Noname manuscript No.
(will be inserted by the editor)

Complete Classifications for the Communication Complexity of Regular Languages

Pascal Tesson, Denis Thérien*

School of Computer Science, McGill University
3480 rue University, Montréal, Québec,
H3A 2A7, Canada.
e-mail: {ptesso,denis}@cs.mcgill.ca

The date of receipt and acceptance will be inserted by the editor

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 similar classifications for the communication complexity of regular languages for the simultaneous, probabilistic, simultaneous probabilistic and Mod_p-counting models of communication.

1 Introduction

Communication complexity has grown to be one of the most ubiquitous tools of theoretical computer science [10] and the work pioneered by Babai, Frankl and Simon in [2] (see also [7]) 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 [2]) for every level of PH^{cc}, PSPACE^{cc} (using a result of [5]) and for $\mathrm{MOD_qP}^{cc}$ (see [7]). Moreover, a majority of the most well-studied examples in the field such as 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 $\{01 \cup 10 \cup 00\}^*$ and $\{\{\{01 \cup 10 \cup 00\}^* 11 \{01 \cup 10 \cup 00\}^*\}^m\}^*$ respectively).

* Research supported in part by NSERC and FCAR grants.

The algebraic theory of finite monoids and semigroups provides a powerful framework and tools to analyze and classify regular languages [11]. It 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 [5,6]. Many recent lower bound techniques for such circuits have relied on communication complexity and, moreover, 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 [17].

It was established in [15] that the class of regular languages having $O(f)$ communication complexity, in some worst-case partition sense, forms a variety and so the question of the communication complexity of regular languages has an algebraic answer. Remarkably, this point of view allows us to determine the complexity of every regular language in five well-known models of communication complexity, the usual two party deterministic model and four important variants of it: bounded-error probabilistic, simultaneous, probabilistic-simultaneous and Mod_p-counting. These complete classifications unexpectedly feature complexity gaps in all five cases: a regular language has either $O(1)$, $\Theta(\log n)$ or $\Theta(n)$ communication complexity in the two-party deterministic, simultaneous, probabilistic-simultaneous and Mod_p-counting (for p prime) models and has communication complexity $O(1)$, $\Theta(\log \log n)$, $\Theta(\log n)$ or $\Theta(n)$ in the bounded-error probabilistic model. The classes of languages thus defined have fairly simple algebraic and combinatorial descriptions and are therefore 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. It also has complexity $O(\log n)$ in the Mod_p- and Mod_q-counting models for distinct primes p and q if and only if it has deterministic complexity $O(\log n)$.

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). Our results give further evidence that regular languages are an important first step in the study of various communication complexity models.

An important part of the technical content of this paper is algebraic in nature. In Section 2, we introduce the relevant background on finite monoids and algebraic automata theory and provide algebraic characterizations for pertinent classes of regular languages. 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. [12] for further background on algebraic automata theory and semigroup theory.

A semigroup S is a set together with a binary associative operation which we usually denote multiplicatively. A monoid M is a semigroup containing a distinguished identity element 1_M for the 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, the *exponent* of M, 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.

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. Examples of monoid-varieties which will be of importance in our context include the variety $\mathbf{Com}$ of commutative monoids, the variety $\mathbf{Ab}$ of Abelian groups and, for a prime p, the variety $\mathbf{G_p}$ of p-groups, that is groups whose order is a power of p. It is easy to show that G is a p-group if and only if $g^{p^\omega} = 1$ for all $g \in G$. We will also be concerned with the variety $\mathbf{G_p} * \mathbf{Ab}$ generated by the semidirect products of a p-group and an Abelian group. One can show that a group G lies in $\mathbf{G_p} * \mathbf{Ab}$ if and only if its *commutator subgroup* $[G, G]$, i.e. the subgroup generated by all *commutators* $[x, y] = x^{-1}y^{-1}xy$ of G, is a p-group. If $\mathbf{H}$ is a variety of finite groups, we denote by $\overline{\mathbf{H}}$ the variety of monoids whose subgroups all belong to $\mathbf{H}$. Finally, we will denote $\mathbf{LG_p}$ the variety of semigroups S such that for every idempotent e in S, the submonoid eSe of S lies in $\mathbf{G_p}$.

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.

A *relational morphism* from a semigroup S to a semigroup T is a mapping $\pi : S \to 2^T$ such that $\pi(x)\pi(y) \subseteq \pi(xy)$ for any $x, y \in S$ and $\pi(x) \neq \emptyset$ for all $x \in S$. Furthermore, if S, T are monoids, we require that $1_T \in \pi(1_S)$.

The *Mal'cev product* $\mathbf{V} \circledM \mathbf{W}$ of the semigroup variety $\mathbf{V}$ and the monoid variety $\mathbf{W}$ is the class of monoids M such that there exists a relational morphism π from M onto a monoid N of $\mathbf{W}$ such that for all idempotents $e \in N$ we have $\pi^{-1}(e) = \{m \in M | e \in \pi(m)\}$ forms a semigroup belonging to $\mathbf{V}$.

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 $M(L) \prec N$. A class of regular languages $\mathcal{L}$ is said to be a *language variety* if it is closed under Boolean operations, left and right quotients and inverse morphisms from A^* to B^*. The Variety Theorem of Eilenberg states that if $\mathbf{V}$ is a class of monoids and $\mathcal{L}$ is the class of regular languages whose syntactic monoids lie in $\mathbf{V}$, then $\mathbf{V}$ is a monoid-variety only if $\mathcal{L}$ is a language-variety. In that case $\mathbf{V}$ is generated by the monoids $M(L)$ for $L \in \mathcal{L}$. The objectives of algebraic automata theory are to exploit this correspondence by explicitly relating the algebraic description of a given variety of monoids with the combinatorial description of its language-variety counterpart.

2.1 Combinatorial Descriptions of Subvarieties of DO

We will almost exclusively be concerned with sub-varieties of the variety $\mathbf{DS}$ 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 B_2 the syntactic monoid of the language $(ab)^*$ and by U the syntactic monoid of $(a \cup b)^* aa (a \cup b)^*$. Both monoids contain the six elements $\{1, a, b, ab, ba, 0\}$ and satisfy $aba = a$, $bab = b$ and $x0 = 0x = 0$ for all monoid elements x. In B_2, we further have $aa = bb = 0$ while we have $aa = 0$ but $bb = b$ in U.

Lemma 1 (see e.g. [1]) *For a finite M, the following are equivalent:*

- *$M \in \mathbf{DS}$;*
- *$((xy)^\omega (yx)^\omega (xy)^\omega)^\omega = (xy)^\omega$ for all $x, y \in M$;*
- *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 $R_x y \subseteq R_x$;*
- *M does not admit B_2 or U as a divisor.*

We further denote by **DO** the subvariety **DS** consisting of monoids satisfying $(xy)^\omega(yx)^\omega(xy)^\omega = (xy)^\omega$ for all x, y. Alternatively, one can characterize **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. Many sub-varieties of **DS**, and **DO** in particular, have played key roles in many computational complexity questions pertaining to finite monoids [19, 18]. The following is a crucial property of monoids in **DO**.

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 $x\phi(w_1) = x\phi(w_2)$.*

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 $euvf = eufevf$.

Let $\phi(w_1) = y_1$ and $\phi(w_2) = y_2$. Since $xy_1\,\mathcal{H}\,xy_2\,\mathcal{R}\,x$ there must exist $s, t \in M$ lying $\mathcal{J}$-below y_1 and y_2 with $xst = x = x(st)^\omega$ and $xy_its = xy_i = xy_i(ts)^\omega$. 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,

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

□

Using the two previous lemmas, it is not hard to show that if M is in **DO** and $w \in M$ lies $\mathcal{J}$-above some idempotent $e \in M$, then $ew^\omega e = e$.

If $L_0, \ldots, L_k$ are languages in A^* and $a_1, \ldots, a_k$ are letters in A, we denote by $\binom{x}{(L_0a_1L_1\ldots a_kL_k)}$ the number of factorizations of the word x as $x = w_0a_1w_1\ldots a_kw_k$ with $w_i \in L_i$. When the a_i and L_i are such that for any x we have either $\binom{x}{(L_0a_1L_1\ldots a_kL_k)} = 0$ or $\binom{x}{(L_0a_1L_1\ldots a_kL_k)} = 1$ then we say that the concatenation $L_0a_1L_1\ldots a_kL_k$ is *unambiguous*. For a language variety $\mathcal{L}$, we denote by $UPol(\mathcal{L})$ the variety of languages generated by the unambiguous concatenations $L_0a_1L_1\ldots a_kL_k$ with $L_i \in \mathcal{L}$ and denote by $M_pPol(\mathcal{L})$ the language-variety generated by the languages $\{x | \binom{x}{(L_0a_1L_1\ldots a_kL_k)} = j \pmod p\}$ for some $0 \le j < p-1$ and $L_i \in \mathcal{L}$.

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 $M(L) \in \mathbf{DO} \cap \overline{\mathbf{H}}$ iff L is the disjoint union of languages of unambiguous concatenations $K = K_0a_1K_1\ldots a_sK_s$ where each K_i can be recognized by the direct product of a group of $\mathbf{H}$ and an idempotent and commutative monoid.*

Building on the work of [20], we will next prove a slight refinement of this theorem by defining a congruence capturing these languages in a convenient way. 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. This congruence captures the languages recognized by monoids in **DO**:

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 $G \in \mathbf{H}$.*

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 $K_0 a_1 \ldots a_s K_s$ where the K_i can be recognized by the direct product of an idempotent and commutative monoid (which suffices to check 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 $u = u_0 a_1 u_1 \ldots a_t u_t$, with $a_i \in A$ and $u_i \in A^*$ such that:

1. $[u_0 a_1 \ldots a_s u_s] \not\mathcal{R} [u_0 a_1 \ldots a_s u_s a_{s+1}]$ for any $s < t$.
2. $[u_0 a_1 \ldots a_s] \mathcal{R} [u_0 a_1 \ldots a_s u_s]$ for any $s \leq t$.
3. $[u_0] \mathcal{H} 1_M$.

Because M lies in **DS** we know by Lemma 1 that $a_i \notin \alpha(u_{i-1})$. In particular, $u = u_0\, a_1\, u_1 a_2 \ldots a_t u_t$ 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 $u = u_0 a_1 u_1 \dots a_t u_t$ be the $\mathcal{R}$-decomposition of u with $t \leq k$. We write w_i for $u_i a_{i+1} \dots a_t u_t$ and, by our earlier remark, have $w_i = u_i a_{i+1} w_{i+1}$ is the a_{i+1}-left-decomposition of w_i, so there must be a decomposition $v = v_0 a_1 \dots a_t v_t$ 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 $[u] \,\mathcal{R}\, [u_0 a_1 \dots u_{t-1} a_t] = [v_0 a_1 \dots v_{t-1} a_t] \geq_{\mathcal{R}} [v]$. 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,

$$[u] = [u_0 a_1 \dots u_{t-1} a_t][u_t] = [u_0 a_1 \dots u_{t-1} a_t][v_t] = [v_0 a_1 \dots v_{t-1} a_t][v_t] = [v]$$

□

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 $n, k \in \mathbb{N}$.

Next, we define $\mathbf{W}$ as the variety of monoids M satisfying:

1. $M \in \mathbf{DO}$;
2. $exwyf = ewxyf = exywf$ for all $w, e, f, x \in M$ such that e, f are idempotents lying $\mathcal{J}$-below w; i.e. M satisfies

$$\begin{aligned}(swt)^\omega wxy(uwv)^\omega &= (swt)^\omega xwy(uwv)^\omega \\ &= (swt)^\omega xyw(uwv)^\omega.\end{aligned}$$

In particular, if M is in $\mathbf{W}$ and we consider some $w \in M$ lying $\mathcal{J}$-above the idempotents e, f, and any $x, y \in M$, we must have $ew^\omega e = e$ (since M is in $\mathbf{DO}$) and so

$$\begin{aligned}exw^\omega yf &= ew^\omega exw^\omega yf \\ &= ew^\omega w^\omega exyf \text{ (By condition 2)} \\ &= exyf\end{aligned}$$

Note also that condition 2) shows that $\mathbf{W} \subseteq \overline{\mathbf{Ab}}$. Indeed, if u, v are elements of a subgroup with identity element e, we have

$$uv = euve = evue = vu.$$

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 occurrences 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 $M = A^*/\approx_{t,p}$ lies in $\mathbf{W}$. It is not hard to show from their definition that the $\approx_{t,p}$ equivalence classes are unambiguous concatenations of languages with syntactic monoids which are the product of an Abelian group and an idempotent commutative monoid. Thus, $M \in \mathbf{DO} \cap \overline{\mathbf{Ab}}$ by Theorem 1. Furthermore, let us consider the words $q = (uwv)^{tp}wxy(swz)^{tp}$ and $r = (uwv)^{tp}xwy(swz)^{tp}$. For any $a \in A$, $|q|_a \equiv |x|_a + |w|_a \equiv |r|_a \pmod{p}$ and

$$\begin{aligned} RED_t((uwv)^{tp}wxy(swz)^{tp}) &= RED_t((uwv)^{tp}xy(swz)^{tp}) \\ &= RED_t((uwv)^{tp}xwy(swz)^{tp}) \end{aligned}$$

since for any letter a occurring in w, the first t occurrences of a lie in $(uwv)^{tp}$, the common prefix of q and r and its last t occurrences lie in the suffix $(swz)^{tp}$. Thus, $q \approx_{t,p} r$ so M satisfies condition 2 in our definition of $\mathbf{W}$.

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 occurrences 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

$$s_i = s_i \phi(a u_{i+1} a \ldots a u_j) = s_i \phi(a u_{i+1} a \ldots a u_j)^\omega.$$

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 $q = u_0 a \ldots u_{t-1} axayav_{t-1} a \ldots av_0$ 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 the definition of **W** to obtain

$$\phi(q) = \phi(u_0 a \ldots u_{t-1} a\, a\, xy\, av_{t-1} a \ldots av_0).$$

Repeating this same process for all occurrences of a in x or y we can get $\phi(q) = \phi(u_0 a \ldots u_{t-1} a\, a^{kp+d} z\, av_{t-1} a \ldots av_0)$ where $a \notin \alpha(z)$ and from condition 1: $\phi(q) = \phi(u_0 a \ldots u_{t-1} a\, a^d z\, av_{t-1} a \ldots av_0)$ where $0 \le d < p$ 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. $\square$

Theorem 4 ([21]) *Let p be prime and* **V** *be a variety of finite monoids with an associated variety of languages $\mathcal{L}$. The variety of monoids associated with $M_p Pol(\mathcal{L})$ is* $\mathbf{LG_p} \circledM \mathbf{V}$.

Because the application of this theorem to the case $\mathbf{V} = \mathbf{Com}$ will bear special importance in Section 4, we need a more precise algebraic description of the variety $\mathbf{LG_p} \circledM \mathbf{Com}$.

Theorem 5 *The monoid M lies in $\mathbf{LG_p} \circledM \mathbf{Com}$ if and only if every subgroup of M lies in $\mathbf{G_p} * \mathbf{Ab}$ and for all $\mathcal{J}$-related idempotents $e, f \in M$ holds $(efe)^{p^\omega} = e$.*

Proof The simplest way to obtain this theorem is through topological results of [13]. Since we do not really need to introduce these sophisticated methods, we sketch here a proof from elementary principles.

Suppose first that there is a relational morphism $\pi : M \to T$ with T commutative and for any idempotent $d \in T$, $\pi^{-1}(d) \in \mathbf{LG_p}$. If e is an idempotent of M, then there is an idempotent d in $\pi(e)$ and $e\pi^{-1}(d)e$ is a p-group. It is easy to show that any idempotent f $\mathcal{J}$-related to e is also in $\pi^{-1}(d)$ so $(efe)^{p^\omega} = e$. Moreover, since T is commutative, for any x, y in the group $\mathcal{H}_e$, holds $d \in \pi(x^{-1})\pi(y^{-1})\pi(x)\pi(y)$. Thus $[x,y] \in \pi^{-1}(d)$ and so $[x,y]^{p^\omega} = e$ and therefore $\mathcal{H}_e$ is a group in $\mathbf{G_p} * \mathbf{Ab}$.

Conversely, let M be such that for all $\mathcal{J}$-related idempotents $e, f \in M$ holds $(efe)^{p^\omega} = e$. In particular, M lies in **DS** and so $x^\omega y^\omega$ is $\mathcal{J}$-related to $(xy)^\omega$. From each $\mathcal{J}$-class J_i of M, we pick a maximal subgroup G_i: the subgroup $[G_i, G_i]$ is normal in G_i and the group $K_i = G_i/[G_i, G_i]$ is Abelian. We also denote by ρ_i the canonical morphism from G_i into K_i and e_i the idempotent in G_i.

We define a commutative monoid T on the set $\bigcup K_i$. For $t_1 \in K_{i_1}$ and $t_2 \in K_{i_2}$ such that $(e_{i_1} e_{i_2})^\omega$ and e_k are $\mathcal{J}$-related idempotents in M, we choose m_1 and m_2 arbitrary pre-images in G_{i_1}, G_{i_2} in $\rho_{i_1}^{-1}(t_1)$ and $\rho_{i_2}^{-1}(t_2)$ respectively and define the multiplication of t_1 and t_2 in T as:

$$t_1 \cdot t_2 = \rho_k(e_k (m_1 m_2)^{\omega+1} e_k).$$

Note that the particular choice of m_1, m_2 is unimportant. It is also easy to verify associativity. We claim that there is a morphism ϕ from M to T such that the inverse image of any idempotent in T lies in $\mathbf{LG_p}$. We define ϕ by:

$$\phi(x) = \rho_i(exe)$$

where x^ω belongs to the $\mathcal{J}$-class J_i and e is the idempotent in group G_i. One must check that this is a well defined morphism. The crucial property we will use is that if e is the idempotent of G_i and f is another idempotent $\mathcal{J}$-related to e, then $\rho_i(efe) = e$. Indeed, since every group in M is in $\mathbf{G_p} * \mathbf{Ab}$ and $(efe)^{p^\omega} = e$ then efe lies in the subgroup $[G_i, G_i]$. Now we get:

$$\begin{aligned}
\phi(x) \cdot \phi(y) &= \rho_x(e_x x e_x) \cdot \rho_y(e_y y e_y) \\
&= \rho_{(e_x y)}(e_{xy}(e_x x e_x)(e_y y e_y)e_{xy}) \\
&= \rho_{(e_{xy})}(e_{xy} x y e_{xy}) \\
&= \phi(xy)
\end{aligned}$$

The same remark allows us to conclude that the inverse image of any idempotent of T lies in $\mathbf{LG_p}$. $\square$

3 Communication Complexity

We only provide here a short introduction to some communication complexity models but refer the reader to the excellent book of Kushilevitz and Nisan [10] 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.

Combining properties of the two previous models, a *simultaneous probabilistic communication protocol* $\mathcal{P}$, is one in which Alice and Bob *simultaneously* send a message depending on their inputs and their random bits to a trusted referee which then outputs $\mathcal{P}(x, y)$ which should equal $f(x, y)$ with probability at least 3/4. We denote $R^{\|}(f)$ the simultaneous probabilistic (or randomized) communication complexity of f.

In contrast to these first four models, the non-deterministic and MOD_p communication models that we present next will only be used to recognize languages, i.e. to compute functions from $S^* \times S^*$ into $\{0, 1\}$. For convenience, we choose a "guess and verify" presentation of non-deterministic protocols although alternative but equivalent definitions can be used [2,10].

A *non-deterministic protocol* $\mathcal{P}$ is a set of deterministic protocols $\{\mathcal{P}_\pi | \pi \in \{0, 1\}^k\}$ indexed by strings of some length k, which we think of as proofs. The language accepted by $\mathcal{P}$ consists of pairs (x, y) such that there exists a proof $\pi \in \{0, 1\}^k$ with the property that $\mathcal{P}_\pi$ accepts (x, y). The cost of $\mathcal{P}$ is the maximum cost of any $\mathcal{P}_\pi$ plus k. It is convenient to think of a non-deterministic protocol as comprising an initial "guessing" stage in which a proof is chosen and a subsequent "verification" stage in which Alice and Bob must agree on the proof's validity.

A *MOD_p-counting communication protocol* $\mathcal{P}$ is syntactically the same as a non-deterministic protocol but it accepts those (x, y) such that the number of proofs that lead Alice and Bob to acceptance is *not* divisible by p [7]. Here also, the cost of the protocol includes both the length of the proof and the subsequent communication exchanged by the two players. We denote by $N^{Mod_p}(L)$ the Mod_p-counting communication complexity of L.

It is not difficult to show that if p is prime then for any languages L_1, L_2:

$$N^{Mod_p}(L_1 \cup L_2) = O(N^{Mod_p}(L_1) + N^{Mod_p}(L_2))$$

and

$$N^{Mod_p}(L_1) = \Theta(N^{Mod_p}(\overline{L_1})).$$

All five of these models have been extensively studied. One can easily show that $\log(D(L)) \le R(L) \le D(L)$ and, similarly, all above models differ by at most an exponential factor. We will use classical complexity bounds for the following functions [10]:

- For $x, y \in \{0, 1\}^n$, we define Disjointness as: $DISJ(x, y) = 1$ if and only if $\bigvee_{1 \le i \le n} x_i y_i = 0$;
- 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 $x \geq y$.
- For $x \in \{0,1\}^n$ and a $\log n$-bit number $s \in [n]$ we define $INDEX(x,s) = x_s$;

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

Theorem 6 *For any prime p and integer q which is not a power of p, we have:*

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

It should be noted that non-trivial work is necessary to prove the GT randomized upper bound (see exercise 3.18 in [10]) and the randomized lower bounds for $DISJ$ [9,16] and IP [4,8]. All the Mod_p-counting bounds are theorems or easy corollaries of [7].

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 ϵ and they want to determine whether $a_1 a_2 \ldots a_{2n}$ belongs to L.

In order to compare the communication complexity of two languages K, L in different models, Babai, Frankl and Simon [2] 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:

Definition 1 *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 $t = O(n)$.

For completeness, we sketch the proofs of two basic facts from [15].

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

Proof (sketch) 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 $u[m]v \in L$ if and only if $u[m_1][m_2]\ldots[m_{2n}]v \in L$. This observation can be used to obtain a rectangular reduction of linear length from M to L. □

In particular the deterministic (resp. simultaneous, randomized, randomized-simultaneous, Mod_p-counting) 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. $D^{\|}(M)$, $R(M)$, $R^{\|}(M)$, $N^{Mod_p}(L)$ for prime p) is $O(f)$ forms a variety.*

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

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. □

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, simultaneous, randomized-simultaneous and Mod_p-counting two-party models. We start with a number of useful upper bounds, the first of which is an easy observation due to [15].

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

Proof By commutativity, we have

$$m = m_1 m_2 \dots m_{2n} = (m_1 m_3 \dots m_{2n-1})(m_2 m_4 \dots m_{2n})$$

so if Alice and Bob respectively send the referee the $\lceil \log |M| \rceil$ bits encoding the product of the odd-indexed and the even-indexed m_i, then he can compute m. $\square$

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 $x = x_1 x_2 \dots x_{2n} \in A^*$, 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 $(|A|\lceil \log p \rceil + 1)$, 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 $x = x_1 \dots x_{i-1} a x_{i+1} \dots x_n$ is the a-left decomposition of x and, by induction, can verify that $x_1 \dots x_{i-1} \sim^G_{|A|-1,k} v$ and $x_{i+1} \dots x_n \sim^G_{|A|,k-1} w$ using only $O(\log n)$ communication. Left-right symmetry completes the proof. $\square$

At least intuitively, we have $M(L)$ lying in $\mathbf{W}$ if and only if membership of a word w in L can be determined by counting, up to threshold t and modulo p, the number of occurrences of each letter in w and determining the relative positions of any of the first and last t occurrences of letters in w. In terms of communication complexity, $\mathbf{W}$ thus forms an "easy" subclass of $\mathbf{DO} \cap \overline{\mathbf{Ab}}$ because these comparisons of $\log n$-bit numbers can be done relatively efficiently. More precisely:

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 6 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. □

We have seen that unambiguous products of languages with commutative syntactic monoids have $O(\log n)$ deterministic communication complexity. It should not come as much of a surprise that in the MOD_p-counting model we can correspondingly obtain:

Lemma 8 *Let $L \subseteq A^*$ be regular with $M(L) \in \mathbf{LG_p} \circledM \mathbf{Com}$. We have $N^{Mod_p}(L) = O(\log n)$.*

Proof We know from the result of [21] cited as Lemma 4 that L is in $M_pPol(\mathcal{L}_{com})$, i.e. is a Boolean combination of languages of the form

$$\{x | \binom{x}{(L_0 a_1 L_1 \ldots a_k L_k)} \equiv j \pmod{p}\}$$

where $M(L_i)$ is commutative for all i. So we only need to exhibit an $O(\log n)$-cost protocol to check if a given word $w \in (A \cup \{\epsilon\})^*$ has a number of factorizations as $u_0a_1u_1 \dots a_ku_k$, with $u_i \in L_i$, that is congruent to j modulo p.

Suppose first that $j = 0$. The protocol we present in the next paragraph will in fact output positively if and only if the number of valid factorizations is *not* congruent to 0 modulo p. This is sufficient as we have mentioned that for p prime $N^{Mod_p}(L) = N^{Mod_p}(\overline{L})$.

The proof chosen in the first step of the protocol consists of k $\log n$-bit integers $t_1 < t_2 < \dots < t_k$. In the next stage of the protocol, Alice and Bob interpret the t_i's as possible locations for the bookmarks $a_1, a_2, \dots, a_k$ in w and accept if they correspond to a valid factorization $u_0a_1u_1 \dots a_ku_k$ with $u_i \in L_i$. This can be done at constant cost since Alice and Bob need only check that position t_i indeed contains letter a_i and that segment u_i belongs to L_i, which requires only $O(1)$ bits since $M(L_i)$ is commutative. The cost of the protocol is dominated by the length of the proof which is $O(\log n)$. The number of proofs accepted by Alice and Bob is thus exactly the number of legal factorizations so the protocol accepts if and only if it is non-zero modulo p. Note that for $j \neq 0$, we need to slightly modify our protocol by creating $(p-j)$ different "special" proofs that always lead Alice and Bob to accept. $\square$

Whereas the upper bounds just presented exploit crucially the combinatorial descriptions of languages with syntactic monoids in particular sub-varieties of **DO**, our lower bounds will all stem from the algebraic descriptions of these same varieties. We give a number of conditions under which a finite M admits a reduction from GREATER THAN, DISJOINTNESS, INNER PRODUCT mod q and INDEX.

Lemma 9 *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. For any prime p, if M lies in **DS** *but contains $\mathcal{J}$-related idempotents e, f such that $(ef)^{p^\alpha}e$ is diffrent from e for any $\alpha \geq 1$ then $IP_q \leq_r M$ for some q not a power of p;*
5. If G is a non-commutative group then $IP_q \leq_r G$ for some integer q;
6. For any prime p, if G is a group outside of $\mathbf{G_p} * \mathbf{Ab}$ then $IP_q \leq_r G$ for some q which is not a power of p;
7. If M is in **DO** *but not in* **W** *then $INDEX \leq_r M$.*

Proof 1- Let $a, b \in M$ be such that $ab \neq ba$. We obtain an exponential length reduction from $GT(x, y)$ by building $m_1m_2 \dots m_{2^n+1}$ 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.
2- If M is not in **DS** then, by Lemma 1, it must admit one of B_2 or U as a divisor. In both cases, 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

$m_{4i-2} = 1_M$ when $y_i = 0$ and $m_{4i} = ab$ for any input. Since in both B_2 and U, we have $abab = ab$ but $aa = 0$, one can easily check that any such four-tuple evaluates to 0 when $x_i = y_i = 1$ and to ab otherwise. Hence, the product of all such tuples will be 0 if $x_i = y_i = 1$ for some i 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. Since M is in **DS**, however, we 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 $m_{2i-1} = e(ef)^\omega = (ef)^\omega$ otherwise and $m_{2i} = fe$ if $y_i = 1$ and $m_{2i} = (ef)^\omega e = e$ 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 $(ef)^{\sum_{1 \le i \le n} x_i y_i \equiv 0 \pmod q} e$ which equals e if and only if $IP_q(x, y) = 1$.

4- The argument is almost the same as 3-. Again, let q be minimal such that $(ef)^q e = e$. We are guaranteed that q is not a power of p and we can reuse the reduction described for 3-.

5- 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 $[g, h]^{\sum_{1 \le i \le n} x_i y_i \pmod q}$.

6- If G is not in $\mathbf{G_p} * \mathbf{Ab}$, there must exist $g, h \in G$ such that the commutator $[g, h]$ has order q which is not a power of p so we can reuse the previous reduction.

7- If M lies in $\mathbf{DO} - \mathbf{W}$, we must consider two cases. Assume first that there exist e, f, u, v, w with e, f idempotent and $\mathcal{J}$-below w such that $euw^\omega vf \neq euvf$. Since M is in **DO** we have $ew^\omega e = e$ and $fw^\omega f = f$.

We obtain a rectangular reduction from $INDEX(x, s)$ by creating $m = m_1 m_2 \dots m_{2n+1}$ as follows:

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

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 $ew^\omega e$ is e and $fw^\omega f$ is f, the product of the m_i's is equal to $euvf$ when x_s is 0 and $euw^\omega vf$ when x_s is 1 which shows the correctness of the reduction.

If on the other hand we do have $euw^\omega vf = euvf$ for all $e, u, v, w, f \in M$ with the relevant properties, then there must exist $e, f, u, w \in M$ with e, f

idempotent $\mathcal{J}$-below w but $ewuf \neq euwf$ for otherwise we have $euwvf = euw^\omega wvf = euw^\omega vwf = euvwf$ and M lies in $\mathbf{W}$. On the other hand, since we assume $euw^\omega vf = euvf$, we have $euwf = ew^\omega uwf$ and $ewuf = ewuw^\omega f$. We now obtain a rectangular reduction from $INDEX(x, s)$ (assuming w.l.o.g. that $x_n = 1$) as follows:

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

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 $w^{k\omega+1}$, for some k, if and only if x_i is 1 and $w^{k\omega}$ if x_i is 0. Similarly, using the fact that $x_n = 1$, the product $m_{2i+2} \dots m_{2n}$ is $w^{k'\omega}$ if and only if $x_i = 1$ and $w^{k'\omega+1}$ otherwise. Using the values assigned by the reduction to the odd-indexed m_i's we have

$$m_1 m_2 \dots m_{2n+1} = e m_2 m_4 \dots m_{2s}\, u\, m_{2s+2} m_{2s+4} \dots m_{2n} f$$

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

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 four following theorems.

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

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

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. $\Box$

Theorem 8 *Let $L \subseteq A^*$ be a regular language with $M = M(L)$. Then*

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

Proof When M is in $\mathbf{W}$ but not commutative we put together Lemma 7 and part 1 of Lemma 9 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. □

Theorem 9 *Let $L \subseteq A^*$ be a regular language with $M = M(L)$. Then*

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

Proof When M is in $\mathbf{W}$ but not commutative we can combine the upper bound of Lemma 7 with the lower bound obtained from part 1 of Lemma 9. When M is not in $\mathbf{W}$, then it admits a linear length reduction from one of $INDEX$, $DISJ$ or IP_q which all have $\Omega(n)$ simultaneous and probabilistic-simultaneous complexity. □

There are non-regular languages for which probabilistic-simultaneous protocols significantly outperform deterministic simultaneous ones. For example, the probabilistic-simultaneous complexity of Equality is $O(\sqrt{n})$ (folklore) and it was established by [3] that this quadratic gain is optimal. Such gaps do not exist for regular languages because they do not exist for GT, IP_q, $INDEX$ or $DISJ$.

Theorem 10 *Let $L \subseteq A^*$ be a regular language with $M = M(L)$ and p be prime. Then*

$$N^{MOD_p}(L) = \begin{cases} O(1) & \textit{if and only if } M \textit{ is commutative;} \\ \Theta(\log n) & \textit{if and only if } M \textit{ is in } \mathbf{LG_p} \circledM \mathbf{Com} \\ & \textit{but is not commutative;} \\ \Theta(n) & \textit{otherwise.} \end{cases}$$

Proof If M is in $\mathbf{LG_p} \circledM \mathbf{Com}$ but not commutative, we use $GT \leq_r^{2^n} M$ for the lower bound and Lemma 8 for the upper bound. If M is not in $\mathbf{LG_p} \circledM \mathbf{Com}$ then it must be either outside of $\mathbf{DS}$, outside of $\overline{\mathbf{G_p} * \mathbf{Ab}}$ or have $\mathcal{J}$-related idempotents e, f with $(efe)^{p^\omega} = (ef)^{p^\omega} e \neq e$ and by Lemma 9 we then have either $DISJ \leq_r M$ or $IP_q \leq_r M$ for some q which is not a prime power of p. In either cases this suffices to get $N^{MOD_p}(M) = \Omega(n)$. □

Corollary 1 *If L is a regular language such that for two distinct primes p and q we have both $N^{MOD_p}(L) = O(\log n)$ and $N^{MOD_q}(L) = O(\log n)$. Then $D(L) = O(\log n)$.*

Proof By our previous theorem, M must lie in the intersection of the varieties $\mathbf{LG_p} \circledM \mathbf{Com}$ and $\mathbf{LG_q} \circledM \mathbf{Com}$. In particular, M lies in $\overline{\mathbf{G_p} * \mathbf{Ab}}$

and $\overline{\mathbf{G_q} * \mathbf{Ab}}$, so any subgroup G of M must be such that $[G,G]$ is both a p-group and a q-group. This implies that $[G,G]$ is trivial and thus, that G is Abelian and that M lies in $\overline{\mathbf{Ab}}$. Similarly, if for all $\mathcal{J}$-related idempotents e,f we have both $(efe)^{p^\omega} = e$ and $(efe)^{q^\omega} = e$ then it must be that $efe = e$ and so M lies in $\mathbf{DO}$. By Theorem 7, all monoids in $\mathbf{DO} \cap \overline{\mathbf{Ab}}$ have $O(\log n)$ deterministic communication complexity. □

It is unclear whether or not such a phenomenon also occurs for non-regular languages. To the best of our knowledge, this question has never been studied.

5 Conclusion

We have proved that in five well-known two-party models of communication complexity: deterministic, simultaneous, probabilistic, probabilistic-simultaneous and Mod_p-counting (for prime p), the complexity of a regular language is determined by simple and decidable algebraic properties of its syntactic monoid and that, in each case, this classification only defines a small number of classes. Because regular languages are so well-behaved in these contexts, they are a natural starting point for further investigations in other models of communication complexity. In particular, one might attempt to obtain a similar classification theorem for the non-deterministic communication complexity of regular languages. Attacking this question from an algebraic angle will require a refinement of our techniques since some regular languages have a non-deterministic communication complexity exponentially smaller than that of their complement. This means that we cannot find tight bounds for $N^1(L)$, the non-deterministic communication complexity of L, by simply looking at the algebraic properties of $M(L)$. However, it is easy to show that the class of regular languages with non-deterministic communication complexity $O(f)$ is closed under positive Boolean operations, inverse homomorphisms and left-right quotients, i.e. that it forms a *positive variety* of languages (see e.g. [12]). Correspondingly, $N^1(L)$ is thus determined by algebraic properties of L's *ordered* syntactic monoid. We conjecture that if L is a regular language, then:

$$N^1(L) = \begin{cases} O(1) & \text{if and only if } M(L) \text{ is commutative;} \\ \Theta(\log n) & \text{if and only if } M(L) \text{ is non-commutative but } L \\ & \text{is the disjoint union of languages } L_0 a_1 L_1 \ldots a_k L_k \\ & \text{with } M(L_i) \text{ commutative;} \\ \Theta(n) & \text{otherwise.} \end{cases}$$

The linear lower bound seems to be a difficult challenge and probably requires both subtle algebraic calculations and non-deterministic communication complexity lower bounds that are of a different nature than the linear lower bounds already known for Inner Product and Disjointness. A curious corollary of this conjecture would be that if L is regular then $D(L) = \Theta(\max\{N^1(L), N^0(L)\})$.

Ongoing work [15, 14, 18] has also focused on the communication complexity of regular languages in the multiparty (number in the forehead) model (see [10]) where one might still hope to prove a generalization of the theorem of Szegedy mentioned in our introduction. This multiparty model is far from being well-understood and the case of regular languages appears to be a simple yet very meaningful first step.

References

1. J. Almeida. *Finite Semigroups and Universal Algebra.* Series in Algebra, Vol 3. World Scientific, 1994.
2. L. Babai, P. Frankl, and J. Simon. Complexity classes in communication complexity theory. In *Proc. 27^{th} IEEE FOCS*, pages 337–347, 1986.
3. L. Babai and P. G. Kimmel. Randomized simultaneous messages. In *Proc. 12^{th} Conf. on Comp. Complexity (CCC '97)*, pages 239–247, 1997.
4. L. Babai, N. Nisan, and M. Szegedy. Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. *J. Comput. Syst. Sci.*, 45(2):204–232, Oct. 1992.
5. 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.
6. 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.
7. 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/.
8. V. Grolmusz. Separating the communication complexities of MOD m and MOD p circuits. In *Proc. 33rd IEEE FOCS*, pages 278–287, 1992.
9. B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM J. on Discrete Mathematics*, 5(4):545–557, 1992.
10. E. Kushilevitz and N. Nisan. *Communication Complexity.* Cambridge University Press, 1997.
11. J.-E. Pin. *Varieties of formal languages.* North Oxford Academic Publishers Ltd, London, 1986.
12. J.-E. Pin. Syntactic semigroups. In G. R. et A. Salomaa, editor, *Handbook of language theory*, volume 1, chapter 10, pages 679–746. Springer Verlag, 1997.
13. J. E. Pin and P. Weil. Profinite semigroups, Mal'cev products, and identities. *Journal of Algebra*, 182:604–626, 1996.
14. P. Pudlák. An application of Hindman's theorem to a problem on communication complexity. To appear in *Combinatorics, Probability and Computing*, 2003.
15. 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.
16. A. Razborov. On the distributed complexity of disjointness. *Theoretical Comput. Sci.*, 106(2):385–390, 14 Dec. 1992. Note.

17. M. Szegedy. Functions with bounded symmetric communication complexity, programs over commutative monoids, and ACC. *J. Comput. Syst. Sci.*, 47(3):405–423, 1993.
18. P. Tesson. *Computational Complexity Questions Related to Finite Monoids and Semigroups.* PhD thesis, McGill University, 2003.
19. P. Tesson and D. Thérien. Diamonds are forever: the variety **DA**. In G. Gomez, P. Silva, and J-E.Pin, editors, *Semigroups, Algorithms, Automata and Languages.* WSP, 2002.
20. D. Thérien and T. Wilke. Over words, two variables are as powerful as one quantifier alternation. In *Proc. 30th ACM Symposium on the Theory of Computing*, pages 256–263, 1998.
21. P. Weil. Closure of varieties of languages under products with counter. *J. Comput. Syst. Sci.*, 45:316–339, 1992.