Report

Algebraic Communication Complexity of Distributed Computation

Author(s):
Bläser, Markus; Vicari, Elias

Publication Date:
2005

Permanent Link:

https://doi.org/10.3929/ethz-a-006788193 →

Algebraic Communication Complexity of Distributed Computation

Markus Bläser, Elias Vicari
{mblaeser,vicariel}@inf.ethz.ch

Institute of Theoretical Computer Science, ETH Zurich

Technical Report 501

8th February 2006

Abstract

In this paper, we present some new results in the field of algebraic communication complexity, the theory dealing with the least number of messages needed between some players, in order to compute the value of a function depending on a input distributed among them. In particular, we introduce 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 completely solve the one-way communication problem with two players, we prove a lower bound for the two-way communication complexity, we show that two important problems, set-disjointness and equality, have high communication complexity and we generalize the former on general networks. The presented lower bound matches the trivial upper bound on trees and even cycles.

We motivate the research by applications in distributed algorithmic mechanism design theory.

1 Introduction

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 $f : k^{n_1} \times \ldots \times k^{n_r} \to k$, we want to know how many messages are necessary between the r players, until one of them has enough information at his/her disposal to be able to compute the function value.

The boolean counterpart, where one deals with boolean functions and bit-strings, is well studied, thanks to its successful application in VLSI-design theory. Knowledge in this more general direction (usually over $\mathbb{R}$) is often embedded in a model, where unrealistically powerful and sometimes non-computable messages (in the sense of standard models) are allowed. Even though lower bounds obtained in this setting evidently apply even in weaker models, we want to develop techniques taking advantages of the algebraic nature of the problem, allowing as permitted operations only additions, multiplications, divisions and possibly comparisons.

Recently Feigenbaum, Papadimitriou and Shenker [8] introduced *distributed algorithmic mechanism design* problems, particularly cost-sharing issues. They characterized some mechanisms by means of their communication complexity and since a lot of them are computable in terms of the basic algebraic operations, together with comparisons, we eventually find then useful to deepen algebraic communication complexity theory. Our goal is to apply it to mechanisms solving a distributed problem, in order to study their intrinsic complexity expressed in terms of number of messages, independently from the kind of encoding.
In this work, we present different kinds of lower bound techniques, mainly using tools from linear algebra, algebra and algebraic geometry.

Most of the presented lower bounds are quite general and can be applied to arbitrary rational functions $f \in k(X, Y)$, where $X = (x_1, \ldots, x_n)$ and $Y = (y_1, \ldots, y_m)$, but they fail, leading to trivial lower bounds if applied in particular to the function solving the *set-disjointness problem*. This is the problem, where the two players try to decide whether the input sets of size n and m, respectively, they have been given are disjoint or not. It aroused our attention, since its high complexity in the boolean model is the key for the proof that group-strategyproof, budget-balanced mechanisms have inherently high communication complexity [7]. Nevertheless, we will show by means of a specific proof that the complexity of the set-disjointness problem in our model is also maximal, lower bound being also equal to $\min(n, m)$. The latter is obviously always a trivial upper bound, since the player with

the least number of variables can transmit all his input with one message per each variable.

2 Related research and new results

As mentioned before, the most developed direction in the field of communication complexity is embedded in a boolean model proposed by Yao [14] and is typically motivated by VLSI design problems. It is mostly of combinatoric nature, whereas our model is based on algebraic structures like fields, where the possible operations are restricted.

Abelson [1, 2] motivated and introduced continuous communication complexity theory on $\mathbb{R}$ assuming some differentiability properties of the involved functions. Luo and Tsitsiklis [12] improved Abelson's results in certain cases making use of algebraic tools, but still under similar strong assumptions.

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 [11].

Our aim is to study 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 function, nevertheless we strongly believe 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 when allowing equality tests, we actually deal with *piecewise rational functions* (in particular not continuous on $\mathbb{R}$) or we can speak about *decision problems*, where the goal is not to compute the function, but decide whether a given input lies in its zero-set.

In particular, in this paper we show the following results:

- We solve completely the one-way case, where only one player is allowed to send messages. In particular, we show that the decision problem is not really easier on algebraically closed fields.
- We prove a general lower bound for the computational problem for the two-way case.
- We show that the set-disjointness problem and the equality problem have both high communication complexity.

- We generalize communication on more than two players, solving in particular the set-disjointness problem on trees and on even cycles.

The motivation for research in this field, as mentioned before, are distributed algorithmic mechanism design problems and, in particular, multicast cost-sharing problems. Applying our result on the set-disjointness problem, we can restate the proof of Feigenbaum et al. [7] about the communication intractability of budget-balanced, group-strategyproof mechanism. The proof there is based on the result in boolean communication complexity which states that for two players, deciding whether their sets chosen from $\{1, \ldots, r\}$ are disjoint or not, takes $\Omega(r)$ bits of communication (for a proof, see [10]). They reduce the mechanism to this problem. But since the input is the characteristic vector of the sets and not the sets itself, the size of the input in this reduction is bigger and therefore, the bound is weaker. In our model, we do not care about the size of the involved numbers, since every field element can be transmitted at unit cost and therefore we get a better bound. In a previous work, Feigenbaum et al. [8] considered a model similar to ours with decision trees but only linear operations on the input.

3 Model and notation

Throughout this paper, k is a field of characteristic zero which we often assume to be algebraically closed. Since our basic results deal with the two-player case, we give them a name: 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 function $f : k^n \times k^m \to k$. Each player may send messages that are rational functions in his input and the messages (s)he has received from the other player. In the end, one of them has to be able to compute f evaluated at the inputs (*computation problem*) or just to decide the value is zero or not (*decision problem*).

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 broadcasting is error-free. In particular we assume that we can send field-elements as they are, that means that every message counts to 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 [12], where the assumptions on the involved functions are different: on one side, we deal uniquely

with polynomial/rational functions, but introducing equality tests (say for $k = \mathbb{R}$) we also allow piecewise continuous functions. Moreover, we also deal with decision problems. The precise algebraic definition of a protocol without equality tests comes next.

Definition 1 (Protocol). A *protocol* P for a multivariate function f is a list of instructions telling the players, the form of the messages they have to send and in which order, in order to let a player to be able to compute $f(X, Y)$. Formally, a two-way protocol P is constituted by:

1. Disjoint input 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 $f \in k(Y, m_1, \ldots, m_r)$

The one-way model is defined in the same way, but forcing $M_{A \leftarrow B} = \emptyset$ and $f \in k(Y, m_1, \ldots, m_r)$.

Definition 2. We define the *two-way communication complexity* for computing f as follows:

$$t_f^{\leftrightarrow} := \min_P r(P)$$

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. A protocol is called *optimal*, if it attains this number of messages. Similarly we define the one-way communication complexity $t_f^{\rightarrow}$.

A message m is said to be *feasible*, if the second property in the definition of a protocol holds. This reflects the intuitive notion of a protocol, where each player forms messages according to what he/she knows at that particular moment.

Example. *Let*

$$f_\star(X, Y) = (y_1 + x_1 y_2 + x_1^2 y_3)x_1 + (y_1 + x_1 y_2 + x_1^2 y_3)^2 x_2 + (y_1 + x_1 y_2 + x_1^2 y_3)^3 x_3$$

then $t_f^{\leftrightarrow} \leq 2$. *Indeed, it is easy to see, that if Alice sends to Bob the value of* x_1, *then he is able to compute the coefficient* $y_1 + x_1 y_2 + x_1^2 y_3$, *which, in turn, enables Alice to compute the whole function. We present this example, since we will show, that in fact* $t_f^{\leftrightarrow} = 2$ *and* $t_f^{\rightarrow} = 3$ *holds.*

4 Transcendence Degree Bound

In this section we deal exclusively with the one-way communication model and we provide a technique to compute exactly the value $t_f^{\rightarrow}$ for every rational function $f \in k(X,Y)$. For such an f we denote with $\mathrm{Coeff}_Y f$ the field extension over k generated by adding the coefficients of f seen as function of Y. Obviously, Bob is able to compute the value of f receiving messages $m_1, \ldots, m_r$ from Alice if and only if $\mathrm{Coeff}_Y f \subset k(m_1, \ldots, m_r)$.

Theorem 1 (Transcendence Degree Bound). *For every field k and rational function $f \in k(X,Y)$, we have*

$$t_f^{\rightarrow} \geq \mathrm{tr\,deg}_k \mathrm{Coeff}_Y f$$

where 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 $q > 0$, consider any message m_i sent by Alice. Then we have for $k' := k(m_1, \ldots, m_{i-1})$:

$$\mathrm{tr\,deg}_{k'(m_i)} \mathrm{Coeff}_Y(f) = \begin{cases} \mathrm{tr\,deg}_{k'} \mathrm{Coeff}_Y f & \text{if } \mathrm{tr\,deg}_{k'} k'(m_i) = 0 \\ \mathrm{tr\,deg}_{k'} \mathrm{Coeff}_Y f - 1 & \text{otherwise} \end{cases}$$

If the ith message m_i, $1 \leq i \leq q$, is such that $\mathrm{tr\,deg}_{k'} k'(m_i) > 0$, then Alice sends at least $q - 1 + i \geq q$ messages because of the induction hypothesis, otherwise the claim is clear. ♢

It is easy to prove with the Primitive Element Theorem from algebra, that $\mathrm{tr\,deg}_k \mathrm{Coeff}_Y f + 1$ is an upper bound for $t_f^{\rightarrow}$: 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. Further, 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:*

$$k(f_1, \ldots, f_r) \subset k(g_1, \ldots, g_q).$$

For the proof, we refer the reader to the appendix.

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 $q = \operatorname{tr\,deg}_k \mathrm{Coeff}_Y f$, so that we don't need the Primitive Element Theorem anymore. This establishes $\operatorname{tr\,deg}_k \mathrm{Coeff}_Y f$ as the correct number of messages for every optimal protocol in the one-way communication model. It can be easily computed for every rational function f: write

$$f(X,Y) = \frac{\sum_{\alpha \in \mathbb{N}^m} g_\alpha^{(1)}(X) y_1^{\alpha_1} \cdot \ldots \cdot y_m^{\alpha_m}}{\sum_{\beta \in \mathbb{N}^m} g_\beta^{(2)}(X) y_1^{\beta_1} \cdot \ldots \cdot y_m^{\beta_m}}$$

where $X = (x_1, \ldots, x_n)$, $Y = (y_1, \ldots, y_m)$ and $g_\alpha^{(i)}$ are almost all zero, $i = 1, 2$. Then

$$\operatorname{tr\,deg}_k \mathrm{Coeff}_Y f = \mathrm{rk} \left[\begin{array}{c|c} \nabla_x \left(g_\alpha^{(1)}\right)_\alpha & \nabla_x \left(g_\beta^{(2)}\right)_\beta \end{array} \right]$$

For a proof of this fact, see [12]. Applying this observation to the function $f_\star$ in the example of third section, one gets $t_{f_\star}^{\rightarrow} = 3$, since the 3×3 identity matrix can be found as submatrix of the gradient-matrix.

At this point, we introduce the ability for the players to send piecewise polynomial messages, enhancing the definition of protocol. In this setting, players are not restricted to compute and send always the same messages independently from the input, but they are allowed to test whether some functions of what they know in a given moment are equal and accordingly decide which message has to be sent. Moreover, we are only interested in checking whether a function vanishes on a given point.

Definition 3 (Protocol with equality test). A *protocol* P for deciding if an input $\{X, Y\}$, distributed over two players A and B, lies in the zero-set of f, is defined on *decision trees*: we consider the straight-line program (cp. Definition 1), where every node expands to a rooted binary tree T, for which the following properties hold:

1. There is a labelling $l : \{\text{roots of trees}\} \to \{A, B\}$, such that for every root v, $l(v)$ denotes the player who goes through the respective tree and sends a message (or decide the function) once a leaf is reached.

2. At every branching α of the tree with root v, a test is performed, i.e.

$l(v)$ checks if

$$g_\alpha^{(1)}(X, m_1, \ldots, m_t) \stackrel{?}{=} g_\alpha^{(2)}(X, m_1, \ldots, m_t), \quad \text{if } l(v) = A$$
$$g_\alpha^{(1)}(Y, m_1, \ldots, m_t) \stackrel{?}{=} g_\alpha^{(2)}(Y, m_1, \ldots, m_t), \quad \text{if } l(v) = B$$

holds, whereby $g_\alpha^{(1)}, g_\alpha^{(2)}$ are rational functions and $m_1, \ldots, m_t$ are the messages sent previously in the protocol (dependent from the input). If the test is positive, he proceeds to the right child, otherwise to the left child.

3. When a leaf β of the tree with root v is reached, player $l(v)$ sends a message $m_\beta(X, Y)$ with the following property:
   - $m_\beta \in k(X, m_1, \ldots, m_t)$, if $l(v) = A$
   - $m_\beta \in k(Y, m_1, \ldots, m_t)$, if $l(v) = B$
4. At the last node of the protocol-tree, $l(v)$ goes through a binary tree as before, but its leaves are now labelled with 0 or 1, with the following meaning:
   - If $\{X, Y\}$ lies in the zero-set of f, then a 0-leaf is reached during the protocol.
   - If $\{X, Y\}$ does not lie in the zero-set of f, then a 1-leaf is reached during the protocol.

The communication complexity $t_f^{\leftrightarrow,=}$ is defined in the same way over a bigger class of protocols than $t_f^{\rightarrow}$ and the decision problem is obviously an easier task than computing a value, therefore it is clear that $t_f^{\star,=} \leq t_f^{\star}$ for all rational functions f, in both the one-way and two-way case.

Intuitively, one would expect that this addition would indeed decrease the number of messages needed between the two players. For the one-way case, it is not decisive when k is algebraically closed, meaning that performing equality tests does not bring any significant help in the communication task.

Theorem 2. *Given an irreducible polynomial $f \in k[X, Y]$ on an algebraically closed field k, we consider protocols P which test whether the input distributed to the players lies in the zero-set of f, where the players are allowed to perform any operation from $\{+, -, \times, /, =\}$ as discussed above. Then*

$$t_f^{\rightarrow,=} \geq \operatorname{tr\,deg}_k \operatorname{Coeff}_Y f - 1$$

when messages are sent by Alice to Bob.

Proof. Consider a protocol P deciding the membership of the inputs in $\Omega := V(f)$, the variety defined by f. Since the possible inputs are infinite (k being 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 call ν the first 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 input following 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 $\subset V(g_r)$. 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 $\mathrm{rad}(\langle f \rangle) = \langle f \rangle$ (since f is irreducible), we have $g = h \cdot f$, for a rational function h, coprime with f. In order to compute g, B then needs to receive at least $t \geq \mathrm{tr\,deg}\,\mathrm{Coeff}_Y\, g \geq \mathrm{tr\,deg}\,\mathrm{Coeff}_Y\, f - 1$ messages, where last inequality is true since h and f are coprime. ♢

5 Rank Bound

In this section, we write the rational function $f \in k(X, Y)$ as the ratio of two coprime polynomials $p, q \in k[X, Y]$. For a protocol P, we denote with h, the rational function that describes the relationship between f and the messages $m_1, \ldots, m_r$ of the protocol, that is, h fulfills one of the following equations:

$$f(X, Y) = h(X, m_1, \ldots, m_r), \quad \text{if } A \text{ computes the value of } f$$

or

$$f(X, Y) = h(Y, m_1, \ldots, m_r), \quad \text{if } B \text{ computes the value of } f$$

for all X, Y.
Given a rational function $f = p/q$, we define its related matrix M_f: $(M_f)_{i,j}$ is the sum of all the coefficients from k obtained in the following way:

- coefficients of all monomials in f containing a product $x_i y_j$;
- product of coefficients of monomials, containing x_i in p and y_j in q, and vice versa.

For example, for $f(X,Y) = x_2(y_1 + x_1y_2)/y_1 = (x_2y_1 + x_1x_2y_2)/y_1$ we get $M_f = \begin{pmatrix} 1 & 1 \\ 3 & 1 \end{pmatrix}$.

In the same spirit, we define the matrix related to a message in a protocol: let $m_1, \ldots, m_r$ be messages from a protocol P for f. Without loss of generality, we assume that it is Alice that retrieves the information for computing f at the end, that is we get a function h depending on X and $m_1, \ldots, m_r$ as explained above.
For a message m_ρ, $1 \leq \rho \leq r$, let M_{m_ρ} be the matrix obtained as follows: consider h as function of m_ρ:

$$\tilde{h}(m_\rho) := h(X, m_1, \ldots, m_n) = \frac{\sum_{u \geq 0} g_u^{(1)} m_\rho^u}{\sum_{v \geq 0} g_v^{(2)} m_\rho^v}$$

whereby $g_u^{(1)}, g_v^{(2)}$ are functions of the other arguments of h. $(M_{m_\rho})_{ij}$ is then the sum of all parameters $\gamma_\kappa \in k$ obtained as follows:

- consider all instances where x_i is contained in m_ρ and y_j is contained in one of its coefficients $g_u^{(1)}, g_v^{(2)}$, $u, v \geq 1$ in $\tilde{h}$, and vice versa. Then every single γ_κ is the product of the k-coefficient of the monomial in m_ρ containing x_i with the k-coefficient of the monomial of the $g_u^{(1)}$s or $g_v^{(2)}$s containing y_j.
- if x_i is contained in numerator (denominator), we also consider y_j contained in $g_0^{(2)}$ ($g_0^{(1)}$), and vice versa.
- we do not consider any of these x_i, y_j instances that have been already treated by any previous message of P involving x_i and y_j yet.

In particular, we consider all messages, even those that do not show up directly in h, but which are only used to build other messages. Note that in this definition, the vice versa is to be taken with respect to the role of x_i and y_j.

For the example from above, we consider following (non-optimal) protocol: messages are $m_1(X) = x_1$, $m_2(Y, m_1) = y_1 + m_1y_2$, $m_3(Y) = y_1$ and the evaluating function h is:

$$f(X,Y) = h(X, m_1, m_2, m_3) = \frac{x_2 \cdot m_2(m_1)}{m_3}$$

so we get:

$$M_{m_1} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}, \quad M_{m_2} = \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix}, \quad M_{m_3} = \begin{pmatrix} 0 & 0 \\ 2 & 0 \end{pmatrix}$$

Note that M_{m_3} does not count the first product x_2y_1 in numerator of h, since it has been already taken into account by M_{m_2}.

Lemma 2. *If $m_1, \ldots, m_r$ are the messages of a protocol P for computing f, then*

$$M_{m_1} + \ldots + M_{m_r} = M_f$$

Proof. We show the contraposition, we therefore assume that $M_{m_1} + \ldots + M_{m_r} \neq M_f$. Therefore we can find a component (i, j), such that the left-hand side is either bigger or smaller than the right-hand side. The first case can not occur, since we do not allow to count a product more than once, while in the second case, this would mean, that we have a combination given by x_i and y_j in f, which has to be (completely) contained in some message or in the coefficients of h. Since a player cannot know both an x-variable and an y-variable without getting one of them from a message, we conclude that protocol P cannot evaluate f. ♢

Lemma 3. *For a protocol P that computes f, we have*

$$\operatorname{rk} M_m \leq 1$$

for every message m in P.

Proof. We assume, that for a message m we have $\operatorname{rk} m \geq 2$. Then we can find a 2×2-submatrix of M_m of rank two, where without loss of generality the involved variables are x_1, x_2, y_1, y_2.
First we note that m has to contain at least one variable x_i and one variable y_j, for $i, j \in \{1, 2\}$, because if m would not contain, say, any y variable, then the x variables in m multiplies with the same factors (in the sense of the definition of M_m) containing y_1, y_2, leading to a matrix of rank 1.
Assume m is sent by A, then, since m is feasible, A has received the information about the y variables of m from some previous messages. This means that the combinations in M_m involving the y variables contained in m have been already counted, leading to a situation as before with a matrix of rank at most 1, a contradiction. ♢

Theorem 3 (Rank Bound). *Every protocol P which computes f needs at least $\operatorname{rk} M_f$ messages, in other words:*

$$t_f^{\leftrightarrow} \geq \operatorname{rk} M_f.$$

Proof. Let $m_1, \ldots, m_r$ be the collection of messages of a protocol P, which computes f. From $M_f = M_{m_1} + \ldots M_{m_r}$ and $\operatorname{rk} m_i \leq 1$ for all i, we get using the subadditivity of the rank function:

$$\operatorname{rk} M_f \leq \operatorname{rk} M_{m_1} + \ldots + \operatorname{rk} M_{m_r} \leq r.$$

So, every protocol needs at least $\operatorname{rk} M_f$ messages. ♢

Example.

- *Scalar product $\langle X, Y \rangle = \sum_i x_i y_i$ has maximal communication complexity, since its matrix $M_{\langle .,. \rangle}$ is the identity matrix.*
- *The rank bound applied to the function $f_\star$ of section 3 leads to a lower bound of 2, indeed we get following matrix:*

$$M_{f_\star} = \begin{pmatrix} 23 & 25 & 25 \\ 5 & 5 & 5 \\ 19 & 19 & 19 \end{pmatrix}$$

6 Rank Bound revisited

In this section we restrict the class of allowed protocols to those not permitting divisions, that is, we deal solely with polynomials. In this case we can give an alternative formulation of the rank bound, giving us the possibility to say something about the complexity of a product of two polynomials.

Definition 4 (Hessian matrix of a polynomial). Given a multivariate polynomial $f \in k[X, Y]$, we define the Hessian-like matrix H_f as follows:

$$\big(H_f\big)_{i,j} = \frac{\partial^2 f}{\partial x_i \partial y_j}\Big|_{X,Y=\mathbf{1}}, \qquad 1 \leq i \leq n, 1 \leq j \leq m$$

whereby $\mathbf{1} = (1, \ldots, 1)^T$

Lemma 4. *Recalling the definition of M_f, we build $\tilde{M}_f$ in the same way, counting now* all *possible monomials $x_i x_j$, that is, we also take the exponents into account. For example $\gamma x_i^{\alpha_1} y_j^{\alpha_2}$, $\gamma \in k$ contains in total $\alpha_1 \cdot \alpha_2$ "copies" of $x_i y_j$ and altogether it brings a term $\alpha_1 \alpha_2 \gamma$ to $(\tilde{M}_f)_{i,j}$, whereby $\gamma \in k$, $\alpha_1, \alpha_2 \in \mathbb{N}$. Then we have:*

$$H_f = \tilde{M}_f$$

Proof. Differentiating f with respect to $x_i y_j$ amounts of dropping all the monomials that do not contain $x_i y_j$. In particular their exponents are now multiplied as coefficients, which we can get by setting all variables to one. This is exactly the definition of $\tilde{M}_f$. ♢

Repeating word-by-word the proofs leading to the rank bound and using Lemma 4, we similarly get that $\operatorname{rk} H_f$ is a lower bound for the communication complexity of f. With the help of this construction, we can now prove following result.

Lemma 5. *If* $\operatorname{rk} M_f = r$ *and* $f(\mathbf{1},\mathbf{1}) = 0$, *then for all* $g \in k[X,Y]$ *with* $g(\mathbf{1},\mathbf{1}) \neq 0$, $f \cdot g$ *has communication complexity of at least* $r - 2$.

Proof. Clearly,

$$H_{fg} = \left(\nabla_X f (\nabla_Y g)^T + \underbrace{g(X,Y) H_f}_{\neq 0} + \underbrace{f(X,Y) H_g}_{=0} + \nabla_X g (\nabla_Y f)^T\right)\Big|_{X,Y=\mathbf{1}}$$

So

$$\begin{aligned} r = \operatorname{rk} H_f &= \operatorname{rk} \underbrace{g(\mathbf{1},\mathbf{1})}_{\neq 0} H_f \\ &= \operatorname{rk}\left(H_{fg} - (\nabla_X f (\nabla_Y g)^T + \nabla_X g (\nabla_Y f)^T)\right)\Big|_{X,Y=\mathbf{1}} \leq \operatorname{rk} H_{fg} + 2 \end{aligned}$$

♢

Note that Luo and Tsitsiklis proved that $\max_{p \in D} \operatorname{rk} H_f\Big|_p$ is a lower bound for the communication complexity involving C^2 functions, where $D \subset D_f$ open [12], so resticting the family of involved functions we get rid of the maximization problem.

7 Variety dimension bound

In this section, we show lower bounds using methods from algebraic geometry and in particular, we apply some known results about the dimension of a variety to our problem. Since a variety is defined by a polynomial functions, we forbid divisions.

The technique presented here works over varieties defined on 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 $a = \lambda b$. It is easy to see, that $\mathbb{P}^{2n} = \left(U_0/\sim\right) \cup U_1$, where

$U_i := \{(x_1, \ldots, x_n, y_1, \ldots, y_n, z) \in k^{2n+1} \setminus \{0\} \mid z = i\}$. Another standard fact about projective spaces is that a polynomial is only then well-defined if it is homogeneous. In order to profit from the nice properties of the dimension of a projective variety, we assume that all polynomials involved in a protocol are homogeneous in the first $2n$ variables, actually we embed k^{2n} into $\mathbb{P}^{2n}(k)$, leaving the $(2n+1)$st variable free.

In this setting, we consider a more general problem: Alice and Bob want to compute a whole set of polynomials $(f_1, \ldots, f_l)$. The definition of the protocol needs only a obvious slight modification and we assume that all polynomials are evaluated by the same player. This assumption can cause a burden of at most $l/2$ messages with respect to an optimal protocol without this restriction. We want to provide a sufficient condition, such that the communication complexity is high. Of course, the trivial upper bound still holds.

Theorem 4. *Let k be an algebraically closed field. Assume $f(X,Y) := (f_1, \ldots, f_l) : k^n \times k^n \to k^l$ vanishes on $(0,0)$ and $\dim V(f(0,Y)) = q < n$. Then, letting Alice compute f has a two-way communication complexity of at least $(n-q)/d$ with polynomial messages of bounded degree d.*

The assumption that $f(0,0) = 0$ is not very strong, since possibly one can consider $\tilde{f}(X,Y) := f(X,Y) - f(0,0)$, which obviously has the same communication complexity as $f(X,Y)$. More limiting is however the fact that $\dim V(f(0,Y))$ has to be small.

Proof. Let P be a protocol with homogeneous messages, where Alice is able to compute the value of f at the end. We embed the messages in $\mathbb{P}^{2n}(k)$ and as explained before, also introducing a fictive variable z.
Furthermore, we assume that $n-q-1$ messages are enough; we denote them with $m_1, \ldots, m_{n-q-1}$. Since the messages are homogeneous, it follows in particular that $m_i(0,0) = 0$, $1 \leq i \leq n-q-1$. Also $M := V(m_1, \ldots, m_{n-q-1})$ (the common zero-set of $m_1, \ldots, m_{n-q-1}$) 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, then it is clear that $\dim M \cap E \geq q+1$. Because $\dim V(f(0,Y)) = q$, we can find $(0,b)$ in the intersection with $f(0,b) \neq 0$. If we run P on $(0,0)$ and $(0,b)$, we notice that A has the same input and she always receives vanishing messages, being then unable to distinguish between $f(0,0) = 0$ and $f(0,b) \neq 0$; obviously this is a contradiction. Since a protocol defined by a set of r arbitrarily polynomials of bounded degree d can be transformed to a protocol with homogeneous messages of size $d \cdot r$, it follows that $r \geq (n-q)/d$. $\diamondsuit$

Example. *This theorem applies in a straightforward manner to every function (or set of functions), which solves the* equality *problem. They are defined to be zero, iff $X = Y$, such has:*

$$\begin{aligned} f_{\mathrm{EQ}} :& k^n \times k^n \to k^n \\ & (X, Y) \mapsto (x_1 - y_1, \ldots, x_n - y_n) \end{aligned}$$

Of course we have $\dim V(f_{\mathrm{EQ}}(0, Y)) = 0$*, since only* $(0, 0)$ *belongs to it.*

Remark. It would be very interesting to prove the same result for the equality decision problem over every field k of characteristic zero not necessarily closed, because this would rule out the existence of an injective polynomial $p \in k[X, Y]$, which is still an open problem for general fields. Indeed, if such a polynomial could be found, then applying it iteratively would lead to an injective polynomial $\tilde{p} : k \times \ldots \times k \to k$, which, in turn could be used to find a protocol with one message for the previous problem: both players compute the value of $\tilde{p}$ on their inputs and check whether the images are equal. The same result would be implied by Theorem 2, if generalized to arbitrarily fields as well. For issue, if the sets are chosen from $\mathbb{N}$, a protocol for the decision problem with one message is induced by the bijection b:

$$\begin{aligned} b :& \mathbb{N} \times \mathbb{N} \to \mathbb{N} \\ & (a, b) \mapsto \frac{1}{2}\big((a + b)^2 + 3a + b\big) \end{aligned}$$

8 Set-disjointness Problem

We consider the following problem: A and B pick a set $X, Y \in \binom{S}{n}$ respectively, $S \subset k$ big enough. They want to exchange messages in order to decide whether there is a common element among their chosen sets, or, in other words, they want to check if $A \cap B \stackrel{?}{=} \emptyset$. For example, a function solving this problem is given by

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

The only bound which seems to apply successfully is the rank bound via Hessian matrix applied to $\tilde{f}_{\mathrm{DISJ}}(X, Y) = \prod_{i,j=1}^{n}(ix_i + jy_j)$, where we add the coefficients i and j in order to break symmetries as much as possible. Of course, if one could prove that $\tilde{f}_{\mathrm{DISJ}}$ has maximal communication complexity, the same would follow for the set-disjointness function. Unfortunately, at

the moment we computed the rank only for $n \leq 13$ using the mathematical software Maple.

Theorem 5. *Set-disjointness problem has full communication complexity, also* $t^{\leftrightarrow}_{f_{\text{DISJ}}} = n$, *for* $n = \min\{|X|, |Y|\}$.

Proof. For this proof, we assume that the players can build more powerful messages. Instead of choosing elements from the field they have access at the moment, they can send any algebraic element of that field. A lower bound proved in this setting holds also in our standard model.
We assume there exists an optimal protocol P which needs at most $n-1$ messages in order to compute f_{DISJ}. Since f_{DISJ} depends on at least $2n$ variables, then at least n of them have to be present in some of the $n-1$ messages. Therefore, at least two variables, say x_1, x_2, cannot be sent alone. Further, we show that it is impossible to find a set of messages $m_{i_1}, \dots m_{i_s}$ and a rational function $\tilde{h}$ such that $x_1 = \tilde{h}(m_{i_1}, \dots m_{i_s})$. If this would be true, in particular, the following obviously holds:

$$k(x_1) \subset k(m_{i_1}, \dots m_{i_s}) =: k'.$$

We want to prove, that in this case we can construct a new feasible protocol P' with the following properties:

- x_1 is sent alone in P';
- The set of messages of P' is a proper subset of the set of messages of P with the exception of the messages in which x_1 is sent alone. In particular, it follows that the protocol is feasible and the number of messages of P' is at most the number of messages of P.

Once we have established this, we get a contradiction by iterating the argument until any variable is sent alone with at most $n-1$ messages, being a contradiction.
In the case where $k(m_{i_1}, \dots m_{i_s})$ is algebraic over $k(x_1)$ we have nothing to prove, since instead of those messages, one can send only the value of x_1 and both the players can recover $k(m_{i_1}, \dots m_{i_s})$. Otherwise, if $\operatorname{tr\,deg}_k k' = r$, then $\operatorname{tr\,deg}_{k(x_1)} k' = r-1$, and hence from the properties of the transcendence basis we can choose $r-1$ elements from $\{m_{i_1}, \dots m_{i_s}\}$ in order to complete a basis for k'. We achieve this by choosing them according to their order; if a message m_{i_k} is not chosen (since it does not increase the transcendence degree), it means that it is algebraic over $k(x_1, M)$, where $M \subset \{m_{i_1}, \dots, m_{i_{k-1}}\}$ is the subset of chosen messages, and therefore useless. This produces the desired basis and hence the new protocol P'.

Further, we denote by $m_1^*(x_1, X_1, Y_1), \ldots, m_s^*(x_1, X_s, Y_s)$, $X_i \subset \{x_2, \ldots, x_n\}$, $Y_i \subset \{y_1, \ldots, y_n\}$ the set of messages of P which contain x_1.
Because of the definition of f_{DISJ}, for every $c \in S$, we have:

$$f_{\text{DISJ}}(X, Y) = 0, \quad \forall X, Y \in \{c\} \times S^{n-1}.$$

This is possible only if at least one of the messages $m_1^*, \ldots, m_s^*$ is of the form $m_i^* = (x_1 - y_1)\tilde{m}_i^*$ or we have $x_1 - m_*(X, Y)$ as a factor in the evaluating function h, where $m_* = y_1$ is a message. Repeating the same argument by changing the role of y_1 to $y_2, \ldots, y_n$, we get that every factor $x_1 - y_i$, $1 \leq i \leq n$ has to be contained in some message or to be present in the evaluation function h. Since x_1 can never be evaluated with the help of other messages, we conclude that in order to build all factors $x_1 - y_i$ or the messages y_k, n messages are necessary, a contradiction. $\diamondsuit$

This result is the algebraic analogous of the mentioned boolean set-disjointness problem. With this is mind, we can restate the proof of Feigenbaum et al. [7] in our model. We abstain ourselves here from introducing the basic notions of the theory of multicast cost-sharing problems, the interested reader may want to refer to the paper by Feigenbaum et al. [8] about this issue or to her survey paper with S. Shenker [9] about distributed algorithmic mechanism design more in general.

Corollary 1. *Any distributed algorithm that exactly computes a budget-balanced, group strategy-proof multicast cost-sharing mechanism by means of the basic operations $\{+, -, \times, /\}$ and that satisfies the four basic properties must send at least $|P|/2$ messages over linearly many links in the worst case, where P is the set of clients residing in the network.*

9 Communication on graphs

A natural generalization of our model is to consider more than two players communicating over an underlying graph. At each node i resides a player having access to a vector $X^i \in k^n$. He can send messages to the players that are his neighbors in the given graph. As usual, messages are sent one after the other and they can be elements of the field extension generated from k adding X^i and the messages received so far. Our previous model is a special case of this one with K_2 as the underlying network. The reduction technique proposed by Abelson [2] works well in our case too, allowing us to reduce the problem to the usual two-way case. We present it again here in a generalization of the set-disjointness problem over an even cycle C_r with

r nodes, r even. The goal is to determine whether all the sets are pairwise disjoint, or in other words, to compute following function:

$$f_{\text{DISJ}}(X^1, \ldots, X^r) = \prod_{\alpha \neq \beta} \prod_{i,j} (X_i^\alpha - X_j^\beta).$$

An upper bound is given by the obvious protocol: players residing in 1 and r send their inputs to player 2 and $r-1$, respectively, with n messages. They proceed then similarly to 3 and $r-2$ with $2n$ messages each, and so on. The protocol ends when these flows of messages have almost met, where then player $r/2$ sends $rn/2$ messages to $r/2+1$, who can compute the function. Altogether, we get

$$t^{C_r}_{f_{\text{DISJ}}} \leq \sum_{i=1}^{r/2} i \cdot n + \sum_{i=1}^{r/2-1} i \cdot n = \frac{r^2}{4} n.$$

To prove a lower bound, we reduce the problem to the two-player case. We denote by $N_{i,i+1}$ the total number of messages flowing from i to $i+1$ and vice versa, setting $r+1 \equiv 1$, for a protocol P. Obviously we have $t^{C_r}_{f_{\text{DISJ}}} = \sum_{i=1}^{r} N_{i,i+1}$. We consider an edge-cut passing through $\{i, i+1\}$ and $\{i + r/2, i + r/2 + 1\}$. If the protocol has to solve the general problem, it has in particular to decide the following subproblem: determine whether the induced set defined by the union of all sets in each partition are disjoint or not. This reduces to solving an usual set-disjointness problem with two players, where each player has a $nr/2$ set and they communicate over the edge-cut. Since we proved the maximal communication complexity of this problem, we get

$$N_{i,i+1} + N_{i+r/2,i+r/2+1} \geq \frac{r}{2} n, \quad i = 1, \ldots, r$$

in P. So, summing over i:

$$t^{C_r}_{f_{\text{DISJ}}} = \sum_{i=1}^{r} N_{i,i+1} = \frac{1}{2} \sum_{i=1}^{r} \frac{r}{2} n \geq \frac{r^2}{4} n$$

An easy upper bound for an arbitrary graph G is always obtainable following the receipt of the example, which then leads to the following result.

Proposition 1. *For every function f over a graph G_r on r nodes, there is a protocol with at most $\overline{R}(G_r)rn$ messages, whereby*

$$\overline{R}(G_r) := \frac{1}{r} \min_{u \in V(G_r)} \sum_{v \in V(G_r)} d(u, v)$$

is the mean radius *of* G_r. *Obviously we have* $\overline{R}(G_r) \leq R(G_r)$, *where* $R(G_r)$ *is the usual radius of the graph.*

Using this proposition and the lower bound technique presented in previous example, we can prove the following results for the set-disjointness problem.

Corollary 2. *For the two-way communication complexity of the set-disjointness problem on general networks, we have:*

$$\begin{aligned} t^{C_r}_{f_{\text{DISJ}}} &= \frac{r^2}{4}n, \quad r \text{ even} \\ \frac{r^2-r}{4}n \leq t^{C_r}_{f_{\text{DISJ}}} &\leq \frac{r^2-1}{4}n, \quad r \text{ odd} \\ t^{T_r}_{f_{\text{DISJ}}} &= \overline{R}(T_r)rn \\ t^{P_r}_{f_{\text{DISJ}}} &= \frac{r^2}{4}n, \quad r \text{ even} \\ t^{P_r}_{f_{\text{DISJ}}} &= \frac{r^2-1}{4}n, \quad r \text{ odd} \\ \frac{r}{2}n \leq t^{K_r}_{f_{\text{DISJ}}} &\leq (r-1)n \end{aligned}$$

where C_r, T_r, P_r *and* K_r *denote a cycle, a tree, a path and a complete graph on* r *vertices, respectively.*

Proof. Almost all the presented bounds follow quite straightforward. An interesting case is:

$$t^{T_r}_{f_{\text{DISJ}}} \geq \overline{R}(T_r)rn.$$

Let $u \in V(T_r)$ be a vertex which leads to $\overline{R}(T_r)$, for every edge $e \in E(T)$, define $T_1^u(e)$ and $T_2^u(e)$ to be the two components of T_r after removing e, where $u \in V(T_1^u(e))$. Further let $w_u(e) := |V(T_2^u(e))|$. Then clearly $r\overline{R}(T_r) = \sum_{e \in E(T_r)} w_u(e)$.
We claim that $|V(T_1^u(e))| \geq |V(T_2^u(e))|$ for all $e \in E(T_r)$, otherwise we can find an edge e', such that the component not containing u is strictly bigger than the other. By moving along the path connecting e' to u, we can assume that e' is such that $u \in e'$. Let be $\overline{u}$ be the vertex incident to e', then we have

$$\begin{aligned} \frac{1}{r}\sum_{v \in V(T_r)} d(\overline{u}, v) &= \frac{1}{r}\sum_{v \in V(T_r)} d(u,v) + |V(T_1^u(e))| - |V(T_2^u(e))| \\ &< \frac{1}{r}\sum_{v \in V(T_r)} d(u,v) = \overline{R}(T_r) \end{aligned}$$

in contradiction with the minimality of u. So $w_u(e)$ always represent the least number of vertices in the components of $T_r \setminus \{e\}$, therefore we have $N(e) \geq nw_u(e)$, for all e, where $N(e)$ is the number of messages over e of a protocol P for DISJ. Hence

$$t^{T_r}_{f_{\text{DISJ}}} = \sum_e N(e) \geq n \sum_e w_u(e) = nr\overline{R}(T_r)$$

for every optimal protocol P. ♢

Remark. By generalizing again the model adding the ability to perform equality tests, it is easy to find a counterexample to the result stated in Theorem 2. We consider the equality decision problem, where the players want to check whether their inputs are equal as vectors. In this case, the function (over $\mathbb{R}$) is the following:

$$f_{\text{EQ}}(X^1, \ldots, X^r) = \sum_{\alpha \neq \beta} \big(\sum_i (X_i^\alpha - X_i^\beta)^2 \big).$$

Since the two-player case has also maximal complexity (as follows from the rank bound), previous bounds are equally true for f_{EQ}, while computing f_{EQ}. Considering this function over P_r, we get a substantially better bound if the players are able to compute equality tests and if we are interested only to verify whether $f_{\text{EQ}} = 0$. An upper bound on the complexity in this case is given by following protocol: player 1 sends his input to 2 who then checks whether $X^1 = X^2$. If it is not the case, then we know that $f_{\text{EQ}} \neq 0$, otherwise it is enough for him to send n messages to 3, since his input and that of 1 are equal. 3 performs the same actions and the protocol stops when it reaches player r. Altogether $(r-1)n$ messages are sent over P_r.

10 Conclusions and open problems

The transcendence degree bound for the one-way communication model exploits the algebraic nature of the problem deeply, whereas the rank bound only takes care of how the x- and y-variables are intertwined, without fully characterizing the field extension leading to the computability of a function in the two-way model. This lack of power is evident when dealing with problems such as the complexity of products of polynomials, where one knows the complexity of each of them, or the complexity in the model with comparisons. So, an important future direction would be to develop other lower bound techniques which turn out to be more appropriate in this context.

One should expect that if a polynomial $f \in k[X,Y]$ has maximal communication complexity, then $g \cdot f$ has maximal communication complexity for every non-zero polynomial $g \in k[X,Y]$. A nice application of this fact would be a direct proof of the communication complexity of the set-disjointness problem, since one can prove easily with the rank bound, that both $f_1(X,Y) := \prod_i (x_i - y_i)$ and $f_2(X,Y) := \prod_{i \neq j}(x_i - y_j)$ have maximal complexity (simply count the products $x_i y_j$ in f_1 and f_2, where one substitutes the difference with an addition) and noting that $f_{\text{DISJ}} = f_1 f_2$.

Possibly, a better knowledge of this case could eventually lead to lower bounds for communication over general networks that do not rely on the simple case (K_2) and takes into account some properties of the underlying graph. Once we have good lower bounds on networks, it is possible to study the hardness of different mechanisms solving distributed issues, possibly pushing the research of canonical hard problem further.

The bound based on the dimension of a variety seems hard to generalize to a wider range of functions, since examples show that the dimension of the induced variety of a function does not seem to be directly linked to its complexity.

It would be interesting to know what happens to the communication complexity if we introduce inequality tests in the decision tree. Proving something in this direction seems to be hard because we cannot rely directly on the properties of the Zariski topology. Parallel results by Bürgisser et al. [5] about the algebraic complexity of deciding the membership in the zero-set of an irreducible polynomial may suggest, that the lower bounds could still hold.

As a last remark, we would like to emphasize the fact a bound general enough would directly answer negatively to the question whether a injective polynomial in $k[X,Y]$ exists for every field of characteristic zero; this indicates strongly how hard the problem actually is.

Acknowledgments

We would like to thank L. Shankar Ram and Andreas Meyer from Institute of Theoretical Computer Science of ETH, Zurich for carefully reading an early draft of this paper and in particular the latter for the help provided in proving Lemma 1.

A Appendix

Here follows the proof of Lemma 1. This could be surely found in the literature, but we present it here for the reader's convenience.

Proof of Lemma 1. We proceed by induction on $d := \operatorname{tr\,deg}_k k(f_1, \ldots, f_r)$ and on r. For $d = r = 1$ the claim is trivial. Assume now $d = 1$ and $r = 2$. Because of the algebraic dependence of f_1, f_2 we can find an non-zero irreducible polynomial $f \in k(X, Y)$ such that $f(f_1, f_2) \equiv 0$. Of course $k(f_1, f_2)$ is then isomorphic to $k(X, Y)/(f)$, since the ideal (f) is the kernel of the map sending f_1 to X and f_2 to Y. It turns out that $V = V(f)$ is a rational curve: choose $\tilde{\varphi}, \tilde{\psi}$ rational functions such that at least one of $f_1(\tilde{\varphi}(t), \tilde{\psi}(t))$ or $f_2(\tilde{\varphi}(t), \tilde{\psi}(t))$ is not a constant, then we obviously have $f(\varphi(t), \psi(t)) \equiv 0$ for $\varphi(t) = f_1(\tilde{\varphi}(t), \tilde{\psi}(t))$, $\psi(t) = f_2(\tilde{\varphi}(t), \tilde{\psi}(t))$.
Therefore we get the following chain of natural isomorphisms:

$$k(f_1, f_2) \simeq k(X, Y)/(f) \simeq k(\varphi(t), \psi(t)) = k(\omega(t)) \simeq k(t)$$

whereby for $k(\varphi(t), \psi(t)) = k(\omega(t))$ (for a suitable $\omega(t)$) we make use of Lüroth's Theorem which states that every subfield of $k(t)$ is of the form $k(\omega(t))$ for some rational function ω. Altogether we get an isomorphism $\sigma : k(f_1, f_2) \xrightarrow{\sim} k(t)$.
Consider $\hat{f}_i := \sigma(f_i)$, then it follows that $\hat{f}_i = \chi_i(t)$ for some $\chi_i \in k(X)$, $i = 1, 2$. Since σ is an isomorphism, the polynomial relation is preserved, and so

$$f_i = \sigma^{-1}(\hat{f}_i) = \sigma^{-1}\big(\chi_i(t)\big) = \chi_i\big(\sigma^{-1}(t)\big) = \chi_i(s), \quad i = 1, 2$$

for $s = \sigma^{-1}(t) \in k(f_1, f_2)$. It clearly implies that $k(f_1, f_2) \subset k(s)$.
For $d = 1$ and an arbitrary $r > 2$, we get the claim inductively as follows: by induction hypothesis we have $k(f_1, \ldots, f_{r-1}) \subset k(s)$ for an s, therefore $k(f_1, \ldots, f_{r-1}, f_r) \subset k(s, f_r) \subset k(s^\star)$, for a suitable $s^\star$, since the transcendence degree of $k(s, f_r)$ over k is 1. Indeed, in order to get $k(f_1, \ldots, f_{r-1}) \subset k(s)$, one applies previous case $r - 2$ times, showing that $f_i = \chi_i(s)$, $1 \leq i \leq r - 1$, being this a stronger condition.
To complete the proof, we show the step from $d-1$ to d with fixed r: consider a subset F of $\{f_1, \ldots, f_r\}$ such that $k(F)$ has transcendence degree $d-1$ over k. By induction hypothesis, we can choose a set $F^\star$ of $d - 1$ elements such that $k(F) \subset k(F^\star)$. Further, by the properties of the transcendence basis, we have that $\operatorname{tr\,deg}_{k(F)} k(F, F^c) = 1$, where $F^c = \{f_1, \ldots, f_r\} \setminus F$, allowing

us to apply the induction assumptions on the extension $k(F, F^c)/k(F)$. We therefore get an element $s^\star \in k(F, F^c)$ such that:

$$k(f_1, \ldots, f_r) \subset k(F, s^\star) \subset k(F^\star, s^\star)$$

with $|F^\star \cup \{s^\star\}| = d$. $\diamondsuit$

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. Bürgisser, M. Clausen, M. Amin Shokrollahi *Algebraic Complexity Theory.* Springer, 1997.

[5] P. Bürgisser, T. Lickteig. Test complexity of generic polynomials. *J. Complexity, 8*, pp 203-215, 1992.

[6] P. Chen. The communication complexity of computing differentiable functions in a multicomputer network. *Theoret. Comput. Sci., 125(2)*, pp 373-383, 1994.

[7] 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.

[8] J. Feigenbaum, C. H. Papadimitriou, S. Shenker. Sharing the cost of a multicast transmission. *J. Comput. Sys. Sci., 63*, pp 21-41, 2001.

[9] 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.

[10] E. Kushilevitz, N. Nisan. *Communication Complexity.* Cambridge University Press, Cambridge, 1997.

[11] Z.-Q. Luo, J. N. Tsitsiklis. Communication complexity of convex optimization. *J. Complexity, 3*, pp 231-243, 1987.

[12] Z.-Q. Luo, J. N. Tsitsiklis. On the communication complexity of distributed algebraic computation. *J. Assoc. Comput. Mach., 40(5)*, pp 1019-1047, 1993.

[13] I. R. Shafarevich. *Basic algebraic geometry 1 - Varieties in projective space.* Second Edition. Springer-Verlag, 1994.

[14] A. C. Yao. Some complexity questions related to distributed computing. *Proc. of 11th ACM Symp. on Th. of Comp.*, pp 209-213, 1979.