Distributed Algorithmic Mechanism Design and Algebraic Communication Complexity

Markus Bläser[1] and Elias Vicari*[2]

[1] FR Informatik, Saarland University
Postfach 151150, 66041 Saarbrücken, Germany
mblaeser@cs.uni-sb.de
[2] Institute of Theoretical Computer Science, ETH Zurich
8092 Zurich, Switzerland
vicariel@inf.ethz.ch

Abstract. In this paper, we introduce and develop the field of algebraic communication complexity, the theory dealing with the least number of messages to be exchanged between two players in order to compute the value of a polynomial or rational function depending on an input distributed between the two players. We define a general algebraic model over an arbitrary field k of characteristic 0, where the involved functions can be computed with the natural operations additions, multiplications and divisions and possibly with comparisons.
We then apply this general theory to problems from distributed mechanism design, in particular to the multicast cost sharing problem, and study the number of messages that need to be exchanged to compute the outcome of the mechanism. This addresses a question raised by Feigenbaum, Papadimitriou, and Shenker [9].

1 Introduction

Distributed algorithmic mechanism design is the art of analyzing, designing, and implementing mechanisms in a distributed setting. Feigenbaum, Papadimitriou, and Shenker [9] were the first to consider distributed aspects of mechanism design by investigating mechanisms for the multicast cost sharing problem in a distributed setting. Many mechanisms, like *marginal costs* and *Shaply value*, can be computed by arithmetic operations ("$+$", "$*$", and "$/$"), equality tests ("$=$"), and comparisons ("$<$"). These are exactly the operations that are considered in *algebraic complexity theory*, a well-developed area of complexity theory. In the multicast cost sharing problem, a provider, the root of a tree, wants to send a transmission to n players, residing at the nodes of the tree. Each player has a utility u_i for receiving the transmission. The provider has to pay some cost c_e when he uses a particular edge e of the tree for the transmission. The cost c_e is a private information that is owned by the players residing at the endpoints of the edge e. The nodes are capable of doing multicast, so edges are shared by the users in the tree. A mechanism is sought that is strategy-proof, fulfills three natural side constraints, and either maximizes welfare or is budget-balanced. We assume that there is no centralized

* The author is partially supported by the Swiss National Science Foundation under grant number $200021 - 105236/1$.

authority that knows all the values. Instead the computation is done at the nodes of the trees and values have to be sent over the links. The mechanisms maximizing welfare have efficient distributed implementations with only a constant number of messages per link.

Feigenbaum, Papadimitriou, and Shenker [9] investigate the communication complexity of computing budget balanced mechanisms with linear operations, that is, only multiplication with scalars but not with variables are allowed. This implies that all the messages sent are linear combinations of the values held by the players. In this setting, they show that the computation of any budget balanced mechanism requires a linear number of messages over a linear number of links, yielding an overall quadratic lower bound. Any function can be computed by this many messages by sending all inputs to all players, so this lower bound is optimal. They leave the extension of their results to non-linear operations, i.e., multiplications and divisions, as an open question.

Later, Feigenbaum et al. [8] used Boolean communication complexity to prove lower bounds for the bit complexity of distributed algorithms for budget balanced mechanisms. In the Boolean model, all numbers that occur are rational numbers (given as the quotient of two natural numbers in binary). The messages can be arbitrary bit strings. In this model, Feigenbaum et al. prove that any distributed Boolean algorithm has to send at least a linear number of *bits* over a linear number of links.

While this seems to solve the problem at a first glance, there is a flaw hidden: If we measure the bit complexity of the messages, we have to relate it to the *bit size* of the input. In particular, in order to prove the mentioned lower bound, Feigenbaum et al. consider the following simple scenario: There is one link from the provider to one node v at which $n/2$ players reside and there is one link from v to another node u at which another $n/2$ players reside. By instantiating the utilities of the players appropriately, it turns out that any distributed algorithm that computes a budget-balanced mechanism on this tree decides whether the sets of utilities at the nodes u and v are disjoint. If we now consider this to be a game between two players, one at u, the other at v, this becomes the well-known set disjointness problem, however with a small twist, the universe is not $\{1, \ldots, n\}$. It is known that deciding the set disjointness problem in the Boolean setting where both players hold the characteristic vector of their subset of $\{1, \ldots, n\}$ requires n bits of communication. The scenario that we get from the multicast cost sharing problem above is a generalization, so we get the lower bound of n bits as well. In this example, the utilities of the players are of the form C/i for some constant C and $1 \leq i \leq n$. Thus the overall input size $\ell = \Theta(n \log n)$, but the lower bound for the communication is only $\Omega(n^2) = \Omega(\ell^2 / \log^2 \ell)$. (Note that we have to send n bits. If we replace the one link between u and v by a path of length n, we get n messages over n links, yielding the quadratic lower bound.) While this is still bad enough, it is not quadratic in the input size.

We will show that in the general algebraic model, we can get a quadratic lower bound. We provide tools to investigate the communication complexity of distributed computations in a general algebraic setting. This means, that the player may perform arbitrary arithmetic operations and not only linear ones, together with comparisons. The messages are now arbitrary rational functions in the inputs of the players. We will show

that in this setting, the computation of any budget balanced mechanism needs a linear number of messages over a linear number of links. This is now a tight lower bound.

The reduction by Feigenbaum et al. to a two-player communcation problem indicates that it is sufficient to understand two-player communication complexity. This is what we do here: *Algebraic communication complexity* deals with the problem of computing a multivariate function, when the input is distributed between different entities referred to as *players*. In particular, given a field k and a rational function $f \in k(X_1, \ldots, X_n, Y_1, \ldots, Y_m)$, we want to know how many messages have to be exchanged between the two players, until one of them has enough information at his/her disposal to be able to compute the function value. Here, one player is holding the X-indeterminates and the other one the Y-indeterminates.[3]

The boolean counterpart, where one deals with boolean functions and bit-strings as inputs, is well studied, thanks to its successful application in VLSI-design theory. [12,13] are excellent textbooks on Boolean communication complexity. Compared to the boolean setting, we work over infinite fields and treat the arising numbers as entities. So our results do not follow from discrete models with larger (but finite) alphabets.

2 Related research and new results

The Boolean model was proposed by Yao [18] and is typically motivated by VLSI design problems, see e.g. [12,13]. It is mostly of combinatoric nature, whereas our model is based on algebraic structures like fields in which the possible operations are the natural arithmetic operations.

Abelson [1,2] motivated and introduced continuous communication complexity theory over $\mathbb{R}$ assuming some differentiability properties of the involved functions. Luo and Tsitsiklis [16] improved Abelson's results in certain cases making use of algebraic tools, but they only consider computional problems and no decision problems. The reason for this is that they only consider smooth messages and smooth problems.

Other than these work, there are also studies leading to more specific directions, like optimization within an error of a sum of two distributed convex functions, where every player has access to a single function [15].

Our aim is to develop a model based on an algebraic structure, in order to take advantage of the powerful tools of algebra and algebraic complexity theory. Even though this restricts the power of the messages the players can exchange and the family of the multivariate objective functions, we strongly believe that this model of computation fits better in the context of possible applications, since the power of the involved messages is realistically bounded. Moreover, in our setting we can introduce new features in a natural way. For instance by allowing equality tests and comparisons, we actually can deal with non-continuous functions. We can speak about *decision problems*, too, where the goal is not to compute the function, but decide whether a given input lies in its zero-set. Or we can introduce *nondeterminism* in a flavor recalling the usual nondeterminism of computational complexity.

[3] One can define the multi-party case accordingly. Due to space limitations, we solely deal with the two player case in this paper.

Feigenbaum, Papadimitriou, and Shenker [9] investigated the multicost cost sharing problem in a restricted algebraic setting. They only allow messages that are linear forms in the inputs (i.e. bids) of the players. They leave the proof of lower bounds in a general model as an open problem. We here prove a quadratic lower bound for the communication complexity of budget-balanced mechanisms for the multicast cost sharing problem in such a general model.

In particular, we show the following results:

- We completely solve the one-way case for the computation problem and give a characterization in terms of the transcendence degree.
- We almost completely solve the one-way case for the decision problem, where "almost" means up to one message.
- We provide several lower bounds techniques for the two-way case:
  - a method that is inspired by the substitution method known from algebraic complexity theory
  - a dimension bound (which is solely suited for homogeneous problems)
- We extend our methods to decision problems with equality tests and with comparisons.
- We apply our methods to distributed mechanism design problems, like the multicast cost sharing problem and auctions with single minded bidders.

In the full version, we also introduce nondeterminism and provide an interesting link to the decision and the computation problem in the one-way case.

Recently and independently, Grigoriev [11] introduced a model similar to ours, but he only deals with the one-way communication complexity. His model is actually an extension of the one-way model presented here: both parties send messages to some referee who then makes the decisions. We will use one of his results, which can be readily generalized to the two-way model, to deal with comparisons. Grigoriev allows randomization, too, but he mainly uses rank-based lower bounds.

3 Model of computation and notation

Throughout this paper, k is a field of characteristic zero which we often assume to be algebraically closed. All our mechanism design problems will be reduced to a two-player algebraic communication problem; so we define and analyze this setting first. We give the players names: Alice (A) and Bob (B). Alice usually holds an input denoted by $X = (x_1, \ldots, x_n)$ and Bob holds $Y = (y_1, \ldots, y_m)$ and their aim is to compute a rational function $f : k^n \times k^m \to k$. Each player may send messages that are rational functions in his/her inputs and the messages (s)he has received from the other player so far. In the end, one of them has to be able to compute $f(X, Y)$ (*computation problem*), to decide if the value is zero or not (*decision problem*), or to prove that indeed the input lies in the zero-set of f (*nondeterminism*).

We denote, as usual in the literature, by $k[X, Y]$ the ring of polynomials with variables $(x_1, \ldots, x_n, y_1, \ldots, y_m)$ and coefficients in k and by $k(X, Y)$ the field of rational functions over the same set of variables. Further, we denote by $M_{A\to B}$ and $M_{A\leftarrow B}$ the index-set of messages sent by Alice to Bob and vice versa, respectively. The network

where messages are sent is completely reliable: there is no data loss and transmissions are error-free. In particular we assume that we can send field-elements as they are, meaning that every message counts as one, no matters how large the number is (for example as real number) and how we encode it. What is really important in this framework is the total number of messages sent during a protocol, neglecting the amount of computation performed by the players. In fact, our model relies on Luo and Tsitsiklis' model [16], but the assumptions on the involved functions are different: we deal solely with polynomial/rational functions. Later, we introduce equality tests and comparisons. In this extended model, we can also compute non-continuous functions and investigate decision problems, too.

A protocol P for computing a rational function $f \in k(X, Y)$ is a list of instructions telling the players the form of the messages they have to send and in which order, so to let a player to be able to compute $f(X, Y)$.

Definition 1 (Protocol). *A* two-way protocol P *for computing* f *consists of:*

1. *Disjoint inputs* X, Y *distributed between player* A *and player* B*, respectively.*
2. *A collection of messages* $m_1, \ldots, m_r$ *belonging to some field extensions of* k*, sent in this order, with the following property: for each* $1 \leq i \leq r$*, we have:*
   - *if* $i \in M_{A \to B}$*, then* $m_i \in k(X, m_1, \ldots, m_{i-1})$
   - *if* $i \in M_{A \leftarrow B}$*, then* $m_i \in k(Y, m_1, \ldots, m_{i-1})$
3. *We have either* $f \in k(X, m_1, \ldots, m_r)$ *or* $f \in k(Y, m_1, \ldots, m_r)$*.*

P *is called* one-way *if in addition* $M_{A \leftarrow B} = \emptyset$ *and* $f \in k(Y, m_1, \ldots, m_r)$*.*

Definition 2. *The* two-way communication complexity *of* f *is defined as*

$$\mathrm{C}(f) := \min_P r(P)$$

where the minimum is taken over the set of all protocols P *for* f *and* $r(P)$ *is the number of messages sent in* P*. Similarly we define the one-way communication complexity* $\mathrm{C}^{\rightarrow}(f)$*.*

A message m is said to be *feasible* in step i, if the second property in the definition of a protocol holds. If a message m is feasible, this exactly means that m can be computed by additions, multiplications, and divisions from the inputs of the particular player and all the messages he received so far.

We will also speak about divisionfree protocols. In this case, m is feasible if $m \in k[X, m_1, \ldots, m_{i-1}]$ or $m \in k[Y, m_1, \ldots, m_{i-1}]$, respectively. With such protocols, we can of course only compute polynomials.

Finally, we can also extend the notion of a protocol and communication complexity to a set of functions $f_1, \ldots, f_\ell$. We just require $f_1, \ldots, f_\ell \in k(X, m_1, \ldots, m_r)$ or $f_1, \ldots, f_\ell \in k(Y, m_1, \ldots, m_r)$ in the third item of the definition of a protocol.

Example 1. Let

$$f_1(X, Y) = (y_1 + x_1 y_2 + x_1^2 y_3)x_1 + (y_1 + x_1 y_2 + x_1^2 y_3)^2 x_2 + (y_1 + x_1 y_2 + x_1^2 y_3)^3 x_3.$$

Then $\mathrm{C}(f_1) \leq 2$. Indeed, it is easy to see, that if Alice sends to Bob the value of x_1, then he is able to compute the polynomial $y_1 + x_1 y_2 + x_1^2 y_3$, which, in turn, enables Alice to compute the whole function. One can easily show with the results of this paper, that in fact $\mathrm{C}(f_1) = 2$ but $\mathrm{C}^{\rightarrow}(f_1) = 3$.

4 One-way communication

In this section we deal exclusively with the one-way communication model. One-way communication corresponds to distributed one-pass algorithms. For the example of multicast pricing, this means that the result has to be computed in one sweep over the multicast tree. We provide a technique to compute exactly the value $\mathrm{C}^{\rightarrow}(f)$ for *every* rational function $f \in k(X, Y)$.

For such an $f \in k(X, Y)$ we denote by $\mathrm{Coeff}_Y f$ the field extension over k generated by adding the coefficients of f seen as function of Y. In particular, $\mathrm{Coeff}_Y f \subseteq k(X)$. Obviously, Bob is able to compute the value of f from the messages $m_1, \ldots, m_r$ received from Alice if and only if $\mathrm{Coeff}_Y f \subseteq k(m_1, \ldots, m_r)$.

Theorem 1 (Transcendence degree bound). *For every field k and rational function $f \in k(X, Y)$, we have $\mathrm{C}^{\rightarrow}(f) \geq \mathrm{tr\,deg}_k \mathrm{Coeff}_Y f$ where $\mathrm{tr\,deg}$ denotes the transcendence degree of the field extension $\mathrm{Coeff}_Y f/k$.*

Proof. We proceed by induction on $q := \mathrm{tr\,deg}_k \mathrm{Coeff}_Y f$. For $q = 0$, we have $f \in k(Y)$, since any X-variable cannot be algebraic over k.
For $q > 0$, let m_i be the first message such that $\mathrm{tr\,deg}_{k(m_1,\ldots,m_{i-1})} k(m_1, \ldots, m_i) = 1$. Then we have $\mathrm{tr\,deg}_{k(m_1,\ldots,m_i)} \mathrm{Coeff}_Y(f) = \mathrm{tr\,deg}_k \mathrm{Coeff}_Y(f) - 1$ and the induction hypothesis applies. Hence Alice sends at least $q - 1 + i \geq q$ messages. □

It is easy to prove with the Primitive Element Theorem from algebra (see [3]), that $\mathrm{tr\,deg}_k \mathrm{Coeff}_Y f + 1$ is an upper bound for $\mathrm{C}^{\rightarrow}(f)$: Alice simply sends a complete transcendence basis of $\mathrm{Coeff}_Y f$ over k to Bob using $\mathrm{tr\,deg}_k \mathrm{Coeff}_Y f$ messages. Furthermore, since fields of characteristic zero are separable, the Primitive Element Theorem for algebraic extensions assures that at most one more message makes the extended field equal to $\mathrm{Coeff}_Y f$. The following lemma helps us to strengthen this result.

Lemma 1. *Let $q = \mathrm{tr\,deg}_k k(f_1, \ldots, f_r)$, then there exist $g_1, \ldots, g_q \in k(f_1, \ldots, f_r)$ with the property that $k(f_1, \ldots, f_r) \subseteq k(g_1, \ldots, g_q)$.*

This lemma assures us that Alice can send a transcendence basis $\{m_1, \ldots, m_q\}$ with the property that $\mathrm{Coeff}_Y f \subset k(m_1, \ldots, m_q)$, with $q = \mathrm{tr\,deg}_k \mathrm{Coeff}_Y f$, so that we do not need the Primitive Element Theorem anymore. This establishes $\mathrm{tr\,deg}_k \mathrm{Coeff}_Y f$ as the correct number of messages for every optimal protocol in the one-way communication model. It can be easily computed as the rank of the matrix, whose columns are given by the gradient of the coefficients of f, see [16].

5 Two-way communication

Next we study the two-way model. Here, we do not have a tight characterization as in the one-way case, but we provide several lower bound techniques that show tight bounds for some specific functions, in particular for the ones arising from our mechanism design problem.

5.1 Substitution method

The idea of the substitution method is to make the first message trivial by adjoining it to the ground field. Since we deal with several ground fields in this section, we write occasionally $\mathrm{C}(f;k)$ instead of $\mathrm{C}(f)$ to stress the underlying ground field.

Lemma 2. *Let P be some protocol for computing a rational function f over some field k that uses r messages. Then there is an extension field $k' \supset k$ with $\operatorname{tr}\deg_k k' \leq 1$ and either $k' \subseteq k(X)$ or $k' \subseteq k(Y)$ and there is a protocol P' for computing f with $r-1$ messages over k'.*

Proof. Let m_1 be the first message of the protocol. Set $k' = k(m_1)$. Then we do not need to send the first message, since it is now known to both players in advance. The claim on the transcendence degree is obvious and also that either $k' \subseteq k(X)$ or $k' \subseteq k(Y)$, since m_1 is the first message. □

Example 2. Assume that Alice and Bob both hold n and m input values. We want to know whether they have at least one value in common. This is called the *set disjointness problem*. It is modelled by the following function

$$\mathrm{Disj}_{n,m}(X,Y) = \prod_{i=1}^{n}\prod_{j=1}^{m}(x_i - y_j).$$

It is now possible to show the following lower bound by induction.

Theorem 2. *For every rational function $g = p/q$ such that q is coprime with $\mathrm{Disj}_{n,m}$, $\mathrm{C}(g \cdot \mathrm{Disj}_{n,m}) \geq \min\{m,n\}$.*

Proof. The proof is by induction on $\min\{m,n\}$. Since q is coprime with $\mathrm{Disj}_{n,m}$, we need at least one message. If this minimum is 1, then the claim is trivial. Otherwise we want to apply Lemma 2. We treat the case $k' \subseteq k(X)$, the other case is symmetric. By renaming variables, we can assume that $x_1,\ldots,x_{n-1},m_1$ are algebraically independent, where m_1 is the message substituted. $x_1,\ldots,x_{n-1}$ and $y_1,\ldots,y_m$ are algebraically independent over $k' = k(m_1)$ We can write $\mathrm{Disj}_{n,m} = g' \cdot \mathrm{Disj}_{n-1,m}$ where g' consists of all terms that contains x_n.

By the inductions hypothesis $\mathrm{C}(gg' \cdot \mathrm{Disj}_{n-1,m}; k') \geq \min\{n-1,m\}$ and by Lemma 2, $\mathrm{C}(g \cdot \mathrm{Disj}_{n,m}; k) \geq \mathrm{C}(gg' \cdot \mathrm{Disj}_{n-1,m}; k') + 1$. This proves the claim. □

5.2 Dimension bound

In this section, we show lower bounds using methods from algebraic geometry and in particular, we relate results about the dimension of a variety to our problem.

We here consider only homogeneous polynomials $f_1,\ldots,f_\ell$ and we want to decide whether the inputs of Alice and Bob lie in the zero set $V(f_1,\ldots,f_\ell)$ of all of them. In this setting, it makes sense to restrict oneself to *projective protocols*, i.e., protocols in which all messages are homogeneous polynomials. Now we can work over the projective space $\mathbb{P}^{2n} := \left(k^{2n+1}\setminus\{0\}\right)/\sim$, where $a \sim b$, iff there exists a $0 \neq \lambda \in k$ such that

$a = \lambda b$. The inputs of the players are embedded canonically into the projective space by adding a 1 as the last coordinate. The input $(x_1, \ldots, x_n, y_1, \ldots, y_n)$ is mapped to the $(x_1, \ldots, x_n, y_1, \ldots, y_n, 1)$. Whenever we speak of the point $(0, 0)$, we mean the corresponding point $(0, \ldots, 0, 0, \ldots, 0, 1)$ in $\mathbb{P}^{2n}$. Since we always work in the affine subspace of $\mathbb{P}^{2n}$ of points with the last coordinate $\neq 0$, we will usually omit the last 1.

Theorem 3. *Let k be an algebraically closed field and $f : k^n \times k^n \to k^\ell$ be a homogeneous polynomial mapping. Assume that $\dim V(f(0, Y)) = q < n$. Then deciding whether $(X, Y) \in V(f) = \{(X, Y) \in k^n \times k^n \mid f(X, Y) = 0\}$ requires at least $n - q$ homogeneous messages even in the two-way model.*

Proof. Let P be a homogeneous protocol where w.l.o.g. Alice is able to eventually decide whether $f(X, Y) = 0$ for $(X, Y) \in k^{2n}$. We embed (X, Y) into $\mathbb{P}^{2n}$ by adding a component 1 as described above. Furthermore, we assume that $n-q-1$ messages are enough for Alice to decide whether $f(X, Y) = 0$; we denote them by $m_1, \ldots, m_{n-q-1}$. Since the messages are homogeneous, it follows in particular that $m_i(0, 0) = 0$, $1 \leq i \leq n - q - 1$. Also $M := V(m_1, \ldots, m_{n-q-1})$ is a projective variety in $\mathbb{P}^{2n}$ with $\dim M \geq n + q + 1$.

Consider the variety $E := V(x_1, \ldots, x_n)$ of dimension n. By the properties of the projective space it follows that $\dim M \cap E \geq q+1$. On the other hand $\dim V(f(0, Y)) = q$, so we can find a point $(0, b)$ in the intersection $M \cap E$ with $f(0, b) \neq 0$. If we run P on the inputs $(0, 0)$ and $(0, b)$, we notice that Alice has the same input and receives vanishing messages in both instances. Hence she is unable to distinguish between $f(0, 0) = 0$ and $f(0, b) \neq 0$; this is a contradiction. □

Example 3. This theorem applies in a straightforward manner to the *equality problem*. We want to decide whether the two inputs are identical. The problem is modelled by the n functions $f_1 = x_1 - y_1, \ldots, f_n = x_n - y_n$. Let $f(X, Y) = (f_1(X, Y), \ldots, f_n(X, Y))$. Of course we have $\dim V(f(0, Y)) = 0$, since only $(0, 0)$ belongs to it. Thus any homogeneous protocol for deciding whether the two inputs are the same requires n messages.

Remark 1. This argument is interesting in another important context as well. If the theorem should extend to the equality problem over every field k of characteristic zero not necessarily closed, then this would rule out the existence of an injective polynomial $p \in k[X, Y]$, which is still an open problem for general fields. On the other hand note that the equality problem can be solved with one message on certain ground sets. For instance, the polynomial $(a, b) \mapsto \frac{1}{2}\big((a + b)^2 + 3a + b\big)$ is a bijection over $\mathbb{N}^2$.

6 Decision problems

Now we allow the players to perform equality tests and comparisons, the latter only over $\mathbb{R}$. In this setting, players do not only compute and send the same messages for every input, but they are allowed to go through a decision tree, where at every node they check whether some two functions of their input are equal. Every internal node in the tree has two successors, one corresponds to the outcome of the fact that the compared functions are equal, one to the outcome that they are not equal. Depending on the result of the

test, the players follow the respective branch in the decision tree. When they reach a leaf, they send a message accordingly. Furthermore, we do not want to compute f but it is sufficient to decide whether the input of the players lies in the zero set $V(f)$ of f or not.

If we allow comparisions, then we can also decide semi-algebraic sets. Every test now has three outcomes, $<$, $=$, or $>$, and every node in the decision tree has three successors. The formal definitions can be found in the full version of this paper.

The communication complexities $\mathrm{C}^{\rightarrow}_{\mathrm{dec}}(f)$ and $\mathrm{C}_{\mathrm{dec}}(f)$ are defined in the same way as $\mathrm{C}^{\rightarrow}(f)$ and $\mathrm{C}(f)$, but over the larger class of protocols with equality tests. Therefore, it is clear that $\mathrm{C}^{\rightarrow}_{\mathrm{dec}}(f) \leq \mathrm{C}^{\rightarrow}(f)$ and $\mathrm{C}_{\mathrm{dec}}(f) \leq \mathrm{C}(f)$ for all rational functions f. In the same way, we define $\mathrm{C}^{\rightarrow}_{\mathrm{dec},<}(S)$ and $\mathrm{C}_{\mathrm{dec},<}(S)$ for semi-algebraic sets S. $\mathrm{C}^{\rightarrow}_{\mathrm{dec},<}(f)$ and $\mathrm{C}_{\mathrm{dec},<}(f)$ shorthand $\mathrm{C}^{\rightarrow}_{\mathrm{dec},<}(V(f))$ and $\mathrm{C}_{\mathrm{dec},<}(V(f))$.

Intuitively, one would expect that these additions would indeed decrease the number of messages needed between the two players. We will see that often, this decrease is very modest.

6.1 Equality tests

Lemma 3. *For every irreducible polynomial $f \in k[X,Y]$ over an algebraically closed field, there is a rational function $h = p/q$ with q and f being coprime, such that $\mathrm{C}^{\rightarrow}_{\mathrm{dec}}(f) \geq \mathrm{C}^{\rightarrow}(hf)$ and $\mathrm{C}_{\mathrm{dec}}(f) \geq \mathrm{C}(hf)$.*

Proof. Consider a protocol P, one-way or two-way, deciding the membership of the inputs in $\Omega := V(f)$, the variety defined by f. Since the possible inputs are infinite (k is closed), almost every input (in the Zariski sense) follows the same path π_0. Let π be the typical path of an element from Ω and let ν be the node where π and π_0 separate for the first time. Following the path π_0 up to ν, we find rational functions $g_1^{(1)}, g_1^{(2)}, \ldots, g_r^{(1)}, g_r^{(2)}$ such that $g_i^{(1)} \stackrel{?}{=} g_i^{(2)}$ is tested by some player. Obviously $g_i := g_i^{(1)} - g_i^{(2)}$ is not identically zero and because we follow the path taken by most inputs, $g_1, \ldots, g_{r-1}$ do not vanish on the given input when we follow π or π_0. Since Ω is a closed set in the Zariski topology, it follows that the elements of Ω reaching ν lie also in $V(g_r)$. Altogether, for an input (X,Y), from $f(X,Y) = 0$, we necessarily have that at least one of $g_1(X,Y), \ldots, g_r(X,Y)$ vanishes, in other words $g := g_1 \cdot \ldots \cdot g_r$ vanishes on (X,Y). Therefore, applying the Nullstellensatz on the numerator of g and noting that $\mathrm{rad}(f) = (f)$ (since f is irreducible), we have $g = h \cdot f$, for a polynomial h, coprime with f. Thus from the protocol P that decides $V(f)$, we get a protocol of the same type that computes hf. □

For the one-way case, we get an almost tight characterisation meaning that performing equality tests does not bring any significant help in the communication task.

Corollary 1. *Given an irreducible polynomial $f \in k[X,Y]$ over an algebraically closed field k, $\mathrm{C}^{\rightarrow}_{\mathrm{dec}}(f) \geq \operatorname{tr\,deg}_k \operatorname{Coeff}_Y f - 1$.*

Proof. We have $\operatorname{tr\,deg} \operatorname{Coeff}_Y g \geq \operatorname{tr\,deg} \operatorname{Coeff}_Y f - 1$, see [6]. □

We cannot apply Lemma 3 to $\mathrm{Disj}_{n,m}$, since $\mathrm{Disj}_{n,m}$ is not irreducible.

Lemma 4. *Let $\ell_1, \ldots, \ell_t$ be linear forms such that any two of them are linearly independent. Let $L = \ell_1 \cdots \ell_t$. Then there is a rational function $h = p/q$ with L and q being coprime, such that $\mathrm{C}_{\mathrm{dec}}(L) \geq \mathrm{C}(hL)$.*

Proof. Consider a protocol P for deciding L and let π_0 be the path taken by almost all inputs. Let π_τ be the path taken by almost all inputs in $V(\ell_\tau)$. Let v_τ be the node where these two paths separate for the first time. Let $g_\tau = p_\tau/q_\tau$ be the rational function tested at v_τ. By Gauss' Lemma, $\ell_\tau | p_\tau$. Thus $g_1 \cdots g_t$ can be written as hL. □

Corollary 2. $\mathrm{C}_{\mathrm{dec}}(\mathrm{Disj}_{n,m}) \geq \min\{n, m\}$.

6.2 Comparisons

Consider the lexicographic monomial ordering on $y_n, \ldots, y_1, x_n, \ldots, x_1$, i.e., a monomal m is smaller than an other one m' if the exponent vector of m is smaller than that of m' in the lexicographic ordering (with variables in the order as above). For a polynomial p, $\mathrm{lt}(p)$ denotes the least term with respect to the chosen monomial ordering. Grigoriev [11] essentially proves the following result.

Lemma 5. *Let $\ell_1, \ldots, \ell_n$ be linear forms such that $x_1, \ldots, x_n, \ell_1, \ldots \ell_n$ are linearly independent. Let V be the union of some hyperplanes, among them $V(\ell_1), \ldots, V(\ell_n)$. Then there is a polynomial f such that $\mathrm{C}_{\mathrm{dec},<}(V) \geq \mathrm{C}(f)$ and $\ell_1, \ldots, \ell_n$ divide $\mathrm{lt}(f)$. The same holds for the seminalgebraic set S defined by $\ell_1 > 0, \ldots, \ell_t > 0$.*

But now the substitution method from Section 5.2 allows us to get rid of one message for each linear form ℓ_i that divides $\mathrm{lt}(f)$ and get a bound for $\mathrm{C}^{\rightarrow}_{\mathrm{dec},<}(\ell_1, \ldots, \ell_n)$ in this way. Note that every indeterminate is substituted by a polynomial, so the $\mathrm{lt}(f)$ does not change and we can perform induction.

Corollary 3. $\mathrm{C}_{\mathrm{dec},<}(\mathrm{Disj}_{n,m}) = \min\{n, m\}$.

Proof. $\mathrm{Disj}_{n,m}$ is a product of linear forms, among them $x_1 - y_1, \ldots, x_n - y_n$. >From the lemma above, we get that these linear forms divide $\mathrm{lt}(f)$. Now we can apply the substitution method to f in the same manner as we did before to $\mathrm{Disj}_{n,m}$. □

In the same way, we can show that the set defined by $x_1 > y_1, \ldots, x_n > y_n$ has communication complexity equal to n. Grigoriev shows the last two bounds using rank based methods.

7 Applications

Using the reduction by Feigenbaum et al. [8] and applying Corollary 3, we get the following lower bound for the multicast cost sharing problem.

Theorem 4. *There is a tree with n players such that every algebraic algorithm with comparisons and equality tests that computes a strategyproof and budget-balanced mechanism for the multicast cost sharing problem sends at least n messages over a linear number of links.*

As a second example, we consider a distributed version of a combinatorial auction with single minded bidders, see e.g. [14,4]. We have n players and a collection of objects. There is a partial order on these objects. Each player bids on exactly one of the objects. A selection algorithm gets the bids of the players and selects those players that get an object. Such an algorithm is called monotone if a selected player that now makes a higher or equal bid on a lesser or equal object still is selected (assuming that all other players bids are the same). >From such an algorithm, one can construct a truthful mechanism via the critical value scheme, see [14,4] for details. We now assume that the players are distributed in a network. The players can send their bids along the links and there has to be one (arbitrary) node at which the selection process takes places. Consider the following scenario: we have a network of two subgraphs connected by one link. Each subgraph contains n players. We have n objects and one player of each subgraph bids on one of them. In particular, such a mechanism can decide whether the bids in one subgraphs are all larger than the corresponding bids in the other subgraph. But this is exactly the semialgebraic set defined by $x_1 > y_1, \ldots, x_n > y_n$, which has communication complexity n.

Theorem 5. *Any distributed truthful critical value scheme for a combinatorial auction with single minded bidders needs to send at least n messages over linearly many links in the worst case.*

8 Conclusions and open problems

>From the point of distributed mechanism design, it would be nice to find more examples to which our techniques can be applied.

>From the point of algebraic complexity theory, the most interesting open problem is to relate the complexity of f to the complexity of $g \cdot f$ in the two-way model. This would directly show that any lower bound for the computation problem is also a lower bound for the decision problem (Lemma 3). For instance, if the rank bound of Abelson [2] yields a lower bound of q for the complexity of f, then it yields a lower bound for $q - 2$ for $g \cdot f$, under some strict assumptions on the functions. However, there are cases where the rank bound is not tight at all. It is not clear to us how to prove the general case. We conclude with the following conjectures.

Conjecture 1. For all irreducible polynomials f over algebraically closed fields,

1. $\mathrm{C}(f) \geq \mathrm{C}(gf) - 1$ for every coprime polynomial $g \neq 0$ and
2. $\mathrm{C}(f) \geq \mathrm{C}(f^j) - 1$ for every $j \geq 1$.

If both conjectures were true, then the decision complexity and the nondetermistic complexity (see the full version for definitions and results) would be closely related to the communicitation complexity of computing f. The case $f = x_1y_1 + x_1^2x_2y_2$, $g = x_2$ shows that $\mathrm{C}(f) \geq \mathrm{C}(gf)$ is in general not true.

Acknowledgments

We would like to thank L. Shankar Ram and Andreas Meyer for carefully reading an early draft of this paper and in particular the latter for the help provided in proving Lemma 1.

References

1. H. ABELSON. Towards a theory of local and global in computation. *Theoret. Comput. Sci. 6(1)*, pp 41-67, 1978.
2. H. ABELSON. Lower bounds on information transfer in distributed computations. *J. Assoc. Comput. Mach., 27(2)*, pp 384-392, 1980.
3. S. BOSCH. *Algebra.* Third Edition. Springer-Verlag, 1999.
4. P. BRIEST, P. KRYSTA, B. VÖCKING, Approximation techniques for utilitarian mechanism design. *Proc. ACM Symp. on Theory of Computing*, 2005.
5. P. BÜRGISSER, M. CLAUSEN, M. AMIN SHOKROLLAHI. *Algebraic Complexity Theory.* Springer, 1997.
6. P. BÜRGISSER, T. LICKTEIG. Test complexity of generic polynomials. *J. Complexity, 8*, pp 203-215, 1992.
7. P. CHEN. The communication complexity of computing differentiable functions in a multicomputer network. *Theoret. Comput. Sci., 125(2)*, pp 373-383, 1994.
8. J. FEIGENBAUM, A. KRISHNAMURTHY, R. SAMI, S. SHENKER. Hardness results for Multicast Cost Sharing. In *Proc. 22nd conf. on Found. of Software Technology and Theor. Comput. Sci. (FSTTCS)*, volume 2556 of *Lecture Notes in Comput. Sci.*, pp 133-144, Springer, 2002.
9. J. FEIGENBAUM, C. H. PAPADIMITRIOU, S. SHENKER. Sharing the cost of a multicast transmission. *J. Comput. Sys. Sci., 63*, pp 21-41, 2001.
10. J. FEIGENBAUM, S. SHENKER. Distributed algorithmic mechanism design: Recent results and future directions. *Proc. 6th Int. Workshop on Discr. Alg. and Methods for Mobile Comput. and Communic.*, pp 1-13, 2002.
11. D. GRIGORIEV. *Probabilistic communication complexity over the reals*. Preprint, 2007.
12. J. HROMKOVIČ. *Communication Complexity and Parallel Computation.* Springer, 1998.
13. E. KUSHILEVITZ, N. NISAN. *Communication Complexity.* Cambridge University Press, Cambridge, 1997.
14. D. LEHMANN, L. O'CALLAGHAN, Y. SHOHAM. Truth revelation in approximately efficient combinatorial auctions. *Proc. ACM Conference on Electronic Commerce*, 2003.
15. Z.-Q. LUO, J. N. TSITSIKLIS. Communication complexity of convex optimization. *J. Complexity, 3*, pp 231-243, 1987.
16. Z.-Q. LUO, J. N. TSITSIKLIS. On the communication complexity of distributed algebraic computation. *J. Assoc. Comput. Mach., 40(5)*, pp 1019-1047, 1993.
17. I. R. SHAFAREVICH. *Basic algebraic geometry 1 – Varieties in projective space.* Second Edition. Springer-Verlag, 1994.
18. A. C. YAO. Some complexity questions related to distributed computing. *Proc. of 11th ACM Symp. on Theory of Comput.*, pp 209-213, 1979.