

Surface Fitting Using Genetic Algorithms

Andrew Mason, Technical Director, Formation Design Systems Pty Ltd, Fremantle, WA, Australia

Surface Fitting

Computer aided hull design systems have been widely adopted and have dramatically reduced the amount of work needed to generate a fair linesplan. However, the rapid change in design technology from drawing board to computer has been responsible for the orphaning of an immense number of existing designs previously stored as lines drawings. Although it is possible to scan these drawings into a computer and create a 3D wireframe model relatively easily, creating a fair surface model suitable for analysis or further design modification is another matter.

In the January 2002 issue of The Naval Architect, Dr Herbert J Koelman gave some background to the various surface types used by marine hull modelling software and mentioned the Genetic Algorithm based surface fitting function added to Maxsurf in 2001. Although the method has now been widely adopted by Maxsurf users, it is not widely known by the broader CAD community. This article is intended to give a better understanding of the use of Genetic Algorithms for surface fitting as implemented in Maxsurf.

As mentioned by Dr Koelman, several methods are available for surface fitting to existing designs. However, these tend to have been created for disciplines that are less fussy about surface fairness and curvature smoothness than the shipbuilding industry. As a result the use of these methods tends to give less than satisfactory results, particularly when it is desirable to use the fitted design as a parent model for future modification. In this case, it is important that the net of surface control points be fair and regular otherwise it will be almost impossible for a designer to manipulate them in a meaningful way.

Existing Fitting Methods

Before describing the Genetic Algorithm approach to surface fitting it is worth looking at existing fitting methods and their problems.

In the past the most common method for surface fitting was the use of the Coons patch. This method simply looks at four boundary curves and blends them to produce a surface that passed through the edge curves. This approach has several drawbacks, primarily the need to create a series of fair boundary curves that cross one another and form a four sided region. This in itself is a time consuming process. Once a set of Coons patches covering the hull are created, they are cumbersome to modify due to the large number of small surfaces that are bonded together to create a larger surface.

A better approach is to fit a small number of NURBS surfaces (Non-Uniform Rational B-splines, as used by Maxsurf, Autoship and Fastship) to the hull offset data. The most common method for the fitting of NURBS surfaces to a set of data points is an approach referred to as “skinning”.

To perform the skinning operation, a series of profiles is first created by fitting curves through offset data for individual sections. Once this has been done, a surface is created that fits through these section profiles. The quality of the surface is very dependent on the spacing and accuracy of the chosen profiles.

For some time the Maxsurf suite of programs has included Prefit, a program that fits a surface to a table of offsets using the skinning method. Over the years, as time was spent developing and using Prefit, the deficiencies of the skinning method became more and more apparent. The problems with Prefit and other skinning methods are significant. In particular, although the time to fit a surface may be short, the time taken to get a good result may be very long indeed. A great deal of trial and error with different surface parameters may be needed before a satisfactory result is achieved and sometimes it is just not possible to produce a surface of sufficient quality to be used for hydrostatic analysis or further design work.

In addition to the time spent trying to achieve an acceptable quality of fit, the time spent preparing data prior to the fitting process taking place can be onerous. Each profile is made up of a series of data points and it is important that these points be in the correct order. If one of the points is created or selected out of order, the fitted profile curve and the resulting fitted surface will be flawed.

The skinning process works best when all of the fitted profiles are transverse sections. It is possible to incline the stem and stern profile, however this can disturb the fitting process and result in a severely distorted network of control points.

It is not possible to use profiles that cross one another or to include curves that intersect these profiles, such as waterlines or diagonals. This means that valuable information about the shape of the hull is often ignored by the fitting process.

In summary, the skinning process suffers from the following problems -

Ordered data points. Profile skinning requires that data points be ordered in consecutive profile from one end of the surface to the other, for example from the stern to the bow. Points within each profile should be in the correct order for example from the hull centreline to the sheerline.

Density of data points. Profile skinning has problems if the data points are unevenly distributed over the surface.

Poor net of control points. The resulting net of control points from the skinning process is often not suitable for modification in normal design work. Columns and rows may be skewed and be irregular in their layout. Typically, the skinning process can generate hundreds or thousands of control points making the net completely unsuitable for future manual modification.

Poor fairness. Even when the skinned surface fits the data closely it is possible for it to contain significant unfairness.

A New Approach

These problems convinced the author to embark on a project to improve on the Prefit program. An approach inspired by some original work by Richard Birmingham and Tim Smith at the University of Newcastle was investigated and adopted.

Birmingham and Smith had investigated the use of Genetic Algorithms to solve a variety of optimisation problems in Naval Architecture. They initially focussed on the problem of parametric transformations of hull surface models (Ref 1). To do this they took a Maxsurf model and manipulated the surface control points using a Genetic Algorithm in order to achieve certain values of LCB, Prismatic Coefficient and Midship Area.

Birmingham and Smith also considered the surface fitting problem, but in their later work chose to investigate the application of Genetic Algorithms to the more complex *ab-initio* design problem, an ambitious undertaking given the number of variables in the design of a ship (Ref 2).

On reviewing Birmingham and Smith's work, it appeared to the author that the surface fitting problem was sufficiently complex to justify the Genetic Algorithm approach and that the number of degrees of freedom was small enough for a solution to be achievable. Despite its applicability, the Genetic Algorithm approach to surface fitting did not appear to have been fully investigated.

The fitting of a NURBS surface to a set of data points can be viewed as an optimisation problem having one or more objectives (minimum error between the surface and the data points, as well as a measure of surface fairness), a number of degrees of freedom (the locations of the surface control points in space) and a number of constraints (such as a flat of bottom or flat of side limit). The solution space of the surface fitting problem is also multi-modal, that is, it contains a large number of solutions that are good but not the best solution. This is very different for a uni-modal problem which has only one optimum solution.

Although a number of different optimisation methods are available, a method was required that would give a good solution with multiple objectives, a large number of degrees of freedom, multiple constraints and a multi-modal solution space. Given the recent success of Genetic Algorithms with problems of this nature it was felt that they warranted investigation as an fitting method.

What is a Genetic Algorithm?

Biological evolution can be viewed as a process that optimises a species to its environment and ecological niche. As an optimisation process it is remarkably effective, although painfully slow. The Genetic Algorithm is simply the use of a computer based model of evolution to find an optimum solution to a problem. Any problem that can be quantified numerically by a set of input parameters and an objective can have a Genetic Algorithm applied to find its optimum solution.

Genetic Algorithms differ from more traditional optimisation techniques in that they involve a search from a "population" of solutions, not from a single point. Each iteration of a Genetic Algorithm involves a competitive selection that weeds out poor solutions. The solutions with high "fitness" are "recombined" with other solutions to produce members of the next generation, which inherit properties from their parents. Solutions are also "mutated" by making small random changes to one or more characteristics. Recombination and mutation are used to generate new solutions that are biased towards regions of the space for which good solutions have already been seen.

Fitting Surfaces with Genetic Algorithms

The Genetic Algorithm applied to the surface fitting problem in Maxsurf performs the following steps–

1. Create an initial population of surfaces each of which varies slightly from the others.
2. Repeat the following steps until a solution is found that meets the fitting criteria sufficiently well.
 - A. Measure the fit of each individual surface to the data points and give it a score based on its quality of fit, fairness and net regularity. Rank all the individuals from best to worst.
 - B. Repeatedly choose two individuals as parents, with individuals being chosen in proportion to how highly they were ranked in step A. This means that the most successful individuals produce the most offspring in the next generation. Create members for the next generation by choosing some surface control points from one parent and the remainder from the other parent. Occasionally include a small amount of random control point movement to provide some mutation from generation to generation.

These steps result in a population that progressively gets closer to an optimum solution of the problem, with poor solutions being eliminated. The highest ranked individual in the final generation is the best solution.

Results

Maxsurf's GA based fitting has proven to be robust and gives good results in most cases. It has been shown to have the following advantages over the skinning algorithm -

Unordered or Random Data. Points can be in any order – it is not necessary to organise data into rows or columns, nor is it necessary for points to be sorted into a particular order. The GA fitting method is able to handle randomly ordered.

Data Density. The density of point data is unimportant, with it being acceptable to concentrate more points in areas of greater curvature or detail.


Multiple Objectives – Maxsurf uses Root Mean Square (RMS) error as a measure of surface fit. RMS error is the square root of the sum of the squares of all of the errors between each data point and its nearest point on the fitted surface. In addition, measures for surface fairness, control point net spacing and perpendicularity of the net rows and columns are used.

Multiple constraints - flat of side and flat of bottom constraints may be applied.

Manual Intervention With a Genetic Algorithm it is possible to interrupt the fitting process at any point and manually modify the surface that has been developed. If fitting is then restarted the manual modifications will be incorporated as long as they are determined by the Genetic Algorithm to be improvements when compared to the prescribed objectives.

Regular Net The net of control points produced tends to be fair and regular, making it possible to use the fitted surface as a basis for further design modifications.

Good Surface Fairness- fitted surfaces tend to be fair with a minimum of surface inflections.



Genetic Algorithm based surface fitting to a table of offsets derived from the NPL series.

The main disadvantage with the Genetic Algorithm is performance, with a reasonable fit to surface data points taking between 20 minutes and 2 hours on a 1GHz Pentium class PC. Although this seems high, it needs to be balanced against the time savings gained in data preparation, and the fact that it is usually possible to get a reasonable result without having to repeat the fitting process over and over.

Conclusion

Maxsurf has contained a Genetic Algorithm based fitting method for more than one year and it has proven to be very effective. Although fairly slow to calculate it is believed that it still results a substantial saving in operator time when the complete process fitting an acceptable surface is considered.

Ongoing research is being done to determine how to use the fitting function to fit multiple surfaces simultaneously, to recognise chine and knuckle lines automatically and to determine the ideal number of rows and columns for a fitted control point net.

In the meantime it is felt that the approach is a significant improvement over previous methods that will greatly assist in the conversion of existing designs from drawings or as built vessels to computer based NURBS surface representations.

References

1. Birmingham, R W , Smith T A G: "A Practical Approach to Form Transformations for Software Generated Designs" Computer Aided Design and Productions for Small Craft Conference, Southampton, 1995.
2. Birmingham, R W , Smith T A G: "Automatic hull form generation, a practical tool for design and research" Practical design of Ships and Mobile Units, The Hague 1998.