On the Design of Optimisers for Surface Reconstruction

Tobias Wagner
Department of Machining Technology(ISF)
University of Dortmund
Dortmund, Germany
wagner@isf.de

Thomas Michelitsch
Department of Machining Technology(ISF)
University of Dortmund
Dortmund, Germany
michelitsch@isf.de

Alexei Sacharow
Department of Machining Technology (ISF)
University of Dortmund
Dortmund, Germany
sacharow@isf.de

ABSTRACT

In many industrial applications the need for an efficient and high-quality reconstruction of free-form surfaces does exist. *Surface Reconstruction* – the generation of CAD models from physical objects – has become an independent area of research. The supplementary modification and the automated manufacturing of workpieces represent typical fields of application. Small tolerances in the desired properties result in a very high number of scan points needed. Thus, modern approaches have to be capable of processing, analysing and modelling these amounts of data.

There are several studies that use evolutionary algorithms (EA) for surface reconstruction tasks. Until now, these studies only describe the general ability of EA to successfully optimise surfaces. Aspects like runtime as well as comparability to other optimisation techniques have not been considered. Since these aspects are of great importance for integration in applicable software tools, in this paper the ability of a state-of-the-art multi-objective EA to be successfully integrated in surface reconstruction software is analysed. Major drawbacks are disclosed and necessary add-on modules are presented.

Categories and Subject Descriptors

J.6 [**Computer-Aided Engineering**]: Computer-aided design (CAD); G.1.6. [**Numerical Analysis**]: Optimization–Least squares methods; Global optimization

General Terms

Algorithms, Design

Keywords

Surface Reconstruction, Hybrid Algorithms, Linear Least Square, Singular Value Decomposition, SMS-EMOA, Computer-Aided Geometric Design

GECCO'07, July 7–11, 2007, London, England, United Kingdom.

1. INTRODUCTION

In common modern development processes, CAD tools are used to generate models of the workpiece. If a physical prototype exists, methods to obtain a CAD model from the prototype are often needed. This process is called *Surface Reconstruction* or *Reverse Engineering*.

Current scanners are capable of providing a very dense and precise set of scan points from the prototype's surface, which can be used as origin of a CAD-compatible representation. In this context the computerised processing of the point set represents a central problem. With growing accuracy of modern scanning systems, the amount of data to be processed increases. Most scanning systems provide only partly structured point sets. Thus, a computer system for surface reconstruction has to be able to identify important regions and efficiently use the information obtained from the calculations. Only in this vein, a reconstruction suitable for real-world applications can be performed.

Most articles on *Surface Reconstruction* make use of triangulations due to the direct reference to the set of scan points. Since the number of valid triangulations drastically increases with the number of scan points used, these triangulations become very large and difficult to handle. The class of *Approximating Triangulations* (AT) tackles this problem. The number of vertices is independent of the size of the point set and considerably smaller than the number of scan points. Weinert, Mehnen, and Schneider [22] document the capabilities of a standard evolutionary strategy using balancing techniques to optimise AT.

Unfortunately, triangulations are often serrated and therefore not convenient for the modelling of most workpiece surfaces. Thus, *Non-Uniform Rational B-Splines* (NURBS) [16] are a commonly used mathematical model to describe free-form surfaces in CAD systems. The advantages of the NURBS representation are the smoothness, the compact mathematical definition, the intuitive local manipulation, as well as the ability to combine NURBS patches to larger structures. However, the direct approximation of scan point data by means of evolutionary optimised NURBS is rarely discussed in literature.

Weinert, Surmann, and Mehnen [23] combine NURBS with Constructive Solid Geometries in a hybrid evolutionary algorithm/genetic programming approach. They report runtimes up to 24 hours for the optimisation of a NURBS to model a simple structure due to expensive function evaluations and the high dimensionality of the decision space. Thus, Beielstein et al. [2] divide the process of optimisation in different steps where only the z-coordinate of the control

points is modified in the beginning. Only afterwards, the other parameters become the subjects of optimisation.

This paper analyses and counteracts the problems experienced in the above papers. Hints on pre-processing the data to allow for an efficient application are given. A multi-objective EA (MOEA) regarding the dominated hypervolume [8] of the population is used within our experiments. This MOEA has shown its superior performance in several benchmarks [8, 15, 21] and real-world applications [14]. Thus, it can be assumed that this MOEA represents a state-of-the-art evolutionary approach.

In the first experiments drawbacks of distance functions often used for the evaluation are disclosed. A comparison to the surfaces calculated by a numerical equation solver that provides solutions, which are optimal in the sense of *Linear Least Square* [24] is performed. The results of the approaches are compared with regard to both quality and runtime. Finally, concepts for the hybrid application of both techniques are presented.

This paper is arranged as follows. In the next three chapters the topics dealt with are briefly introduced. First, the problem of surface reconstruction is described and NURBS are analytically defined. The approaches used for optimisation are explained in sections 3 and 4. In the following section the effect of the objectives is analysed and promising objective functions for both approaches are selected. The inferred hybrid approach applied in section 7 is assisted by some data structures that organise the point data for a faster evaluation. In section 6 these structures are described and their improvement is empirically proven.

2. SURFACE RECONSTRUCTION

The intention of surface reconstruction or reverse engineering is to generate a CAD model from a physical representation of an object [13]. In the first step the object is scanned by a tactile or digital scanner resulting in a set $\mathcal{P}$ of scan points. Based on $\mathcal{P}$ the CAD model of the object's surface can be generated. Since we restrict the CAD model to the *Non-Uniform Rational B-Splines* (NURBS) representation, we assume a parametric NURBS surface $S(u, v)$ that has to be optimised to approximate $\mathcal{P}$. This initial solution can either be randomly generated, constructed by a heuristic, or based on a primary representation of the object.

NURBS surfaces [16] are defined as parametric bivariate tensor products that map $\mathbb{R}^2$ to $\mathbb{R}^3$. The definition of a NURBS surface of order (p, q) comprises the following components:

- the control point net with $n \times m$ vertices
$$\mathbf{C} = \{\mathbf{c}_{i,j} \in \mathbb{R}^3, i = 1, \ldots, n, j = 1, \ldots, m\},$$
- the knot vector $\mathbf{u}$ and $\mathbf{v}$, where
$$\mathbf{u} = (\underbrace{0, \ldots, 0}_{p+1}, u_{p+1}, \ldots, u_{r-p-1}, \underbrace{1, \ldots, 1}_{p+1})^T,$$
$$\mathbf{v} = (\underbrace{0, \ldots, 0}_{q+1}, u_{q+1}, \ldots, u_{s-q-1}, \underbrace{1, \ldots, 1}_{q+1})^T,$$
$r = n + p + 1$, and $s = m + q + 1$,
- and the $n \times m$ weight matrix $\mathbf{W}$ providing weights $w_{i,j}$ for each control point $\mathbf{c}_{i,j}$.

The function describing the parametric surface is defined on $(u, v) \in [0, 1] \times [0, 1]$ and expressed by

$$S(u, v) = \frac{\sum_{i=0}^{n} \sum_{j=0}^{m} N_{i,p}(u) N_{j,q}(v) w_{i,j} \mathbf{c}_{i,j}}{\sum_{i=0}^{n} \sum_{j=0}^{m} N_{i,p}(u) N_{j,q}(v) w_{i,j}}, \qquad (1)$$

where $N_{i,p}(u) : \mathbb{R} \to \mathbb{R}, N_{j,q}(v) : \mathbb{R} \to \mathbb{R}$ are the ith basis functions of order p and q respectively [16].

NURBS can be locally manipulated by a variation of the control points' weights and positions. The influence of a single control point is determined by the degrees p and q. If $p = q = 1$ holds, $S(u, v)$ corresponds to piecewise linear models that can be roughly compared with *Approximating Triangulations*. Typically, degrees $p = q = 3$ are used. Thus, our investigations are restricted to this case.

3. MULTI-OBJECTIVE OPTIMISATION

Solutions to real-world problems often cannot be accurately evaluated using a single decision criterion. If more objectives are considered, a negative correlation among each other is the typical case. Therefore, in the field of multi-objective optimisation, the dominance relation is used to order solutions with respect to multiple objectives. An objective vector is said to *dominate* another if no objective is worse and at least one objective is better. Thus, no single optimum does exist.

The non-dominated solutions are called (Pareto) optimal. With many multi-objective problems, knowledge about the structure of the Pareto optimal objective vectors (also called Pareto front) assists the decision maker to find a sophisticated compromise solution, e. g. by identifying knee regions. Multi-objective evolutionary algorithms (MOEA) are designed to discretely approximate the Pareto front of a given problem. In contrast to single-objective EA where convergence towards the optimum is the main performance criterion, MOEA shall additionally obtain a diverse and well-distributed set of objective vectors.

3.1 SMS-EMOA

The *$\mathcal{S}$-metric Selection Evolutionary Multi-Objective Algorithm* (SMS-EMOA) by Emmerich et al. [8] supplements the often used non-dominated sorting procedure [9] by concerning the exclusive $\mathcal{S}$-metric (often referred to as hypervolume) contribution of an individual as secondary selection criterion. This measure considers convergence as well as a wide and well-distributed spread of individuals and, thus, is suggested as one of the best indicators to measure multi-objective performance by Zitzler et al. [25].

Since a steady-state selection scheme is applied, the hypervolume, which is covered by the population regarding a fixed reference point, monotonically increases during the optimisation. Effects of deterioration, often observed when only clustering measures are concerned, are avoided. Numerous studies (e. g. [8, 14, 15, 21]) affirm the superior performance relative to the known MOEA.

The most important drawback of this approach constitutes the runtime complexity, which is exponential in the number of objectives m. In the given application, runtime improvements [4, 15], the restriction to $m \leq 2$, and long evaluation-times compared to the time needed for the selection process allow for the use of the SMS-EMOA. Actually, faster convergence rates can significantly improve the runtime until satisfying solutions are obtained.

Although the superior performance of the SMS-EMOA mainly has to be ascribed to the selection process, the variation operators used are important for the progress of an evolutionary algorithm. Like most known MOEA, SMS-EMOA uses, the genetic algorithm inspired, simulated binary crossover (SBX) with polynomial mutation (PM) [5, pp. 106 et seqq.]. The adaptive skills of the SBX operator show several desired properties [6], hence an additional adaptation of the distribution, which is used for the generation of the random numbers, is not needed.

4. SINGULAR VALUE DECOMPOSITION

Singular Value Decomposition (SVD) [7] represents a stable approach for solving systems of linear equations (SLE) $\mathbf{Ax} = \mathbf{b}$. The problem of finding a solution for the SLE is reduced to the calculation of the inverse matrix of A and obtaining the solution $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$.

Contrary to many other numerical methods, which directly solve the equations by transformation, the accumulation of rounding errors is avoided. Press et al. [17] provide an implementation of SVD to solve over-determined SLE, which minimises the sum of squared errors over all equations. With this implementation, even ill-conditioned SLE can be robustly solved.

5. ON OBJECTIVE FUNCTIONS

The vector of objective functions $\mathbf{f}$ represents the only information usable for a MOEA during the optimisation of the genome. Thus, a successful optimisation is strongly related to a sophisticated selection of $\mathbf{f}$. If the needs of the decision maker are not completely encoded, a solution that is optimal in sense of $\mathbf{f}$ may not be satisfying in the given application.

In the process of surface reconstruction, the distance between the NURBS $S(u, v)$ and the point data $\mathcal{P}$, and vice versa respectively, surely are the most important subjects of optimisation. Due to the continuous definition of $S(u, v)$, only a discrete approximation of its distance to $\mathcal{P}$ is possible. Therefore, discrete point samples $\mathbf{s}_{u,v}$ of $S(u, v)$ are calculated by uniformly dividing the u- and v-intervals. For each sample point the point $\mathbf{p}_{uv}$ in $\mathcal{P}$ with the minimal squared distance is determined. The first objective function f_1 of the MOEA corresponds to the average over all minimal distances.

For fixed parameters u and v, the basis functions in equation (1) have constant values. Thus, the position vector $\mathbf{s}_{u,v}$ can be expressed as linear-weighted combination of the control points. In an optimal approximation

$$\forall u, v : \mathbf{s}_{uv} = \mathbf{p}_{uv} \tag{2}$$

holds. The basic functions and the weights form a coefficient matrix $\mathbf{C}$, which is applied to the vector $\mathbf{x}$ of control points to interpolate the vector of corresponding scan points $\mathbf{s}$. In this vein, a linear system of equations

$$\mathbf{Cx} = \mathbf{s}$$

is established. If the number of sample points is higher than the number of control points, this system is over-determined. Thus, no exact solution exists. SVD obtains a solution $\mathbf{x}$ that minimises the squared distance between the corresponding points in equation (2). Hence, $\mathbf{x}$ gives a discrete approximation of the optimal control point coordinates with regard to the chosen corresponding points.

The other direction, where the average squared distance of $\mathcal{P}$ to $S(u, v)$ is determined, is realised by projecting each $\mathbf{p}_k \in \mathbf{P}$ on $S(u, v)$ [20] leading to pairs $(\mathbf{p}_k, \mathbf{p}'_k)$ of correspondent points, where $\mathbf{p}'_k$ denotes the projection of $\mathbf{p}_k$ on the surface. For the evaluation within the MOEA, the distance of $\mathbf{p}_k$ to $\mathbf{p}'_k$ is determined and the averaged distance over all scan points is denominated f_2. Since the projection of a point on a NURBS surface is a complex task and the number of scan points is very high in most cases, the use of f_2 is very expensive.

In order to process f_2 with SVD, for each pair the following equation is set up:

$$\mathbf{p}'_k = \mathbf{p}_k. \tag{3}$$

Since $\mathbf{p}'_k$ is a point on the surface, it can be described as a linear combination of control points. According to f_1 we obtain a set of linear equations. Usually, this system of equations is over-determined due to the high number of scan points in $\mathcal{P}$. SVD obtains the solution that minimises the squared distance between the corresponding points.

Although it seems that the direction of the distance evaluation has no significant effect, the surfaces obtained with regard to only one of the distance functions show contradicting shapes. Experiment 1 provides insight at interesting effects of both objectives. All experiments performed in the paper are described and organised using a scheme suggested by Preuss [18].

5.1 Experiment 1: Effect of the Objectives

Pre-experimental planning. Due to the evaluation runtime of f_2, in this case the scan point set has to be intensively filtered. For the given amount of scan points, the runtime of a single evaluation amounts to over 20 seconds leading to a optimisation runtime of several days. Thus, we implemented a method to filter the point data while preserving its characteristics, which is introduced in section 6. With this method the scan data has been filtered to 823 points decreasing the runtime to nine hours for 20000 function evaluations.

Task. The experiment aims at gaining insight into the general effects of both distance functions with regard to the MOEA and the numerical approach repsectively.

Setup. The experiments are performed using a self-implemented Visual C++ surface reconstruction software on an AMD Athlon 64 2GHz processor with the NVidia GeForce 6200 Turbo Cache graphics card. The initial surface and the 17307 scan points to be approximated are shown in Figure 1. This surface combines challenging characteristics such as smooth parts with changing curvature, squared edges, and precipitous vertical areas. The genome represents the $x-$, y-, and z-coordinates of the control points where all weights are assumed to be equal to one. The given surface is defined by a 5×10 control net resulting in a dimensionality of 150.

Three EA runs over 10000 generations with different random seeds have been performed. A low number of parents and offspring $\mu = \lambda = 2$ is chosen because only a single objective is considered and function evaluations have to be sparingly utilised. In this vein, the concept of greedily using information obtained during selection, applied in the SMS-EMOA, is transferred to the single-objective case. For SBX and PM, $\eta_c = 15$ and $\eta_m = 20$ are used. More runs and an

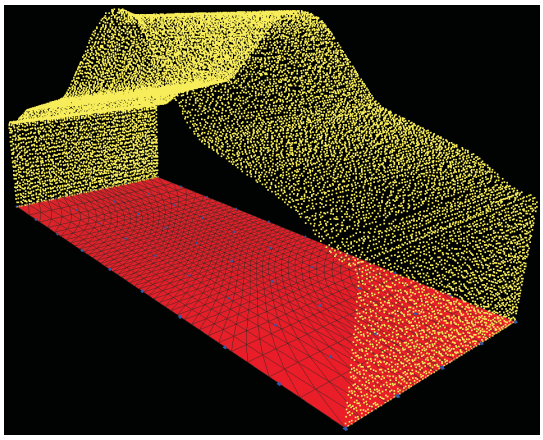

Figure 1: The test surface and the point data used within the experiments.

optimised parameter setting [1] would be desired, but cannot be provided due to running times of several hours per run. In the evaluation of f_1, a 60×120 grid of sample points is applied to guarantee an accurate representation of the surface and establish an overdetermined system of equations.

The effects are assessed by the software-provided visualisation of the genome and the objectives. Here, the distance of the parts of the NURBS to the scan points is encoded by colours from green, alternatively light gray, (close) to red, alternatively dark gray (far). The scan points are coloured according to their distance to the surface.

The SVD approach is also iteratively applied. After one step of the reconstruction process, new equations are set up with regard to the new shape of the surface and the optimisation is repeated. A total of 100 iterations is performed per objective. This approach is closely related to the *iterative closest point* (ICP) algorithm of Besl and McKay [3], which is a standard approach for the registration of a point cloud to a given CAD model by translation. The improvement, which is obtained by the refined assignment of corresponding points, is also analysed.

In both objective functions the assignment of corresponding points presumes the same spread for the surface and the point set. To assure this condition, the corners of the surface are positioned on the corner points of the scan data. A method to find these points is introduced in section 6.

Experimentation/Visualisation. For the progression of the objective values obtained by the EA, box and whiskers plots are shown in Figure 2. Figure 3 visualises the trend of the objective values over the iterations of the SVD approach. The final solutions in decision space are shown.

Observations. The EA can rapidly reduce both distance functions, whereby the optimisation of f_1 proceeds faster, shows less variance, and reaches better final median values (0.074 compared to 0.658). With regard to the concerned objective, the application of the EA is successful.

However, the surface which results from the genome does not represent a satisfying solution. On both objectives the minimisation of one distance causes a degradation of the other. In addition the general trend of changing the spread of the surface is different. Objective f_1 constricts the surface in order to simplify the convergence to some of the points while f_2 enlarges the surface to cover the complete scan data.

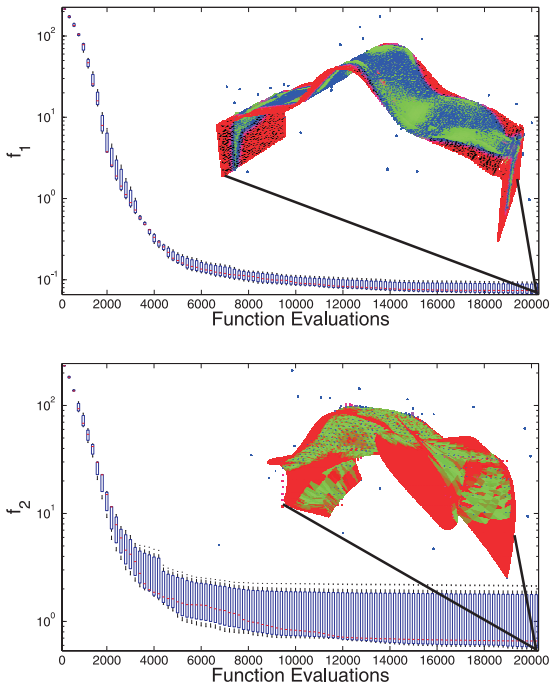

Figure 2: Box and whiskers plots of the objective values over the generations performed. Objective f_1 is depicted in the upper and f_2 is depicted in the lower figure. The solutions ultimately obtained are visualised in the top right corner.

The SVD approach can further improve its objective values if it is repeated because the assignment of corresponding pairs can be refined with a better approximation. The applicability of the ICP algorithm can be transferred. In contrast to the EA, objective f_2 can be faster improved and reaches a much lower value (0.042). As per the evolutionary approach, the surface is enlarged, which results in parts that are far from the the scan data and, thus, not desired. The value of f_1 is higher than the one obtained by the EA (0.118). However, the visual impression of the surface resulting from optimising f_1 is the best in this experiment. The surface also reaches a low value of 0.121 with regard to f_2, indicating that only good comprise solutions provide the desired shape.

Discussion. Both surfaces found by the EA are very irregular, giving a reason to fix the x- and y- coordinates of the control points as done by Beielstein et al. [2]. However, a modelling of the sharp edges as well as of the vertical planes will not be possible with this restriction.

The solution provided by SVD when optimising only f_1 is very promising because both objectives are equally small. In addition the objective functions have shown a negative correlation when applied within the EA. Thus, a multi-objective optimisation of both distance functions promises to lead to the desired compromise results and is therefore analysed in experiment 2.

However, the optimisation using the SVD approach also took much longer than one hour. The time for the calculation of the equations should be optimised by restricting the points that are candidates to be the nearest neighbour of a given sample point. We propose such a method in section 6.

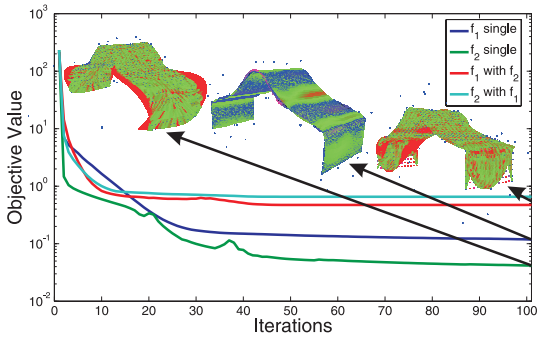

Figure 3: The run of the objective values for different setups of the SVD approach.

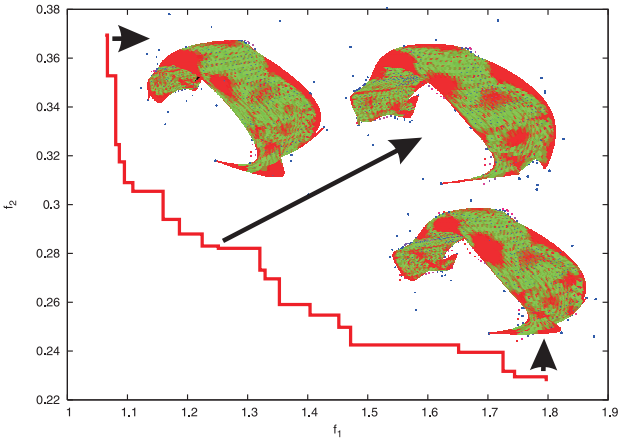

Figure 4: The maximum attainment surface of the three runs of the $(20 + 1)$ SMS-EMOA. Exemplary, the genomes of some solutions are visualised.

The slower convergence of the SVD optimisation with respect to the first objective is due to the effect that the surface initially does not cover the upper part of the scan points because the sample points in the middle of the surface are projected to undesired scan points that are closer. By adjusting the surface the projections of these points gradually get better and the final solution can be provided.

5.2 Experiment 2: Combining the Objectives

Task. In this experiment the combination of both objectives is analysed to find out whether the shape of the reconstructed surface can be improved.

Setup. The basic settings are the same as in Experiment 1. The SMS-EMOA is applied in the suggested steady-state scheme using a population size of 20 to provide a set of compromise solutions. The number of parents is limited to two as in the single-objective case. The number of generations is set to 20000 to admit the same number of function evaluations. The SVD approach solves the combined system of linear equations for both objectives.

The progression of the objective values in the combined optimisation with SVD are depicted in Figure 3 (f_1 with f_2 and f_2 with $f1$ respectively). This Figure also provides the visualisation of the resulting surface.

Observations. The multi-objective optimisation of f_1 and f_2 using SMS-EMOA provides compromise solutions that dominate the solutions found by the EA optimising only f_2. Unfortunately, this improvement has to be compensated with values of f_1 that are much higher than in the single-objective optimisation (1.0–1.8 compared to 0.074). As depicted in Figure 4, the genomes of the solutions found are very close to each other resulting in surfaces with similar shapes.

When both systems of equations are solved together, the improvement of the objective functions within an iteration is much smaller. After 10 iterations no further improvement can be achieved. The solutions are worse compared to the results obtained by optimising only one function. The shape of the surface is not regular and some data points are not covered.

Experimentation/Visualisation. The maximum attainment surface [25] of all runs is shown in Figure 4.

Discussion. The far distance to the minimal values of each objective found in Experiment 1 and the similarity of the genomes indicate that the EA has been trapped in a local extremum. The solution $(0.12, 0.12)^T$ found by SVD only with respect to f_1 dominates the whole maximum attainment surface, indicating that the Pareto front has not been found. Since the other runs performed achieved even worse individuals, the results obviously demonstrate that the optimisation of the control points should be performed using SVD. Here, only objective f_1 should be considered because f_2 does engage an undesired spread of the surface and the projection of a scan point on the surface is computationally expensive. The multi-objective approach scales bad between the different kinds of equations resulting in objective values that are nearly the same.

Right now, only the position of the control points has been optimised. The weights have been assumed to equally emphasise each control point. For a more accurate approximation of the square edges to the vertical areas, the weights of some control points have to be shifted. Nevertheless, the optimisation of the control point's weights cannot be expressed via linear equations. Thus, a hybrid approach, which uses SVD for the positioning of the control points and the MOEA to adjust the weights, is analysed in Experiment 4.

6. RUNTIME OPTIMISATION

The first experiments point out that a pre-processing of the data is necessary before the optimisation of the surface can be efficiently applied. Since most scanners can only provide disconnected and partly unstructured data with an unnecessary accuracy in most of the parts, a method to organise and filter the given information has been developed as a module within our software. In this method the rectangle defined by the minimal and maximal coordinates of the scan points is divided in cubes (voxels) with an approximated edge length of ε specified by the user. Afterwards, each scan point is assigned to the corresponding voxel and the voxel is stored in a list of occupied voxels.

Based on the grid of now organised scan points, the filtering can be implemented by only retaining the scan point with the minimal distance to the centre of the corresponding voxel. This technique is closely related to the distribution scheme in the ϵ-MOEA [12], which is known to produce well-distributed sets. Figure 5 shows the filtered point data from Figure 1. The amount of scan points has been decreased

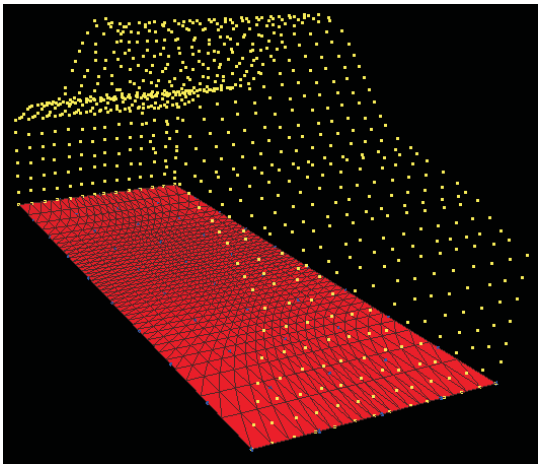

Figure 5: The filtered point data of Figure 1.

from 17307 to 823 scan points. The distance between the points is nearly uniform in parts with same curvature and the density increases in the areas of square edges.

Furthermore, it is possible to detect the corners of the scan point set by simply counting the voxels in the neighbourhood that are occupied by a scan point[1]. The four points with the minimal count are assumed to be the corner points. The corners of the NURBS can be assigned to the corners with regard to their geometric relation.

The calculation of the distance between a point on the surface $\mathbf{s}$ and a point $\mathbf{p}$ of the data set is the most frequently applied operation within the calculation of f_1. To reduce the number of scan points to be taken into account and allow for a faster evaluation, again, the grid of voxels is used. Each voxel shall contain the scan points that are candidates to be the nearest point $\mathbf{p}$ to a point $\mathbf{s}$ in this voxel. To this end, for each unoccupied voxel the minimal distance of a point to this voxel is determined. For occupied voxels the minimal distance is set to zero. Afterwards, all points, whose distance is smaller than the sum of the determined minimal distance and the length of the voxel's diagonal, are added to the candidate set. This procedure guarantees that no possible candidate point is excluded, but often produces unnecessary large candidate sets. Thus, an elimination step is performed to filter this unnecessary data. Here, for each point $\mathbf{p}$ in the candidate set, the projection $\mathbf{p}'$ to the corresponding voxel is determined. Afterwards, the distance of all other points in the candidate set to $\mathbf{p}'$ is calculated. If a point is closer to $\mathbf{p}'$ than $\mathbf{p}$, $\mathbf{p}$ is removed. For the calculation of the candidate set, the grid should be slightly expanded. In this manner sample points $\mathbf{s}$ that lie above or beneath the scan points can also be processed.

6.1 Experiment 3: Runtime

Pre-experimental Planning. It can be expected that the runtime of f_1 is significantly improved by restricting the points that have to be concerned during the determination of the minimal distances. The effect of different grid resolutions to provide an adequate setting has also to be analysed.

[1]The size of the neighbourhood must be chosen large enough to avoid gaps in the occupied areas of the grid.

Table 1: Run-time of the given settings in ms

Objective	Size of grid	Surface_{init}	Surface_{opt}
$f_{1,unimpr.}$	arbitrary	2219	2250
$f_{1,impr.}$	952 voxels	47	47
	1680 voxels	31	46
	2750 voxels	31	46

Task. Our hypothesis is

$$T_{exp}(f_{1,unimpr.}) - T_{exp}(f_{1,impr.}) > 0,$$

where T_{exp} denotes the expected running time. The indices *impr.* and *unimpr.* denote the improved and unimproved version of the evaluation algorithm respectively.

Setup. If no other jobs are running on the CPU, the runtime shows no stochastic effects. Thus, each parameterisation has to be only performed a single time. The runtime is determined for two different kinds of surfaces. Surface_{init} is the starting solution shown in Figure 1 situated far away from the point data. Surface_{opt} is taken from the SVD optimisation and provides a well approximation of the whole point data. The hypotheses have to hold for all surfaces and different resolutions of the grid of candidates. The maximal amount of voxels is increased until no further improvement can be measured, starting from 1000 in steps of 1000 voxels[2]. Again, the sampling of the surface is performed in a regular 60×120 grid of the u-v-space.

Experimentation/Visualisation. The results of the experiment are reported in Table 1.

Observations. The results approve the first hypothesis. The use of the grid accelerates the evaluation regarding both surfaces and all improved objective functions. In the best case, an increase in efficiency of factor seventy is realised.

Discussion. From the results can be inferred that a grid resolution smaller than 3000 voxels should be chosen because no further improvement can be gained. A size of 2000 voxels is recommended.

7. RESULTS OF THE HYBRID MOEA

This section reports the results obtained with the hybrid approach and the pre-processing steps that are emerged from the experiments in the section 5 and 6. To evaluate the performance of the algorithm, the objective values, the shape of the surface, as well as runtime aspects are considered.

7.1 Experiment 4: The hybrid approach

Task. In the experiment, we analyse the hybrid approach with regard to all important aspects for integration into surface reconstruction software.

Setup. Sequentially, the control points are positioned by SVD and the weights are adjusted using the EA, both only with respect to f_1. This objective is chosen because of good results of the SVD approach in Experiment 1 and the fast improved evaluation (Experiment 3). The effects of constriction shown by the EA in Experiment 1 are avoided because only the weights of fixed control points are optimised. This way only modifications of the inner shape of the surface are possible. The EA parameters are chosen according to Ex-

[2]The implemented grid sizes are smaller because the dimensions of the grid are adapted to provide approximately cubic voxels

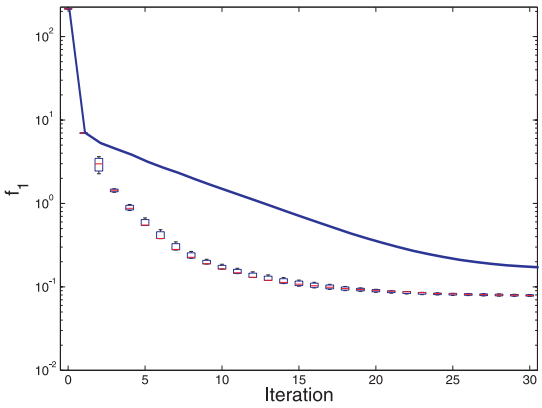

Figure 6: Box and whiskers plot of the objective value over the iterations performed with the hybrid approach. Also the run of the single SVD approach is depicted.

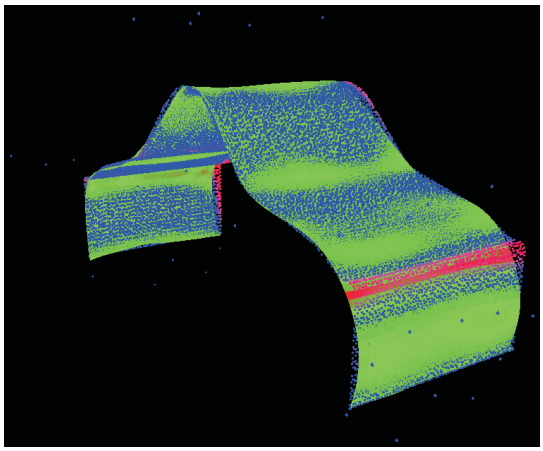

Figure 7: The final surface obtained by the hybrid optimisation approach.

periment 1. With regard to the observations from Figure 3, only 30 iterations are performed.

Experimentation/Visualisation. The objective values obtained after each SVD iteration are shown in Figure 6 using a box and whiskers plot. For comparison, the run of f_1 using only SVD is also depicted. The final surface after 30 iterations is shown in Figure 7.

Observations. The convergence speed can be significantly increased. Just approximately 15 iterations are necessary to reach the objective values obtained after 100 iterations when only SVD is applied. The runtime is about 1 minute per iteration providing very good solutions within less than 10 minutes. Thus, the use of the grid of candidate points provides acceptable evaluation times. The shape of the final surface is very regular. The best objective vector that can be obtained $(0.078, 0.125)$ dominates the best solution found in Experiment 1 requiring only 5 percent of the runtime.

Discussion. The hybridisation of both approaches can improve the results. The refinement of the corresponding pairs when setting up the equations in combination with an iterative optimisation of the weights provides high-quality results in short time.

8. SUMMARY AND OUTLOOK

The paper analyses the general ability of EA to be utilised in surface reconstruction. It is shown that standard EA only guided by distance functions rapidly get trapped in local optima, which are found by constricting and expanding the surface respectively. The multi-objective optimisation of both objectives just slightly evades this problem. Only a narrow set of compromise solutions is found due to a small diversity in decision space. The identification of the border points and a new objective function with regard to the borders of both the scan points and the surface is necessary to avoid these problems. Methods to provide and use this information will be presented in the near future.

However, the runtime of the EA under investigation ia still too long to be reasonable applied. In the experiments, evolutionary operators and algorithms are analysed, which are known to be the state-of-the-art. Nevertheless, it cannot be inferred that no EA exists, which is able to successfully optimise surfaces in acurate time due to a lower number of required function evalutions. The CMA operator by Hansen and Ostermeier [10] and its multi-objective variant [11] may be candidates. Beyond, Rudolph, Naujoks, and Preuss [19] give hints on heuristics, which also provide diversity in decision space. Maybe with these techniques, the problem of only finding parts of the Pareto front can be tackled.

An alternative to evolutionary techniques is provided by expressing the problem in terms of equations. The iterative application of SVD, a robust solving method for systems of linear equations, obtains solutions of a better quality about 10 times faster considering the distance between the surface and the scan data. Right now, only a method to find the optimal control point positions by means of SVD is established. The optimal setting of the weight vector has to be found by more general types of optimisers. The experiment with a hybrid evolutionary approach provide the desired results. The use of additional data structures allows for an efficient evaluation and an acceptable runtime.

Acknowledgements

This work was supported by the German Research Foundation (DFG) as project T3 in the collaborative research center SFB 531 and project C4 in the collaborative research center SFB 708 at the University of Dortmund, Germany.

9. REFERENCES

[1] T. Bartz-Beielstein. *Experimental Research in Evolutionary Computation – The New Experimentalism.* Natural Computing Series. Springer, Berlin, 2006.

[2] T. Beielstein, J. Mehnen, L. Schönemann, H.-P. Schwefel, T. Surmann, K. Weinert, and D. Wiesmann. Design of evolutionary algorithms and applications in surface reconstruction. In H.-P. Schwefel, I. Wegener, and K. Weinert, editors, *Advances in Computational Intelligence - Theory and Practice*, pages 164–193. Springer, Berlin, 2003.

[3] P. J. Besl and N. D. McKay. A method for registration of 3D shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 8:239–256, 1992.

[4] N. Beume and G. Rudolph. Faster S-Metric Calculation by Considering Dominated Hypervolume as Klee's Measure Problem. In *International Conference on Computational Intelligence (CI 2006)*, 2006. (in print).

[5] K. Deb. *Multi-Objective Optimization using Evolutionary Algorithms*. Wiley, Chichester, UK, 2001.

[6] K. Deb and H.-G. Beyer. Self-adaptation in real-parameter genetic algorithms with simulated binary crossover. In W. Banzhaf et al., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999)*, pages 172–179. Morgan Kaufmann, 1999.

[7] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.

[8] M. Emmerich, N. Beume, and B. Naujoks. An EMO algorithm using the hypervolume measure as selection criterion. In C. A. C. Coello et al., editors, *Evolutionary Multi-Criterion Optimization: 3rd Int'l Conf. (EMO 2005)*, pages 62–76. Springer, Berlin, 2005.

[9] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley, Reading, MA, 1989.

[10] N. Hansen and A. Ostermeier. Completely Derandomized Self-Adaptation in Evolution Strategies. *IEEE Computational Intelligence Magazine*, 9(2):159–195, 2001.

[11] C. Igel, T. Suttorp, and N. Hansen. Steady-state selection and efficient covariance matrix update in the multi-objective CMA-ES. In S. Obayashi et al., editors, *EMO2007, LNCS 4403*. Springer, Berlin, 2007.

[12] M. Laumanns, L. Thiele, K. Deb, and E. Zitzler. Combining convergence and diversity in evolutionary multi-objective optimization. *Evolutionary Computation*, 10(3):263–282, 2002.

[13] J. Mehnen. *Evolutionäre Flächenrekonstruktion*. Dr.-Ing. Dissertation, University of Dortmund, 2000.

[14] B. Naujoks, N. Beume, and M. Emmerich. Metamodel-assisted SMS-EMOA applied to airfoil optimization tasks. In R. Schilling et al., editors, *Proceedings EUROGEN'05* (CD-ROM).

[15] B. Naujoks, N. Beume, and M. Emmerich. Multi-objective optimisation using S-metric selection: Application to three-dimensional solution spaces. In *Evolutionary Computation Congress (CEC'05), Edinburgh, UK*, volume 2, pages 1282–1289, Piscataway NJ, 2005. IEEE Press.

[16] L. Piegl and W. Tiller. *The NURBS book*. Springer, Berlin, 1997.

[17] W. H. Press, S. A. Teukolsky, W. T. Vettering, and B. P. Flannery. *Numerical recipes in C++. The Art of Scientific Computing*. Cambridge University Press, 2002.

[18] M. Preuss. Reporting on experiments in evolutionary computation. Reihe CI 221/07, SFB 531, Universität Dortmund, Januar 2007.

[19] G. Rudolph, B. Naujoks, and M. Preuss. Capabilities of EMOA to detect and preserve equivalent Pareto subsets. In S. Obayashi et al., editors, *EMO2007, LNCS 4403*. Springer, Berlin Heidelberg, 2007.

[20] I. Selimovic. Improved algorithms for the projection of points on NURBS curves and surfaces. *Computer Aided Geometric Design*, 23:439–445, 2006.

[21] T. Wagner, N. Beume, and B. Naujoks. Pareto-, aggregation-, and indicator-based methods in many-objective optimization. In S. Obayashi et al., editors, *EMO2007, LNCS 4403*, pages 742–756. Springer, Berlin Heidelberg, 2007.

[22] K. Weinert, J. Mehnen, and M. Schneider. Evolutionary optimization of approximating triangulations for surface reconstruction from unstructured 3D data. In H. J. Caulfield and M. Grana, editors, *Proceedings of the 6th Joint Conference on Information Science (JCIS 2002)*, pages 578–581. Duke University, 2002.

[23] K. Weinert, T. Surmann, and J. Mehnen. Evolutionary surface reconstruction using CSG-NURBS-hybrids. In L. Spector et al., editors, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2001*, pages 1456–1463. Morgan Kaufmann Publishers, 2001.

[24] D. York. Least-square fitting of a straight line. *Canadian Journal of Physics*, 44:1079–1086, 1966.

[25] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. Grunert da Fonseca. Performance Assessment of Multiobjective Optimizers: An Analysis and Review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003.