Conservative Methods for Structural Optimization

J. A. Bennett* and R. V. Lust†
General Motors Research Laboratories, Warren, Michigan 48090-9055

Several methods can be used to generate conservative designs when using structural optimization techniques. A method is introduced which adds a padding term to each constraint that is proportional to the gradient of the constraint. One way to implement this approach is to determine the gradients of the constraints at the unpadded optimum design, to modify the constraint allowables, and then to restart the optimization. This approach assumes that these gradients will not change significantly when the new optimum is located. A second implementation is proposed which imbeds the padding calculation into the optimization process by updating the padding term periodically during the optimization process. Numerical experiments indicated that the additional nonlinearity introduced by this method did not affect the solution progress. Finally an implementation which uses nonlinear constraint approximations and a second order update method is described. It was shown that all implementations produced more mass efficient designs than does the more traditional method of adding a constant safety factor to each constraint.

Introduction

MATHEMATICAL optimization techniques require that the user explicitly specify limiting values on the constraints of the design problem. These limits bound the space in which acceptable designs are located. The optimization algorithm will not converge to a solution outside this region. In a typical structural design problem, limits are placed on the stresses in all elements, and displacement limits are imposed at certain points in the structure. In addition limits may be placed on other performance variables such as frequencies or on design variables such as thicknesses, widths, or heights. Since the optimal design will be limited by one or more of these constraints (e.g., some stresses will be at the maximum allowable value), a question frequently arises as to the sensitivity of the design to unforeseen changes in either the values of the limiting constraint allowables or the values of the design variables as manufactured. This concern generally is obviated in more traditional design methods because the constraints are rarely so clearly controlling the design. In optimal design, however, it would be desirable to provide for some amount of robustness in the final design by providing ways to explicitly handle a safety margin on the constraints.

A number of methods have been proposed to alleviate this problem.[1-3] A traditional approach is to use the concept of a safety factor on the values of the constraints in order to provide a margin of safety. This approach is particularly valid when large safety factors are used to provide a significant safety margin. Alternatively, it might be possible to scale up some or all of the member dimensions to make them stronger than required. Schmit[4] has observed that interior penalty function methods can automatically produce a sequence of designs, which tend to funnel down the center of the design space as opposed to tightly following the constraints. In addition, techniques based on assuming that the design parameters are random variables with some distribution of values have also been developed.[3] Though this latter approach may be quite powerful, to be effective it requires detailed knowledge about the statistical variation of the design parameters. Therefore this paper will concentrate on deterministic approaches; although it is clear the motivation for robust design philosophies is based on the assumption of uncertainty in either the model or the data. This paper will introduce the concept of generating robust designs by providing constraint padding which is proportional to the constraint gradients. Further, based on both linear and nonlinear approximations of the structural response, an implementation will be developed which updates this effect as the design progresses.

Received April 1989; revision received June 1989.

*Principal Research Engineer. Member AIAA.
†Staff Research Engineer. Member AIAA.

Measures of Robustness

Consider the standard optimization problem in which W is the objective function, G_i are the constraints, and x_j are the design variables. There are two rather traditional ways of generating a design which contains some measure of conservativeness. A simple expedient approach which is often used is to increase all structural design variables at the optimum (assuming they all lead to an increase in structural integrity) by a predetermined amount, i.e.,

$$x_j = x_j(1 + k_j) \tag{1}$$

This approach will automatically modify the values of the constraints but does not directly provide for additional protection from variability in the constraint allowable itself (such as yield stress).

In contrast there is the approach of providing a safety margin to each of the constraints. Then the problem becomes

$$\begin{aligned} &\text{Min} \quad W(x) \\ &\text{Subject to } G_i(x) + \bar{\kappa}_i \le 0 \\ &x_j^{\ell} \le x_j \le x_j^{u} \end{aligned} \tag{2}$$

where $\bar{\kappa}_i$ is the constraint tolerance or padding. The optimization can then be executed in the normal fashion.

Some other approaches to robust design have been presented in the structural optimization literature. Schmit[4] has suggested that the penalty function methods can automatically provide a set of designs which are pushed away from the constraints. These methods convert a constrained problem into a sequence of unconstrained problems [see Eq. (3)] in which the penalty is gradually changed so that the solutions converge to the solution of the constrained problem.

$$\underset{r \to 0}{\text{Min}} \quad W(x) - r\Sigma[1/G_i(x)] \tag{3}$$

The penalty term r forces the design away from the constraints until the final steps of the iterative process. Thus a series of progressively less robust designs is produced. The measure of robustness can be controlled at each step by the choice of the penalty parameter r; however, the actual margin is not clearly determined since it is an implicit function of r. At the optimum, of course, r approaches zero, and therefore no padding is guaranteed.

Another approach that is often put forward is to formulate the problem in the following fashion

$$\text{Max } \beta \tag{4}$$

$$\text{Subject to } G_i(x) + \kappa_i \beta \leq 0$$

$$\Delta W < \Delta W_{\max}$$

where $\Delta W_{\max}$ is a maximum allowable mass penalty. It should be noted that this is in some sense a dual to the problem stated in Eqs. (2) in that a maximum robustness is determined for a given mass penalty; whereas Eqs. (2) determine a minimum mass for a given robustness. The choice of which form to use may be based on the specific problem at hand or the specific software that is available.

It is also interesting to note that if the formulation of Eqs. (2) is used, and the solution is started from the infeasible region (corresponding to a solution for $\bar{\kappa}_i = 0$) and if a feasible directions algorithm is used, that the initial direction finding problem has the following form[5]

$$\text{Max } -\nabla W(x) \cdot s + \Psi \beta \tag{5}$$

$$\text{Subject to } \nabla G_i(x) \cdot s + \beta \kappa_i \leq 0$$

$$s \cdot s \leq 1$$

The search direction determined by Eqs. (5) is one that maximizes the distance ($\beta \kappa_i$) of the critical constraints from their boundaries with a minimum increase in the objective function. The interplay between increasing the objective function and pushing away from the boundaries is controlled by the parameter Ψ. If Ψ is large, the search direction determined by Eqs. (5) is similar to the search direction that would be produced by a feasible directions algorithm acting on Eqs. (4). This again suggests that the formulations of Eqs. (2) and (4) will give similar answers for many problems; it is merely a question of whether the robustness parameters or the mass penalty is to be specified.

Now let us reconsider the formulation in Eq. (2). This approach explicitly addresses the uncertainty in the constraint values themselves, but to represent the design variable uncertainty choosing a constant value for all $\bar{\kappa}_i$ will probably produce an overly conservative design. What is needed is some way to include in $\bar{\kappa}_i$ the direct effect of the design variables on the constraints. One possible way is to split $\bar{\kappa}_i$ into two terms. One term, κ_{i0}, would contain the padding relative to the constraint itself. The second term, κ_i, is taken to be proportional to the constraint gradient and thereby provides a larger padding term for constraints which are more sensitive to variations in the design variables. In order to calculate the effect on the constraint, the gradients must be multiplied by an assumed variation, k_j, in each design variable. This takes the form

$$\bar{\kappa}_i = \kappa_{i0} + \kappa_i \tag{6}$$

$$\kappa_i = \Sigma \left| (\partial G_i / \partial x_j) \right| k_j x_j$$

The absolute value of the derivative is used because the sign of the variation in x is undetermined. If k_j is equal for all j in Eq. (6), then the protection against change will be equal for all design variables. Some measure of equal robustness can then be generated by choosing $k_j = k$ and calculating κ_i. Alternatively a maximum κ_i can be selected, which will determine k, and subsequently the remaining κ_i can be calculated. In what follows we will neglect κ_{i0} and concentrate on κ_i. It could be pointed out, however, that if similar sensitivity information describing the behavior of the constraints with respect to other parameters of the design is available, it might be handled in a similar fashion in κ_{i0}.

Implementation Strategies

A simple way to implement this approach using existing structural optimization software is to first execute the optimization in the normal fashion. Then a padding term κ_i is determined for each constraint based on the values of the constraint derivatives at the optimum and an assumed value of k. The constraints are then modified by these quantities, and the optimization is restarted. Although the design is now infeasible, if the κ_i are relatively small, it should take only a few steps to converge. The assumption is made here that the gradients of the constraints are nearly constant and will not be significantly different at the new padded optimum. This assumption is consistent with the linearization assumptions made in many structural optimization problems. The major advantage of this implementation is that it does not require modification of existing structural optimization programs.

If the constraint derivatives are not nearly invariant, then the approach described above does not insure that the final design will provide precisely the padding specified since, at the new design point, the derivatives will have changed. In addition since the design is but a small modification of the original final design, there is little opportunity for the design process to identify an entirely different portion of the design space that might produce a superior design. These difficulties can only be addressed by embedding the padded constraint calculation of Eq. (4) into the optimization process itself. The immediate concern is that this would then require the derivatives of the padding term, which would then imply that second derivatives of the performance functions would be required. However, in the linearized approximate problem format,[6] during each design step, the constraint derivatives are assumed to be constant. Therefore the second derivatives are not required, and the constraint and its derivatives take the following form

$$\tilde{G}_i = G_{0i} + \Sigma(\partial G_i/\partial x_j)(x - x_0)_j + k \Sigma \left| (\partial G_i/\partial x_j) \right| x_j \tag{7}$$

$$\partial \tilde{G}_i / \partial x_n = (\partial G_i / \partial x_n) + k \left| (\partial G_i / \partial x_n) \right|$$

Thus a simple modification to the approximate constraint calculation allows a continuous update of the padding during the optimization. This will guarantee (to first order) that the converged design does provide the required padding, and no additional optimization runs are required. However, this does introduce an additional nonlinearity into the sequence of approximate problems since the effect of the padding may change from step to step due to changes in the constraint derivatives. These problems can be resolved by an appropriate choice of move limits. Also the parameter which defines the potentially active constraints for each step must be chosen carefully since the padding term will alter the value of the constraint after these constraints have been selected. In the examples shown later these concerns were easily resolved.

Now consider a slightly different formulation. Recent work in structural optimization[7] has shown that economies can be achieved by approximating only the force terms in the stress calculations and retaining the explicit nonlinearities in the stress calculations. Thus for each approximate problem, the forces are approximated by a first-order Taylor series, and the stresses are calculated, from these approximate forces, using the updated design variables (section dimensions). Therefore the derivatives of the constraints are not constant within the approximate problem and must be continually evaluated during the approximate problem solution. This can be easily done by a finite difference calculation on the constraint values since these are relatively inexpensive computations. For consistency then, the second derivative term should be included when the derivative of the padding term is calculated. In this case the approximate constraint is given by

$$\tilde{G}_i = G_{\text{non}l(i)} + k\Sigma |(\partial G_{\text{non}l(i)}/\partial x_j)| x_j \tag{8}$$

and the derivative is

$$\partial\tilde{G}_i/\partial x_n = \partial G_{\text{non}l(i)}/\partial x_n + k|(\partial G_{\text{non}l(i)}/\partial x_n)|$$

$$+ k\Sigma \text{sgn}(\partial G_{\text{non}l(i)}/\partial x_j)(\partial^2 G_{\text{non}l(i)}/\partial x_n \partial x_j)x_j$$

The derivative of $G_{\text{non}l}$ is calculated by finite differencing the values of the constraint obtained from the nonlinear expansion as described previously. One approach would be to neglect the second derivative term entirely under the assumption that it would have a negligible effect on the direction finding problem. Though this is attractive from an efficiency standpoint, it leads to a situation in which incomplete information about how the derivatives affect the padding term is supplied to the direction finding problem. Conceptually this could cause the optimizer to be unable to find the optimal solution. Two alternatives are available to obtain this second derivative information. It could be obtained by a full finite difference calculation, or one of the update techniques could be used to construct an approximate Hessian matrix. For this work the latter approach was taken.

The symmetric secant or Powell-symmetric-Broyden[8] update was chosen. This produces a symmetric but not necessarily positive definite Hessian matrix. The update formulas generate Hessian approximations, which are most accurate in a particular direction. Here we seek an approximation which is accurate in a direction parallel to the perturbed design vector x [see Eq. (8)]. The updates to the Hessian are calculated during the one-dimensional search of the optimization. The first point is taken to be the current point in the one-dimensional search. The second point is generated a small distance away from this point and in such a way that the vector between the two points is parallel to x.

Examples

In order to illustrate these concepts, they will be applied to a simple example problem, a seven-bar frame, which is described in Fig. 1. All solutions are obtained with the ODYSSEY structural optimization program.[9] The baseline optimum solution is shown in Table 1. The active and potentially active constraints and the type of the constraint are shown. All constraints are normalized to 1. All constraints are upper bounds except for 2, and 18 which is the lower bound on ω_2.

For the first example, the concept of moving the design variables by a fixed amount as in Eq. (1) is used. For this case $k = .05$ was chosen. In this case the mass becomes 9.3 kg, and the new constraint values are shown in Table 2 for the active constraints. One problem with this approach is that if the design is close to a side constraint, it may be impossible to move the size variables by the required amount without violating the original side constraint. For this problem this occurred for the height variables on elements 2 and 3 (see Fig. 1). For purposes of comparison, the variables were allowed to exceed the original side constraints, but in practice this may be impossible. As discussed previously, the sign of the constraint derivatives may vary, and this may affect the robustness of this design. In particular, for this design, the frequency constraint derivatives have varying signs. To evaluate this effect, the design, which had been obtained by increasing all size variables by 5%, was modified by assuming that all variables were perturbed by 5% in the direction to give a worst-case effect on the frequency constraint. As can be seen from Table 2, this violates the original frequency constraint by 1.9%, which originally had been satisfied by 1.3%.

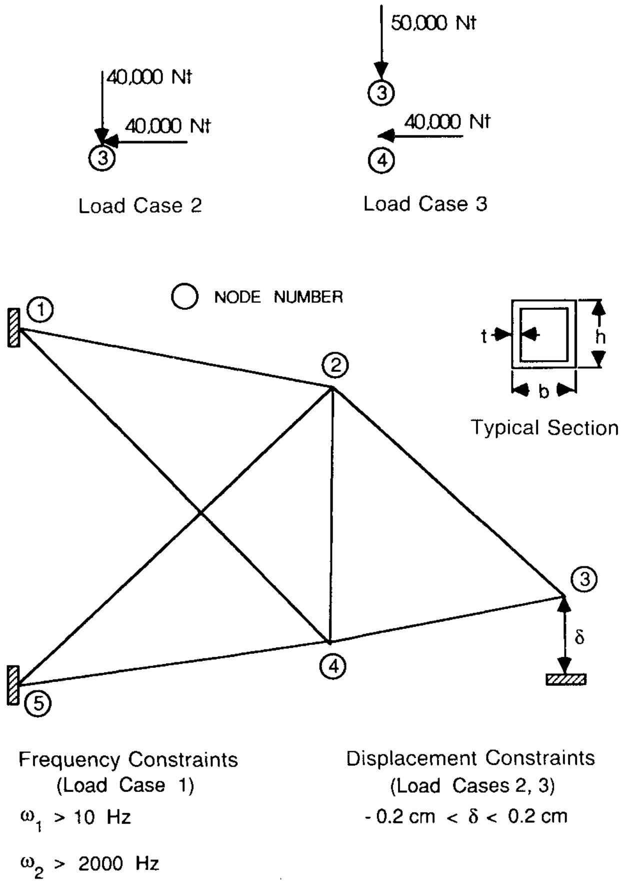

Fig. 1 Example problem.

Table 1 Values for base optimum design

Constraint	Value	Type
2	−0.997	Displacement
5	0.998	Stress
10	0.980	Stress
11	1.004	Stress
13	0.998	Stress
16	0.913	Stress
18	1.013	Frequency

Optimum mass = 8.47 kg

Now let us consider the measures of robustness based on adjusting the values of the constraints described by Eqs. (2). First a constant value of κ of .1 or 10% was assumed. This produced a design with a mass of 11.42 kg. This is a significantly higher mass than that produced by the first example, but all constraints have a 10% margin of safety. For the next example, the κ_i were chosen to be proportional to the constraint gradients as described in Eq. (6). The values of κ_i are given in Table 3 as well as the value neglecting the absolute value calculation and are for $k_j = 1$. As can be seen, the absolute value computation has a significant effect on the value of κ for the frequency constraint, which is due to the sign differences in the components of the gradient as described earlier. To execute this example, a maximum κ of 10% was chosen. Constraint 2 has the greatest k; therefore a corresponding k was calculated from this constraint ($.1*1 = 2.032*k$; $k = .05$). The remaining κ_i are shown in Table 4. With the constraints modified by these values, the optimization was run which produced a mass of 9.40 kg. The similarity of the results of this run and the results produced by merely increasing the design variables are interesting. The final masses are virtually identical, and the amount of constraint padding is also very similar except for the frequency constraint, which is affected by the absolute values of the derivatives as described earlier (compare Tables 2 and 4). In this latter approach, there is no guarantee

Table 2 Constraint values for variable scaling

Constraint	Original design (Fig. 1)	Size variables increased by 5%	Worst case variation
2	−0.997	−0.910	−0.991
5	0.998	0.906	0.833
10	0.980	0.894	0.979
11	1.004	0.926	0.854
13	0.998	0.920	0.975
18	1.013	1.019	0.981

Table 3 Proportional padded constraint κ

Constraint	κ	κ Neglecting absolute value
2	2.032	−2.032
5	1.907	1.907
10	1.371	1.371
11	1.998	1.998
13	1.770	1.770
16	1.168	1.168
18	0.733	−0.151

Table 4 Proportional padded constraint moves for 10% maximum move

Constraint	κ	Constraint value
2	0.100	−0.895
5	0.095	0.901
10	0.069	0.906
11	0.100	0.900
13	0.098	0.901
16	0.058	0.835
18	0.037	1.037

Table 5 Constraint values for variation in size variables from 10% proportional padded optimum design

Constraint	5% Decrease	Worst case
2	−1.008	−0.944
5	0.992	0.817
10	1.018	1.024
11	1.003	0.835
13	1.016	1.015
16	0.955	0.955
18	1.025	0.995

that the process described will insure the original constraints will be satisfied if the design variables are moved the anticipated amount. When the optimization is carried out, the material may be redistributed such that if all the size variables are decreased by the anticipated percentage, the design may be infeasible relative to the original constraints.

For the example problem, this was tested by working with the solution to the problem having proportional constraint padding as shown in Table 4. From this final design, all size variables were changed by the apppropriate amount, and this design was checked against all constraints. Again both a decrease of all size variables and a worst-case scenario based on the frequency derivative signs are created. As can be seen in Table 5, for the 5% decrease scenario, 4 constraints were slightly violated (2,10,11,13) and 3 constraints were violated for the worst-case scenario (10,13,18). For a highly nonlinear problem, this could be a significant difficulty but is probably not significant here.

In Fig. 2 optimization results have been plotted for both constant and proportional constraint padding. For this range of values, constant moves of the design variables [see Eq. (1)] would almost exactly correspond to the proportional constraint padding as far as final mass is concerned. For the constant constraint padding, the specific size variable protection is undetermined. It is clear that moving the constraints proportional to the gradients is a more mass effective way to provide this performance. For instance, to provide a maximum of 7.5% protection on the size variables (corresponds to $\kappa = .15$) requires a 72% mass increase using the constant strategy and requiring a 16% increase using the proportional strategy.

Finally, this approach has been applied to a 33-bar automotive frame (see Ref. 9 and Fig. 3). For this particular problem, a k of .05 led to a maximum κ_i of .18 rather than the .1 of the 7-bar problem. As a result constant κ of both .10 and .18 were used. The results of these optimizations are shown in Table 6. Once again similar trends were observed with the proportional constraint padding producing a smaller mass penalty than does the constant constraint padding. However, the contrast is not as dramatic as in the 7-bar since a smaller proportion of the constraints is active.

Examples of Alternate Implementations

We have proposed three ways to perform the optimization taking into account the constraint padding during the optimization 1) linear constraint approximations, 2) nonlinear constraint approximations with no second derivative approximation for the padding term, and 3) nonlinear constraint approximations with approximate Hessians for the padding terms. In order to evaluate these methods, runs were made for the 7-bar problem using $k = .05$. The results are shown in Table 7. This figure shows the mass and constraint values for each of these approaches as well as the case where padding is provided at the end of the run, labeled "proportional fixed." It is clear that for this problem that significantly different optimum masses are not produced; although a slight (but probably not significant) mass penalty was observed for the linear approximation approach. A similar result was obtained for linear approximation in the 33-bar problem (see Table 6); although the mass penalty in this case was more significant. However, all of the new approaches do guarantee (to first

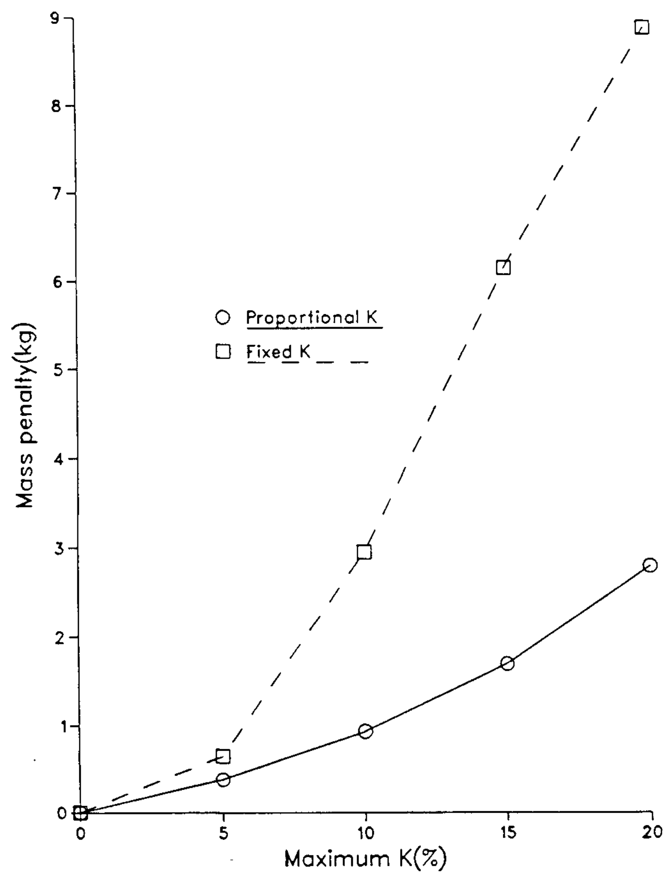

Fig. 2 Mass penalty as constraints are padded.

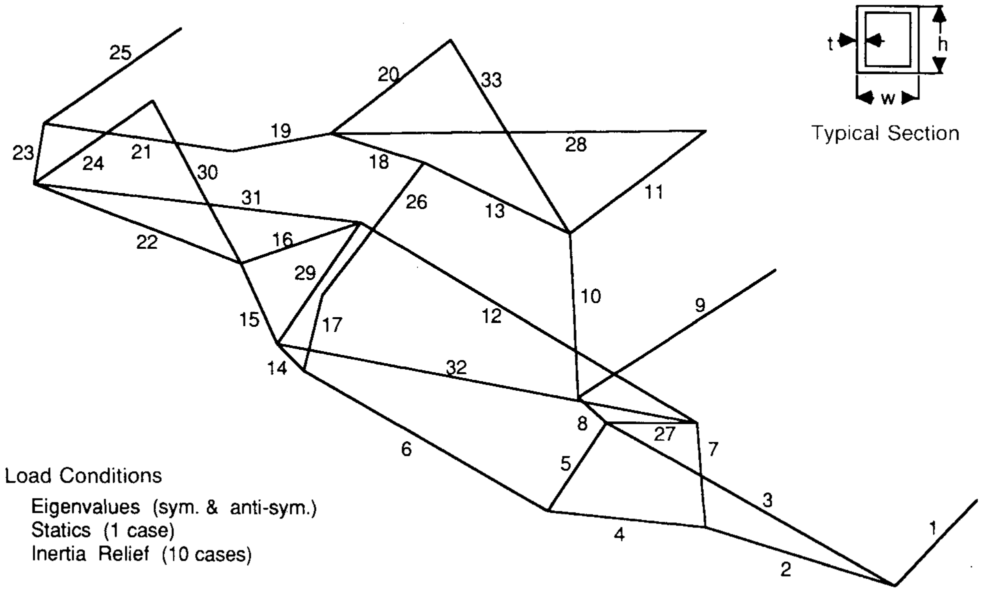

Fig. 3 33-bar frame example.

Table 6 Results for 33-bar problem

	Mass
Baseline Design	94.9
Fixed κ, .10	110.7
Fixed κ, .18	113.5
Proportional Padding, $k = .05$	102.3
Proportional Padding, $k = .05$ Linear approximation	104.9

Table 7 Alternative strategies for constraint padding $k = .05$, 7-bar problem

	Proportional fixed	Padding in approximate problem Linear	Nonlinear	Nonlinear/Hessian
Mass	9.40	9.54	9.41	9.42
Constraint				
2	−0.90	−0.90	−0.90	−0.90
5	0.90	0.89	0.85	0.87
10	0.93	0.95	0.92	0.95
11	0.90	0.93	0.91	0.92
13	0.91	0.95	0.96	0.96
16	0.85	0.84	0.85	0.85
18	1.04	1.04	1.03	1.03

Table 8 Alternative strategies for constraint padding $k = .10$, 7-bar problem

	Proportional fixed	Padding in approximate problem Linear	Nonlinear	Nonlinear/Hessian
Mass	11.25	10.63	10.43	10.48
Constraint				
5	0.80	0.83	0.84	0.83
10	0.84	0.90	0.84	0.90
11	0.80	0.88	0.82	0.88
13	0.82	0.86	0.85	0.87
16	0.70	0.75	0.77	0.75
18	1.07	1.07	1.07	1.07

order) that the final design will protect against a 5% change in all of the design variables. As was shown previously, this is not necessarily true for the case where the padding is accounted for at the end of the run. In addition both the linear approach and the nonlinear method with the Hessian approximation produce designs in which the padding term $k\Sigma|(\partial G_i/\partial x_j)|\,x_j$ is smaller (i.e., the constraint values are closer to 1) for constraints 10, 11, and 13 than for the padding term calculated at the end of the original optimization (see Table 7). This appears to indicate that the optimization process is searching for a design which tends to minimize the padding term.

This last observation suggests that this effect might be more pronounced for a larger value of k. Therefore a similar set of runs was executed for $k = .1$, and the results are shown in Table 8. With the exception of constraint 16, all of the constraints are active at all of the final designs. The two nonlinear update designs are only marginally lighter than the linear update design. The designs found by including the padding during the approximate steps clearly identified a place in the design space where the padding term was smaller than that for the fixed proportional case, which resulted in a lower mass design for the same protection against design variable variation. The designs obtained in this fashion exhibited changes in the design variables and sensitivities that were significantly different (> 10%) from the values for the proportionally fixed design. This effect should be more pronounced for more highly nonlinear problems.

Summary

It is clear from the simple examples shown that the technique of providing robustness against variations in the design variables by padding all constraints by the same amount is not a mass effective way to provide robustness because different constraints will have different dependence on the design variables. Therefore it is important that one carefully assess any tolerances that are imposed on constraints to insure that they accurately reflect the uncertainty in the design since they may significantly affect the final mass of the structure. A rational way of selecting these tolerances was proposed based on the gradients of the constraints. It was shown that this was a more mass effective approach to providing this measure of robustness. Two classes of methods to implement this concept were presented. One approach does not require modification of existing programs but does not guarantee a specified tolerance to design variable variation. The second approach does require program modification but appears to produce conceptually more robust designs.

Acknowledgment

The authors would like to thank P. A. Fenyes for his helpful discussions concerning the Hessian updates.

References

[1]Sandgren, E., Gim, G., and Ragsdell, K. M., "Optimal Design of a Class of Welded Structures Based on Design for Latitude," *ASME Transactions: Journal of Mechanisms, Transmissions, and Automation in Design,* Vol. 107, Dec. 1985, pp. 482–487.

[2]Balling, R. J., Free, J. C., and Parkinson, A. R., "Consideration of Worst-Case Manufacturing Tolerances in Design Optimization," *ASME Transactions: Journal of Mechanisms, Transmissions, and Automation in Design,* Vol. 108, Dec. 1986, pp. 438–441.

[3]Moses, F., "Structural System Reliability and Optimization," *Computers and Structures,* Vol. 7, 1977, pp. 283–290.

[4]Schmit, L. A., "Structural Optimization—Some Key Ideas and Insights," *New Directions in Optimum Structural Design,* edited by E. Atrek et al., John Wiley & Sons, 1984.

[5]Vanderplaats, G. N., *Numerical Optimization Techniques for Engineering Design: with Applications,* McGraw-Hill, 1984.

[6]Schmit, L. A., and Miura, H., "A New Structural Analysis/Synthesis Capability—ACCESS 1," *AIAA Journal,* Vol. 14, No. 5, 1978, pp. 661–671.

[7]Vanderplaats, G. N., and Salajegheh, E., "A New Approximation Method for Stress Constraints in Structural Synthesis," *Proceedings of AIAA/ASME/ASCE/AHS 28th Structures, Structural Dynamics and Materials Conference,* AIAA, New York, 1987.

[8]Dennis, J. E., Jr., and Schnabel, R. B., *Numerical Methods for Unconstrained Optimization and Nonlinear Equations,* Prentice-Hall, 1983.

[9]Botkin, M. E., Lust, R. V., Yang, R. J., Katnik, R. B., and Song, J. O., "Structural Sizing Optimization Using an External Finite Element Program," *Proceedings of AIAA/ASME/ASCE/AHS 28th Structures, Structural Dynamics and Materials Conference,* AIAA, New York, 1987.