the number of $\mathcal{L}$-classes of M. Let $B \subseteq A$ be a sub-alphabet with $|B| = m$ and suppose u, v are strings in B^*. We claim that if $u \sim^G_{m,mk} v$ then $[u] = [v]$ and prove this by induction on m.

For $m = 0$, the claim is trivially true. For the inductive step, assume now $m \geq 1$ and suppose $u \sim^G_{m,mk} v$, with $\alpha(u) = \alpha(v) \neq \emptyset$. Let $u = u_0a_1u_1 \dots a_tu_t$ be the $\mathcal{R}$-decomposition of u with $t \leq k$. We write w_i for $u_ia_{i+1} \dots a_tu_t$ and, by our earlier remark, have $w_i = u_ia_{i+1}w_{i+1}$ is the a_{i+1}-left-decomposition of w_i, so there must be a decomposition $v = v_0a_1 \dots a_tv_t$ such that $u_i \sim^G_{m-1,mk-i} v_i$ for $i < t$ (Note that this actually implies $u_i \sim^G_{m-1,(m-1)k} v_i$). By the induction hypothesis, we have $[u_i] = [v_i]$ for all $i < t$ so $[u] \mathcal{R} [u_0a_1 \dots u_{t-1}a_t] = [v_0a_1 \dots v_{t-1}a_t] \geq_{\mathcal{R}} [v]$. Symmetrically, we get $[v] \geq_{\mathcal{R}} [u]$ and thus $[u] \mathcal{R} [v]$. With the symmetric argument, we can also establish $[u] \mathcal{L} [v]$ so we have $[u] \mathcal{H} [v]$.

It remains to show that in fact $[u] = [v]$. Note that by definition of $\sim^G$, we have that u_t and v_t are G-equivalent. So by Lemma 2 we obtain,

$$[u] = [u_0a_1 \dots u_{t-1}a_t][u_t] = [u_0a_1 \dots u_{t-1}a_t][v_t] = [v_0a_1 \dots v_{t-1}a_t][v_t] = [v]$$

□

Communication Complexity

In defining the communication complexity of regular languages, we allowed the a_i's given to Alice and Bob to be the empty word ϵ. This might seem at first less natural than a simple worst-case partition definition but it is in fact more sound since it limits the amount of information that the players might gain from seeing letter a_j in position i. Consider for example the set L of words starting with a. In the worst-case partition framework, membership in L can always be determined by the player which has access to the first letter in the input without any communication, making the whole task trivial and uninteresting. In our alternative definition, however, testing membership in L boils down to determining which of the players first see a "real" letter (not an ϵ) and whether that letter is an a. This requires $\Theta(\log n)$ bits of communication deterministically. These issues are discussed in some more detail in [8].

We now sketch proofs for two basic Lemmas about the communication complexity of regular languages and monoids.

Lemma 4. *3 Let $L \subseteq A^*$ be regular with $M(L) = M$. We have $D(M) = \Theta(D(L))$ and similarly for $D^{\|}$ and R.*

Proof. It is quite clear that the recognizing morphism $\phi : A^* \to M$ allows us to build a rectangular reduction from L to M of linear length which shows that $D(L) = O(D(M))$.

Perhaps not as clear is the converse. The syntactic congruence of L on A^* is given by $x \equiv_L y$ if for all $u, v \in A^*$ we have $uxy \in L$ if and only if $uyv \in L$, and, by definition, $M(L) = A^*/\equiv_L$. There must exist a k such that every element $m \in M$ has a minimal representative $[m] \in A^*$ of length at most k. It is easy to verify that the product $m_1 m_2 \ldots m_{2n}$ is equal to m if and only if for all $u, v \in A^{\leq k}$ we have $u[m]v \in A$ if and only if $u[m_1][m_2]\ldots[m_{2n}]v \in A^*$. This observation can be used to obtain a rectangular reduction of linear length from M to L.

Lemma 5. *For any increasing $f : \mathbb{N} \to \mathbb{N}$ the class of monoids such that $D(M)$ (resp. $D^{\|}(M)$, $R(M)$) is $O(f)$ forms a variety.*

Proof. Suppose M and T are monoids of communication complexity $O(f)$. Then $M \times N$ has complexity $O(f)$ since evaluating the product $(m_1, t_1)\ldots(m_{2n}, t_{2n})$ can be done by computing separately $m_1 \ldots m_{2n}$ and $t_1 \ldots t_{2n}$.

If T is a submonoid of M then any protocol for M is a protocol for T and if $T = \phi(M)$ then $t_1 \ldots t_{2n} = \phi(\phi^{-1}(t_1)\ldots\phi^{-1}(t_{2n}))$ where $\phi^{-1}(t)$ denotes some arbitrary preimage of t. In both cases the complexity of T is bounded by the one of M.

Lemma 9. *2 If M is not in* **DS** *then $DISJ \leq_r M$.*

4 If G is a non-Abelian group then $IP_q \leq_r G$ for some integer q.

Proof. 2- It can be shown (see e.g. [6]) that if M is not in **DS** then it is divided by the syntactic monoid of $(ab)^*$, called BA_2 or by U the syntactic monoid of (A^*aaA^*). The monoid BA_2 has six elements $\{1_M, a, b, ab, ba, 0\}$ and its multiplication is given by $m0 = 0m = 0$ for all $m \in BA_2$, $aa = bb = 0$, $bab = b$ and $aba = a$. In this case, the reduction from $DISJ$ builds for every pair x_i, y_i a four-tuple $m_{4i-3} \dots m_{4i}$ where $m_{4i-3} = a$ and $m_{4i-1} = b$ when $x_i = 1$ and $m_{4i-3} = m_{4i-1} = 1_M$ when $x_i = 0$, $m_{4i-2} = ab$ when $y_i = 1$ but $m_{4i-2} = m_{4i} = 1_M$ when $y_i = 0$ and $m_{4i} = ab$ for any input. One can check that any such four-tuple evaluates to 0 when $x_i = y_i = 1$ and to ab otherwise so the product of all of them is 0 if $x_i = y_i = 1$ for some i and is ab otherwise.

The elements of U are also $\{1_M, a, b, ab, ba, 0\}$ but its multiplication is slightly different since $b^2 = b$. Nevertheless one can verify that the reduction described for BA_2 also constitutes a reduction to U.

4- If G is not Abelian, there must exist $g, h \in G$ such that the commutator $[g, h] = g^{-1}h^{-1}gh$ is not the identity and thus has order $q \neq 1$. We obtain a reduction from IP_q by creating for each pair x_i, y_i a four-tuple of monoid elements $m_{4i-3}m_{4i-2}m_{4i-1}m_{4i}$ where $m_{4i-3} = g^{-1}$ and $m_{4i-1} = g$ when $x_i = 1$ and $m_{4i-3} = m_{4i-1} = 1_G$ when $x_i = 0$ and where $m_{4i-2} = h^{-1}$ and $m_{4i} = h$ when $y_i = 1$ and $m_{4i-2} = m_{4i} = 1_G$ when $y_i = 0$. This four-tuple thus evaluates to $[g, h]$ if and only if $x_i = y_i = 1$ and to 1_G otherwise and the product of all such tuples is $[g, h]^{\sum_{1 \leq i \leq n} x_i y_i \pmod q}$. □