

Efficient Projection Orders for CAD

Andreas Dolzmann*

Fakultät für Mathematik und Informatik
Universität Passau
Germany

Andreas Seidl†

Departamento de Matemáticas, Estadística y Computación
Universidad de Cantabria
Spain

Thomas Sturm‡

Fakultät für Mathematik und Informatik
Universität Passau
Germany

MIP-0401
January 22, 2004

Abstract

We introduce an efficient algorithm for determining a suitable projection order for performing cylindrical algebraic decomposition. Our algorithm is motivated by a statistical analysis of comprehensive test set computations. This analysis introduces several measures on both the projection sets and the entire computation, which turn out to be highly correlated. The statistical data also shows that the orders generated by our algorithm are significantly close to optimal.

*<http://www.fmi.uni-passau.de/~dolzmann/>

†<http://www.fmi.uni-passau.de/~seidl/>

‡<http://www.fmi.uni-passau.de/~sturm/>

1 Introduction

Cylindrical algebraic decomposition (CAD) [ACM84a] is an important and powerful tool in real algebraic geometry. One major application is real quantifier elimination [Col75]; one prominent other application is adjacency algorithms [ACM84b].

The pure decomposition algorithm constructs for a finite set $A \subset \mathbb{R}[x_1, \dots, x_r]$ of multivariate polynomials in r variables a decomposition of real r -space into finitely many *cells* such that over each cell all polynomials in A are sign-invariant. With each cell there is simultaneously a *test point* inside this cell constructed in such a form that the sign of the polynomials in A can be effectively determined.

The CAD construction is commonly described as consisting of three phases:

1. The *projection phase* starts out with the original set A of polynomials in r variables. In $r - 1$ projection steps there are $r - 1$ further finite sets A_{r-1}, \dots, A_1 of polynomials with one fewer variable in each step generated.
2. In the *base phase*, the finitely many zeros of the set A_1 of univariate polynomials obtained in the last projection step are computed in some exact symbolic representation. Both the zeros and the open intervals between these zeros are cells that yield a decomposition D_1 of 1-space. The zeros are at the same time test points for themselves. For the intervals there are *test points* explicitly computed.
3. The *extension phase* starts out with the decomposition D_1 of 1-space and successively constructs decompositions D_2, \dots, D_r of higher-dimensional spaces. For obtaining D_{i+1} there are the test points of D_i plugged into the polynomials of A_{i+1} , which yields univariate polynomials. Then one proceeds like in the base phase.

One crucial result, which makes this technique work, is that the projection generates polynomials in such a way that the polynomials in A_2 are sign-invariant over the cells obtained from A_1 in the base phase, and that this remains true for higher dimensions when plugging in test points from the cells.

Consider now a prenex first-order formula φ containing only polynomials from A as terms, and let x_{k+1}, \dots, x_r denote the quantified variables:

$$\varphi = Q_{k+1}x_{k+1} \dots Q_r x_r \psi, \quad Q_i \in \{\exists, \forall\}.$$

It is easy to see that ψ has an invariant truth value over each cell in D_n , and using the test points, this truth value can be effectively determined. Moreover, it is possible to produce quantifier-free descriptions for the cells in D_k . By means of a rather straightforward recursive algorithm, which checks according to the




Figure 1: CAD wrt. the variable orders $y \rightarrow x$ (left) and $x \rightarrow y$ (right)..

quantifiers Q_{k+1}, \dots, Q_r either *at least one* or *all* cells in D_{k+1}, \dots, D_r , a CAD can be exploited for eliminating the quantifiers from φ .

During the past 30 years there have been considerable research and publications on optimizing CAD. For the application to real quantifier elimination, the introduction of *partial* CAD (PCAD) [CH91] has been one major progress, which affects the extension phase.

The vast majority of improvements, however, extremely focused on improving the projection phase [McC84, McC88, Hon90, Laz94, Bro01, SS03]. Most surprisingly, all these contributions concentrating on improved projection operators never examined the relevance of the *order* in which the variables are projected. If one is only interested in pure CAD, then this order can be chosen completely arbitrarily. For quantifier elimination there are restrictions imposed by projecting unquantified variables last and not interchanging \exists with \forall . There is, however, still a most considerable degree of freedom.

We are going to demonstrate by means of a small example that the projection order is highly relevant for the practical complexity of the overall procedure: We consider two circles of radius 2,

$$c_1 = (x + 3)^2 + (y + 1)^2 - 4, \quad c_2 = (x - 3)^2 + (y - 1)^2 - 4,$$

one located at $(-3, -1)$ and the other one located at $(3, 1)$:

Figure 1 shows on the left hand side a CAD for these circles choosing the projection order $y \rightarrow x$. This CAD contains $1 + 3 + 5 + 3 + 1 + 3 + 5 + 3 + 1 = 25$ cells.

On the right hand side, Figure 1 shows, by the way of contrast, a corresponding CAD using the projection order $x \rightarrow y$. Note that the y -axis is drawn horizontally here. We obtain considerably more cells: $1 + 3 + 5 + 7 + 9 + 7 + 5 + 3 + 1 = 41$.

It is obvious that the computation of this second CAD requires more computational resources while delivering a result of equal quality for most purposes.

This paper provides results for determining good variable orders from the input beforehand, i.e., without actually constructing any CAD, in order to then construct the desired CAD wrt. such an order.

The plan of the paper is as follows: In Section 2, we introduce time and space measures for characterizing the complexity of a particular CAD computation. These measures apply partly to the projection phase and partly to the overall computation. We discuss a comprehensive example set of CAD's wrt. all relevant orders and apply all our measures to all results. We then show on a precise formal basis that all our measures are statistically strongly correlated. In particular there are measures on the projection that are suitable for predicting the complexity of the overall computation.

In Section 3, we introduce a heuristic algorithm for efficiently constructing one good projection order wrt. to the relevant measures on the projection phase. This is done *without* trying all relevant projection orders. We reuse our example database from Section 2 to show two facts: First, our heuristic algorithm yields projection orders that are statistically significantly close to optimal. Second, the overhead originating from constructing the good projection order is negligible while the gain is immense.

In Section 4 we finally summarize and evaluate our results.

2 Measures on CAD Computations

In this section we consider the following situation: There is a set A of polynomials in r variables given, which possibly origin from a prenex formula φ . For a variable order $X = (x_r, \dots, x_1)$, which we also write $x_r \rightarrow \dots \rightarrow x_1$, the projection results in projection sets A_r, \dots, A_1 of projection factors; for convenience we set here and in the sequel $A_r := A$. Based on these sets, the CAD tree $D = (D_1, \dots, D_r)$ is constructed. The levels D_1, \dots, D_r of D are CAD's for $\mathbb{R}^1, \dots, \mathbb{R}^r$, respectively. This situation is completely specified by (A, X) or (φ, X) , respectively. We have developed our methods on the basis of the classical Collins–Hong projection. They are, however, also applicable to the other projection operators discussed in the literature.

Our goal is to find for given A or φ a favorable projection order X at the earliest possible stage. For this we want to be able to draw conclusions on the size of the CAD from properties of the intermediate projection sets A_r, A_{r-1}, \dots, A_1 .

Therefore, we are going to systematically investigate numerous complete CAD's wrt. all relevant projection orders, consider certain *measures* on the projection sets as well as on the CAD's, and examine statistical correlations between these measures.

2.1 Measures

There are six measures that we take into account:

1. The number of projection factors in all the projection sets A_r, \dots, A_1 :

$$\text{card}(A, X) = \sum_{i=1}^r |A_i|.$$

2. The sum of total degrees of *all* monomials of *all* polynomials in *all* the projection sets A_r, \dots, A_1 :

$$\text{sotd}(A, X) = \sum_{i=1}^r \sum_{f \in A_i} \sigma(f),$$

where, using the convention $\mathbf{e} = (e_1, \dots, e_r)$,

$$\sigma\left(\sum_{\mathbf{e} \in E} a_{\mathbf{e}} x_1^{e_1} \cdots x_r^{e_r}\right) = \sum_{\mathbf{e} \in E} \sum_{i=1}^r e_i.$$

3. The number of cells in the resulting full CAD:

$$\text{ncad}(A, X) = |D_r|.$$

4. The overall computation time of the full CAD computation in seconds:

$$\text{tcad}(A, X).$$

5. The number of leaves in the partial CAD tree that is generated for quantifier elimination:

$$\text{npcad}(\varphi, X) = \left| \left\{ c \in D_k \cup \dots \cup D_r \mid c \text{ is a leaf} \right\} \right|.$$

6. The overall computation time of the quantifier elimination by partial CAD in seconds:

$$\text{tqe}(\varphi, X).$$

The time measures tcad and tqe depend on the implementation and the machine. We have used the CAD implementation of the second author in the REDLOG package [DS97] of the widespread computer algebra system REDUCE. We have carried out all our computations on a 2.0 GHz Intel Pentium IV using 128 MB of RAM.

The first four measures card , sotd , ncad , and tcad are defined for sets A of polynomials. They can be as well applied to formulas φ by considering the set of polynomials occurring in these formulas. The last two measures npcad and tqe , in contrast, do not make sense outside a quantifier elimination context.

It might appear more natural to consider instead sums of total degrees of monomials in our sotd the total degrees of the corresponding polynomials. Experiments have shown that these measures are highly correlated. Our choice has the advantage to favor sparse polynomials.

For the definition of npcad we have to recall the basic idea of the partial CAD procedure for quantifier elimination [CH91]: When using full CAD, there is the CAD tree $D = (D_1, \dots, D_r)$ computed, and then the matrix formula is evaluated at the test points of the cells in D_r . The truth values thus obtained are propagated down to the corresponding root cell in D_k according to the types of quantifiers. Hence full CAD computation and its use for quantifier elimination are two isolated subsequent steps. For partial CAD, in contrast, one tries to determine truth values for the matrix formulas *during extension* already for cells in D_k, \dots, D_{r-1} , i.e., without fixing all variables to real numbers. For instance, $x_{k+1} - 42 > 0 \wedge x_{k+2}^3 - 4711x_{k+3} = x_{k+4}$ is false for $x_{k+1} = \sqrt{2}$, no matter what x_{k+2}, \dots, x_r are. Whenever one succeeds like this for a cell $c \in D_i$, where $i \in \{k, \dots, r\}$, this cell c need not be further extended. In other words, the partial CAD tree is pruned at this point, and c becomes a leaf. From this discussion, it is not hard to see that

$$|D_k| \leq \text{npcad}(\varphi, X) \leq \text{ncad}(\varphi, X) = |D_r|.$$

It is now also clear, that it is not reasonable to consider the time for the partial CAD construction as a measure: This construction is not isolated from the quantifier elimination but contains a considerable part of the quantifier elimination work, viz. hard computations with algebraic numbers for trial evaluation of the matrix formula. The other part of the quantifier elimination work, viz. solution formula construction, is, in contrast, still an isolated subsequent step. From that point of view, we consider tqe an appropriate counterpart for tcad .

The first two measures card and sotd can be applied after or even during projection. They are candidates for suitable criteria for determining projection orders. The latter four measures ncad , tcad , npcad , and tqe , in contrast, can be applied only after a complete CAD computation or quantifier elimination, respectively. They are going to be used for evaluating the significance of our candidate criteria.

2.2 Computation of the Test Set

We are going to discuss a test set consisting of six CAD examples, mostly from the literature. Each example has been computed wrt. a significant number of projection orders. In fact, we have computed a much more comprehensive example set comprising 48 examples, and selected these six examples as a representative subset for this paper.

All our examples are in fact quantifier elimination examples, which allows us to apply all our measures. Note that also from the point of view of pure CAD, it is not at all a restriction to consider only such examples; they can be considered deliverers of interesting sets of polynomials.

2.2.1 Admissible Projection Orders

On the other hand, we consider the quantifier block structure of the examples in order to restrict the set possible orders to a reasonable subset: Strictly speaking, quantifier elimination by CAD for a prenex formula

$$\varphi = Q_{k+1}x_{k+1} \dots Q_r x_r \psi$$

requires a projection order with two properties: First, all quantified variables x_r, \dots, x_{k+1} are projected before all unquantified variables x_k, \dots, x_1 .

Second, the quantified variables have to be projected *essentially* in the order $x_{k+1} \rightarrow \dots \rightarrow x_r$. This requirement for a fixed projection order is weakened by the fact that in φ like neighbored quantifiers can be equivalently interchanged.

For instance, $\exists x \exists y \forall z \psi$ is equivalent to $\exists y \exists x \forall z \psi$ but *not* generally equivalent to $\exists x \forall z \exists y \psi$. Consequently $x \rightarrow y \rightarrow z$ and $y \rightarrow x \rightarrow z$ are both possible projection orders in this example, while $x \rightarrow z \rightarrow y$ is not. This observation suggests to rewrite $\exists \{x, y\} \forall \{z\} \psi$ thus making visible the *quantifier blocks*.

So returning to the general discussion, we can rewrite our prenex formula φ as

$$\varphi = Q_1 B_1 \dots Q_n B_n \psi,$$

where $Q_i \neq Q_{i+1}$ for $i \in \{1, \dots, n-1\}$ and B_1, \dots, B_n are finite sets of variables. This is the unique *block representation* of a prenex formula. For convenience, let B_0 denote the set of unquantified variables.

From that point of view, *admissible* projection orders for quantifier elimination are characterized by projecting

$$B_n \rightarrow \dots \rightarrow B_1 \rightarrow B_0,$$

while within each block B_i the order can be freely chosen. Obviously, the number

of admissible projection orders is given by

$$\prod_{i=0}^n |B_i|!$$

2.2.2 The Quartic Problem

The `quartic` problem has been suggested by Lazard [Laz88]. It asks for necessary and sufficient conditions on the coefficients of a quartic polynomial to be positive semidefinite:

$$\text{quartic} = \forall x (px^2 + qx + r + x^4 \geq 0).$$

There are $1! \cdot 3! = 6$ admissible orders. The following table presents the computation results:

	order	card	sotd	ncad	tcad	npcad	tqe
1	(x, r, q, p)	7	54	445	4.71	251	7.04
2	(x, r, p, q)	7	54	445	83.39	251	138.18
3	(x, q, r, p)	7	50	417	0.54	235	0.89
4	(x, q, p, r)	7	50	417	1.64	239	2.55
5	(x, p, r, q)	9	66	\perp	>600	\perp	>600
6	(x, p, q, r)	9	66	\perp	>600	\perp	>600

We see that `card` cannot predict differences in `ncad`, where `sotd` can. Note that `ncad` and `tcad` are surprisingly unrelated here. On the other hand, there is one order, viz. no. 3, that is optimal wrt. all criteria. We have automatically aborted all our computations after 10 minutes. Measures that are unknown due to such unfinished computations are marked with \perp .

2.2.3 A Real Implicitization Problem

This example is an exercise on complex implicitization in a textbook [CLO92]. Our formulation asks for a corresponding real implicitization:

$$\text{c1s7} = \exists u \exists v (x = uv \wedge y = uv^2 \wedge z = u^2).$$

The number of admissible orders is $2! \cdot 3! = 12$.

	order	card	sotd	ncad	tcad	npcad	tqe
1	(v, u, z, y, x)	9	25	889	0.09	266	0.16
2	(v, u, z, x, y)	9	25	889	0.09	266	0.15
3	(v, u, y, z, x)	9	25	889	0.15	268	0.20
4	(v, u, y, x, z)	9	25	889	0.14	266	0.22
5	(v, u, x, z, y)	9	25	889	0.14	268	0.19
6	(v, u, x, y, z)	9	25	889	0.15	266	0.19
7	(u, v, z, y, x)	11	36	1571	0.19	508	0.28
8	(u, v, z, x, y)	11	36	1571	0.18	508	0.28
9	(u, v, y, z, x)	11	36	1571	0.17	582	0.29
10	(u, v, y, x, z)	11	36	1571	0.16	580	0.29
11	(u, v, x, z, y)	11	36	1571	0.17	582	0.29
12	(u, v, x, y, z)	11	36	1571	0.17	580	0.28

Compared to the previous example there is much less variation here. Note that a choice of projection order according to card or sotd yields significantly good ncad, tcad, npcad, and tqe.

2.2.4 Range of Lower Bounds

The following formula asks for the possible range of strict lower bounds on the values of a parabola that has no real zeros.

$$\text{as6} = \forall x \forall a \forall b \forall c \exists x' ((a > 0 \wedge ax'^2 + bx' + c \neq 0) \longrightarrow y < ax^2 + bx + c).$$

There are $1! \cdot 4! \cdot 1! = 24$ admissible orders:

	order	card	sotd	ncad	tcad	npcad	tqe
1	(x', c, b, a, x, y)	11	42	4199	0.39	283	0.07
2	(x', c, b, x, a, y)	11	42	4199	0.48	307	0.09
3	(x', c, a, b, x, y)	12	44	5231	0.55	487	0.15
4	(x', c, a, x, b, y)	12	44	5231	0.55	487	0.14
5	(x', c, x, b, a, y)	13	53	6389	1.37	341	0.13
6	(x', c, x, a, b, y)	12	49	6389	0.38	357	0.09
7	(x', b, c, a, x, y)	11	50	4007	0.44	241	0.06
8	(x', b, c, x, a, y)	11	50	4007	0.55	255	0.08
9	(x', b, a, c, x, y)	12	50	5027	0.62	395	0.12
10	(x', b, a, x, c, y)	12	50	5027	0.62	395	0.11
11	(x', b, x, c, a, y)	12	50	5027	0.58	305	0.10
12	(x', b, x, a, c, y)	12	50	5027	0.59	321	0.10
13	(x', a, c, b, x, y)	12	43	5007	0.46	523	0.18

continued on the next page

continued from the previous page

	order	card	sotd	ncad	tcad	npcad	tqe
14	(x', a, c, x, b, y)	12	43	5007	0.40	523	0.17
15	(x', a, b, c, x, y)	11	39	3975	0.38	423	0.16
16	(x', a, b, x, c, y)	11	39	3975	0.40	423	0.15
17	(x', a, x, c, b, y)	11	39	3975	0.28	415	0.16
18	(x', a, x, b, c, y)	11	39	3975	0.29	415	0.14
19	(x', x, c, b, a, y)	11	36	3709	0.52	348	0.16
20	(x', x, c, a, b, y)	11	36	3709	0.21	358	0.11
21	(x', x, b, c, a, y)	9	28	2365	0.27	272	0.11
22	(x', x, b, a, c, y)	9	28	2365	0.27	288	0.11
23	(x', x, a, c, b, y)	9	28	2365	0.13	290	0.08
24	(x', x, a, b, c, y)	9	28	2365	0.14	290	0.08

This is another example with little variation. Here card and sotd do not discover optimal orders wrt. npcad or tqe. The orders that they point at are, however, absolutely acceptable.

2.2.5 Consistency in Strict Inequalities

This problem decides whether the intersection of the open ball with radius 1 centered at the origin and the open cylinder with radius 1 and axis the line $x = 0$, $y + z = 2$ is nonempty. It has been introduced by McCallum and used by Collins and Hong for demonstrating PCAD [McC87, CH91]:

$$\text{con} = \exists z \exists x \exists y (x^2 + y^2 + z^2 < 1 \wedge x^2 + (y + z - 2)^2 < 1).$$

There are $3! = 6$ admissible orders:

	order	card	sotd	ncad	tcad	npcad	tqe
1	(y, x, z)	12	46	251	0.02	51	0.01
2	(y, z, x)	8	43	365	2.24	43	0.04
3	(x, y, z)	11	33	193	0.02	29	0.01
4	(x, z, y)	11	33	193	0.02	37	<0.01
5	(z, y, x)	8	43	365	2.90	47	0.07
6	(z, x, y)	12	46	251	0.02	51	0.01

Comparing the lines 2 and 5 with 3 and 4, we observe that card and sotd contradict each other. Following sotd in these cases yields the best values for ncad, tcad, npcad, and tqe.

2.2.6 Parametrized Collision Problem

The following formula asks if two moving objects, a circle and a square, are going to collide at some time t in the future. The circle is moving with constant velocity $(1, 0)$, while the velocity of the square is parameterized with (v_x, v_y) . This example has been used by Collins and Hong for several fixed choices of (v_x, v_y) [CH91].

$$\text{pcol} = \exists t \exists x \exists y (t > 0 \wedge -1 \leq x - v_x t \leq 1 \wedge \\ -9 \leq y - v_y t \leq -7 \wedge (x - t)^2 + y^2 \leq 1).$$

The number of admissible orders is $3! \cdot 2! = 12$:

	order	card	sotd	ncad	tcad	npcad	tqe
1	(y, x, t, v_y, v_x)	62	250	144971	47.86	6969	3.78
2	(y, x, t, v_x, v_y)	62	250	144971	121.13	6969	4.91
3	(y, t, x, v_y, v_x)	\perp	\perp	\perp	>600	\perp	>600
4	(y, t, x, v_x, v_y)	4678	227337	\perp	>600	\perp	>600
5	(x, y, t, v_y, v_x)	62	248	149925	58.95	13310	4.53
6	(x, y, t, v_x, v_y)	62	248	149925	151.50	13310	7.51
7	(x, t, y, v_y, v_x)	57	323	\perp	>600	\perp	>600
8	(x, t, y, v_x, v_y)	57	323	\perp	>600	\perp	>600
9	(t, y, x, v_y, v_x)	\perp	\perp	\perp	>600	\perp	>600
10	(t, y, x, v_x, v_y)	\perp	\perp	\perp	>600	\perp	>600
11	(t, x, y, v_y, v_x)	\perp	\perp	\perp	>600	\perp	>600
12	(t, x, y, v_x, v_y)	\perp	\perp	\perp	>600	\perp	>600

In this example we observe an immense variety in all measures. In particular, the orders in the lines 3, 9–12 do not even allow to finish projection within the time limit. The probability of failing to finish full as well as partial CAD computation for a random order is $2/3$. Minimal card misleadingly points at such failing orders. Minimal sotd does not point at the optimal order but still at acceptable ones.

2.2.7 The X-Axis Ellipse Problem

The X-Axis Ellipse Problem has been firstly stated by Kahan [Kah75]. It has been formulated as a quantifier elimination problem by Lazard [Laz88]. The problem is to write down conditions such that the ellipse

$$\frac{(x - c)^2}{a^2} + \frac{(y - d)^2}{b^2} = 1$$

is inside the circle $x^2 + y^2 = 1$. We treat the special case $d = 0$:

$$\text{e11} = \forall x \forall y (b^2(x - c)^2 + a^2y^2 = a^2b^2 \longrightarrow x^2 + y^2 \leq 1).$$

There are $2! \cdot 3! = 12$ admissible orders:


	order	card	sotd	ncad	tcad	npcad	tqe
1	(y, x, c, b, a)	32	107	114541	46.64	37883	29.68
2	(y, x, c, a, b)	28	103	51477	45.54	11635	28.70
3	(y, x, b, c, a)	26	89	64625	17.10	21059	15.59
4	(y, x, b, a, c)	27	91	96833	29.71	20431	17.70
5	(y, x, a, c, b)	36	119	⊥	>600	74587	260.70
6	(y, x, a, b, c)	43	129	⊥	>600	⊥	>600
7	(x, y, c, b, a)	122	522	⊥	>600	⊥	>600
8	(x, y, c, a, b)	109	537	⊥	>600	⊥	>600
9	(x, y, b, c, a)	77	345	⊥	>600	⊥	>600
10	(x, y, b, a, c)	74	331	⊥	>600	⊥	>600
11	(x, y, a, c, b)	136	761	⊥	>600	⊥	>600
12	(x, y, a, b, c)	143	751	⊥	>600	⊥	>600

Similar to the previous example, there is probability of only 1/3 to finish full CAD, and there is only a slightly higher probability to finish partial CAD for a random order. Both minimal card and sotd point at an optimal order wrt. tcad, npcad, and tqe and this order is almost optimal wrt. ncad.

2.3 Statistical Correlations

In the informal remarks after presenting for each of our examples the resulting table, we have collected some positive and negative observations concerning the possible correlations between our measures.

We are now going to systematically examine the correlations on a precise formal basis. To motivate this, consider the x -axis ellipse problem from Section 2.2.7. For the 12 admissible projection orders the card values (lower line) and the corresponding sotd values (upper line) differ pretty much as can be seen from the following interpolated plot:



Nevertheless, it is our impression that there is a correlation between these measures. To substantiate this, we introduce a formal notion of correspondence: Consider lists $p = (p_1, \dots, p_l)$ and $q = (q_1, \dots, q_l)$ in $\mathbb{R} \cup \{\perp\}$ like the card column and the sord column for the ellipse example. The value \perp is used to encode that the corresponding value is unknown. Define

$$I = \{ \{i, j\} \mid 1 \leq i < j \leq l \text{ and } p_i, p_j, q_i, q_j \in \mathbb{R} \};$$

this is the set of all unordered pairs of different indices (orders) for which all values are known. On this basis, we define

$$C = \{ \{i, j\} \in A \mid \text{sign}(p_i - p_j) = \text{sign}(q_i - q_j) \},$$

the subset of all pairs of indices where the corresponding values for p are ordered in exactly the same way as the ones for q . Finally, we define

$$C' = \{ \{i, j\} \in A \mid |\text{sign}(p_i - p_j)| \neq |\text{sign}(q_i - q_j)| \},$$

where the corresponding orders between the values are at least not completely opposite.

If we observe $C = I$, then this would suggest very good correspondence. If $\{i, j\} \notin C$ for some $\{i, j\} \in I$, then we would consider it good correspondence to at least have $\{i, j\} \in C'$. This gives rise to the following definition of the *degree of correspondence* between p and q . It is defined for $|I| > 0$:

$$\text{doc}(p, q) = \frac{|C| + 0.5|C'|}{|I|} \in [0, 1].$$

Note that doc is commutative.

It is not hard to see that for fixed list length l , the average over the degrees of correspondences for all possible choices of pairs of lists is 0.5. A value greater than 0.5 thus indicates an above-average correspondence. For our motivating ellipse example above, we obtain, e.g., $\text{doc}(\text{card}, \text{sord}) = 0.97$.

The following table collects for all our examples the doc values for the relevant combinations of card, sord, ncard, tcard, npcad, and tqe. For convenience, the doc's are multiplied by 100 thus rescaling them to percentages:

	quartic	cls7	as6	con	pcol	ell	\emptyset
card–ncad	67	100	89	46	67	67	72
sord–ncad	100	100	88	74	33	67	77
card–tcard	50	81	78	30	50	100	64
sord–tcard	83	81	85	56	33	100	73
card–npcad	58	84	67	67	67	70	68
sord–npcad	91	84	54	93	33	70	70
card–tqe	50	82	62	33	50	100	62
sord–tqe	83	82	47	60	33	100	67

The last column gives the averages over the corresponding lines. These averages indicates that `sotd` has a significantly higher `doc` with all measures on the complete computation than `card`.

We thus consider `sotd` in contrast to `card` to be a suitable indicator after projection for the time and space to be expected for the overall computation.

3 Constructing Good Orders

We recall from the previous section that $\text{sotd}(A, X)$ for a set A of polynomials and a projection order X is the sum of the total degrees of all monomials in all polynomials in all projection sets A_r, \dots, A_1 . For a prenex first-order formula φ we may also speak of $\text{sotd}(\varphi, X)$ referring to the set of polynomials occurring in φ .

The main result of the previous section is the fact that `sotd` is a suitable measure that is correlated with a high degree of correspondence to all measures that one possibly wishes to optimize in order to save computational resources:

1. small size of the full CAD (`ncad`),
2. fast computation time for the full CAD (`tcad`),
3. small size of the partial CAD (`npcad`),
4. fast computation time for quantifier elimination (`tqe`).

On the basis of this result we can conclude from the knowledge of *all* projection sets for *all* admissible orders on the interesting time and space measures listed above without actually performing any base phase or extension phase.

The remaining problem is the following: In order to get a basis for the decision, there are still projection phases wrt. *all* admissible orders to be performed. The worst-case number of admissible orders is the factorial of the number of variables and hence exponential in the word length of the input.

In this section we are going to suggest a heuristic algorithm for finding a good projection order wrt. `sotd`. This algorithm will require quadratically many projection steps in the number of variables and thus in the word length.

We are going to evaluate the quality of the orders determined by our algorithm on the basis of our example set introduced in the previous section. It is going to turn out that the practical computation times absolutely negligible, while the gain obtained from using the computed orders is immense.

3.1 Greedy Projection

We suggest a *greedy* algorithm for finding a good admissible projection order wrt. *sotd*. By greedy, we refer to roughly the following idea: We perform the first projection step wrt. to all possible variables. Then we determine the sum of total degrees for each single obtained sets. We greedily take the best one, throw away all others, and repeat like that until there are no variables left to order.

Algorithm 1 (Greedy Projection) *Input:* a finite set $A \subset \mathbb{R}[x_1, \dots, x_r]$, $B \subseteq \{x_1, \dots, x_r\}$. *Output:* A projection order ω , an intermediate projection set A .

```

1  begin
2       $\omega := ()$ 
3      while  $B \neq \emptyset$  do
4           $A' := \perp$ 
5          for each  $x \in B$  do
6               $A'' := \text{project}(\{f \in A \mid x \text{ not in } f\}, x)$ 
7               $s'' := \sum_{f \in A''} \sigma(f)$ 
8              if  $A' = \perp$  or  $s'' < s'$  then
9                   $A' := A'' \cup \{f \in A \mid x \text{ not in } f\}$ 
10                  $s' := s''$ 
11                  $x' := x$ 
12             fi
13         od
14          $\omega := \omega \circ (x')$ 
15          $A := A'$ 
16          $B := B \setminus \{x'\}$ 
17     od
18 end

```

For sets A of polynomials we can straightforwardly apply our algorithm to the given set A and the set of all variables occurring in the polynomials in A .

For formulas φ , we have to recall our discussion of admissible orders in Section 2.2.1: For a formula $\varphi = Q_1 B_1 \dots Q_n B_n \psi$ in prenex block representation, the admissible projection orders are characterized by projecting the blocks in the order

$$B_n \rightarrow \dots \rightarrow B_1 \rightarrow B_0,$$

where B_0 denotes the set of all unquantified variables. For each of these blocks, the order can be freely chosen. The input A of our algorithm for the first block B_n is the original set of polynomials in φ . In the sequel, for each block B_i with $i \in \{0, \dots, n-1\}$ the input A is the A returned from the preceding call for B_{i+1} .

Note that the algorithm does for several reasons not necessarily find a projection order yielding really minimal *sotd*:

1. We apply a local criterion on projection levels, while `sotd` is a global criterion on entire projections. In other words, the projection order with smallest `sotd` need not necessarily have small sums of total degrees in the set obtained after the first projection step.
2. When judging in Line 7 the quality of a trial projection wrt. some variable x , the set A'' is not perfectly the set that has to be considered: Some of the polynomials sorted out in Line 6 would belong into A'' while others would not. This depends on the next variable to be projected, which we do not yet know.

Nevertheless, we shall see that our greedy algorithm provides orders of very high quality in the following sense: First, they are as a rule close to optimal. We are going to substantiate this in the following section by means of a statistical analysis. Second, and even more important, the computed orders never exceed the time limit in any of our examples, even when there is a high probability of failing.

3.2 CAD Performance with Greedy Projection

The following tables provide a statistical analysis of the performance of the orders generated by our greedy algorithm. It shows that the performance of these orders is considerably above-average. The extra computing time is negligible.

	quartic	cls7	as6
order no. by greedy	3	1	22
time for greedy	0.01	<0.01	0.01
ncad by greedy	417	889	2365
rank w/i ncad	1 of 6	1 of 12	1 of 24
median of ncad	445.00	1230.00	4103.00
mean of ncad	⊥	1230.00	4273.00
tcad by greedy	0.54	0.09	0.27
rank w/i tcad	1 of 6	1 of 12	4 of 24
median of tcad	44.05	0.16	0.42
mean of tcad	>215.04	0.15	0.45
npcad by greedy	235	266	288
rank w/i npcad	1 of 6	1 of 12	5 of 24
median of npcad	251.00	388.00	352.50
mean of npcad	⊥	411.67	364.25
tqe by greedy	0.89	0.16	0.11
rank w/i tqe	1 of 6	2 of 12	10 of 24
median of tqe	72.61	0.25	0.11
mean of tqe	>224.78	0.24	0.11

	con	pcol	ell
order no. by greedy	3	5	3
time for greedy	<0.01	0.13	0.01
ncad by greedy	193	149925	64625
rank w/i ncad	1 of 6	3 of 12	2 of 12
median of ncad	251.00	⊥	⊥
mean of ncad	269.67	⊥	⊥
tcad by greedy	0.02	58.95	17.10
rank w/i tcad	1 of 6	2 of 12	1 of 12
median of tcad	0.02	⊥	⊥
mean of tcad	0.87	>431.62	>411.59
npcad by greedy	29	13310	21059
rank w/i npcad	1 of 6	3 of 12	3 of 12
median of npcad	45.00	⊥	⊥
mean of npcad	43.00	⊥	⊥
tqe by greedy	0.01	4.53	15.59
rank w/i tqe	2 of 6	2 of 12	1 of 12
median of tqe	0.01	⊥	⊥
mean of tqe	0.02	>401.72	>379.36

4 Conclusions

We have shown that the projection order is of crucial importance for the success of CAD computations. For determining the quality of a given projection order, we have shown that there is one single measure, viz. *sotd*, on the projection sets that is correlated to all interesting time and space measures on the computation of a full CAD as well as on quantifier elimination by partial CAD. We have introduced a greedy algorithm for efficiently constructing a good projection order wrt. *sotd*. Our work closes a considerable gap within the CAD framework.

References

- [ACM84a] Dennis S. Arnon, George E. Collins, and Scott McCallum. Cylindrical algebraic decomposition I: The basic algorithm. *SIAM Journal on Computing*, 13(4):865–877, November 1984.
- [ACM84b] Dennis S. Arnon, George E. Collins, and Scott McCallum. Cylindrical algebraic decomposition II: An adjacency algorithm for the plane. *SIAM Journal on Computing*, 13(4):878–889, November 1984.

- [Bro01] Christopher W. Brown. Improved projection for cylindrical algebraic decomposition. *Journal of Symbolic Computation*, 32(5):447–465, November 2001.
- [CH91] George E. Collins and Hoon Hong. Partial cylindrical algebraic decomposition for quantifier elimination. *Journal of Symbolic Computation*, 12(3):299–328, September 1991.
- [CLO92] David Cox, John Little, and Donald O’Shea. *Ideals, Varieties and Algorithms*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, Berlin, Heidelberg, 1992.
- [Co175] George E. Collins. Quantifier elimination for the elementary theory of real closed fields by cylindrical algebraic decomposition. In H. Brakhage, editor, *Automata Theory and Formal Languages. 2nd GI Conference*, volume 33 of *Lecture Notes in Computer Science*, pages 134–183. Springer-Verlag, Berlin, Heidelberg, New York, 1975.
- [DS97] Andreas Dolzmann and Thomas Sturm. Redlog: Computer algebra meets computer logic. *ACM SIGSAM Bulletin*, 31(2):2–9, June 1997.
- [Hon90] Hoon Hong. An improvement of the projection operator in cylindrical algebraic decomposition. In Shunro Watanabe and Morio Nagata, editors, *Proceedings of the International Symposium on Symbolic and Algebraic Computation (ISSAC 90)*, pages 261–264. ACM Press, New York, August 1990.
- [Kah75] William Kahan. An ellipse problem. *ACM SIGSAM Bulletin*, 9(3):11, August 1975. SIGSAM Problem #9.
- [Laz88] Daniel Lazard. Quantifier elimination: Optimal solution for two classical examples. *Journal of Symbolic Computation*, 5(1&2):261–266, February 1988.
- [Laz94] Daniel Lazard. An improved projection for cylindrical algebraic decomposition. In C. L. Bajaj, editor, *Algebraic Geometry and its Applications: Collections of Papers from Shreeram S. Abhyankar’s 60th Birthday Conference*. Springer, Berlin, 1994.
- [McC84] Scott McCallum. *An Improved Projection Operation for Cylindrical Algebraic Decomposition*. Ph.D. dissertation, Computer Sciences Department, University of Wisconsin, Madison, 1984.

- [McC87] Scott McCallum. Solving polynomial strict inequalities using cylindrical algebraic decomposition. Technical Report 87-25.0, RISC, Johannes Kepler University, Linz, Austria, 1987.
- [McC88] Scott McCallum. An improved projection operation for cylindrical algebraic decomposition of three-dimensional space. *Journal of Symbolic Computation*, 5(1–2):141–161, February–April 1988.
- [SS03] Andreas Seidl and Thomas Sturm. A generic projection operator for partial cylindrical algebraic decomposition. In Rafael Sendra, editor, *Proceedings of the 2003 International Symposium on Symbolic and Algebraic Computation (ISSAC 03)*, Philadelphia, Pennsylvania, pages 240–247. ACM Press, New York, NY, 2003.