Accuracy in Sail Simulation: Wrinkling and growing fast sails.

Peter Heppel[1], peterheppel@compuserve.com, peter@peterheppel.com

Abstract. This paper describes some techniques for improving accuracy in the computation of structural membranes. Of importance in membrane computations is the modelling of the wrinkling of the surface that occurs to relieve compression. The paper describes an improvement to current theory by consideration of optimum-shift in the stiffness matrix derivation. The paper also describes the way that wrinkling interacts with the kinematics of the element grid. Results are presented demonstrating the effects of varying grids, with and without the modelling of wrinkling. The paper discusses error propagation in linked procedures that include membrane structural analysis and describes an automatic membrane design scheme using optimisation techniques.

NOMENCLATURE

_		Wrinkling angle
C	{cx,cy,cxy}	curvature
D	(3x3) symmetric	material stress-strain matrix.
Dw	(3X3)	wrinkled stiffness matrix
E	(ex,ey,gxy}	plane strain
em		material strain
et	{0, wt_2,0}	wrinkling strain at angle _
j	√(-1)	
J(_)		stress transformation matrix
	| Cos^2(_), Sin^2(_), 2 Cos(_)Sin(_) |	
	| Sin^2(_), Cos^2(_), -2 Cos(_)Sin(_) |	
	| -Cos(_)Sin(_), Cos(_)Sin(_),Cos^2(_)-Sin^2(_)|	
k		change in Gaussian curvature
P	scalar	pressure
s	{sx, sy, txy}	plane stress
s_	{s_1,0,0}	wrinkled stress, _ axes
sw	D em	Wrinkled stress, material axes

1. INTRODUCTION

This paper describes some techniques for the accurate structural modelling of membranes, which are particularly relevant to aeroelastic studies. Aeroelastic modelling is done using coupled structural and aerodynamic computations. This can be done by combining the stiffness, damping, mass and loading from both structure and aerodynamics into one system of equations. More usually, it is done by iterating between separate steady-state structural and aerodynamic codes, using heuristic means of ensuring the coupled system's numerical stability.

The paper concentrates on the structural computation, and on issues affecting the accuracy of the deflected form passed to the aerodynamics code.

2. MEMBRANES AS STRUCTURES

A shell structure is usually defined as a surface structure that supports tensile and compressive stress only and in which bending moment is negligible. Such structural action is termed 'membrane action'. A shell carries normal load by the product of curvature and membrane stress, related through the *membrane equation:*

$$P + C.\, s = 0$$

C.R. Calladine [Ref 3] provides a particularly clear account of shell behaviour and theory.

Although bending stiffness plays little part in the shell's global structural action, it determines its ability to resist buckling under compressive stress. The buckling strength of the shell is a complex function of the shape, stress state and boundary conditions. The buckling mode may be global or local. When the buckling strength is exceeded locally, the shell deforms locally. If the stress is tensile in one direction and compressive in the other, the form of the deformation is a pattern of waves with fronts approximately parallel to the direction of upper principle stress. The wavelength, amplitude and direction depend on both membrane and bending stiffness as well as the shell curvature and the local stress.

In the limiting case where the thickness tends to zero, a shell is termed a membrane. In this idealization, the material can only resist tension. Many classes of extremely thin-shell structures may be considered as membranes. Examples are textile and thin laminate structures such as sails, as well as natural structures such as cell walls and heart valves.

An unloaded object with bending stiffness (and without hinges) has an invariant three-dimensional shape. Any variation of this shape requires strain in the material. The object has no kinematic degrees of freedom apart from its rigid-body modes.

An unloaded membrane does not have an invariant three-dimensional shape, but it does have a corresponding invariant, its distribution of Gaussian curvature. Gaussian curvature is related to membrane strain by:

$$-k = \frac{\partial^2 ex}{\partial y^2} + \frac{\partial^2 ey}{\partial x^2} - \frac{\partial^2 gxy}{\partial x\, \partial y}$$

A membrane has an infinite number of kinematic degrees of freedom, representing, in addition to its rigid-body modes, the set of deformations with corresponding zero

1 Independent engineer

change in strain. These are termed the *kinematic modes*. An *inextensional deformation* is defined as one in which the change in membrane strain is at least second-order with respect to displacement. Thus, the kinematic modes are inextensional deformations.

All other deformations imply a first-order change of strain. They are termed *extensional deformations*.

Note that there exist extensional deformations that do not cause a change of Gaussian curvature. Because the change of Gaussian curvature depends only on the second spatial derivative of strain, an extensional deformation consisting of a linear distribution of strain does not change Gaussian curvature.

Extensional and inextensional deformations have different stiffness characteristics.

An extensional deformation implies a change in stress. It has an associated stiffness that depends only on the shape and material properties and is effectively constant for small deformations. This is termed *direct stiffness*.

Inextensional deformations imply a change of curvature but only high-order change of stress. If the membrane is unstressed, inextensional deformation does not affect equilibrium. If the membrane is already stressed, the membrane equation shows that its equilibrium alters proportionally to the change in curvature. Thus an inextensional deformation has an associated stiffness that depends only on the shape and state of stress. The stiffness is constant for small deformations. It is termed *geometric stiffness*.

A membrane structure of elastic material may deform extensionally or inextensionally, depending on the patterns of loading and constraints. If the deformation is purely extensional OR purely inextensional, the response is linear.

In the general case the deformation is a combination of extensional and inextensional and the response is non-linear. The extensional deformations change the stress, so modifying the stiffness of some inextensional modes. The inextensional deformations change the shape, so modifying the stiffness of some extensional modes.

Thus, any calculation procedure for solving general membrane structures must be non-linear. It should take into account any geometric stiffness due to prestress as well as the change of geometric stiffness due to extensional deformations and the change of direct stiffness due to the inextensional deformations.

3. WRINKLING OF MEMBRANES

3.1 Introduction

We have seen that a membrane wrinkles when subject to uniaxial compression and slackens when subject to biaxial compression. These effects represent an important non-linearity that is not accounted in the convention linear membrane stress-strain formulation. In the case of linear isotropic stress-strain relations, the principle stress and principle strain directions coincide and this simplifies the formulations for wrinkled stress and material stiffness. For anisotropic materials the formulation is more complex. Wrinkling models have been extensively studied. The following development parallels the development due to Lu, Accorsi and Leonard [Ref 4] with a different treatment of some components of the wrinkled stiffness matrix.

3.2 Application to sails

Most sails have such low bending stiffness that wrinkling relieves all significant compressive stress. Notable exceptions are small, stiff sails such as competition dinghy foresails (when new) and sailboard sails.

It is important to model wrinkling in finite element sail analysis to increase the grid kinematic freedom (as will be shown in the next section) and also because racing sails are usually designed for selective wrinkling.

It is a frequent design goal to make a racing sail cover a wide wind range. The optimum sail lift coefficient tends to reduce with increasing windspeed. As sails have sharp leading edges, to change the lift coefficient it is desirable to change the ideal lift coefficient, which implies changing the camber.

It is usual to shape a racing sail to provide Gaussian curvature appropriate to the required camber at the upper wind limit. To increase the Gaussian curvature at lower windspeeds, the sail might be engineered to generate the appropriate (large) strain distribution, but this tends to limit the sail's life. It is usually done by allowing wrinkling perpendicular to the luff.

3.3 Theory

Engineering notation for stress and strain is used throughout.

Consider a patch of membrane subject to a fixed strain e, with linear orthotropic stress-strain relation

$$s = D\,e$$

where s and e are the engineering stress and strain in material axes and D is the symmetrical stiffness matrix. The membrane may be unwrinkled in a state of biaxial tension, singly wrinkled in a state of uniaxial tension, or slack. The membrane will be in biaxial tension if both the principle strain values of e are non-negative. In this case the stiffness matrix is unchanged and the stress is D e.

In the membrane idealization, any wrinkles have infinitesimal wavelength. Their generators are parallel to the upper principle stress. Thus the real deformed shape of the membrane is undefined but the wrinkled surface may be mapped onto a smooth pseudo-surface. The difference in membrane strain between these two surfaces is termed the *wrinkling strain*.

$$e_w(\_) = J^T(\_)\, e_t$$

where J(_) is the stress transformation matrix, following Tsai [Ref 5] and et is the wrinkling strain transformed to angle _.

The patch is subject to a fixed strain e. If it wrinkles or slackens, e remains the strain in the pseudo-surface, while the material strain em changes.

If wrinkling occurs at some angle _ to the material axis, then:

Direct strain in direction _ is unchanged.

Shear strain in direction _ is unchanged.
Direct strain in direction _+p/2 is increased
Direct stress in direction _ is finite and +ve
Shear stress in direction _ is zero
Direct stress in direction _+p/2 is zero.

Thus

$$em = e + J^T(\_)et$$

where

$$et = \{0, wt_2, 0\}$$

and

$$sw = D\ em = J^{-1}(\_)\ s\_$$

where

$$s\_ = \{s\_1, 0, 0\}$$

Rearranging and substituting {sxxo, syyo, sxyo} for D e we obtain the condition on wt_2 for s_2 to be zero.

$$wt2(\_) = ((syyo\ Cos(\_)^2 - 2\ sxyo\ Cos(\_)\ Sin(\_) + sxxo\ Sin(\_)^2)/ (d22\ Cos(\_)^4 - 4\ d23\ Cos(\_)^3\ Sin(\_) + 2\ d12\ Cos(\_)^2\ Sin(\_)^2 + 4\ d33\ Cos(\_)^2\ Sin(\_)^2 - 4\ d13\ Cos(\_)\ Sin(\_)^3 + d11\ Sin(\_)^4))$$

where $wt_2(\_)$ is the (scalar) wrinkling strain.
The energy per unit area is

$$E(\_) = [\ D.\ em\ ]\ .\ em$$

and the wrinkling angle _ will be found at a local minimum of E(_) .
$wt_2(\_)$ has four roots :

-j Log(-√ ((-(syyo +sxxo) - 2 √(syyo sxxo - $sxyo^2$))/ (syyo - sxxo + 2 **j** sxyo)))
-j Log(√ ((-(syyo +sxxo) - 2 √(syyo sxxo - $sxyo^2$)) /(syyo - sxxo + 2 **j** sxyo)))
-j Log(-√ ((-(syyo +sxxo) + 2 √(syyo sxxo - $sxyo^2$)) /(syyo - sxxo + 2 **j** sxyo)))
-j Log(√ ((-(syyo +sxxo) + 2 √(syyo sxxo - $sxyo^2$)) /(syyo - sxxo + 2 **j** sxyo)))

where **j** = √(-1)
$wt_2(\_)$ may have 0 or 4 real roots .
If $wt_2(\_)$ has 4 real roots, the membrane is wrinkled and the wrinkling angle _ will be the unique minimum values of E(_) found in the domains where $wt_2(\_)$ is positive. Because in E(_) there are no terms in _, but only terms in 2_ and higher, these stationary values are p apart. If $wt_2(\_)$ has no real roots and is everywhere positive, the membrane is slack. If $wt_2(\_)$ has no real roots and is everywhere negative, the membrane may be wrinkled or taut. It is wrinkled if at least one of the principle strain values of e is negative. In this case E(_) has just two minima in {0,2p} separated by p.
Because E(_) is a 6^{th} –order equation, it is necessary to find its stationery values numerically, using the above to isolate the relevant domain. Then sw may be determined.
The wrinkled stiffness matrix Dw may be found from

$$Dw = \partial sw/\partial e$$

on the assumption that the wrinkling state (taut, wrinkled or slack) is constant, and that the wrinkling angle _ varies with e so that E(_ + d_) remains a minimum.

$$\partial sw/\partial e = D\ [I + (\partial J^T/\partial\_\ et\ \partial\_/\partial e) + J^T\ (\partial\_/\partial e\ \partial et/\partial\_ + \partial et/\partial e)]$$

The terms $\partial JT/\partial\_$, $\partial et/\partial\_$ and $\partial et/\partial e$ may be found by direct differentiation. The term $\partial\_/\partial e$ is found by consideration of optimum-shift. Writing E (_) as E (_,e) and expanding in the region of _ gives:

$$E(\_+d\_,e) = E(\_,e) + d\_\ E^-(\_,e) + \_\ d\_^2\ E^{--}(\_,e) + O(d\_^3)$$

Differentiating with respect to d_ and setting to zero gives.

$$E^-(\_+d\_,e) = 0 = E^-(\_,e) + d\_\ E^{--}(\_,e) + O(d\_^2)$$

or

$$d\_ = -\frac{E^-(\_,e)}{E^{--}(\_,e)}$$

Differentiating with respect to each of the components e_i of e in turn, gives:

$$\frac{\partial\_}{\partial e_i} = -\frac{E^{-e^i}(\_,e)}{E^{--}(\_,e)} + \frac{E^-(\_,e)\ E^{--e^i}(\_,e)}{E^{--}(\_,e)^2}$$

Noting that $E^-(\_,e) = 0$, this yields $\partial\_/\partial e$,

3.4 Implementing in the Finite Element scheme

The wrinkling method has been implemented in a Dynamic Relaxation [Ref 1] analysis scheme. Dynamic relaxation achieves full geometric non-linearity by the separation of the equations of equilibrium and compatibility. Implementation of the wrinkling method involves simply testing for wrinkling and updating each element's stiffness matrix during the residual force calculation. The wrinkling method (in a debug implementation) slows the analysis by a factor of about four, but requires no changes to the convergence parameter algorithms.

4. GRID KINEMATICS

Most standard works on the Finite Element Method discuss criteria for selecting grid size and pattern. See for instance Ref 7 . In the finite element method, within any one element a relatively simple displacement distribution is modelled. To model a rapid spatial

variation, the analyst chooses a grid size so that the elements provide an acceptable piecewise approximation to the actual displacement. This is a reasonable approach for resolving extensional deformation, but to resolve inextensional deformation there are some additional requirements on the grid. The finite element solution consists of (non-linearly) superposed extensional and inextensional modes of deflection. The accuracy with which the extensional modes are resolved depends on the accuracy with which the strain is resolved. The accuracy of resolution of the inextensional modes depends on the grid's kinematic degrees of freedom.

A grid has a finite number of kinematic degrees of freedom. In the case of a grid of three-node triangular elements, the number can be estimated [Ref 6] as:

$$3j - b - f$$

Here b is the number of element corners, j is the number of element edges and f is the number of constrained degrees of freedom. There are exceptions to this rule. An assembly may be redundant in some regions and under-restrained in others. There are other exceptions for certain axisymmetric cases.

The analysis can only resolve inextensional modes that are combinations of the grid's kinematic modes. The analysis will be insensitive to all other inextensional modes of deflection. Thus if the grid does not provide the appropriate kinematic modes, deflection will be under-estimated.

To choose a grid therefore requires judgement about which inextensional modes are significant as well as knowledge of the grid's kinematic modes. The shapes of the kinematic modes might be calculated, but in practice it is sufficient to provide fold lines that correspond to the direction of the lower principle change of curvature determined from a test analysis [Figure 1]. The deformed lower principle curvature is often a reasonable approximation.. Note that even though [Figure 1]. was generated from a 3500-element grid [Figure 2], the shape in the tack region is still poorly captured.

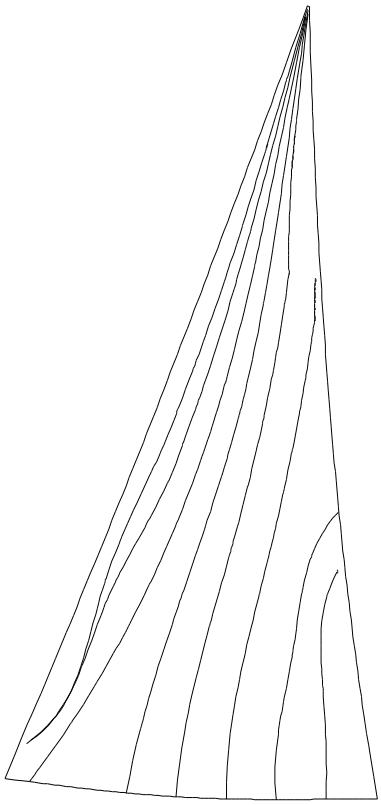

Figure 1 Lower principle curvature directions from figure 2

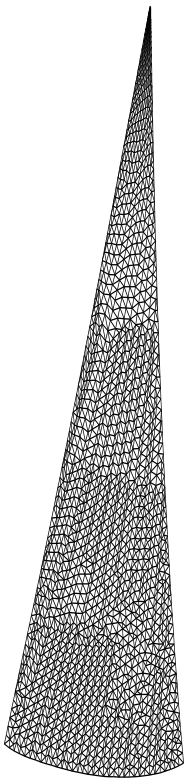

Figure 2 a 0.4m triangular grid on 10^o equiangular guides still has noticeable kinematic restrictions.

5. ELEMENT CHOICE

Three-node triangular elements are popular in membrane analysis because they provide the most kinematic freedom of common element types, and because they simplify the implementation of non-linear material stress-strain relations. A better element for membrane analysis would be a multi-node element that can freely deform into any conical shape. If this hypothetical element also modelled the relation between strain and Gaussian curvature instead of implicitly applying the Plane Strain Compatibility Equations, it could be assembled into a model free of kinematic restrictions.

6. KINEMATICS AND THE RELATION TO WRINKLING

In practice, the use of three-noded triangular wrinkling elements makes the task of choosing the grid pattern less critical. When a two adjacent elements wrinkle, one grid degree of freedom is added. This increases the probability (but does not guarantee) that the grid's kinematic modes can combine to represent the desired extensional modes. However, if such an additional degree of freedom is to cause a large change in displacement, this implies a large wrinkling strain., This tends to create a large local change in Gaussian curvature. Although the overall shape may be more accurate, such regions of 'kinematic' wrinkling will become bumpy. For racing headsails this region often coincides with regions where the boundary layer is sensitive and in an aeroelastic analysis this can cause an artefact in the boundary layer separation model.

The closer the grid kinematic freedoms model the inextensional deformations, the smaller these 'kinematic' wrinkling strains. This can usually be achieved by choosing a grid to follow the direction of principle change of curvature from a previous analysis, and then surveying the result to check that its principle change of curvature is oriented with its grid. This process usually converges if repeated.

7. AN EXPERIMENT ON THE EFFECT OF GRID SIZE AND ORIENTATION ON A GENOA

To investigate the effects of grid orientation and wrinkling on sail analysis results, the same sail was modelled with seven different grids of approximately the same density. Each model was run with and without wrinkling. Finally, one of the models was run using a finer grid. The results were post-processed into the 'digitised perspective up-shot' form commonly used by sailboat racing teams.

The sail mould shape was derived from North Sails' proprietary design software as a NURBS surface. The shape represents a generic IACC headsail.

This is a tape-laid sail. Individual groups of fibres were represented in CAD as space curves. During mesh generation their elastic properties were interpolated onto the elements and combined using the theory of composite laminates.

The seven models have primary grid-lines which are equiangular spirals of –30, -20, -10,0,10,20 and 30 degrees to the radius [Figure 3] The spiral's centre is approximately at the intersection of the extensions of the luff and leech. The primary grid-lines divide the sail into zones. Each zone is gridded using Delauney triangulation with characteristic edge-length of 0.75m.

The models were prepared using Rhino™ based on Desman™ data and meshed in GAUSS, the Relax pre-processor.

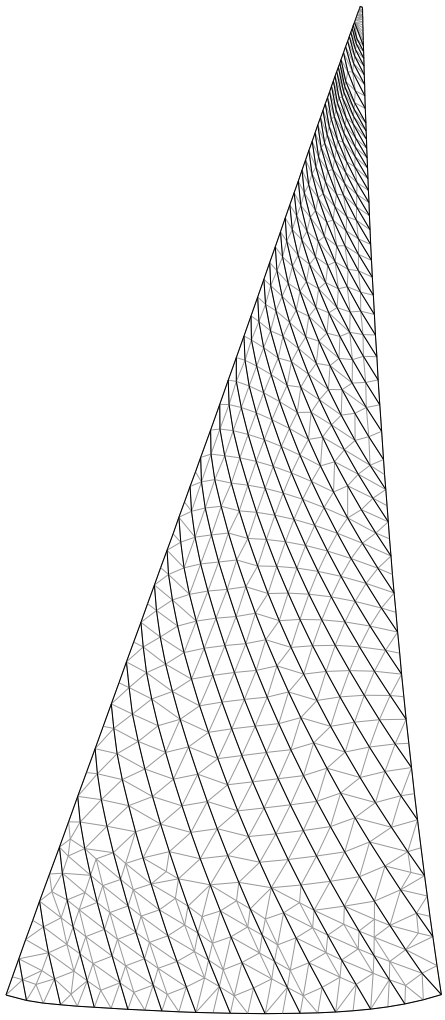

Figure 3 equiangular spiral grid fold lines with 0.75m grid

7.1 Analysis

The models were run using Relax/Pansail [Ref 8] to find the solution where aerodynamic and structural forces are in equilibrium. All models were run with the same wind,

boat attitude and trim. The headstay was set to a fixed length and so its tension increased with loading, but the sag was more constant than if it had been set to constant tension.

7.2 Results

See:
Figure 5 Grid - 30 wrinkled
Figure 6 Grid 0 wrinkled
Figure 7 Fine grid Wrinkled
Figure 8 Grid -30 not wrinkled
Figure 9 Grid 0 Not Wrinkled
Figure 10 Fine grid wrinkled uprange
Table 1 Depth as function of grid angle
Table 2 Draft position as function of grid angle
Table 3 Twist as function of grid angle
Table 4 Entry Angle as function of grid angle

Tables 1,2,3 and 4 show the variation of upshot results with grid angle. The stripes are numbered from the bottom, with stripe 4 being at 87.5% height. Each table shows the results for the unwrinkled and wrinkled models, with the fine-grid wrinkled model for comparison
The most striking aspect of the results is the variation. Table 1 shows a spread of 4% on upper depth although wrinkling improves this a little.
The wrinkled models show less variation and are closer to the fine model, especially for grid –10 and 0. The unwrinkled models are on average about _% flat, but the mid stripe is some 1 _% flat.
Tables 2 and 3 show a wide scatter on draft position and entry angle. The variation on the mid-stripe varies smoothly with grid angle between 25% and 35% chord.
Table 4 shows some consistency of twist, except for the $+20^{o}$ grid.

Figures 5,6,8 and 9 show the -30^{o} and 0^{o} grids with and without wrinkling. The non-wrinkled models show more variation and are noticeably flatter, which is consistent with the tables.
Figs 5,6,7 are the three wrinkled cases. They show less variation that figs 7 and 8.
Figure 7 is probably the most accurate result.
Figure 10 is included for comparison. The sail has been set up in a flatter trim.

7.3 Experiment Conclusions.

Omitting wrinkling made the genoa appear _ - 1 _% too flat. The –10 and 0 degree results are closest to the 'fine' result, which is consistent with their grid orientation being closest to Figure 1 .
All the results, even the 'fine' ones, show errors due to grid kinematics.

8. AN OPTIMISATION SCHEME

8.1 Membranes in optimisation procedures

Membrane computations are frequently linked to other procedures. Sail simulation requires a combined solution for both the structure and the aerodynamics. Membrane computations linked to optimisation schemes may be used to find equilibrium shapes for certain structures, especially those containing post-buckled components.
Accuracy is especially important in linked procedures because noise tends to propagate unnoticed. When procedures are linked into a numerical optimisation scheme, any noise is amplified through its effect on the derivatives of the goal with respect to the optimisation parameters.
Numerical optimisation schemes are frequently used in studies of optimal yacht performance. In such procedures, the goal value is often a relatively flat function of the optimisation parameters. The sail design is usually considered as fixed, with a simple, noise-free trim parameterisation that represents crew choices.
This approach to the VPP sail model uses heuristic functions that are smooth. This guarantees the absence of noise, but may obscure important detail. Another approach is to use results from numerical simulation or wind tunnel tests. Such results may better represent the sail performance, but unless they are carefully obtained, the detail is often hidden in the noise.
However, a sufficiently noise-free aeroelastic simulation would improve the accuracy of such procedures. If the sail model were to respond to design as well as trim parameters, such a procedure would further serve to optimise sail designs for performance.

8.2 Growing Sails

Once the sail simulation is reasonably free from artefacts, it can be used as part of the process of choosing fast sail shapes. It may be used as an aid to the manual design process, using analysis to inform decisions. Alternatively, it may be part of an automated optimisation scheme, seeking out favourable regions of the design space.
Such an optimisation scheme is a generalization of the standard VPP process, in which it is usual for the sail model to be parameterised using factors such as Reef, Flat, Sheeting angle and Twist. In the classic VPP method, these parameters represent the choices made by the crew: sail selection, running rigging settings, appendage settings, weight disposition and heading. They are the set of state variables solved by the performance optimiser.
For a sail-growing scheme, a second set of optimisation parameters is added. This set represents the choices made by the sailmaker: area, edge curves, unstrained Gaussian curvature distribution and construction. The overall scheme is represented in [Figure 4]

8.3 Components of the scheme

8.3 a) Variables

At its simplest, the scheme could be used to find maximal forward force subject to a given heeling moment. The minimum useful set of state parameters is:
Attitude and trim parameters

- Running rigging settings: Two per sail
- Sail design parameters
- A parameter varying the total subtended solid angle
- A parameter to scale the luff curve.

Goal Function

- Forward force

Constraints

- Heeling moment

For a more general search of the design space, the set of state parameters might be:
Attitude and trim parameters

- The hull attitude
- Appendage settings
- Rig settings, standing and running
- Moveable ballast configuration

Rig design parameters

The choice would depend on the design task, but a minimal set might be:

- Parameters varying each sails' total subtended solid angle
- Two parameters causing smooth variations of the luff curves: high and low.
- Two or three parameters varying the vertical and horizontal distributions of the sails' Gaussian curvature
- Two to four parameters varying the sails' material properties
- At least one parameter describing mast stiffness

Goal Function
The choice of the goal function depends on the design task.

- For steady-state upwind design, the goal function would be the weighted mean of a set of boatspeed values at different headings, true windspeeds and sea states. Using a single sailing point results in designs with narrow 'grooves'. Sea state can be modelled simply as an added drag term, but a closer approach to upwind performance would be to take the time-averaged mean net driving force given a cyclic hull motion and an analysis that includes unsteady effects.

Constraints

- Only as required for numerical stability

8.3 b) Procedure 1: Optimiser/goal seeker

The Optimiser/goal seeker searches for favourable regions of design parameter space by varying the state parameters. For successive sets of state parameters, it invokes the goal function to determine the corresponding performance.
A characteristic of aeroelastic simulation is that the goal function is floating-point intensive. A practical computing scheme should be designed to require the minimum number of goal-function samples.
For the task of refining a good existing design, an uphill technique may be used. The author has had reasonable results using the Microsoft Excel solver. This is flexible and reliable for single-peak problems, but somewhat slow for expensive goal functions because it forms the derivatives numerically. A polynomial-interpolation method such as Algy is more efficient because it determines the shape of the solution surface with the least number of goal-function samples.
For wider exploration of the design space, a genetic-algorithm method could be used. These methods have the drawbacks of requiring a large number of goal-function samples as well as of the granularity inherent in the integer description of the parameter space. However, given sufficient computing resources, this could lead to new design spaces. Once a promising design space has been identified, it would be more efficient to refine it using a polynomial-interpolation method.

8.3 c) Procedure 2: Database

The database logs each point in parameter space and provides communication between the components. Compared with most database applications, the amount of data is small and so performance is not critical. The authors have used the Relax internal database, but most customisable commercial packages would be suitable.

8.3 d) Procedure 3: Parametric modeller

The parametric modeller converts the Attitude and Trim parameters into the appropriate analysis actions such as rotations, constraints and forced displacements.
More importantly, it constructs analysis models from the sail design state parameters. The author uses GAUSS, the Relax pre-processor. Alternatively, a programmable parametric design system such as Pro/Engineer would be suitable.

8.3 e) Procedure 4: Structural analysis

The structural analysis calculates the displaced shape given the undeformed geometry, the loading and the constraints. The method should model geometric stiffness and non-linearity because membrane structures carry load by variation of curvature as much as by variation of stress. In some of the degrees of freedom the geometric stiffness dominates the direct stiffness. For the sail membranes, the geometric stiffness is positive but for the mast it is negative. For certain degrees of freedom, it may nearly cancel the direct terms.

Most racing sails use materials that are reasonably linear in their tensile stress-strain behaviour. Crimp interchange gives some small but complex interaction effects between the different layers of the laminate, but satisfactory results may be obtained using conventional laminates theory, based on tests where available. However the wrinkling behaviour is a significant material non-linearity, which should be modelled.
At higher windspeeds, the aerodynamic stiffness approaches and exceeds the structural stiffness. It is advantageous in terms of analysis time to include this in the structural solution. Otherwise negative feedback is needed to maintain numerical stability.
The Relax solver, based on M.R. Barnes' Dynamic Relaxation method [Ref 1], has full geometric non-linearity and certain material non-linearity, including wrinkling. An alternative would be to use a general-purpose non-linear code such as MSC.MARC [Ref 2], with the wrinkling behaviour provided by a user subroutine.

8.3 f) Procedure 5: Grid transformation

The aerodynamic and structural analyses use different element grids. The grid transformation process converts between them, generating the aerodynamic grid from the deflected structural grid, and generating the structural loading from the aerodynamic surface pressures.
Typically, the structural analysis uses a triangular grid of (o) 0.5m edge dimension. The grid relates to boundaries between material types and is refined in regions of greatest curvature. Typically, aerodynamic panel methods use a rectangular grid, with a larger spanwise than chordwise dimension. The grid is often refines at the leading edge. Navier/Stokes methods use a significantly finer grid.
The grid kinematics of the structural model often causes artefacts leading to local unfairness, although the use of a wrinkling algorithm reduces this. These artefacts cause local unfairness in the streamwise pressure distribution that can promote tripping of the boundary layer.
On the return path, a simple linear interpolation, sampling at the structural grid centroids, will often lose pressure transients at the leading edge because the structural grid is coarser than the aerodynamic grid. This can lead to errors in force integration. It is advisable to integrate the pressures on the structural grid over each element by determining its union with each aerodynamic element.

8.3 g) Procedure 6: Aerodynamic analysis

With currently available computing power, it would be impractical to run a Navier/Stokes method as part of an optimising scheme.
However, if the optimising scheme is using aerodynamic forces as goals or constraints, limitations in the aerodynamics method can bring magnified errors in the optimisation result. A particular limitation of panel methods is their assessment of the penalty of leading edge separation. It would be advisable to apply a high penalty to the leading edge singularity in order to drive the scheme towards solutions close to ideal incidence, even though these may be sub-optimal.
It would also be advisable to check any optimisation result obtained using panel methods by using a Navier/Stokes code for a refined search in the neighbourhood of the solution.

8.3 h) Procedure 7: Aeroelastic coupling.

For an undamped iteration between structural and aerodynamics codes to provide linear convergence, the aerodynamic stiffness would have to be negligible compared to the structural stiffness. This is a reasonable assumption for a sail at low windspeeds and increasingly invalid as the windspeed increases. Convergence may be improved by introducing negative feedback: loading the structural model with a weighted mean of the current and the previous aerodynamic results.
If the aerodynamic code were to provide the aerodynamic stiffness as well as loading, this iteration would approach second-order convergence.

8.3 i) Procedure 8: Post-processor

The function of the post-processor is to extract the relevant results from the two analyses and transfer them in the database. Relevant results include all the measures of aerodynamic forces, rig tensions, hull attitude and surveys of the sail shapes. It is useful to survey the sails using a method identical to those used for on-the-water records, representing the process of photographic 'up-shots' and digitising. This provides a direct check on the analysis results. Such a survey must model perspective in the same way as the physical camera.

9. CONCLUSIONS

The finite element grid in a membrane analysis should be chosen to ensure appropriate kinematic freedoms. This is best done via a trial analysis. An inappropriate grid will suppress certain inextensional deformations resulting in erroneous shape predictions.
The incorporation of a wrinkling model reduces the sensitivity of the analysis to the grid orientation and permits sail analysis that better approaches racing experience.
An analysis of an IACC Genoa without wrinkling will under-predict depth by _ - 1%.
Accuracy in sail structural analysis is especially important when the analysis is linked to other procedures such as aerodynamics codes, and critical if the aeroelastic analysis is used as the goal function in an optimisation procedure.

Acknowledgements
The author would like to thank Mickey Ickert of Oracle BMW Racing for providing the sample IACC sail designs.

References

Ref 1 Form and stress engineering of tension structures. Michael Barnes. Structural Engineering Review Vol **6** No 3-4 175-202

Ref 2 MSC.MARC reference manual. MSC Software corporation

Ref 3 C.R. Calladine. Theory of Shell Structures CUP . ISBN 0 521 36945 2

Ref 4 Finite element analysis of membrane wrinkling K. Lu, M. Accorsi and J. Leonard. Int. J. Numer. Meth. Engng. 2001; **50**:1017-1038

Ref 5 Composites Design. Stephen Tsai. ISBN 0-9618090-2-7

Ref 6 Maxwell, J.C. (1864) On the calculation of the equilibrium and stiffness of frames. Philosophical Magazine, **27**, 294-99.

Ref 7 Zienkiewicz & Taylor. The Finite Element Method. Fifth edition. Butterworth-Heinemann. ISBN 0 7506 5049 4

Ref 8 S.P. Fiddes, J.H. Gordon. A new vortex lattice method for calculating the flow past yacht sails. J. Wind. Eng. Ind. Aerodyn **6 3** (1996) 35-39

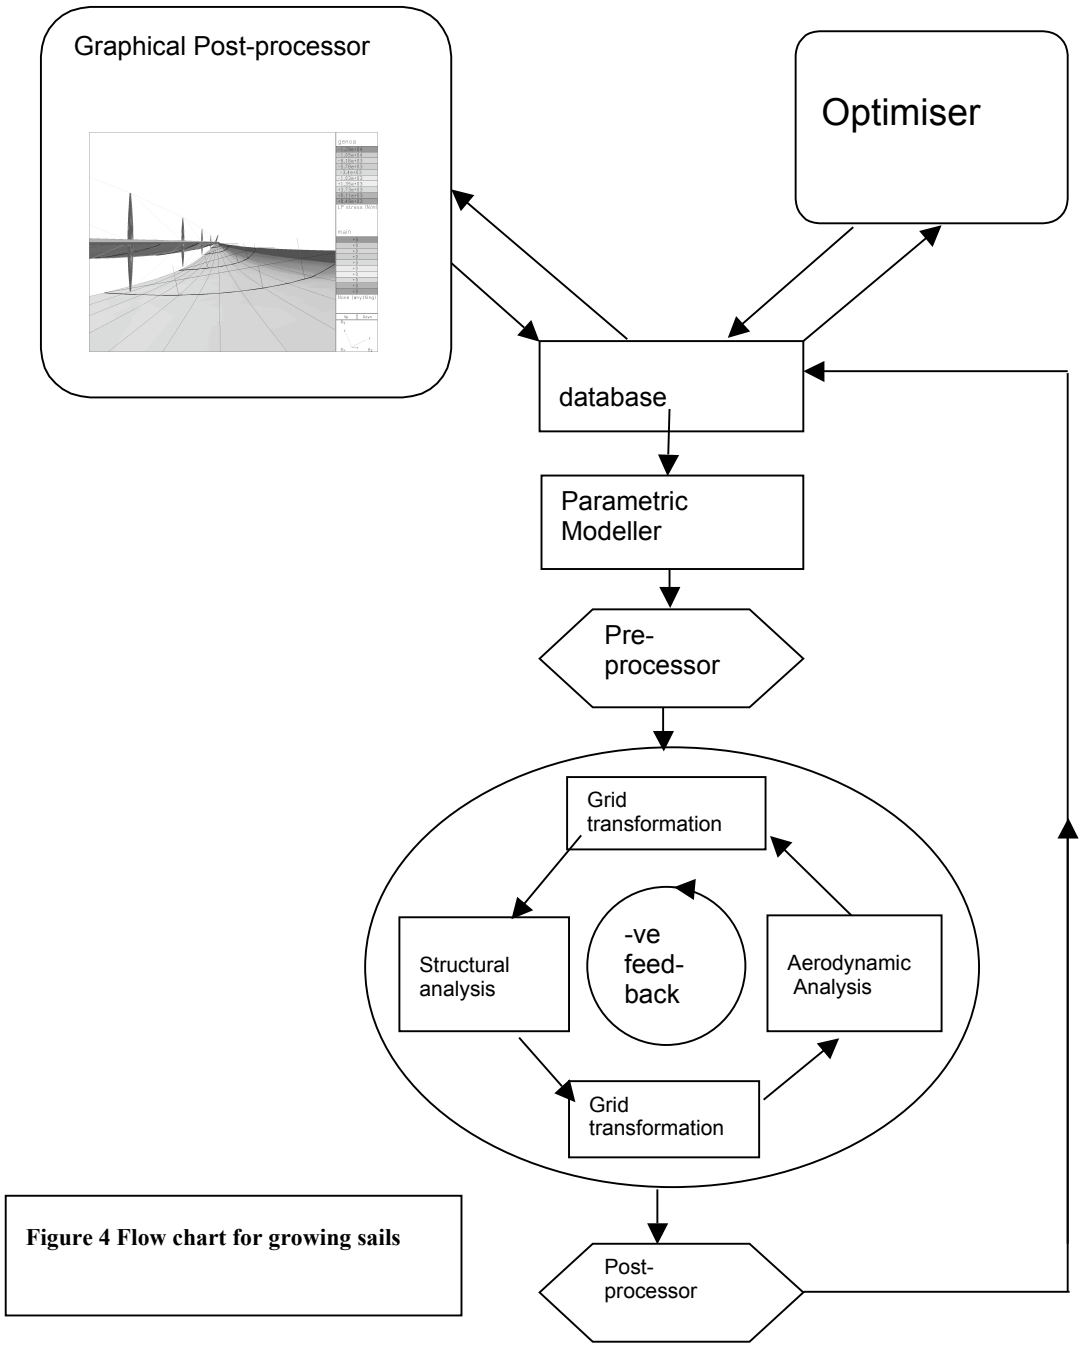

Figure 4 Flow chart for growing sails

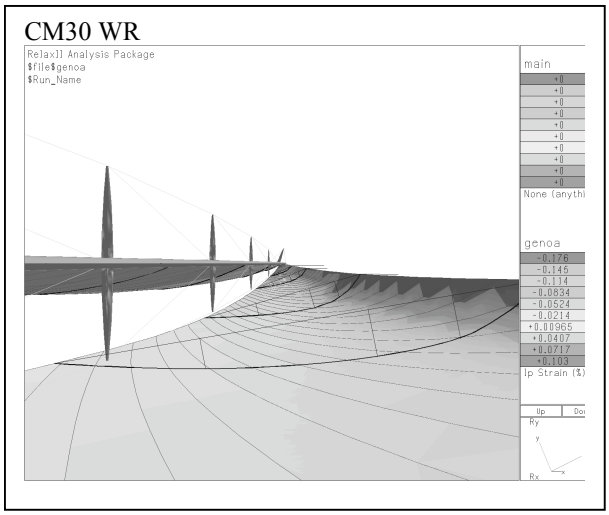

Figure 5 Grid - 30 wrinkled

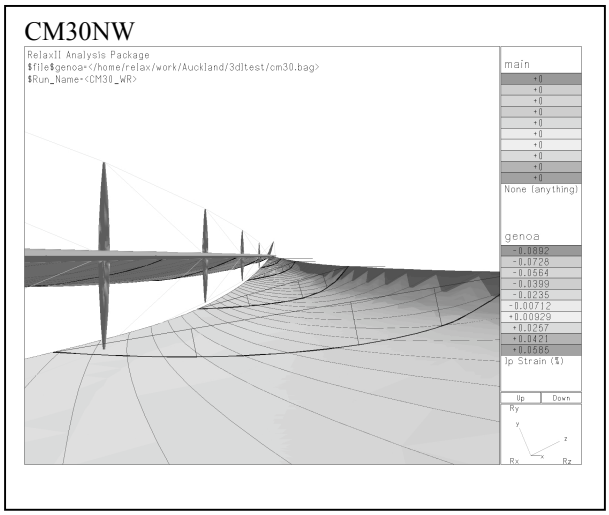

Figure 8 Grid -30 not wrinkled

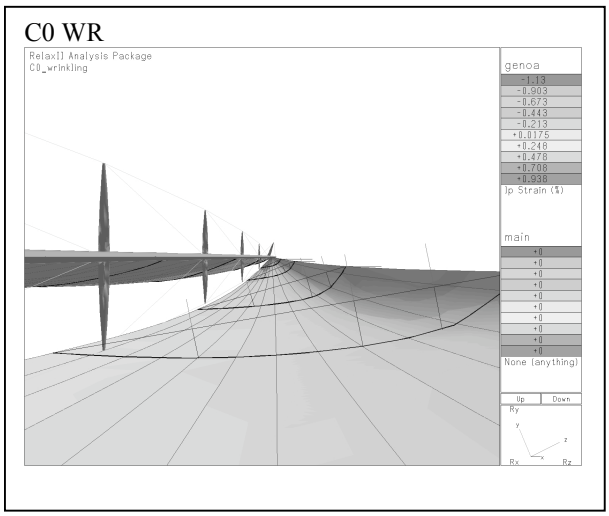

Figure 6 Grid 0 wrinkled

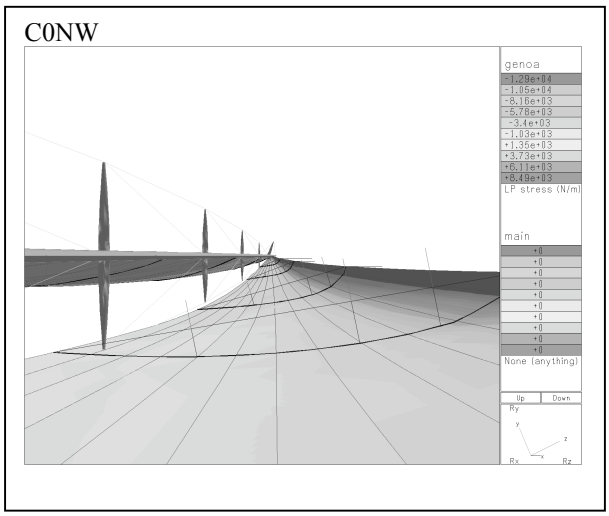

Figure 9 Grid 0 Not Wrinkled

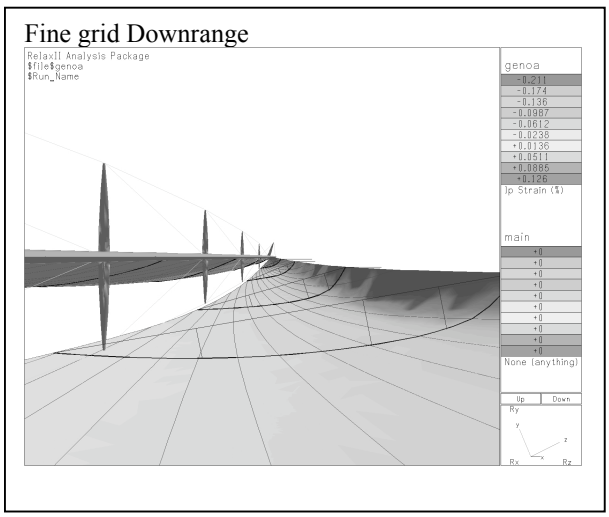

Figure 7 Fine grid Wrinkled

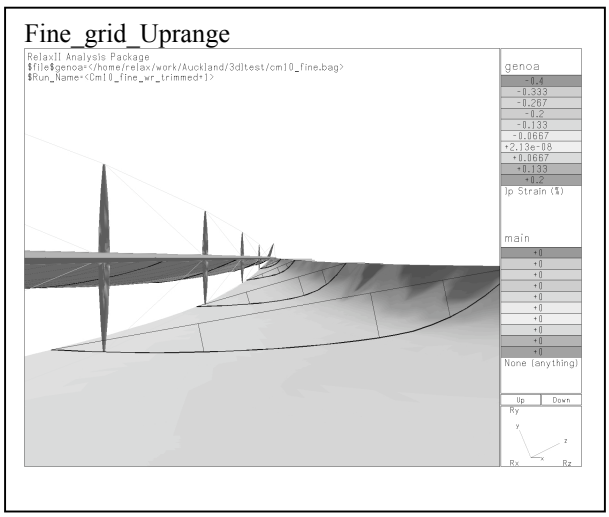

Figure 10 Fine grid wrinkled uprange

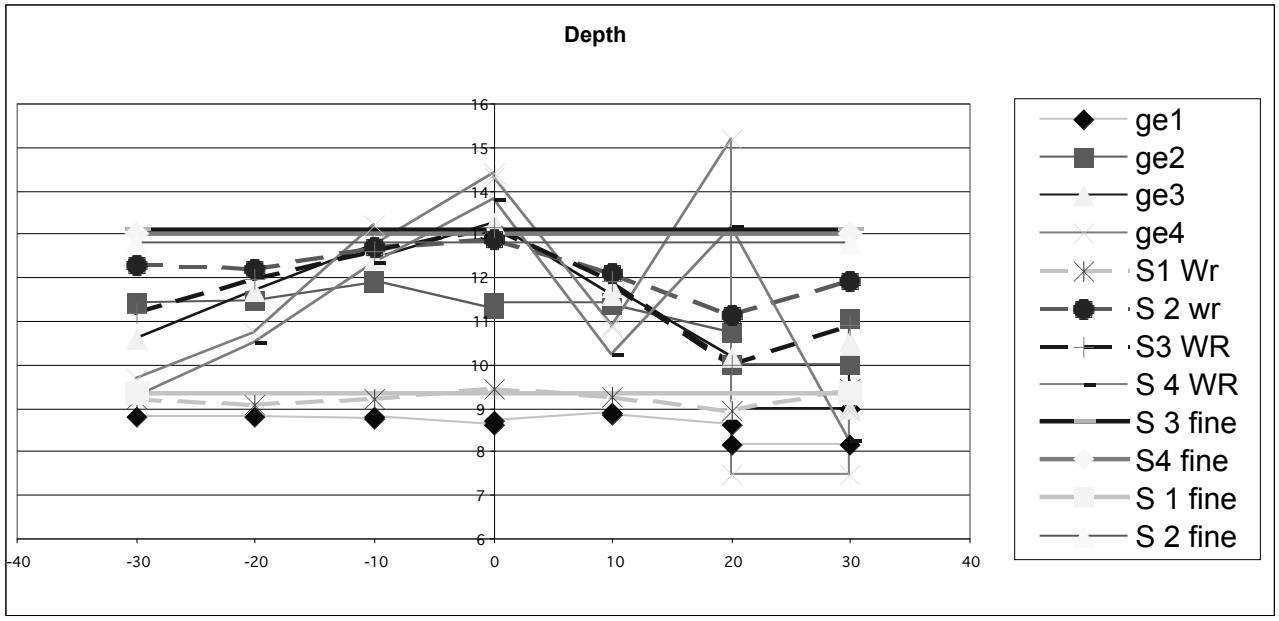

Table 1 Depth as function of grid angle

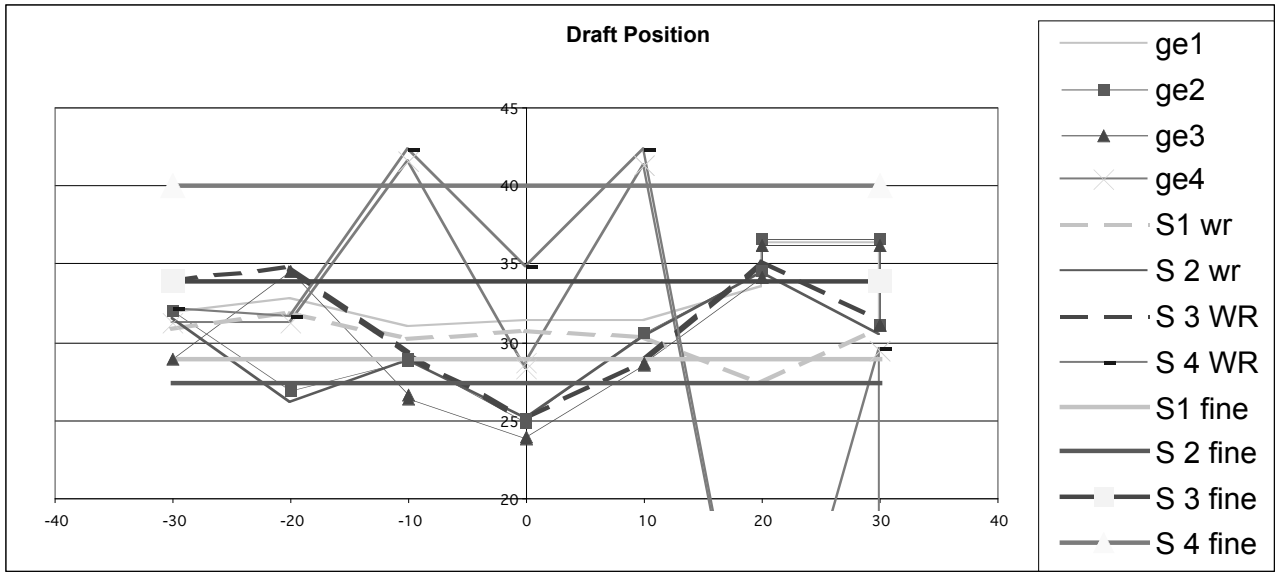

Table 2 Draft position as function of grid angle

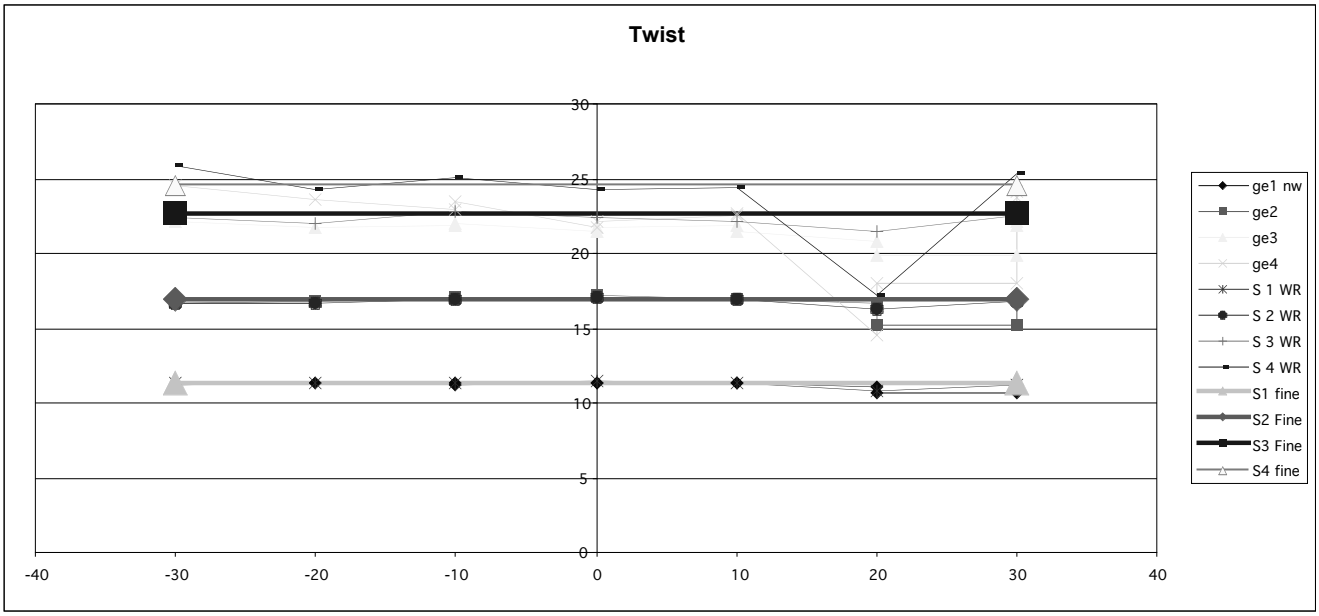

Table 3 Twist as function of grid angle

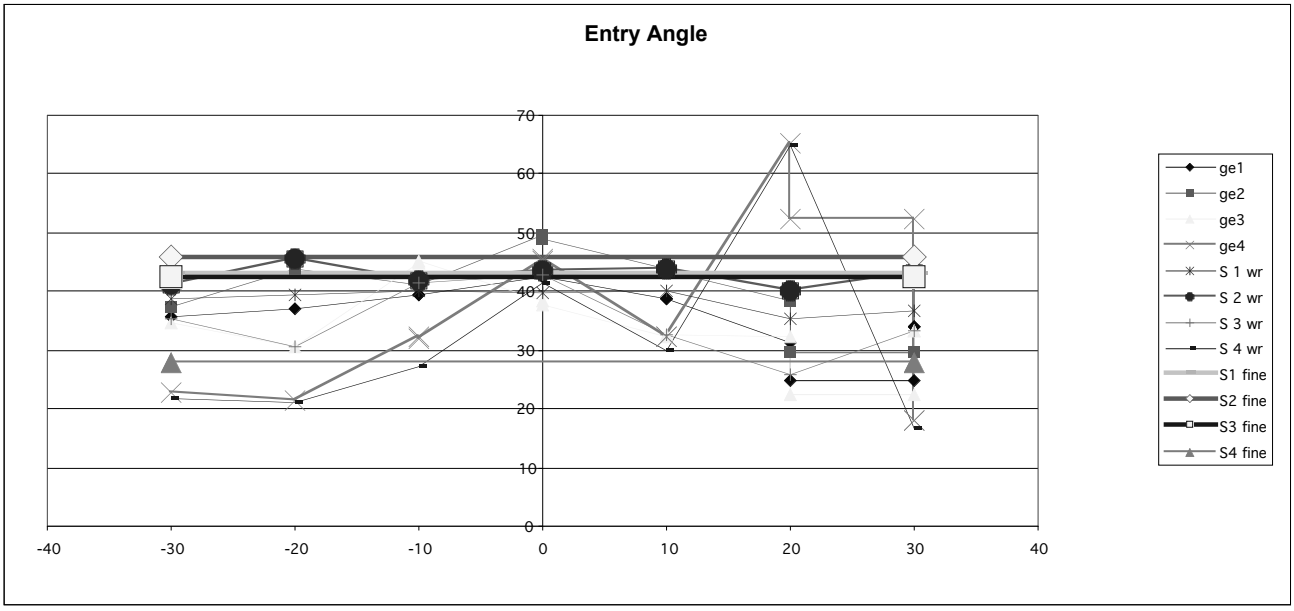

Table 4 Entry Angle as function of grid angle