Reconstructing Crystalline Structures from Few Images Under High Resolution Transmission Electron Microscopy

Peter Gritzmann and Sven de Vries

Zentrum Mathematik, Technische Universität München,
80290 München, Germany
E-mail: gritzman@ma.tum.de, devries@ma.tum.de

1 Introduction

The present paper is motivated by the demand from material sciences to reconstruct crystalline structures given through their images under high resolution transmission electron microscopy (HRTEM) in a certain limited number of directions. In particular, [31] and [22] show how a quantitative analysis of images from high resolution transmission electron microscopy can be used to determine the number of atoms on atomic columns in certain directions; see Sects. 2 and 3. Mathematically, this leads to the problem of reconstructing finite lattice sets from certain of their marginal sums; see Sect. 4.

A full 3D-reconstruction of interfaces at the atomic scale may in particular be used to fine-tune the production processes in semiconductor industry. In fact, as silicon wafers are coated with an amorphous layer of oxidized silicon, methods based on external surface scanning – like scanning tunnel microscopy – cannot be used.

There are many other applications and potential applications of the mathematical techniques developed in this project in image processing, graph theory, scheduling, statistical data security, game theory, to name just a few areas (see e.g. [32], [30], [8], [21], [19], [11], [12]). In fact, discrete tomography can be regarded as a prime paradigm for a whole new field dealing with (typically ill-posed) discrete inverse problems.

Acknowledgement: The project was carried out in cooperation with the Institute for Semiconductor Physics, Frankfurt/Oder and Tietz Video & Image Processing Systems GmbH, Gauting. The authors are deeply indebted to Peter Schwander for his constant support, particularly for helping with the physics and the simulation of HRTEM and QUANTITEM.

2 High Resolution Transmission Electron Microscopy

In the early 1930's conventional light-microscopy was pushed to its limits to provide a resolution of about 0.5 μm. In 1931/32, Ernst Ruska (*Nobel Prize in Physics* in 1986) build the first imaging device utilizing electron beams. The

W. Jäger et al. (eds.), *Mathematics - Key Technology for the Future*

first transmission electron microscopes *(TEM)* were mainly used to study biological specimens. A TEM works essentially like a light microscope except that the light beam is replaced by an electron beam and the optical elements are replaced by electron-optical analogues (like electromagnetic or electrostatic lenses). Images are produced by measuring the electron intensities after the beam passed through the specimen; regions with heavier atoms or greater thickness diffract more electrons away from the detector and can hence be detected.

High resolution transmission electron microscopy further improved the resolution possible for very thin specimen. Since for thin samples only marginal absorption occurs within the HRTEM imaging with high acceleration potential of about 200 keV the imaging process is driven by other principles. As there is no loss of energy the resulting wave function Ψ of the interaction between the electron beam and the electrostatic potential of the crystal is governed by the time independent *Schrödinger equation*

$$\nabla^2\Psi(r) + \frac{8\pi^2 me}{\hbar^2}[E + V(r)]\Psi(r) = 0,$$

where e denotes the electronic charge, E is the acceleration potential of the microscope, $\hbar$ is Planck's constant, m is the mass of the electron, and $V(r)$ denotes the crystal potential at position r. The interaction creates an electron wave emanating from the 'bottom side' of the crystal (if the electron beam hits the specimen from 'above'). It is this wave that is magnified by electrostatic lenses to obtain the final image. Unfortunately, at this high magnification the lens distorts the image rather strongly; to worsen matters, the distortion depends very sensitively on imaging conditions that cannot be fully assessed. So traditionally, given an HRTEM image, the art is to do simulations of different objects under different imaging conditions until a simulated phantom is created that matches the obtained image. In the next subsection we will sketch one particular simulation technique, the *Multislice Method* due to Cowley and Moodie [5].

3 Simulation and Analysis of HRTEM Images

We will sketch how the electron wave at the exit face of a specimen can be calculated from the specimen's crystal potential and the imaging conditions by using the multislice method. This allows us to simulate HRTEM and compute realistic (though small) phantoms under full control of the imaging process. Then we describe a vector analysis technique that can be applied to recover height information. It is called QUANTITEM (quantitative analysis of the information provided by transmission electron microscopy) and was developed by Schwander, Kisielowski, Baumann, Seibt, Kim, and Ourmazd [31], [22].

The main issue with the simulation is to solve the Schrödinger equation in order to obtain a description of the wave at the exit face of the specimen.

For an appropriate simulation it is important to take into account that the involved energies are of quite different orders of magnitude: on the one hand the low energy potential field of the crystal and on the other hand the electron beam of about 200 keV. This difference in scale leads to an almost complete absence of backscattering, i.e. most electrons of the beam are scattered only forward, with the scattering occurring only within a very small angular range. Hence the specific conditions allow it to describe the interaction in terms of a wave propagation, permitting a slice-by-slice simulation with slices orthogonal to the direction of the beam. The full potential of each slice is projected onto the plane of the slice closer to the source of the electron beam (entry face). It turns out that the interaction of the beam with the projected potential on the entry face can be described approximately by a phase shift for the beam's wave. The propagation of the resulting wave from one entry face to the next is then done by using Fresnel approximation.

The leaving wave is modeled as the sum of many spherical (actually, in this approximation, paraboloidal) waves. Their joint effect on the next entry plane is modeled by a convolution integral. Computationally, the convolution is performed in Fourier space; for details see [5]. The magnifying electrostatic lens is modeled by a similar procedure.

For the simulation of HRTEM we utilized the package *EMS* by Stadelmann [33]. As specimen we always used a wedge of pure silicon with some cells at the top of atomic columns (viewed in direction $\langle 001 \rangle$) removed at random. The heights of the final probe varied between 0 and 14 atoms; see Fig. 1a for the depiction of such a height-field (the tiny numbers give the number of atoms in the corresponding column).

Using EMS we then computed the simulated image for the resolutions 2048×2048 subdivided into 30 slices orthogonal to the beam. (Higher resolutions were used to verify that the simulation stabilized.) The simulated image is given in Fig. 1b.

It turns out that for important specimen examples like silicon (under viewing direction $\langle 100 \rangle$) or Germanium (under viewing directions $\langle 100 \rangle$ or $\langle 110 \rangle$) the wave function can be approximated very well at the exit face as the superposition of only two Bloch waves. The vectors whose coordinates are the intensities of all fundamental solutions of the PDE that satisfy the necessary boundary and initial value conditions all lie on an ellipse. Furthermore, the angle (to a fixed reference point on the ellipse) corresponds to the height of the sample.

For a view in a main lattice direction an atomic column 'influences' only a small region, in effect, the height information is localized. It is observed in [22] that after being magnified and distorted by the lens the images can still be approximated well by two basic images. In [22] this is utilized by segmenting the picture into small rectangles that are influenced essentially only by a single atomic column. Each of these rectangles is considered an image vector on which a principal component analysis is based.

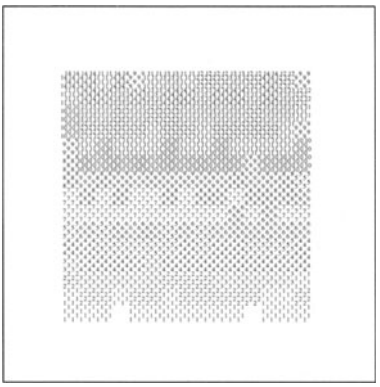

(a) Height of the wedge.

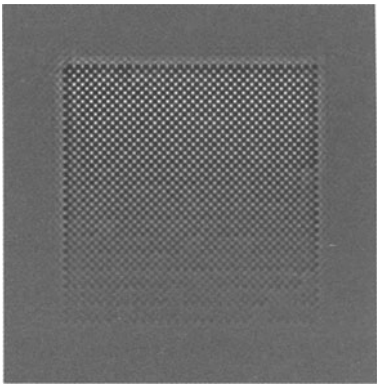

(b) HRTEM Simulation.

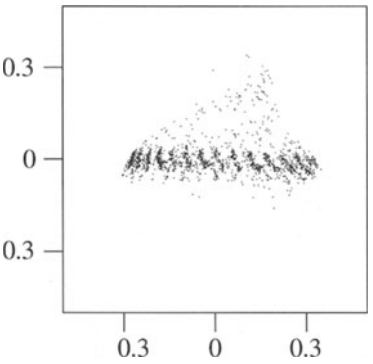

(c) Coordinates of the image-cells with respect to the largest and second largest eigenimage.

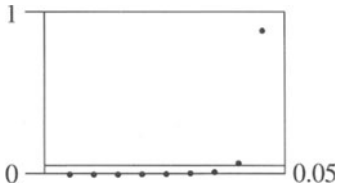

(d) Sizes of the largest eigenvalues.

Fig. 1. HRTEM-simulation of a $15 \times 15 \times 15$-wedge at resolution 4096×4096. The tiny(!) numbers in (a) show the heights of the corresponding columns

In our experiments we used the cells of the Voronoi diagram of the projection of the object to determine the (in this case: rhombic) cells. Then the pixels within a cell were used to construct the image vectors. We computed the principal components for the resolution 2048×2048; their projections onto the two eigenimages with the largest eigenvalues are plotted in Fig. 1c.

The quality of the simulation and analysis can be inferred from the values of the largest eigenvalues. Theory says, that there should be two significant eigenvalues and all others should be smaller than 0.05; see Fig. 1d. For a resolution of 2048×2048 this turned out to be true; resolutions smaller than 2048×2048 did not suffice to show this behavior. Actually, one would expect the largest two eigenvalues to be similar in magnitude; but as the wedges

we studied are very thin we see only a small piece of the ellipse which looks locally like a line; so one eigenvalue dominates the other.

The projection onto the main eigenimages in Fig. 1c reveals a lot about the atomic structure under consideration. Note that from left to right the 15 clusters correspond to columns of height 0 through 14 with only heights 0 and 1 not being separated. The band structure of the clusters is partly due to errors at the boundary of the sample and to the fact that the sample is still very thin compared to the extinction length that lies somewhere between 150 and 200 atoms. But it may also exhibit some additional information related to the relative position of the atoms that may be utilized in the subsequent reconstruction phase.

Since there was no substantial change between resolutions 2048×2048 and 4096×4096 the simulation seems to be stable. This is also supported by the striking similarities between the two eigenimages at medium resolution (Figs. 2b, e) to those at higher resolution (Figs. 2c, f), respectively. Even though the second most important eigenimages (Figs. 2e, f) are of very little influence here, the slight asymmetry between their top and bottom part can be nicely explained by the fact that the bottom-top direction is the direction along which the wedge rises.

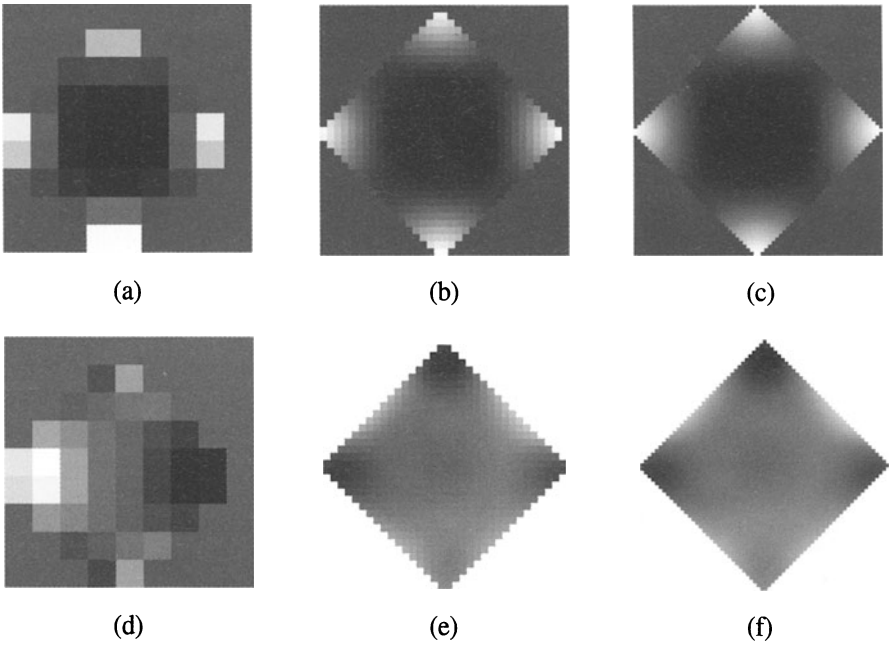

Fig. 2. Comparison of the most *(top)* and second most *(bottom)* dominating eigenimages for HRTEM-simulations of a $15 \times 15 \times 15$-wedge at resolution 512×512, 2048×2048, and 4096×4096 *(from left to right)*

4 The Mathematical Model

In principle, the imaging process described in the previous section essentially provides us with the information how many atoms of a given object interact with given 'sharply focussed electron beams' in a certain small number of viewing directions. In a simple but already highly relevant model suggested by Peter Schwander and Larry Shepp the atoms are identified with lattice points in 3-space while the electron beams are modeled as lines parallel to a given direction. Due to the crystalline structure of the samples and the affine invariance of the basic problem it turns out that it is mostly enough to consider subsets of the integer lattice.

Since in practice, one degree of freedom for moving the imaging device is used to control the position of the object, the view directions for which data are provided all lie in one plane (i.e., virtually, the microscope is rotated around one axis). Hence the 3D-problem lends itself to a 2D-slice-by-slice reconstruction.

Now, let us be more precise. Let F be a finite subset of $\mathbb{Z}^n$, let S be a line through the origin, and let $\mathcal{A}(S)$ denote the set of all lines of Euclidean n-space $\mathbb{E}^n$ that are parallel to S. Then the *(discrete) X-ray of F parallel to S* is the function $X_S F : \mathcal{A}(S) \to \mathbb{N}_0 = \mathbb{N} \cup \{0\}$ defined by

$$X_S F(T) = |F \cap T| = \sum_{x \in T} \chi_F(x), \qquad \text{for } T \in \mathcal{A}(S),$$

where χ_F denotes the characteristic function of F. Note that the integral in the definition of the standard X-ray transform of some function in computerized tomography reduces to a finite sum here; see [27] for an exposition of the mathematics of computerized tomography. Of course, for the practical applications outlined before only the cases $n = 2$ and $n = 3$ are relevant. However, many of the following results hold also in higher dimensions, and for some of the other applications of discrete tomography mentioned in the introduction higher dimensions are important.

In the following let $\mathcal{F}^n = \{F : F \subset \mathbb{Z}^n \wedge F \text{ is finite}\}$ and $\mathcal{S}^n = \{\lin\{u\} : u \in \mathbb{E}^n \setminus \{0\}\}$. The elements of $\mathcal{F}^n$ are called *lattice sets*. Quite typically, we have some additional a priori information available. This is modeled by considering a suitable subset $\mathcal{G}$ of $\mathcal{F}^n$. For instance, $\mathcal{G}$ may incorporate some contiguity condition that reflects that the crystalline structures that are to be reconstructed do not consist of 'scattered' atoms but are highly connected. In most cases it will also be necessary to consider (highly restricted) subsets $\mathcal{T}$ of $\mathcal{S}^n$ since electron microscopic images of high enough resolution can only be obtained in certain directions. An important subset of $\mathcal{S}^n$ is the set $\mathcal{L}^n = \{\lin\{u\} : u \in \mathbb{Z}^n \setminus \{0\}\}$ of *lattice lines*.

Suppose now that $S_1, \ldots, S_m \in \mathcal{L}^n$ are m different lines specified beforehand. Typically, in practice, m is at most 5 since the high energy needed to produce the images leads to distortions already after a few images are taken.

In the inverse problem RECONSTRUCTION$_{\mathcal{G}}(S_1, \ldots, S_m)$, we are given *data* functions

$$f_i : \mathcal{A}(S_i) \to \mathbb{N}_0, \qquad i = 1, \ldots, m$$

with finite support and the task is to find a set $F \in \mathcal{G}$ whose X-rays $X_{S_i}F$ coincide with f_i for $i = 1, \ldots, m$. Of course, due to data errors, in practice it is in general only reasonable to ask for approximate solutions.

To exclude trivial cases we will in the following always assume that $m, n \in \mathbb{N}$ with $m, n \geq 2$.

It should be mentioned that the whole model can be rephrased in a purely combinatorial form that relies only on incidences. Hence many of our results hold for query sets much more general than lines. Further, it is straight forward to extend the model to 'polyatomic' structures so as to deal with the reconstruction of compounds.

The following sections describe the main results some of which are based on the additional assumption of exact data. Of course, this assumption is not realistic for our prime application. It is all right though for some other applications indicated above. Also, the uniqueness theorems, results on the computational complexity of the problem, and the development of exact algorithms under this additional assumption lay the ground for the subsequent parts that deal with the treatment of data errors in practice.

We pointed out already that simulations were performed. For applications in practice, visualization of the results is another important issue. In fact, since typically, samples consist of 10^8 to 10^9 atoms, visualization tools are needed that can cope with many orders of magnitude, viewing the object from some distance in order to detect 'interesting parts' on which has to be zoomed in then.

The main emphasis in the following will be on structural and algorithmic aspects of the problem. Part of what is following is based on [13]. Figure 3 shows a (small) 3D-sample that has been reconstructed from three of its X-rays by methods outlined in Sect. 8.

5 Uniqueness Results

A most fundamental question is that of whether the given information (measurements plus a priori knowledge) determines the underlying object uniquely. Of course, in practice data errors may corrupt solutions hence one has to resort to approximate solutions. But how sensitive are solutions under data errors? Clearly, this question, while looking quite innocent, touches all basic practical problems related to the ill-posedness of the task.

In the present section we are dealing with the issue of uniqueness under the assumption of exact data. This will give a first glimpse of the underlying difficulty of the problem; see also [10].

We say that a subset $\mathcal{G}$ of the set $\mathcal{F}^n$ of all finite lattice sets of $\mathbb{R}^n$ is *determined by m X-rays* parallel to the lines in a subset $\mathcal{T}$ of $\mathcal{S}^n$ if there

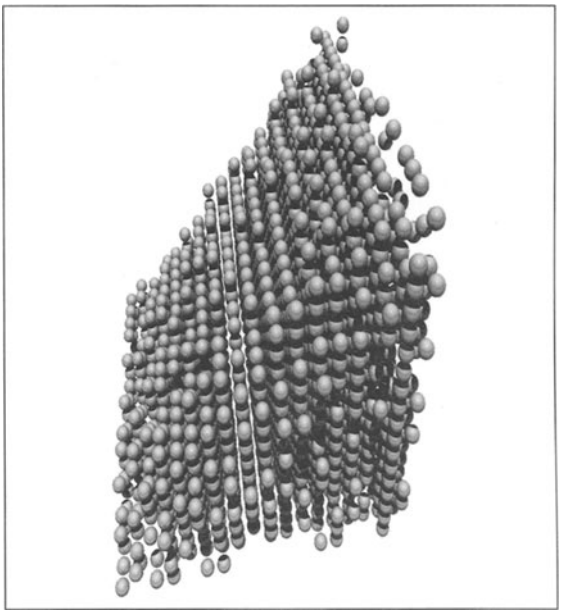

Fig. 3. Reconstruction of a 3D-phantom based on X-ray data in directions $(1,0,0)$, $(0,1,0)$ and $(1,1,0)$. (The front plane is the xy-plane)

exist $S_1, \ldots, S_m \in \mathcal{T}$ such that the following holds: When $F_1, F_2 \in \mathcal{G}$ and $X_{S_j} F_1 = X_{S_j} F_2$ for $j = 1, \ldots, m$, then $F_1 = F_2$.

Here is an utterly trivial result.

With respect to $\mathcal{S}^n$, the class $\mathcal{F}^n$ is determined by one X-ray.

Of course, a line in a non-lattice direction either misses $\mathbb{Z}^n$ or, if it contains a lattice point, this lattice point is the unique lattice point on this line. As simple as it is, this result indicates already the fundamental difference between discrete and continuous tomography. In fact, it is well-known that compact sets are not determined by their (continuous) X-rays in finitely many directions. Hence discrete tomography is not just the discretization of continuous tomography in the sense that the latter comes from a limiting process of the former using lattice refinements.

Since sufficiently high resolution in the practical imaging process can only be achieved in certain main directions of the lattice, for all practical purposes we have to restrict ourselves to $\mathcal{L}^n$, actually to small subsets thereof. Then the situation is less promising.

Let $\mathcal{L}$ be a finite subset of $\mathcal{L}^n$. Then there are sets in $\mathcal{F}^n$ that are not determined by the their X-rays parallel to the lines in $\mathcal{L}$.

There are examples that resemble crystals with just 'a few impurities', see [10]. Hence this negative result captures already some of the problems that can actually be seen in practice. Hence, at best, one can only hope for somewhat weaker uniqueness results. Here is a simple one due to [28] and [18].

Let $\mathcal{F}^n(m)$ be the class of sets in $\mathcal{F}^n$ of cardinality less than or equal to m. Let $\mathcal{L} \subset \mathcal{L}^n$ with $|\mathcal{L}| \geq m+1$. Then the sets in $\mathcal{F}^n(m)$ are determined by their X-rays parallel to the lines in $\mathcal{L}$.

The problem with this result in practice is that the typical atomic structures that have to be reconstructed comprise about 10^8 to 10^9 atoms. Hence an extremely large number of X-ray images would be needed to guarantee uniqueness. Since the energy needed for the imaging process is very high, after just very few images the object is, however, corrupted by the radiation.

Here is a uniqueness result of [9] for the restricted class $\mathcal{C}^n$ of *convex* lattice sets, i.e., finite subsets F of $\mathbb{Z}^n$ such that $F = \mathbb{Z}^n \cap \operatorname{conv}(F)$, see [6] for an extension.

Let $\mathcal{L} \subset \mathcal{L}^n$ with $|\mathcal{L}| \geq 7$ and all lines in $\mathcal{L}$ being coplanar. Then the sets in $\mathcal{C}^n$ are determined by X-rays parallel to the lines in $\mathcal{L}$. Further, there is a set $\mathcal{L} \subset \mathcal{L}^n$ with $|\mathcal{L}| = 4$ such that sets in $\mathcal{C}^n$ are determined by X-rays parallel to the lines in $\mathcal{L}$.

This result says that the class $\mathcal{C}^n$ is determined by X-rays parallel to suitable 4 or any 7 different coplanar lattice lines. As pointed out before, in practice the coplanarity assumption is satisfied since in essence the microscope is only rotated about one axis.

While this result is quite reassuring, it is only partly practical. There may be some applications, for instance in colloid physics, but the main demand for mathematical methods for solving the inverse problems of discrete tomography comes from applications that involve the reconstruction of nonconvex objects. In particular, quality control in certain stages of chip production involves the detection of 'bumps' on the interfacial surface of silicon chips, hence convexity is not an appropriate condition in this situation.

Another possibility to outsmart the general nonuniquess result is to change the 'experimental environment'. The above concept of unique determination was based on an a priori choice of the lines for taking X-rays. What if we take the first X-ray in an arbitrary direction but then use the information gained from analyzing the image in order to determine the next line for taking an X-ray? For the third direction, then use the complete information given by the first two X-rays and so on. This approach of *successive determination* leads to strong uniqueness results even for higher dimensional X-rays and even for sets more general than lattice sets, see [9]. In particular

$\mathcal{F}^n$ can be successively determined by 2 X-rays.

Again, while seeming to be satisfactory, this result is not practical at all. In essence it resembles the first uniqueness result that utilized irrational slopes. In fact, the second X-ray has to be so 'skew' that one cannot produce HRTEM images of high enough resolution in this direction.

The lack of practically satisfactory uniqueness results even in the absence of noise indicates already that we need to settle for less. We may ask for the 'core' of all solutions, the set of all *invariant points* that must belong to all solutions. Equally, we are interested in as large as possible sets of points that do not belong to any solution. We might also wish to determine a 'typical' solution. Of course, such weaker concepts of 'solution' do make immediate sense even in more realistic models that allow data errors.

6 Computational Complexity

Another realistic relaxation of the concept of uniqueness in the previous section is to ask only for a certificate of uniqueness for a given concrete instance of the problem. I.e., rather than requiring general a priori uniqueness results we want to check uniqueness of a solution for each instance separately. Hence, in effect we are asking for an efficient procedure for checking uniqueness algorithmically. This brings up the algorithmic aspect of discrete tomography in the context of uniqueness. But it is certainly clear that efficient procedures for reconstruction are needed anyway.

In this section we begin the discussion of algorithmic features of discrete tomography by stating results on the computational complexity of the questions of checking consistency of X-ray data, of determining uniqueness of given solutions and, of course, of finally reconstructing the objects. In the following we will focus on the full family $\mathcal{F}^n$, and the results will only be stated for that case. However, most of the results hold for a great variety of other subclasses $\mathcal{G}$ as well, [10].

Suppose that $S_1, \ldots, S_m \in \mathcal{L}^n$ are $m \geq 2$ lines specified beforehand. The inverse problem RECONSTRUCTION$(S_1, \ldots, S_m)$ was already defined in Sect. 4. (Here and in the sequel we will omit the subscript $\mathcal{F}^n$.) In the realm of computational complexity theory decision problems are more appropriate than reconstruction problems. Hence we consider also the problem CONSISTENCY$(S_1, \ldots, S_m)$ whose instances are just the same, i.e, consist of given data functions

$$f_i : \mathcal{A}(S_i) \to \mathbb{N}_0, \qquad i = 1, \ldots, m$$

with finite support, but whose task is restricted to the decision whether a solution exists (without being obliged to produce one). Similarly, the problem UNIQUENESS$(S_1, \ldots, S_m)$ asks whether, given a solution F, there exists another one.

Here are some basic tractability and intractability results. (Of course, one needs to specify the data structures for the problems more precisely. But this is easily done; see [11].)

Let $S_1, \ldots, S_m \in \mathcal{L}^n$ be m different lines. If $m = 2$, the problems CONSISTENCY(S_1, S_2), UNIQUENESS(S_1, S_2) *and* RECONSTRUCTION (S_1, S_2) *can be solved in polynomial time. For $m \geq 3$,* CONSISTENCY $(S_1, \ldots, S_m)$ *and* UNIQUENESS$(S_1, \ldots, S_m)$ *are* NP*-complete in the strong sense while* RECONSTRUCTION$(S_1, \ldots, S_m)$ *is* NP*-hard.*

Proofs of the tractability part of the above result, i.e. for the statement for $m = n = 2$, can be found in [4], [32], [29], [30] or [1] while the intractability result is due to [11]. The particular case of the NP-completeness of CONSISTENCY(S_1, S_2, S_3) when S_1, S_2, S_3 are the coordinate axes in $\mathbb{E}^3$ was previously dealt with in [21, Section 4.1] in the context of statistical contingency tables.

Let us emphasize that the constructions exhibit intractability already for solid crystals with 'just a few impurities', i.e. for physically reasonable objects, whence explaining the algorithmic difficulties observed in practice. Some extensions focusing on nonapproximability results are given in [16]. Some related complexity results can be found in [2], [25] and [10].

7 Integer Programming, LP-Relaxation and Randomized Rounding

Various approaches have been suggested for dealing with the basic NP-hard problems in discrete tomography; see [16] for an account of the success and failure of many of those techniques. In the following we will outline some basic standard approaches for RECONSTRUCTION$(S_1, \ldots, S_m)$ that are based on an integer programming formulation of this problem whose variables correspond to the possible positions of elements of a solution. The *grid* G of a given instance of the problem consists of all (finitely many) lattice points that arise as points of intersection of m lines parallel to $S_1, \ldots, S_m$, respectively, whose data function value is nonzero, i.e., ,

$$G = \mathbb{Z}^n \cap \bigcap_{i=1}^{m} \bigcup_{T \in \mathcal{T}_i} T,$$

where $\mathcal{T}_1, \ldots, \mathcal{T}_m$ denote the supports of the given data functions $f_1, \ldots, f_m$, respectively. The incidences of G and $\mathcal{T}_i$ can be encoded by an incidence matrix A_i. If G consist of, say, N points, $M_i = |\mathcal{T}_i|$ for $i = 1, \ldots, m$, and $M = M_1 + \cdots + M_m$, then the incidence matrices A_i are in $\{0,1\}^{M_i \times N}$, and can be joined together to form a matrix $A \in \{0,1\}^{M \times N}$. Identifying a subset F of G with its characteristic vector $x_F \in \{0,1\}^N$, the reconstruction problem amounts to solving the integer linear feasibility problem

$$Ax = b, \text{ s.t. } x \in \{0,1\}^N,$$

where $b^T = (b_1^T, \ldots, b_m^T)$ contains the corresponding values of the data functions $f_1, \ldots, f_m$ as the right hand sides of $A_1, \ldots, A_m$, respectively. Of course,

the equality constraints model the situation of exact data. In practice, this problem will be infeasible and should be replaced by some relaxation, e.g.

$$\max c^T x \quad \text{s.t. } Ax \leq b,\ x \in \{0,1\}^N,$$

where the objective function $x \mapsto c^T x$ models some specific goal. The all-ones vector $\mathbb{1} = (1,\ldots,1)^T$, for instance, is used to maximize the number of points that can be placed under the data constraints, but other objectives like the 'closeness' to some template etc. can also be facilitated. Approximation algorithms will be dealt with in Sect. 9.

Since linear programming problems can be solved in polynomial time the first natural approach is to consider the *LP-relaxation*

$$\max c^T x \quad \text{s.t. } Ax \leq b,\ 0 \leq x \leq \mathbb{1};$$

see [7]. Since linear programming codes are available for solving these problems very efficiently for all sizes of crystalline structures that are relevant in practice, computation time is not so much of an issue for this heuristic. However, the solution is usually far from being integer. N. Young studied a randomized rounding strategy where in some stochastic experiment an atom is placed at some lattice point with the probability coming from the fractional solution produced by the LP-solver. This way an approximative solution is produced. Compared to the heuristics described in Sect. 9 known bounds for such solutions are, however, in general rather weak.

8 Polytopes in Discrete Tomography

Of course, one cannot expect too tight general a priori error bounds for polynomial-time approximations in general. For most practical purposes, however, good bounds generated in the course of the computation for a given specific instance are good enough. One can run an improvement algorithm on the given data until the gap between upper and lower bounds is small enough (or until the allotted time elapsed) and then terminate with an approximate solution including a performance guarantee for exactly that solution. This is the underlying idea of branch-and-cut strategies. These are branch-and-bound methods augmented by cutting plane techniques that we are going to sketch now.

The ultimate object for studying cutting planes is of course the convex hull

$$P(A,b) = \operatorname{conv}\{x \in \{0,1\}^N : Ax = b\}$$

of all solutions of the problem. In general, even computing the dimension of the polytope $P(A,b)$ is $\mathbb{NP}$-hard. Also, $P(A,b)$ is only relevant in case of consistency whence typically irrelevant in the presence of noise in the imaging process. Hence the submissive

$$T(A,b) = \operatorname{conv}\{x \in \{0,1\}^N : Ax \leq b\}$$

is much more appropriate and we study the facial structure of that polytope. As we know that the reconstruction problem is NP-hard, we cannot aim at a 'concise' description of $T(A,b)$ by means of a system of linear inequalities. However, appropriate local information is all that is needed; see [3] for a nontechnical introduction to polyhedral combinatorics.

The algorithm begins by maximizing a suitable objective function, say $x \mapsto \mathbb{1}^T x$, over the polytope

$$T^{\mathrm{LP}}(A,b) = \mathrm{conv}\,\{x \in [0,1]^N : Ax \leq b\}.$$

This is a task of linear programming that can be handled efficiently. If the produced vertex x^* at which the optimum is attained is integral then the problem is solved. Otherwise, we try to improve the relaxation $T^{\mathrm{LP}}(A,b)$ of $T(A,b)$ by adding a suitable inequality that separates x^* from $T(A,b)$. Geometrically we intersect $T^{\mathrm{LP}}(A,b)$ with an appropriate closed halfspace that contains $T(A,b)$ but not x^*. Such a halfspace is called a *cut* and the approach of successively solving the current LP-relaxation, and adding a cut if still necessary is called a *cutting-plane* algorithm. In general, it is not clear how to find reasonably deep cuts efficiently, and typically some lurking NP-hardness is connected to this. (Of course, if the original problem was NP-hard, its intractability cannot magically disappear.) Hence it may be possible that the current LP-optimum is not integer, yet we do not know any cut that can be added. Then we need to resort to a branching strategy, splitting the problem into various others by forbidding certain properties for some and forcing them for other subproblems. For instance, we may split the tomographic reconstruction task into two, one where a certain point must be present, the other, where this point must not be present.

This is the *branch-and-cut* approach, performed within the branch-and-bound framework, with cutting planes used for improving the upper bound at each node of the branching tree.

Of course, it is similarly important to obtain good lower bounds, i.e. approximate solutions close to the optimum. They can be achieved by running heuristics in each branching node. We report about theoretical and practical results for such approximation algorithms in Sect. 9. Primal improvement strategies based on tests sets were studied in [35], where a general algebraic framework for improvement strategies in discrete tomography was developed.

The *tomography polytopes* $T(A,b)$ have some quite special structure. In particular, all submatrices of A corresponding to just two directions are *totally unimodular*. Since all tomography polytopes are 0-1-polytopes it follows from [26] that the combinatorial *diameter* of a tomography polytope is at most N. This means that, in principle, an edge-path could be found leading from 0 to a solution of the problem that is rather short.

Further, by means of various preprocessing techniques, the practically relevant dimensions of these polytopes can be reduced to about 10^4. Judged on the base of the sizes of successfully solved instances of the traveling salesman problem this seems encouragingly small. While the traveling salesman

polytopes correspond to the complete graph on the given number of 'cities' and are hence universal, tomography polytopes depend, on the other hand, on the right-hand side b. So the most important goal of polytopal investigations is to find large systems of valid inequalities that are facet-defining under week conditions on the right-hand side. In [34] various classes of facet-defining inequalities are determined under weak assumptions on b. Usually it is only necessary to require that all components of b are at least 2 or 3 and that no X-ray line is completely filled with atoms.

9 Approximation Algorithms

In view of the computational complexity of RECONSTRUCTION$(S_1, \ldots, S_m)$ for $m \geq 3$ and in view of the presence of noise in the data b it is reasonable to study approximation algorithms for this problem. It turns out that in spite of the underlying $\mathbb{NP}$-hardness some simple heuristics already yield good a priori approximation guarantees and behave surprisingly excellent in practice.

The problem BEST-INNER-FIT$(S_1, \ldots, S_m)$ [BIF] accepts as input data functions $f_1, \ldots, f_m$; its task is to find a set $F \in \mathcal{F}^n$ of maximal cardinality such that

$$X_{S_i} F(T) \leq f_i(T) \qquad \text{for all } T \in \mathcal{T}_i \text{ and } i = 1, \ldots, m.$$

This problem and its outer counterpart were studied in [15]. In particular, various greedy and improvement strategies were fully analyzed. The main results will now be stated.

The greedy algorithm considers the positions of the grid G of candidate points in some order and successively fills in points as long as this is possible without violating the constraints. A simple observation shows that any such algorithm produces a solution V such that

$$|V|/|F| \geq 1/m,$$

where F is an optimal solution; cf. [24]. The bound is sharp and reflects the fact that the more data are given, the harder it is for a greedy strategy to satisfy them. In the experiments it turns out, that $|V|/|F|$ is typically greater than 0.9 and for large instances greater than 0.96 even for $m = 5$.

The greedy strategy is very flexible and allows various specifications for breaking the ties between different choices for points to be placed next. For example, the X-ray data can be used in a way that is very similar to back-projection techniques to express preferences. Also connectivity of the solution (in a sense that is justified by the physical structure of the analyzed material) can be rewarded. Similarly, information of neighboring layers can be taken into account in a layer-wise reconstruction of a 3D-object.

In the practical experiments, an algorithm 'GreedyC' was particularly good. Here the insertion order is based on weights that are dynamically

assigned to the candidate points and represent the 'changing importance' of a point to be included in a solution. The computational part of [15] shows that GreedyC actually produces solutions with small *absolute error*. The average for GreedyC when applied to instances with 250000 variables and point density of 50% were 21.62, 64.13 and 111.88 missing atoms for 3, 4, 5 directions, respectively.

In order to further improve the performance of such iterative insertion algorithms, one can apply *r-improvements* for $r \in \mathbb{N}_0$ that replace an r-point subset of a current feasible solution $F \subset G$ by $r+1$ points of $(G \setminus F)$ while maintaining feasibility. (Of course, 0-improvements are just insertions.) A feasible set $F \subset G$ is called *t-optimal* for the given instance of [BIF] if no r-improvement is possible for any $r \leq t$.

The power of r-improvements can be seen in the next result; (asymptotically) the a priori performance bound is improved by another factor of 2.

Let $t \in \mathbb{N}$, let F be a solution of a given instance of [BIF] and let $V \subset G$ be t-optimal. Then

$$\frac{|V|}{|F|} \geq \frac{2}{m} - \epsilon_m(t),$$

where $\epsilon_m(t)$ approaches 0 exponentially fast.

In [15] $\epsilon_m(t)$ is given explicitly and examples are constructed that show that the bounds are tight. Computationally, it turns out that performing 1-improvements after GreedyC yields substantial improvements. In fact, in our computational study the *absolute* errors go down to 1.07, 23.28, 64.58 for 3, 4, 5 directions, respectively; see [15].

Another approximation algorithm is based on LP-relaxation. It turns out that by rounding down all fractional components of an optimal vertex of T^{LP} we obtain a solution of [BIF] with objective value at most M worse than the optimum. This permits to devise polynomial approximation schemes for certain classes of instances of [BIF]; see [14] for details.

10 Dealing with Ill-Posedness

The results of Sect. 9 are extremely encouraging in that they show that the 'combinatorial optimization part' of the reconstruction problem of discrete tomography can be handled very efficiently in spite of its computational complexity. However, there is more to be taken care of. In fact, the relevant measure for the quality of an approximation to a binary image would of course be the deviation from this image. Hence in order to devise the most appropriate objective function one would have to know the underlying solution of the given inverse problem. However, the whole point is of course to find this unknown solution. Hence one can only consider objective functions

for measuring the quality of approximation that are based on the given input data. While a good approximation in this sense is close to a solution in that its X-ray images parallel to the given lines are close to those of the original set, the approximating set itself may be off quite substantially. In fact, the inverse discrete problem is ill-posed and it is precisely this property that causes additional difficulties. In particular, if the input data do not uniquely determine the image even a 'perfect' solution that is completely consistent with all given data may be quite different from the unknown real object.

There are various approaches for dealing with nonuniqueness. If we use the fact, that in our prime application the lines parallel to which the X-ray images are taken are coplanar, then the objects can be reconstructed in a layer-by-layer wise process. Suppose that the first layer has been reconstructed and that it is uniquely determined by the given information (or known beforehand). Then it may be reasonable to assume, that the second layer does not vary too much from the first. As pointed out already this can be modeled easily by an objective function for the second layer that is given through the incidence vector of the solution for the first layer. In case of just two X-rays such an approach was suggested by [32].

In [7] an interior point heuristic is proposed for identifying positions that are uniquely determined by the given data. While the problem of detecting whether a given subset of G belongs to all solutions for a given instance is again $\mathbb{NP}$-hard, [11], [16], it is demonstrated in [7] that for certain phantoms such an approach produces quite a large number of fixed variables. Note that the problem of deciding whether a subset of a possible solution can actually be extended to a full solution is again $\mathbb{NP}$-complete, see [11].

Various other techniques for reducing ambiguity are proposed in the study [17]. In any case, it is still important to utilize additional physical knowledge and experience in order to be able to produce solutions that are close to the actual physical objects.

The main problem, however, is that of instability under changes of the data b. Of course, it is easy to see that the cardinality of a maximum solution of [BIF] is stable. Also, if b_1 and b_2 are close and F_1 is a solution for the measurement b_1, then there is always an approximate solution F_2 for b_2 that is close to F_1 in terms of their symmetric difference. However, one can use constructions of [23] to produce for any given $m \geq 3$ and $k \in \mathbb{N}$ two vectors b_1 and b_2 with $\|b_1 - b_2\|_{(1)} \leq 2(m-1)$ such that there are unique but disjoint sets F_1, F_2 of cardinality at least k whose X-rays coincide with b_1, b_2, respectively. Hence the data differs only by a constant, while the symmetric difference of the optimal solutions can be arbitrarily large. This shows the ultimate limitations of general theoretic stability results. However, the examples are very specific and in practice the results are typically quite reasonable in that they exhibit the main features of the original object. In any case, a general theory of discrete inverse problem is still to be developed.

11 Software

The software is build around a module that encapsulates all data-structures that are specific to questions of discrete tomography. In particular, they permit to treat the case of arbitrary three-dimensional lattices.

We implemented all the approximation algorithms mentioned in Sect. 9 and several others. Further, we included a branch-and-bound module based on CPLEX [20].

On top of these modules with data-structures and solver algorithms resides the user-interface. Besides being the command center for the solvers it provides means for file-input/output and different methods for generating instances of various types. One such feature is line-convexity in certain directions, used in a first experimental study of large instances that mimic the etching process applied in silicon wafer production.

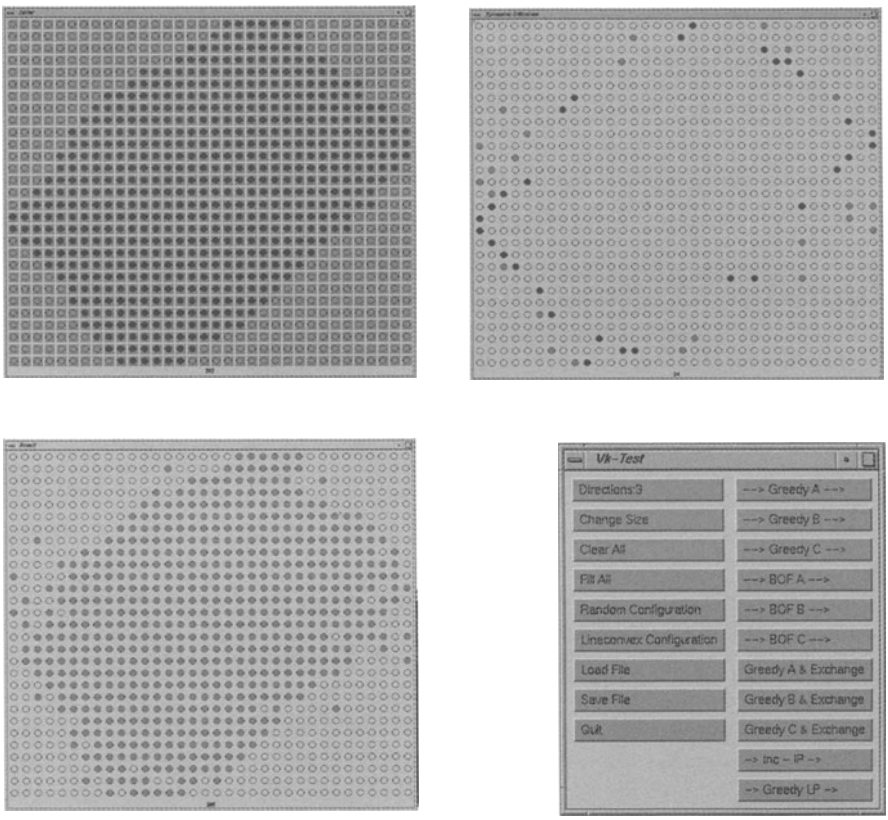

Fig. 4. Screenshots of the user interface for a small example. In the upper left the underlying phantom is depicted. The lower left shows a reconstruction from images parallel to the lines spanned by $(1,0),(0,1)$, and $(1,-2)$. The symmetric difference between the phantom (black) and its reconstruction (grey) is printed in the upper right corner. The main menu of the tool is shown in the lower right

References

1. R.P. Anstee, The network flows approach for matrices with given row and column sums, Discrete Math. **44** (1983), 125–138.
2. E. Barcucci, A. Del Lungo, M. Nivat, and R. Pinzani, Reconstructing convex polyominoes from their horizontal and vertical projections, Theor. Comput. Sci. **155** (1996), no. 2, 321–347.
3. T. Burger and P. Gritzmann, Polytopes in combinatorial optimization, Geometry at work, (ed. by C.A. Gorini), Cambridge University Press, 2000, pp. 11–33.
4. S.K. Chang, The reconstruction of binary patterns from their projections, Commun. ACM **14** (1971), no. 1, 21–25.
5. J.M. Cowley and A.F. Moodie, The scattering of electrons by atoms and crystals. 1. A new theoretical approach, Acta Cryst. **10** (1957), 609–619.
6. A. Daurat, Uniqueness of the reconstruction of q-convexes from their projections, submitted.
7. P. Fishburn, P. Schwander, L. Shepp, and R.J. Vanderbei, The discrete Radon transform and its approximate inversion via linear programming, Discrete Appl. Math. **75** (1997), no. 1, 39–61.
8. P.C. Fishburn, J.C. Lagarias, J.A. Reeds, and L.A. Shepp, Sets uniquely determined by projections on axes II: Discrete case, Discrete Math. **91** (1991), no. 2, 149–159.
9. R.J. Gardner and P. Gritzmann, Discrete tomography: Determination of finite sets by X-rays, Trans. Am. Math. Soc. **349** (1997), no. 6, 2271–2295.
10. ———, Uniqueness and complexity in discrete tomography, Discrete Tomography: Foundations, Algorithms, and Applications, (ed. by G.T. Herman and A. Kuba), Birkhäuser, Bosten, 1999, pp. 85–113.
11. R.J. Gardner, P. Gritzmann, and D. Prangenberg, On the computational complexity of reconstructing lattice sets from their X-rays, Discrete Math. **202** (1999), 45–71.
12. R.J. Gardner, P. Gritzmann, and D. Prangenberg, On the computational complexity of determining polyatomic structures by X-rays, Th. Comput. Sci. **233** (2000), 91–106.
13. P. Gritzmann, On the reconstruction of finite lattice sets from their X-rays, Discrete Geometry for Computer Imagery, (ed. by E. Ahronovitz and C. Fiorio), Springer, Berlin, 1997, pp. 19–32.
14. P. Gritzmann and S. de Vries, LP-based approximation for discrete tomography, in preparation.
15. P. Gritzmann, S. de Vries, and M. Wiegelmann, On the approximate reconstruction of binary images from their discrete X-rays, SIAM J. Opt. (to appear), 25 pages.
16. P. Gritzmann, D. Prangenberg, S. de Vries, and M. Wiegelmann, Success and failure of certain reconstruction and uniqueness algorithms in discrete tomography, Int. J. Imaging Syst. Technol. **9** (1998), 101–109.
17. P. Gritzmann, P. Schwander, and S. de Vries, On the size of switching components, in preparation.
18. A. Heppes, On the determination of probability distributions of more dimensions by their projections, Acta Math. Acad. Sci. Hung. **7** (1956), 403–410.
19. G.T. Herman and A. Kuba (eds.), Discrete Tomography: Foundations, Algorithms, and Applications, Birkhäuser, Bosten, 1999.

20. ILOG Inc. CPLEX Division, Using the CPLEX callable library, 1997, 930 Tahoe Blvd. # 802-279, Incline Village, NV 89451-9436, USA, Information available at URL `http://www.cplex.com`.
21. R.W. Irving and M.R. Jerrum, Three-dimensional statistical data security problems, SIAM J. Comput. **23** (1994), 170–184.
22. C. Kisielowski, P. Schwander, F.H. Baumann, M. Seibt, Y. Kim, and A. Ourmazd, An approach to quantitative high-resolution transmission electron microscopy of crystalline materials, Ultramicroscopy **58** (1995), 131–155.
23. T.Y. Kong and G.T. Herman, On which grids can tomographic equivalence of binary pictures be characterized in terms of elementary switching operations?, Int. J. Imaging Syst. Technol. **9** (1998), 118–125.
24. B. Korte and D. Hausmann, An analysis of the greedy heuristic for independence systems, Ann. Discrete Math. **2** (1978), 65–74.
25. A. Del Lungo and M. Nivat, Reconstruction of connected sets from two projections, Discrete Tomography: Foundations, Algorithms, and Applications, (ed. by G.T. Herman and A. Kuba), Birkhäuser, Bosten, 1999, pp. 163–188.
26. D. Naddef, The Hirsch conjecture is true for (0, 1)-polytopes, Math. Program. **45** (1989), 109–110.
27. F. Natterer, The Mathematics of Computerized Tomography, Teubner, Stuttgart, 1986.
28. A. Rényi, On projections of probability distributions, Acta Math. Acad. Sci. Hungar. **3** (1952), 131–142.
29. H.J. Ryser, Combinatorial properties of matrices of zeros and ones, Canad. J. Math. **9** (1957), 371–377.
30. ———, Combinatorial mathematics, ch. 6, Matrices of zeros and ones, Mathematical Association of America and Quinn & Boden, Rahway, NJ, 1963, pp. 61–78.
31. P. Schwander, C. Kisielowski, F.H. Baumann, Y. Kim, and A. Ourmazd, Mapping projected potential, interfacial roughness, and composition in general crystalline solids by quantitative transmission electron microscopy, Phys. Rev. Lett. **71** (1993), no. 25, 4150–4153.
32. C.H. Slump and J.J. Gerbrands, A network flow approach to reconstruction of the left ventricle from two projections, Comput. Graphics Image Process. **18** (1982), 18–36.
33. P.A. Stadelmann, EMS – a software package for electron diffraction analysis and HREM image simulation in material science, Ultramicroscopy **21** (1987), 131–146.
34. S. de Vries, Discrete tomography, packing and covering, and stable set problems: Polytopes and algorithms, Ph.D. thesis, Technische Universität München, 1999.
35. M. Wiegelmann, Gröbner bases and primal algorithms in discrete tomography, Ph.D. thesis, Technische Universität München, 1998.