

Massively parallel computations in algebraic geometry – not a contradiction^a

Janko Böhm (TU Kaiserslautern)

Anne Frühbis-Krüger (Leibniz Universität Hannover)

Mirko Rahn (Fraunhofer ITWM Kaiserslautern)

boehm@mathematik.uni-kl.de

anne@math.uni-hannover.de

mirko.rahn@itwm.fraunhofer.de



^a This work has been partially supported through SPP 1489 and through Project II.5 of SFB-TRR 195 “Symbolic Tools in Mathematics and their Application” of the German Research Foundation (DFG) and was developed with active and financial support from the “R&D Lab : BIG Data Analysis and High Performance Computing” in the Fraunhofer “High Performance Center Simulation and Software Based Innovation”

The design and implementation of parallel algorithms is a fundamental task in computer algebra. The combination of the computer algebra system SINGULAR and the workflow management system GPI-Space provides an infrastructure for massively parallel computations in commutative algebra and algebraic geometry, which has already shown very promising first results. In this note, we give an overview on the current capabilities of this framework by looking into three sample applications: determining smoothness of algebraic varieties, computing GIT-fans in geometric invariant theory, and determining tropicalizations. These applications employ algorithmic methods originating from commutative algebra, sheaf structures on manifolds, local geometry, convex geometry, group theory, and combinatorics, illustrating the potential of the framework in further problems in computer algebra.

This work grew out of a cooperation of the project II.5 “SINGULAR: A new level of abstraction and performance” within SFB/TRR 195 and the “R&D Lab: BIG Data Analysis and High Performance Computing” within Fraunhofer ITWM. The joint project is led by Janko Böhm, Wolfram Decker (both TU Kaiserslautern), Anne Frühbis-Krüger (Leibniz Universität Hannover), Franz-Josef Pfreundt and Mirko Rahn (both Fraunhofer ITWM). In this note, we discuss the general infrastructure and sample applications which were studied in the PhD thesis of Lukas Ristau [23] (TU Kaiserslautern and Fraunhofer ITWM; supervised by Wolfram Decker and Anne Frühbis-Krüger), the Master’s thesis of Christian Reinbold [21] (Leibniz Univer-

sität Hannover and Fraunhofer ITWM; supervised by Janko Böhm and Anne Frühbis-Krüger) and the Bachelor’s thesis of Dominik Bendle [1] (TU Kaiserslautern; supervised by Janko Böhm). The general infrastructure and the smoothness test were implemented in the context of Ristau’s thesis, the fan traversal and GIT-fan computation in Reinbold’s thesis and the computation of tropical varieties in Bendle’s thesis.

Introduction

Parallel computations play an increasingly important role in the development of computer infrastructure. In particular, high-performance clusters have the potential for a game-changing increase in computing power. This raises the question of development and efficient modeling of parallel algorithms in any field relying on large scale calculations. In the context of numerical simulations, massively parallel computations on clusters are nowadays an indispensable tool used for example in weather forecasts or seismic data analysis. In symbolic computing and, in particular, computational algebraic geometry, parallelism also starts to take a center stage [6, 5, 4, 18]. Massively parallel computations have been applied to problems in combinatorics, e.g. in [19]. However, in algebraic geometry, due to the unpredictability of time and memory consumption of key algorithmic tools like Buchberger’s Algorithm for computing Gröbner bases, the use of parallel algorithms has been limited to rather specific contexts. With the goal of a widespread use of parallelism in computer algebra,

an effort to adopt the approach of separating the actual computation from the coordination layer has been started. From an approach of this kind, fields like numerical simulation have already benefited significantly in recent years. In our context, the computer algebra system SINGULAR provides the computational back-end, for the coordination layer the workflow management system GPI-space is used, which employs Petri nets to model the respective algorithms.

Parallel Computing

From the technical point of view, there is more than one paradigm to achieve parallelism. In general, computer scientists identify two different models: in a shared memory based approach several threads have access to the same data in memory, whereas distributed models run independent processes, possibly on different computers, and communicate their results to the other processes as needed. Not being subject to the obvious restriction of running on the same machine, the latter approach is the suitable one for massively parallel computations, that is for applications scaling up to several hundreds or thousands of parallel computations. On the other hand, it also holds many technical challenges to be met by the expertise of computer scientists, in particular, the need for a clever coordination of the computations to ensure scaling of the performance with the number of cores.

Framework Based on GPI-space and SINGULAR

For the coordination layer we use the workflow management system GPI-space, which is developed by Fraunhofer ITWM (Kaiserslautern) [20], and allows for automated parallel execution of algorithms. GPI-space provides a scalable runtime system suitable for huge dynamic environments like scientific computing clusters, but works equally well on sets of computers with heterogeneous hardware, or an individual server. It manages available resources and is tolerant to failure of a node during computation. GPI-Space introduces an application independent global memory layer, that is used to manage the data-flow across a parallel machine during the computation. This memory driven computing approach gives applications access to very large memories of different character (DRAM NVMe). A key feature is the Petri net based workflow engine, which allows for automatic parallelization and dependency tracking, and will be addressed in the subsequent section.

As back-end we use SINGULAR [15], which is a computer algebra system developed for polynomial computation in commutative algebra, algebraic geometry and singularity theory. Its main workhorse is the Groebner basis engine, around which most of its core algorithms are built. SINGULAR exists as a stand-alone software with user interface and as a library version `libSingular`, which is designed for use within other

systems like SAGE or OSCAR. For use in conjunction with GPI-space, we rely on `libSingular` in its existing form, which for this purpose did not undergo any specific changes.

Petri Nets and Parallelism

While sequential algorithms are usually described step by step, Petri nets intrinsically reflect the structure of the algorithms and the state of the computation. As a result they allow for exploiting the parallel structure automatically. A Petri net looks like a directed graph with two kinds of vertices, places (represented as circles in the figures) and transitions (shown as rectangular boxes). The latter contain the elementary functional units, the places can hold marked tokens. These should be understood to represent pieces of data (in the sense of a so-called *colored* Petri net). Transitions link places, consuming one token from each input place and putting one token onto each output place. So a transition can only fire, if all input places hold a token. Figure 1 shows a small example of a Petri net.




Figure 1: Example of a Petri net.

A token placed in place i causes transition s to fire which produces one token in each output place. These tokens are then treated independently by transitions f and g , which place their output tokens on l and r respectively. The transitions f and g can be executed in parallel, that is, the Petri net allows for task parallelism. The last transition j needs to consume one token from each of l and r . Note that the transition j can only fire, if there is at least one token in each of the places l and r . If several tokens are available on the places l and r the Petri net does not determine which of them will be consumed first by the transition j . By meeting this rule, we can allow several parallel instances of each transition to run in parallel without destroying the integrity of the computation (data parallelism). By the use of so-called *conditions*, it is possible to ensure that, depending on properties of the tokens, only certain transitions can fire.

Applications

Smoothness Test

One of the central tasks of computational algebraic geometry is to explicitly construct objects which exhibit prescribed properties, e.g., general members of moduli

spaces or counterexamples to conjectures. In this context, often the question of deciding smoothness of geometric objects arises, since singularities change intrinsic properties of these objects. Thinking in pictures, an algebraic set is smooth, if it looks in a sufficiently small neighborhood of each of its points like an affine space (that is, like a \mathbb{C}^k), see Figure 2.




Figure 2: A singular and a smooth quadric.

It is classically known that smoothness can be decided by considering the Jacobian ideal generated by all minors of the Jacobian matrix of the correct (that is, codimension) size. Unfortunately, for many problems arising in current research, the Jacobian ideal or even a single minor can be too large to be handled with current computational means; for instance, a recently constructed numerical Godeaux surface (a minimal surface of *general type* with geometric genus and irregularity zero and $K^2 = 1$, constructed in the PhD-thesis of Isabel Stenger [24]) is minimally described by 38 generators in 13 variables, which leads to billions of minors.

Taking a completely different approach to deciding smoothness (or in this case more precisely regularity), one can follow ideas of Hironaka’s resolution of singularities and use his termination criterion for the resolution process. This approach has been described by the first and second author in [9]. In a nutshell, given an affine algebraic set X , this criterion locally uses a descending induction on the dimension of ambient spaces $W_i \supset X$ and checks at each level and point $p \in W_i$ the value of a certain invariant, the order of the defining ideal of the given affine algebraic set X in $\mathcal{O}_{W_i,p}$. If at some level i and point p the order exceeds one, the point p is singular. From an algorithmic point of view, we can then terminate the computation with a certificate of non-smoothness; if the order at every level and point has the value one until we reach the lowest level where the ideal of X equals the whole ring $\mathcal{O}_{W_k,p}$, we obtain a certificate of smoothness. Transferring this approach from the local analytic setting at individual points to the algebraic setting making use of Zariski open charts, the descending induction requires the computation of a suitable finite open covering at each level. This leads to a tree of charts with a potentially large number of leaves growing with the level i . Moreover, the deeper the level, the more difficult becomes the computation of the open covering of the current ambient space.

In practice, it is most useful to first proceed with the descending induction in order to pass from a com-

putationally very hard global problem to many computationally simpler local problems. However, before the combinatorial complexity becomes too large, one rather passes in each open set to a relative version of the Jacobian criterion when reaching a given codimension c . Since then both the number of generators of the ideal of X and the codimension are smaller, the computation of the ideal of minors (in each of the open sets) is a significantly easier task than in the beginning.

A simplified version of the Petri net used for modeling the algorithm is shown in Figure 3. Here tokens represent tuples consisting of X , a principal open subset U of affine space and a local ambient algebraic set $W \supset X \cap U$. The input token is placed on i . Depending on whether the codimension c has been reached for a given token (which is tested for by using conditions), either the relative Jacobian criterion (transition Jac) or the descent in dimension (transition $desc$) fires. These transitions add to each of their output tokens the knowledge whether a singularity has been detected. If so, the respective transition $sing$ fires and places a token on the output place o , indicating that the variety is singular. Otherwise the token is deleted by sm or replaced on i , in the respective cases. If there are no tokens left in the Petri net, we know that the input algebraic set in the input principal open subset was smooth. For details on the implementation see [8].

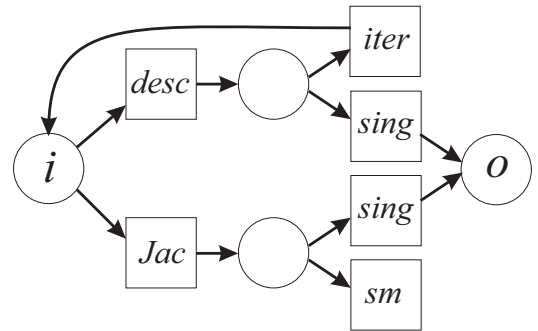


Figure 3: Basic structure of the Petri net for the smoothness test.

It is important to note that, due to the massively parallel execution of the Petri net, the computation will automatically determine, from all possible choices of charts, an open covering of X which leads to the (non-) smoothness certificate in the fastest possible way. This in fact leads to a super-linear speedup with the number of cores: While the computation on 16 cores takes about 53000 seconds, the computation on 128 cores finishes after 3100 seconds, i.e. using the 8-fold number of cores we achieve a speedup by a factor of 17 (with timings measured on a cluster of ITWM with 192 nodes, each with 16 Intel Xeon E5-2670 cores and 64 GB of RAM). The maximal speedup is limited by the total number of leaves of the tree of charts produced by the algorithm.

Another big advantage of the new approach is its lower memory consumption: For the above mentioned Godeaux surface, no individual transition consumed

more than 3,1 GB, whereas the Groebner basis calculation at the end of the classical sequential algorithm (global application of the Jacobian criterion) easily grows beyond 350 GB.

Computing GIT-fans

Geometric Invariant Theory (GIT) aims at associating a reasonable quotient $X//G$ to an algebraic variety X on which an algebraic group G acts. This setup occurs, e.g., when a parameterizing space for classes of geometric objects is to be constructed, where the action of the group G on X arises from isomorphisms between the objects. The homogeneous space X/G is not a good candidate for $X//G$, as it may not be an algebraic variety. For an affine variety X , one can rather define the quotient $X//G$ via the ring of invariant functions of X . However, this quotient may show very little structure: For example, the quotient $\mathbb{C}^2//\mathbb{C}^*$ with the action given by component-wise multiplication is just a point. If we allow for choosing an open subset of X , we obtain a much richer geometry, e.g., $(\mathbb{C}^2 \setminus \{(0, 0)\})//\mathbb{C}^* = \mathbb{P}^1$. For given X , there are usually many choices of open subsets $U \subseteq X$ such that different choices lead to different birationally equivalent quotients $U//G$.




Figure 4: Scaling with the number of cores for computing the Mori chamber decomposition of $\overline{M}_{0,6}$.

To describe this behaviour, Dolgachev and Hu [17] introduced the *GIT-fan*, which is a polyhedral fan¹ describing this variation of GIT-quotients. In [12] a parallel algorithm for computing GIT-fans, which is also designed to make use of symmetries of the setup, has been described. The algorithm is based on symbolic methods from commutative algebra (Gröbner bases), convex geometry (double description) and group theory (orbit decomposition according to the finite symmetry group). Its most important substructure is a traversal of a complete fan which starts at some maximal dimensional cone of the GIT-fan, passes through codimension one faces to all its neighbors, as far as they are not known yet, and iterates. For a short account on the algorithm

¹A polyhedral fan is a finite set of strongly convex rational polyhedral cones such that their faces are again elements of the set and the intersection of any two cones is a face of both.

see [7]. In the Master's thesis of Christian Reinhold [21], this algorithm has been modeled in terms of a Petri net and has been implemented using our framework. We have applied this implementation to compute the Mori chamber decomposition of the cone of movable divisors of the Deligne-Mumford moduli space $\overline{M}_{0,6}$ of 6-pointed stable curves of genus 0. Figure 5 shows the computation time plotted against the number of cores in use (on the cluster described above). We observe an impressive linear scaling up to 640 cores, the maximum number we have tried so far, see Figure 4.




Figure 5: Computation times (in minutes) for the Mori chamber decomposition of $\overline{M}_{0,6}$.

Computing tropical varieties

Tropical geometry is a piecewise linear combinatorial version of algebraic geometry. From an element f of a polynomial ring $R = K[x_1, \dots, x_n]$ over a field K with valuation (usually the Puiseux series field in a variable t) one obtains a piecewise linear function by replacing the field operations with those of the tropical semi-ring

$$a \oplus b := \min(a, b), \quad f \odot g := a + b,$$

and coefficients by their valuation. Given an ideal $I \subset R$ and hence an algebraic variety $V(I)$, one can associate to I the tropical variety $T(I)$ defined as the common corner locus of the tropical polynomials corresponding to elements of I . By the Bieri-Groves theorem [2], the tropical variety corresponds to the set of valuation-tuples of K -points of I . See Figure 6 for a visualization of (fibers in) the family of plane elliptic curves defined by

$$I = \langle t \cdot (x^3 + y^3 + 1) + xy \rangle \in K[x, y]$$

and the associated tropical variety.

Computation of tropical varieties is usually a difficult task, both due to the combinatorial complexity of the resulting tropical variety and the computations of the individual faces, which are based on Buchberger’s algorithm for finding a Gröbner basis. The standard algorithm for this problem starts with one face of the tropical variety and passes to all neighbors by making use of the fact that tropical varieties of prime ideals are irreducible and connected in codimension one [3].




Figure 6: Tropicalization of a family of elliptic curves.

Our massively parallel algorithm for computing tropical varieties is based on the fan traversal through codimension one faces (i.e., on a graph expansion problem), see [1]. Neighboring faces are computed by determining the tropical link via reduction to the curve case and considering Puiseux expansions, see [22]. The traversal is derived from the fan traversal discussed in the subsection on GIT-fans. As in the setting there, the algorithm can make use of symmetries. However, the algorithm is more involved since, although tropical varieties can also be described as fans, they are not of maximal dimension. As a result, when passing through a codimension one face, there will typically exist more than one neighboring cone of maximal dimension.

Our implementation computes the previously unknown tropical Grassmannian $\mathcal{G}_{3,8}$ on 768 cores in less than 20 minutes with a parallel efficiency comparable to the GIT-fan algorithm. This result opens up, for example, the possibility of investigating the relation of the tropical Grassmannian to the Dressian using methods from matroid theory, and the study of the positive Grassmannian.

Conclusion and Outlook

Using GPI-Space in conjunction with SINGULAR and possibly other computer algebra systems with C-library interface provides an efficient infrastructure for massively parallel computations in computer algebra. By modeling algorithms in terms of Petri nets, parallel structures of the problem can be exploited in a transparent and straight-forward manner. Our infrastructure can make use of the potential of large clusters, but is also suitable for use on smaller multi-core servers and personal computers. Based on the example applications considered so far, we believe that a multitude of algorithmic problems in computer algebra can benefit significantly from our approach. Current research addresses,

for example, the computation of Hironaka resolutions of singularities [16], Zeta-functions, positive tropical varieties, integration-by-parts identities for Feynman integrals in high-energy physics via the algorithm developed in [13, 14], and generating functions for tropical numerical invariants via the algorithms introduced in [10, 11].

References

- [1] Bendle, D.: *Massively Parallel Computation of Tropical Varieties*. Bachelor’s Thesis, TU Kaiserslautern (2018).
- [2] Bieri, R.; Groves, J.R.J.: *The geometry of the set of characters induced by valuations*. *Journal für die reine und angewandte Mathematik* **347** (1984), 168 – 195.
- [3] Bogart, T.; Jensen, A.N.; Speyer, D.; Sturmfels, B.; Thomas, R.R.: *Computing tropical varieties*, *J. Symbolic Comput.* **42** (2007), 54 – 73.
- [4] Böhm, J.; Decker, W.; Fieker, C.; Pfister, G.: *The use of bad primes in rational reconstruction*, *Math. Comp.* **84**, 3013–3027 (2015).
- [5] Böhm, J.; Decker, W.; Laplagne, S.; Pfister, G.: *Local to global algorithms for the Gorenstein adjoint ideal of a curve* in Böckle et al. (ed.) *Algorithmic and Experimental Methods in Algebra, Geometry, and Number Theory*, Springer (2018), 51–96.
- [6] Böhm, J.; Decker, W.; Laplagne, S.; Pfister, G.; Steenpaß, A.; Steidel, S.: *Parallel Algorithms for Normalization*. *J. Symbolic Comput.* **51** (2013), 99–114.
- [7] Böhm, J.; Decker, W.; Keicher, S.; Ren, Y.: *Current Challenges in Developing Open Source Computer Algebra Systems*. LNCS **9582** (2016).
- [8] Böhm, J.; Decker, W.; Frühbis-Krüger, A.; Pfreundt, F.-J.; Rahn, M.; Ristau, L.: *Towards Massively Parallel Computations in Algebraic Geometry*. Preprint (2018). <https://arxiv.org/abs/1808.09727>
- [9] Böhm, J.; Frühbis-Krüger, A.: *A smoothness test for higher codimensions*. *J. Symbolic Comput.* **86** (2018), 153–165.
- [10] Böhm, J.; Goldner, Ch.; Markwig, H.: *Tropical mirror symmetry in dimension one*. Preprint (2018). <https://arxiv.org/abs/1809.10659>
- [11] Böhm, J., Bringmann, K., Buchholz, A.; Markwig, H.: *Tropical mirror symmetry for elliptic curves*, *J. Reine Angew. Math.* **732** (2017), 211–246.
- [12] Böhm, J.; Keicher, S.; Ren, Y.: *Computing GIT-fans with symmetry and the Mori chamber decomposition of $\overline{M}_{0,6}$* . Preprint (2016). <https://arxiv.org/abs/1603.09241>
- [13] Böhm, J.; Georgoudis, A.; Larsen, K.J.; Schöenmann, H.; Zhang, Y.: *Complete integration-by-parts reductions of the non-planar hexagon-box via*

- module intersections*. J. High Energ. Phys. **2018:24** (2018).
- [14] Böhm, J.; Georgoudis, A.; Larsen, K.J.; Schulze, M.; Zhang, Y.: *Complete sets of logarithmic vector fields for integration-by-parts identities of Feynman integrals*, Phys. Rev. D **98** (2018) 025023.
- [15] Decker, W.; Greuel, G.-M.; Pfister, G.; Schönemann, H.: *SINGULAR 4-1-0 — A computer algebra system for polynomial computations*. <http://www.singular.uni-kl.de>.
- [16] Frühbis-Krüger, A.; Ristau, L.; Schober, B.: *Embedded desingularization for arithmetic surfaces – toward a parallel implementation*. Preprint (2017). <https://arxiv.org/abs/1712.08131>
- [17] Dolgachev, I.V.; Hu, Y. and Yi Hu, *Variation of geometric invariant theory quotients. (With an appendix: “An example of a thick wall” by Nicolas Ressayre)*., Publ. Math., Inst. Hautes Étud. Sci. **87** (1998), 5–56.
- [18] Jensen, A.; Sommars, J.; Verschelde, J., *Computing Tropical Prevarieties in Parallel*, Preprint (2017), <https://arxiv.org/abs/1705.00720>
- [19] Jordan, C.; Joswig, M.; Kastner, L., *Parallel Enumeration of Triangulations*, Electronic Journal of Combinatorics **25**, (2018)
- [20] Pfreundt, F.-J. ; Rahn, M.; et al.: *GPI-space*, Fraunhofer ITWM Kaiserslautern, <http://www.gpi-space.de/>.
- [21] Reinbold, Ch.: *Computation of the GIT-fan using a massively parallel implementation*. Master’s Thesis, Leibniz Universität Hannover (2018).
- [22] Ren, Y; Hofmann, T.: *Computing Tropical Points and Tropical Links*. Discrete and Computational Geometry **60** (2018), 627–645.
- [23] Ristau, L.: *Using Petri nets to parallelize algebraic algorithms*, Dissertation, TU Kaiserslautern (2019)
- [24] Stenger, I.: *A Homological Approach to Numerical Godeaux Surfaces*, Dissertation, TU Kaiserslautern (2018)

Workshop-Förderung der Fachgruppe:

Sie veranstalten einen Workshop zu einem Thema aus dem Bereich der Computeralgebra und könnten mit einer kleinen finanziellen Unterstützung den Workshop deutlich interessanter oder effektiver gestalten? Die Fachgruppe Computeralgebra unterstützt Workshops mit bis zu 1000,- Euro.

Anträge können bis mit einer kurzen Beschreibung des Workshops (ca. 1 DIN A4 Seite; kurze Beschreibung des Gebiets, Thema des Workshops, Zielgruppe, Budget-Planung) und einer Darstellung, inwiefern diese Förderung einen deutlich erkennbaren Beitrag zum Gelingen des Workshops und zur Nachwuchsförderung liefert, an den Sprecher der Fachgruppe gerichtet werden: **kemper@ma.tum.de**, bitte „**Workshop-Förderung**“ im Betreff angeben.

