

Shape Modification of Ship Hulls in H-rep

Bastiaan N. Veelo, Norwegian University of Science and Technology¹

1 Introduction

The transfinite interpolation of an irregular network of curves is an effective modelling methodology for the design of non-trivial free-form shapes. Recent literature refers to this methodology with the term hybrid representation or H-rep, to indicate that it is based on the merger of wire-frame modelling and solid modelling. In practice, the term H-rep also implies the integration of a curve fitting/fairing algorithm, which is essential for the discussions in this article. *Koelman (1999)* gives a detailed description of the conception of the H-rep concept, and of an implementation called 'Fairway', which is a module in the 'PIAS' package and tailored to (but not limited to) the geometric design of ship hulls. Introductions to the H-rep concept and its value for the maritime industry include *Koelman et al. (2001)*, *Koelman (2003a,b)*, *Veelo (2004)*.

In essence, Koelman's implementation restores the traditional way of lines plan draughting in a computer method. To the user, the system presents itself as a curve modeller. The surface generation, which consists of filling the mesh cells of the curve network with transfinite surface patches, is completely hidden for the user. The patches may have an arbitrary number of sides and are tangent-plane continuous across shared boundaries. Advantages of this modelling concept are the absence of topological restrictions on surface features, the independence of curves and their details, and tight control over the exact shape of the composite surface.

Typically, when starting a design from scratch using this system, the designer starts with an initial model defined by a centre line contour, a deck line and a mid-ship ordinate. These curves are computer-generated, based on user-defined main dimensions. The shape of the surface patches is completely derived from the shape of the curves, so the only means to control the shape of the model is through manipulation of curves. Thus, the first step in the design process is to modify the existing curves to their correct shape, by traditional control point manipulation. When the surface is visualised at this stage, the designer will be probably not satisfied with the composite shape of surface patches. The patches are still large and the defining curves are too far apart to describe every detail that is envisioned for the design. The solution is to add more curves to the network, effectively subdividing patches into smaller ones. New curves can be generated automatically by intersecting the model with a user-defined plane, or by projecting a separate curve onto the surface. After addition of a new curve, the shape of newly subdivided patches can be modified by manipulating the curve. With this process, the shape of a sculpted model evolves from a coarse definition to a detailed definition, until the designer is satisfied with the result.

There is a downside to this modelling methodology. As the design progresses and more curves are added to the model, more of its shape is rigidly defined. The more curves present, the smaller the surface patches, and the more local shape manipulations become. More curves also mean more curve intersections. As a result, during curve manipulation, there is a higher risk that a curve is pulled away from these intersections, rendering the network invalid as a surface description. Such inconsistencies appear as gross surface defects. Currently, no mechanism is implemented that prevents this, or that resolves the incompatibilities. Although it is possible to restore the intersections manually, by adjusting all affected curves to the changes, this is a lengthy, iterative task.

Consequently, making design changes that affect larger surface areas of the model, is discouraged at late design stages. One could say that designing sculpted shapes this way, in practice is a one-dimensional process because the design has a preference to evolve only in one direction. This paper proposes a simple and efficient method for shape manipulation of a dense curve network that is not strictly local and does not destroy the consistency and the geometric continuity of the network. This

¹NTNU, Dept. Eng. Design and Materials, N-7491 Trondheim, Norway, Bastiaan.N.Veelo@ntnu.no

makes it possible to migrate directly from one shape variation to the other, by which the design process becomes, in our way of speaking, multi-dimensional.

2 Background

In the computer-aided design (CAD) of free-form or sculpted shapes, ‘Non-Uniform Rational B-Spline’ (NURBS) surfaces enjoy great popularity. NURBS surface patches derive their geometry from control points that they approximate. Whether their popularity is justified, is debatable. The challenge of designing sculpted shapes boils down to two main problems: geometric continuity and control.

A single NURBS surface patch without degenerate sides or corners fits only well to deformations of the square, the cylinder and the torus – of which only the torus can describe the boundary of a solid. All other geometries, thus including most subjects of design, are commonly denoted as having arbitrary topology. For a composition of (non-degenerate) patches to describe shapes of arbitrary topology, the patches must be allowed to be organised in an arbitrary manner, where the number of patches that mutually connect with one of their corners is not always four. Maintaining geometric continuity, i.e. a smooth composite surface, is especially difficult at these irregular points.

The problem of control is implied by the fact that NURBS surface patches approximate a regular grid of control points. The problem becomes eminent, e.g. when more control points are needed locally in order to define some local detail in a surface patch. Since extra control points can only be added in complete rows or columns, they also appear in regions where they are not wanted, because they make achieving surface fairness more difficult.

Transfinite surface patches, which derive their geometry from curves that they interpolate, do not suffer from a control problem, as the patch does not care how its bounding curves are defined. For adjacent transfinite patches to connect with tangent plane continuity, tangent information is required along the curves, which is represented by so-called tangent ribbons. In order to model arbitrary topology, the curve network must also be arbitrary, i.e. without regularity requirements. The only requirement on curves is that they intersect and not cross each other (within some tolerance) and that they start and end at other curves. *Jensen et al. (1991)* were the first to develop a technique for the generation of tangent ribbons on such networks by using a boundary representation (B-rep), which is a data structure used primarily in solid modelling. A B-rep data structure consists of topological elements of type ‘node’, ‘edge’ and ‘face’. References exist in the data structure such that for each element, its neighbouring elements can be determined. *Jensen et al.* applied their technique to automotive styling. *Van Dijk (1994)* took this to conceptual industrial design and *Michelsen (1995)* to naval architecture.

Although having developed an alternative to NURBS surface modelling without the associated problems, it remained a challenge to keep the curve network simultaneously fair and consistent as a surface representation. By integrating a curve fitting and fairing algorithm, *Koelman (1999)* was able to improve that situation, and produced the described implementation for the geometric design of ship hulls, at production quality. In addition, he removed the need for the user to worry about surface patches, by following a suggestion of *Michelsen (1995)* to use the B-rep to its full potential as a solid representation.

Fig.1 shows the hybrid nature of the H-rep. The nodes in the B-rep refer to intersection points in the wire-frame for their geometry. The edges refer to the curve sections in-between the intersection points and the faces refer to the n -sided patches that can be generated to fill the openings in the wire-frame. Tangent ribbons are also partly indicated.




Fig.1: The hybrid representation with references between topology and geometry

3 Related Research

The problem of the global shape of a model getting fixated by the definition of details is not specific to the H-rep and its precedents. It also exists in systems that are based on approximation of control points such as NURBS surfaces, although the consequences are not as dramatic as the surface defects that can arise in an H-rep. If a surface region is to be modified for which a larger number of control points need to be shifted, this must be done in a way that preserves the coherence between the control points, so that both the global shape remains fair and the detailed surface features are not damaged. This has been addressed by the integration of physics based properties, e.g. by *Terzopoulos and Qin (1994)* and *Leon and Trompette (1995)*, and hierarchical refinement, by *Forsey and Bartels (1988)*. These references do not explicitly consider arbitrary topology however. *Gonzalez-Ochoa and Peters (1999)* proposed the hierarchical refinement principle for so-called surface splines, which approximate an arbitrary mesh of control points and thus support models with arbitrary topology. Contrary to plain NURBS surfaces, their contribution may be a viable alternative to the H-rep.

Free-form deformation (FFD), popularised by *Sederberg and Parry (1986)*, is a technique to reshape a geometric model indirectly by warping the space in which it is defined. FFD is independent of the model definition, and thus competes with the method presented here.

4 Manipulation of Sets of Data Points

We will state our problem as follows: “Given a certain region on a surface that interpolates a network of curves, manipulate all curves in that region simultaneously, in a way that does not destroy the consistency of the network and does not introduce unwanted geometric discontinuities.”

This statement can be symplified after the following observation. The details behind the process of adding a new curve consist of tracing a string of data points over the surface, as a sampling of the intersection curve or the projected curve. Then the fitting/fairing algorithm is invoked to generate a curve through these points, which is added to the model. During manipulation of curves (and thus the surface) these data points can be made to move with the curve, so that they indeed remain positioned on the surface. If we assume for the moment that all data points are persistent, meaning that they remain in existence throughout the modelling process, then the complete model can be regenerated from the data points and the B-rep alone, with the help of the fitting/fairing algorithm. Thus, we can

reformulate the problem as: “Given a point set belonging to a consistent H-rep, shift a selection of points to a new position, so that the distance and the direction of the shift of each individual point varies smoothly over the set.”

We will now assert our assumption. Data points that are associated with intersections between curves are persistent, because they represent the geometry of node elements in the B-rep. Currently, other data points are not persistent, as they serve no purpose after the creation of a curve. Nevertheless, in a dense curve network, there will likely be enough intersections (and thus persistent data points) to record the shape of the curves. A simple heuristic can verify this, e.g. by checking whether the number and distribution of data points belonging to a curve stands in proportion to the number and distribution of control points of the curve. If the verification fails, extra data points can be inserted at low computational cost.

4.1 Shift Vectors

Let us declare \mathbf{s}_i to be the ‘shift vector’ for a data point i , i.e. the difference between the position of that point after and before the shape modification. We will define this shift vector as the vector sum of the sample of one or more three-dimensional vector fields. A vector field is primarily defined by a ‘selection field’ j of varying intensity, which is concentrated around a point, a curve or a surface, which we will call the ‘base’ of the selection field, Fig.2. This base may be part of the model, or be dedicated to support the selection field. The intensity f_j of the field will be unity at its base, decreasing smoothly with increasing distance d to the base, and level off to zero at a distance r_j to the base, which we will call the ‘extent’ of the selection field. If the base is singular, r_j is a constant; but if the selection field is based on a curve (or surface), r_j may be a function of the curve parameter (or surface parameters). In a similar fashion, we will define a vector on the base, whose length and direction may be a function of the base parameters. We will call this vector the ‘typical shift vector’ of the selection field, denoted by \mathbf{S}_j . For the vectors in the field to vary smoothly, it is important that selection fields do not self-intersect.




Fig.2: Schematic presentation of how the shift (\mathbf{s}_i) of a point i is derived from a selection field of smoothly decaying intensity f_j and a given shift (\mathbf{S}_j) at the selection base. In this case, the base is singular and $\mathbf{s}_j = f_j(d_{i,j}/r_j)\mathbf{S}_j$.

In addition to a selection field, one or more ‘deselection fields’, enumerated by k , may take part in the definition of a vector field. Deselection fields reverse the effect of the selection field. Their definition is similar to the definition of selection fields, except that they lack a typical shift vector and their intensity g_k is opposite to the intensity of selection fields: unity outside their extent, smoothly decreasing in inverse proportion to the distance d to the base inside their extent, and levelling off to zero at their base.

The vector field is then defined as the typical shift vector, evaluated on the closest point on the base, multiplied by the selection field intensity and the deselection field intensities. Especially for deselection fields it is interesting to have them act differently on the x , y and z coordinates of the

vectors in the field, and thus we will redefine field intensities as diagonal matrix functions:

$$\begin{aligned} \mathbf{f}_j \left(\frac{d_{i,j}}{r_j} \right) &\equiv \begin{bmatrix} f_{j,x} \left(\frac{d_{i,j}}{r_j} \right) & 0 & 0 \\ 0 & f_{j,y} \left(\frac{d_{i,j}}{r_j} \right) & 0 \\ 0 & 0 & f_{j,z} \left(\frac{d_{i,j}}{r_j} \right) \end{bmatrix} \\ \mathbf{g}_k \left(\frac{d_{i,k}}{r_k} \right) &\equiv \begin{bmatrix} g_{k,x} \left(\frac{d_{i,k}}{r_k} \right) & 0 & 0 \\ 0 & g_{k,y} \left(\frac{d_{i,k}}{r_k} \right) & 0 \\ 0 & 0 & g_{k,z} \left(\frac{d_{i,k}}{r_k} \right) \end{bmatrix} \end{aligned} \quad (1)$$

If we then say $g_{k,y} \equiv 0$ for a deselection field based on the plane $y = 0$, data points in that plane will only shift in that plane and not away from it, regardless of the direction of the typical shift vector. This is advantageous if the design is symmetrical around $y = 0$ and only one half of it is modelled. Note that in this specific case, an alternative is to mirror the selection field in the symmetry plane. Deselection fields however are capable of enforcing more general constraints.

The definition of the shift vector can now be formalised as

$$\mathbf{s}_i \equiv \sum_j \left(\prod_k \left(\mathbf{g}_k(d_{i,k}/r_k) \right) \mathbf{f}_j(d_{i,j}/r_j) \mathbf{S}_j \right), \quad (2)$$

where $d_{i,j}$ denotes the shortest distance through space between data point i and the closest point on the base of selection field j , and \mathbf{S}_j and \mathbf{f}_j are evaluated at that position on the base. Analogously, $d_{i,k}$ denotes the shortest distance through space between data point i and the closest point on the base of deselection field k , and \mathbf{g}_k is evaluated at that position on the base. What remains is to find suitable definitions for the selection functions f and g , and for the typical shift vector \mathbf{S} .

4.2 Selection Functions

Any function that behaves as described will give useful results. For more control of the shape of the resulting modification, one may want to vary the shape of the selection function over the base of the selection, as a function of the base parameters. This is possible in the following definition of a cubic piecewise polynomial, in which a parameter $\kappa \in [0, 1]$ defines how fast the function falls off.

$$f \left(\frac{d_{i,j}}{r_j} \right) \equiv f(\delta) \equiv \begin{cases} \frac{\kappa^2 + \kappa\delta^2(\delta - 3) + \delta^3}{\kappa^2} & \text{if } 0 \leq \delta < \kappa \\ \frac{(\delta - 1)^3}{\kappa - 1} & \text{if } \kappa \leq \delta < 1 \\ 0 & \text{if } 1 \leq \delta \end{cases} \quad (3)$$

in which δ has been substituted for $d_{i,j}/r_j$ for simplicity. This function, Fig.3, is derived from the Cox-deBoor recursive definition of B-spline basis functions. The selection function of deselection fields can simply be defined as $g \equiv 1 - f$.




Fig.3: The selection function defined by Eq.(3): light grey for $0 \leq \delta < \kappa$, dark grey for $\kappa \leq \delta < 1$, black for $1 \leq \delta$. Smaller values of κ make the function fall off faster. Plotted are $\kappa = 1.0$, which is point-symmetric about $(0.5, 0.5)$, $\kappa = 0.75$, $\kappa = 0.5$, $\kappa = 0.25$ and $\kappa = 0.0$. The latter has a cusp at the base, which is not beneficial for surface fairness and therefore not of much interest to us.

4.3 Typical Shift Vector

A selection field with a singular base can very well be based on a data point on the surface. It will be natural to take the surface normal at that point as the typical shift vector, scaled up or down if necessary.

For selection fields that are based on a curve, a powerful modelling tool results if the typical shift vector can be varied along the curve. Put simply, the shift vector can be defined as the difference between the base curve, say $\mathbf{c}(t)$, and an other curve, say $\hat{\mathbf{c}}(\hat{t})$. If $\mathbf{c}(t)$ is a curve on the surface prior to the shape modification, then $\hat{\mathbf{c}}(\hat{t})$ is exactly what the model will look like at this location, after the modification. Thus, designers will be able to manipulate feature curves, or even completely redesign them, while they will be able to control how the other curves (and thus the surface) in their vicinity adapts to the changes with the parameters r and κ . In addition, they will be able to protect other feature curves during the modification, by basing a deselection field on them.

To make this principle work as expected, it needs to be refined. As $\hat{\mathbf{c}}(\hat{t})$ may be completely different from $\mathbf{c}(t)$, their parameterisation may be different. In other words, when two particles are considered, one travelling down each curve at proportional increments of t and \hat{t} , the variation in velocity of the two particles may not be parallel. The effect on the typical shift vector will be that it changes direction more often than necessary. We remedy this by evaluating the curves with respect to arc length. In addition the designer may not want to change the complete curve, and $\hat{\mathbf{c}}(\hat{t})$ may partly coincide with $\mathbf{c}(t)$. But due to the different lengths of $\mathbf{c}(t)$ and $\hat{\mathbf{c}}(\hat{t})$, the typical shift vector may still have non-zero length in these parts, which is not intended. To counter this, we must evaluate the curves only over the curve sections that actually have different geometries.

Let the curves $\mathbf{c}(t)$ and $\hat{\mathbf{c}}(\hat{t})$ differ from each other for $t \in [t_a, t_b]$ and $\hat{t} \in [\hat{t}_c, \hat{t}_d]$, with $t_{\text{begin}} \leq t_a < t_b \leq t_{\text{end}}$ and $\hat{t}_{\text{begin}} \leq \hat{t}_c < \hat{t}_d \leq \hat{t}_{\text{end}}$. For other parameter values, the curves coincide, although not necessarily for equal parameter values. The exact value of t_a , t_b , \hat{t}_c and \hat{t}_d can be determined by analysis of the control points and knot vectors of the curves. For a formal definition of the shift vector, we need a mapping $m: t \mapsto \hat{t}$. For a certain parameter value t_i , $m(t_i)$ must produce a \hat{t}_j so that the arc lengths of the curve sections on either side of these parameter values are proportional:

$$\frac{\int_{t_a}^{t_i} |\dot{\mathbf{c}}(t)| dt}{\int_{t_i}^{t_b} |\dot{\mathbf{c}}(t)| dt} = \frac{\int_{\hat{t}_c}^{\hat{t}_j} |\dot{\hat{\mathbf{c}}}(\hat{t})| d\hat{t}}{\int_{\hat{t}_j}^{\hat{t}_d} |\dot{\hat{\mathbf{c}}}(\hat{t})| d\hat{t}} \quad (4)$$

in which the arc length of a curve is defined as the integral of the length of the first derivative of that curve with respect to the curve parameter. Because m is inefficient to be evaluated directly, one should first assess whether arc length evaluation is at all worthwhile, by comparing internal knot spacings and control point distances of $\mathbf{c}(t)$ and $\hat{\mathbf{c}}(\hat{t})$, looking for large discrepancies. If so, the map m may be approximated by evaluating $m(t)$ at distinct values of t_i , and fitting a polynomial, say $\hat{m}(t)$, through the mapped values \hat{t}_j .

Now we are able to define the typical shift vector \mathbf{S} as the difference between corresponding positions on the two curves according to arc length, expressed as a function of the curve parameter t :

$$\mathbf{S}(t) \equiv \hat{\mathbf{c}}(\hat{m}(t)) - \mathbf{c}(t) \quad (5)$$

Fig.4 illustrates the necessity of going through the trouble of considering only curve sections of dissimilar geometry and computing arc lengths.




Fig.4: Upper left: original stem curve $\mathbf{c}(t)$ and re-designed stem curve $\hat{\mathbf{c}}(\hat{t})$. Middle: straight-forward definition of typical shift vector $\mathbf{S}(t) \equiv \hat{\mathbf{c}}(\hat{t}) - \mathbf{c}(t)$. Data points are shifted even where geometry is not changed, and their spacing becomes unbalanced. Lower right: considering arc lengths over sections of changed geometry, $\mathbf{S}(t) \equiv \hat{\mathbf{c}}(\hat{m}(t)) - \mathbf{c}(t)$, gives a much cleaner result.

4.4 Unintended Selections

The proposed method for shape modification is simple, as we do not regard the surface of the model at all, and only consider data points and their shortest distance to selection bases. The advantage is speed. Shift vectors can be computed quickly enough to visualise them in real time, while the designer manipulates the modification parameters. They give a sufficient indication of how the shape will be modified once the parameters are accepted. Therefore, the presented method for shape modification is highly interactive.

However, in some situations this approach can be too simple. For instance, when modifying an area on the upper side of a thin wing. Because the data points on the lower side are close to the upper side, they may be selected unintentionally. Even though this may be prevented by careful definition of deselection fields, there is an alternative that can be automated, which involves putting the B-rep to good use.

Let us define the ‘root node’ of a selection field as the node that references the data point closest to the base of the selection field. If there are several nodes that qualify, any one of these will do. The algorithm in Fig.5 will only shift data points that form a contiguous selection that is rooted at the base, and prune away isolated sets of selected data points that are separated from the main selection by more than one surface patch.

```

Let  $\mathcal{A}$  be an empty set, capable of containing node references
Mark all nodes and faces in the B-rep as unconsidered
Mark the root node as considered and add it to  $\mathcal{A}$ 
While  $\mathcal{A}$  is not empty {
    Compute  $\mathbf{s}_i$  for a node  $i \in \mathcal{A}$ 
    If  $|\mathbf{s}_i| > 0$ , then {
        For all faces  $j$  that are adjacent to node  $i$  and that are still unconsidered {
            Mark  $j$  as considered
            For all nodes  $k$  that are adjacent to  $j$  and that are still unconsidered {
                Mark  $k$  as considered and add it to  $\mathcal{A}$ 
            }
        }
        shift node  $i$ 
    }
    remove  $i$  from  $\mathcal{A}$ 
}

```

Fig.5: Pseudo-code of an algorithm that only shifts contiguous sets of data points.

5 Finishing Up

Once the fitting/fairing algorithm has re-interpolated the curves over the shifted data points, the H-rep has become a consistent and smooth modification from the original, by which we have succeeded in our objective. However, if the original primarily consisted of planar curves, such as is customary in the design of ship hulls, these may no longer be planar after the modification.

Planar curves can be restored by intersecting the modified model with the planes in which the curves were originally defined, and adding the intersection curves to the model. These new curves take over the definition of the modified shape, by which the old curves become redundant and may be removed.

6 Applications

6.1 Example of Shape Variation

Fig.6 shows how the fore ship of a frigate has been made slightly narrower. The selection base was a point on the hull, and the shift was in transverse direction. The system has no difficulties with the knuckle line that is present in this region. A critique on this particular shape variation may be that the shell near the stem contains too much of the original shape, resulting in an extra inflection. This is due to the shape of the deselection function that was chosen to constrain the shell to the plane of symmetry, as the value of κ in Eq.(1) was set to 1.0, Fig.3. A value of 0 would have been better in this regard. Optimal would have been not to base the deselection field on the plane of symmetry, but on the contour line itself. Then κ could have been varied along it, 0 where the stem is sharp and non-zero elsewhere.




Fig.6: Modification of the hull of a frigate (shaded surface and light grey wire-frame) together with the original shape (black wire-frame overlay).




Fig.7: Three successive curve-based selections turn a plain bow (white wire-frame) into a bulbous bow.

6.2 Example of Bulb Design

Fig.7 shows shape variations resulting from a curve-based selection, illustrating that the method provides a powerful modelling tool. It shows all three steps in the process of designing a bulb from scratch. Starting off with the same model as in Fig.6, the first step was to re-design the stem curve, Fig.7a. An auxiliary waterline was added ending in the fore-most position of the bulb, just below the second ordinary waterline counted from below. On this line the second selection was based, giving the bulb more body, Fig.7b. Finally, the lowest waterline was dragged slightly outward, to improve the shape of the frames in the lower region, Fig.7c. The dent at the top of the bulb is due to scarce geometric data. This was corrected with two extra frames and one extra waterline, Fig.8.




Fig.8: Extra defining data added to correct dent in upper part of bulb.

6.3 Performance

Shape variations with selections based on points and planes are fast. However, curve-based selections can be time consuming, because computing the global closest distance of a point to a curve is expensive. In the current implementation, which is not for production and serves proof of concept only, all data points in the entire model are processed. Depending on the length and complexity of the curve, the time needed for distance computations in the examples of Fig.7 took up to several minutes on a 1.5 GHz PC. Once the distances are computed, shift vectors can be previewed interactively while r and κ are varied.

This performance can be improved. Firstly, it may be possible to omit several iterations in the closest point finding algorithm per data point, if the loop termination is not based on the accuracy of the closest point, but on the accuracy of the resulting shift vector. Secondly, it may be possible to disregard data points on curves that were never manipulated since they were added to the model. These curves do not actually contribute to the definition of the shape, and serve only for surface

visualisation, surface quality interrogation and/or manufacturing (e.g. frame contours and the butts and seams of shell plating). In late stages of the design, there may be many of these curves. Unless the resulting shape has so much detail that the extra curves actually contribute to the definition of the new geometry, they can safely be deleted before the shape variation and restored afterwards. A heuristic based on the typical shift vector can determine this possibility. Finally, the shape variations that are discussed here rarely involve every single data point in the model; in case it does, an ordinary affine transformation probably performs better. So computing the distance for all data points is a waste of time. It would be much better to consider data points on demand, based on the extent of the selection. For this a graph search is required, similar in nature to the algorithm discussed in Ch.4.4.

With the above in mind, it is advisable for an implementation of the H-rep concept that when new curves are interpolated, their addition to the B-rep data structure be deferred until they are selected by the designer for explicit manipulation, which is the point at which they start taking part in actual shape definition. Curves that are not in the data structure should then be regenerated automatically whenever any of their underlying patches change.

7 Conclusion

A simple method for the modification of a network of intersecting curves was presented, which preserves the consistency of the network, the fairness of the surface and local surface features. Curves may be redesigned explicitly, regardless of the detail in a design, with intuitive control over how the surface in their vicinity adapts to the changes. This can be regarded as an advantage over the competing method of free-form deformation (FFD).

References

- FORSEY, D.R.; BARTELS, R.H. (1988), *Hierarchical B-spline refinement*, 15th Annual Conf. Computer Graphics and Interactive Techniques, ACM Press, pp.205-212
- GONZALEZ-OCCHOA, C.; PETERS, J. (1999), *Localized-hierarchy surface splines (LeSS)*, Symp. Interactive 3D graphics, Atlanta, pp.7-15
- JENSEN, T.W.; PETERSEN, C.S.; WATKINS, M.A. (1991), *Practical curves and surfaces for a geometric modeler*, Computer Aided Geometric Design 8/5, pp.357-369
- KOELMAN, H.J. (1999), *Computer support for design, engineering and prototyping of the shape of ship hulls*, TU Delft, ISBN 90-901-2888-3
- KOELMAN, H.J. (2003a), *Exploring the H-rep ship hull modelling concept*, 2nd Conf. Comp. and IT Appl. in the Maritime Ind., COMPIT, Hamburg, pp.140-153, <http://www.sarc.nl/compit03.pdf>
- KOELMAN, H.J. (2003b), *Application of the H-rep ship hull modelling concept*, Ship Techn. Res. 50, pp.172-181
- KOELMAN, H.J.; HORVÁTH, I.; AALBERS, A. (2001), *Hybrid representation of the shape of ship hulls*, Int. Shipb. Progr. 48, pp.247-269
- LEON, J.C.; TROMPETTE, P. (1995), *A new approach towards free-form surfaces control*, Computer Aided Geometric Design 12/4, pp.395-416
- MICHELSSEN, J. (1995), *A free-form geometric modelling approach with ship design applications*, PhD thesis, NAOE, TU Denmark, Lyngby
- SEDERBERG, T.; PARRY, S. (1986), *Free-form deformation of solid geometric models*, ACM SIGGRAPH 20/4, pp.151-160
- TERZOPOULOS, D.; QIN, H. (1994), *Dynamic NURBS with geometric constraints for interactive sculpting*, ACM Trans. on Graphics 13/2, pp.103-136, <http://mr1.nyu.edu/~dt/papers/tog94/tog94.pdf>
- VAN DIJK, C.G.C. (1994), *Interactive modeling of transfinite surfaces with sketched design curves*, PhD thesis, TU Delft
- VEELO, B.N. (2004), *Shape modification of hull models in H-rep*, 3rd Conf. Comp. and IT Appl. in the Maritime Ind., COMPIT, (Ed. Bertram), Siguënza, pp.212-222