Conf-950715--1

SAND95-1364C

APPLICATION OF OPTIMIZATION TO THE INVERSE PROBLEM OF FINDING THE WORST-CASE HEATING CONFIGURATION IN A FIRE

V.J. Romero, M.S. Eldred, W.J. Bohnhoff, and D.E. Outka
Sandia National Laboratories
Albuquerque, New Mexico, 87185, USA

ABSTRACT

Formal optimization procedures have been applied to determine the worst-case heating boundary conditions that a safety device can be credibly subjected to. There are many interesting aspects of this work in the areas of thermal transport, optimization, discrete modeling, and computing. The forward problem involves transient simulations with a nonlinear 3-D finite element model solving a coupled conduction/radiation problem. Coupling to the optimizer requires that boundary conditions in the thermal model be parameterized in terms of the optimization variables. The optimization is carried out over a diverse multi-dimensional parameter space where the forward evaluations are computationally expensive and of unknown duration *a priori*. The optimization problem is complicated by numerical artifacts resulting from discrete approximation and finite computer precision, as well as theoretical difficulties associated with navigating to a global minimum on a nonconvex objective function having a fold and several local minima. In this paper we report on the solution of the optimization problem, discuss implications of some of the features of this problem on selection of a suitable and efficient optimization algorithm, and share lessons learned, fixes implemented, and research issues identified along the way.

1. INTRODUCTION

Thermally induced failures and indeterminacies in critical structures and systems such as aircraft, weapon systems, naval vessels, and petrochemical processing plants can put people and engineered systems at risk. It is highly desirable to design such systems so that fire-triggered catastrophies are avoided. This requires probing the thermal robustness of candidate designs in various credible thermal environments. Alternatively or in conjunction, it would be useful to approach the problem from an inverse sense, where the heating scenario that a design is most vulnerable

This work was supported by the United States Department of Energy under Contract DE-AC04-94AL85000.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

to is identified. Then, the relative safety of the design can be quantified under "worst-case" conditions and design trade-off studies can be conducted.

Computer models validated by physical testing can be combined with formal optimization procedures to determine worst-case heating scenarios. In this paper we describe the use of optimization to solve the 2-parameter inverse problem of finding the size and location of a circular region of fire exposure that most severely threatens the integrity of a weapon safing component. We define the problem in an optimization context, study the forward problem and probe the parameter space, describe the interfaces between the optimizer module and the thermal model, report on the optimization process and results obtained, reiterate lessons learned and fixes implemented, and cite areas where improvements and advances are most critically needed.

2. PROBLEM DEFINITION

2.1 Thermal Model

Figure 1 shows an exploded view of a discretized model of the safing device. (An assembled unit is shown in Figure 2.) The model was created with the finite-element modeler PATRAN 2.5 [1]. It turns out that only half of the safing device had to be modeled due to a plane of symmetry in the problem (see Section 2.3). This

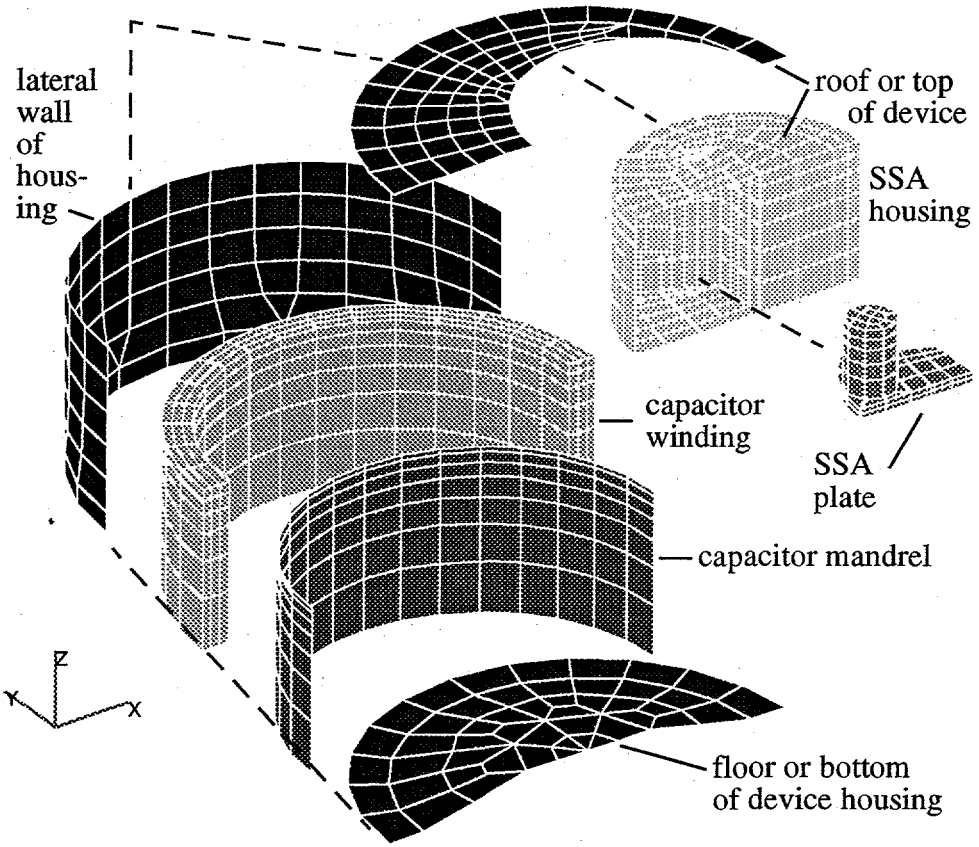

Figure 1 Safing Device Model Geometry (exploded view)

is very fortunate because seven radiation cavities exist in the model, and halving their sizes by symmetry decreases the size of the numerical radiation problem by a factor of about 3 for the present geometry and discretization.

Essentially, the safing-device housing is a cylindrical thin-walled stainless steel can with a diameter of roughly 6 inches and a height of about 2.6 inches. It is predominantly modeled with linear isoparametric quadrilateral and triangular shell[1] elements. The stainless steel single stronglink assembly (SSA) mates to a hole in the roof of the safing device via a perimeter weld. All solid (non-shell) finite elements in the model are linear isoparametric elements, either 8-node hexahedra or 6-node wedges. The SSA plate, to which several devices are attached, fits inside a cavity in the SSA. The physical corners of the plate (as opposed to apparent corners created by the symmetry cut through the model) are bolted to shoulders inside the cavity of the SSA housing. Perfect thermal conductance is assumed across all intimate (bolted, welded, mounted, wound) interfaces in the model. A Mylar-and-foil laminate is wound around the stainless-steel shell-element mandrel. This capacitor winding has highly anisotropic properties because of its layered structure. Thus, the finite elements making up the winding are assigned individually oriented orthotropic property tensors. The top end of the mandrel, which extends just slightly beyond the winding, is welded to the roof of the safing-device housing. There are nearly 2000 finite elements in the model. The highly temperature-dependent properties of stainless steel and the large temperature excursions involved make the conduction problem a fairly nonlinear one.

Because of the high temperatures involved in this problem thermal radiation is extremely important. However, radiative transport is a very nonlinear and computationally demanding problem to solve, adding a component of difficulty to this problem that is just now beginning to be surmountable with todays fast computers and advanced thermal modeling and simulation packages. A system of 7 enclosures (1046 surfaces total) is used to account for internal radiant exchange within the safing device. Upper and lower enclosures exist inside the stronglink housing, separated by the plate. The main enclosure includes the surfaces on: the exterior lateral and bottom faces of the SSA, inside wall of the mandrel, bottom face of the winding, bottom plate of the safing device, and ceiling of the device between the mandrel and SSA. The enclosure where the top and outside faces of the winding exchange with the device housing is modeled as four 90-degree arcing enclosures to reduce the number of view factors generated (while giving up only incremental accuracy).

External forcing conditions (boundary conditions) are described in Section 2.3. Approximately four man-weeks were required to construct the model from drawings and apply boundary conditions. The thermal solver used in the analyses was QTRAN [2].

---

1. Thermal shell elements do not support temperature variation across the element in the thickness direction but do support in-plane temperature gradients. They are used to avoid numerical stiffness that can accompany the resolution of very small temperature gradients through relatively thin, conductive walls.

2.2 Device Operation and Relation to Objective Function

The safing device is intended to prevent operation or triggering of a weapon by unauthorized personnel or unintended occurances. The single stronglink assembly, which we will also call the "stronglink", prevents the transfer of uncleared electrical signals to critical operation and control components in the system. In any operational or abnormal environment it must serve its function until other components in the system critical for operation of the weapon are irreversibly neutralized. In particular, the capacitor winding, which we will also refer to as the "weaklink", must become incapable of holding an electrical charge before the stronglink succumbs.

In abnormal environments a race to time-of-failure between the weaklink and stronglink is set up in the local confines of the safing device. This begs interesting and important questions: What, if any, abnormal environments cause the race to be lost by the weaklink (and thus cause the safing device to be defeated)? Can the worst-case environment be identified so that the device can be designed to survive it? These questions must be answered in an inverse manner, *i.e.* by trial and iteration. Though iterative in nature, these problems can be systematically approached within the context of optimization.

From a thermal perspective, we define failure criteria for the weak and strong links in terms of "failure temperatures" that can in general be complex functions of temperature history, heating rate, geometry, boundary conditions, etc. By experimental testing it has been determined that, under thermal conditions in the neighborhood of those to be probed in the optimization run, the status of the stronglink becomes indeterminate when certain critical elements on the underside (bottom face in Figure 1) of the SSA plate reach 1100° F (600° C). For the weaklink capacitor it has been determined that Mylar melts at about 480° F (250° C), and begins to shrink at significantly lower temperatures. Flow of the Mylar or separation at inclusions during the shrinking process allows coalescence of the oppositely charged aluminum foils in the winding, which shorts the capacitor, rendering it irreversibly inoperable. Testing has shown that, in the neighborhood over which the optimization is to be performed, a conservative failure criterion for the weaklink is a temperature of 500° F (260° C).

In our simulations we periodically scan the temperatures of all nodes on the underside of the SSA plate and all nodes of the winding. We define $t\_fail_{stronglink}$ as the elapsed time (from time zero at the beginning of the simulation) required for the hottest node on the underside of the plate to reach a temperature of 1100° F (600° C). Similarly, $t\_fail_{weaklink}$ is the time required for the hottest node on the capacitor winding to reach a temperature of 500° F (260° C). The difference $t\_fail_{stronglink}$ - $t\_fail_{weaklink}$ is a measure of the safety margin, and will constitute the objective function in our study. We seek a global minimum for this value over a range of heating scenarios to be discussed next. By minimizing the value of this function we identify increasingly threatening heating scenarios, finally arriving at the worst-case scenario at the objective- function minimum. Thus, we solve the inverse problem.

2.3 Two-Parameter Heating Function

The bottom of the safing device is mounted to a bulkhead and is largely isolated

vibrationally and thermally via rubber retainers. Thus, an adiabatic boundary condition is applied to the bottom surface of the device. However, the lateral and top surfaces of the device can be exposed to heat sources. Given that the weaklink has a much lower failure threshold than the stronglink, and has much less heat capacitance, uniform heating of the safing device will result in weaklink failure that precedes stronglink failure. Therefore, for our purposes, we only need concern ourselves with those environments that heat the stronglink preferentially relative to the weaklink. Indeed, we seek by formal optimization to maximize the degree to which the stronglink is preferentially heated. Any element of lateral heating is to be avoided because this heats the weaklink more directly than the stronglink. Hence, the lateral surfaces of the safing device are kept adiabatic in our study. Localized intense heating from the top, however, does meet our criteria. In particular, the device appears potentially vulnerable to application of heat to the roof of the safing device, localized to a region directly above the SSA plate.

We now bring the problem out of the abstract by imposing a form to the heating environment. We postulate that in an accident a hydrocarbon fuel fire erupts and irradiates a small circular region on the top of the firing set, the rest of the firing set being completely shaded from the fire. We model the fire as a blackbody radiator at a temperature of 1000°C (1832°F) (see [3]). We assume the device is completely insulated except at the said irradiated region, which is assumed to fully view the fire (view factor = 1). From geometrical and heat transfer considerations it can be concluded that, if limited to a single circular window of exposure on the roof of the device, the circular region should be centered on the diametral line corresponding to the plane of symmetry of the device. (Thus a plane of symmetry exists in the total thermal problem {geometry + boundary conditions}.)

Initial temperature has a nonlinear effect on the stronglink/weaklink race and is a determining factor in the worst-case heating configuration. However, it is not included as a free parameter in the current investigation. In our simulations the model is initially at 25°C and the fire temperature is ramped from 25°C to 1000°C over the first 10 seconds and held constant thereafter.

We therefore have a two-parameter description over the set of relevant and allowable heating configurations. Radius of the circular irradiated region is one parameter and location of this "spot" along the said diameter is the other. We proceed to find the optimum (worst-case) heating configuration with respect to these two parameters.

3. PARAMETER-SPACE SAMPLING

3.1 The Forward Problem (An Example)

The forward problem is solved to obtain the value of the objective function at given values of the heating parameters. Figures 2 and 3 show results of a simulation run with the thermal model for the parameter values r = spot radius = 1.020 inches and x = 0.142 inches = distance from the center of the safing device in the positive x direction as shown in Figure 2. These values define a region on the roof of the device pointed to by the white arrows in the figure. Though difficult to verify by looking at the gray-scale plot, high temperatures are concentrated about the stron-

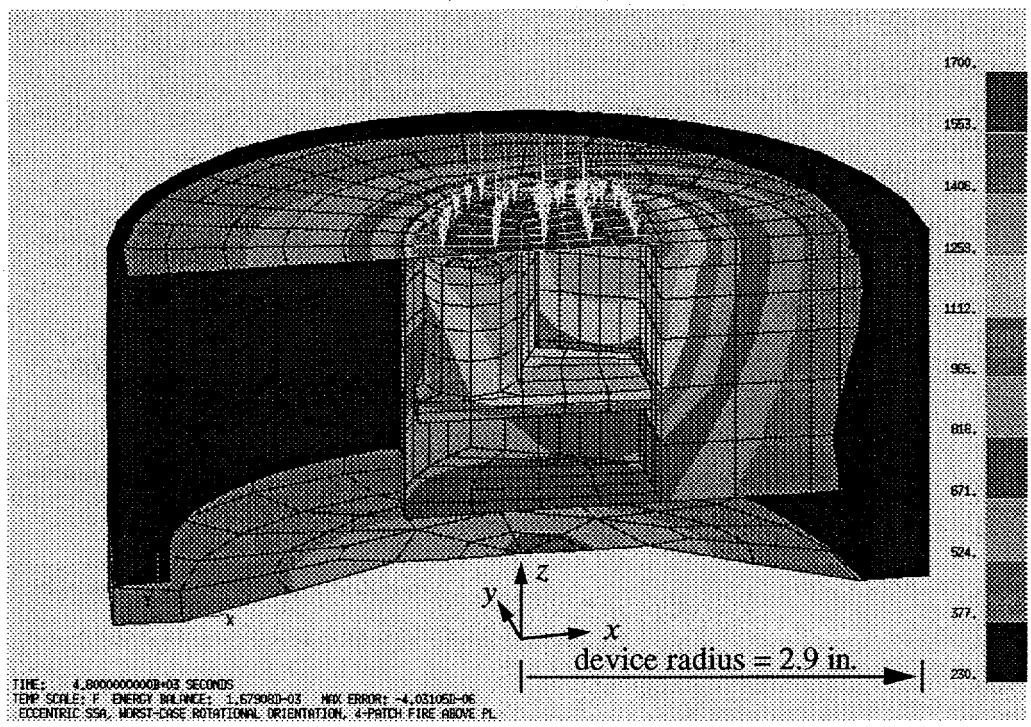

Figure 2 Safing device temperature distribution (°F) 80 minutes after the start of the fire, heating parameters r=1.02 in., x=0.142 in.

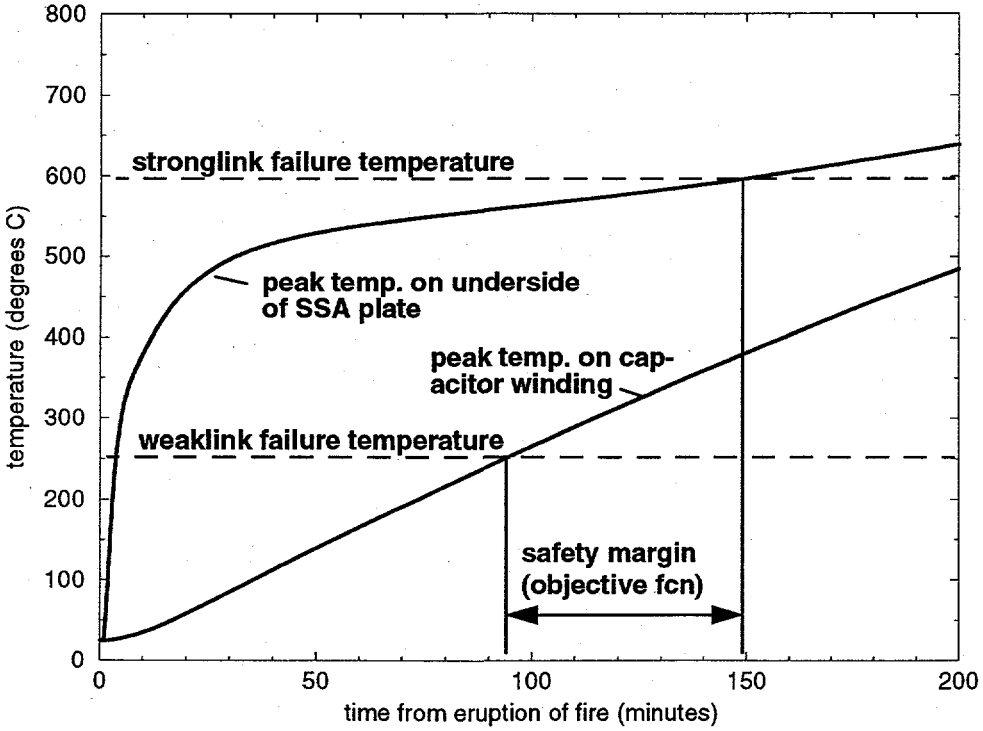

Figure 3 Pertinent temperature histories for r=1.02 in., x=0.142 in.

glink plate, with the capacitor winding being relatively cool. Thus, this combination of parameters results in a highly localized heating configuration that preferentially heats the stronglink to a high degree. This is what we are after. Figure 3 shows the relevant weaklink and stronglink temperature responses over time. For the stated parameters the value of the objective function is $O(x = 0.142 \text{ in.}, r = 1.020 \text{ in.}) \cong 56$ minutes.

3.2 Objective-Function Character: 1-D Parameter Studies

We can get a local indication of the variation of the objective function over the optimization space by performing 1-D parameter studies in the optimization variables r and x.

With regard to location x of the irradiated spot, it would seem that a region centered somewhere over the SSA plate (e.g. the region indicated by white arrows in Figure 2) is most likely to constitute worst-case positioning. A parameter study in x bears this out. Figure 4 shows the variation of the objective function as x is varied over a particularly relevant span of its allowable range when the radius is held constant at $r = 1.89$ inches. As the irradiated spot proceeds away from the center of the safing device a minimum point is reached somewhere over the exposed portion of the SSA plate. As the spot proceeds further away, the objective function rises rapidly as the capacitor winding begins to "feel" the heat much more directly and the SSA plate simultaneously experiences a decrease in heating. Significantly (from an optimization perspective), a fold in the objective function exists as indicated by the cusp point at the minimum in Figure 4 where the slope becomes discontinuous. This occurs because an abrupt shift occurs in the location on the winding where the failure temperature is first experienced. When the center of the irradiated spot is to the left of the cusp point, the peak temperature on the capacitor always occurs in (-)x half-space at the top of the winding where the plane of symmetry intersects the inner cylinder of the winding. On the right of the cusp, failure always first occurs at the mirror image in (+) x half-space. Additionally, the objective function in the viscinity of the cusp exhibits negative curvature, which is poorly captured by the positive-definite Hessian approximations used in many second-order optimizers.

Assuming that the irradiated spot is centered somewhere above the SSA plate, an interesting interplay exists between the localization of heating and the amount of thermal energy deposited to the device. As the circular region that sees the fire decreases in size, the heating becomes more localized to the stronglink (and remote from the weaklink), tending to worsen the severity of the event. However, a counteracting effect exists in that the magnitude of heat applied to the firing set diminishes, allowing the transport mechanisms (conduction and radiation) to diffuse heat more effectively (*i.e.* fast enough to avert localized concentration of heat). Thus, we would expect that a minimum in the objective function would occur for a medium-sized spot. This is verified by Figure 5, which presents the results of a parameter study over spot radius r for a spot centered at $x = 0.8$ in. Though almost imperceptible because of the convexity (positive curvature) in the viscinity of the minimum, the slope of the objective function is discontinuous at the minimum due to a switch in the nodal location where weaklink failure occurs. The figure also verifies the presence of multimodality in the objective function, which presents a diffi-

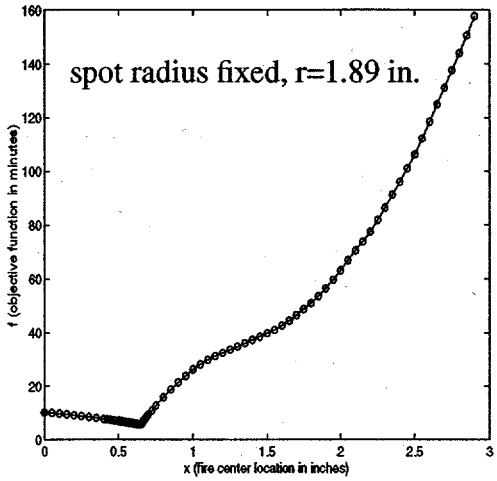

Figure 4 Objective fcn. variation w.r.t. spot center location x, for fixed r

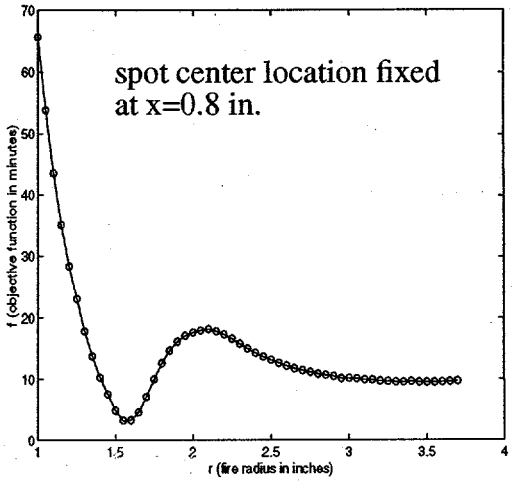

Figure 5 Objective fcn. variation w.r.t. spot radius r, for fixed x

culty in optimization because a local minimum might mistakenly be identified as the global minimum.

3.3 Effect of Model Parameters on Optimization Process

The finite precision of numerical solutions to the forward problem can introduce difficulties into the optimization process. Figure 6 is a magnification of a portion of Figure 4 showing small-scale nonsmoothness that can occur for improperly set error tolerance parameters EPSIT and EPSIT2 (see [2]) in the thermal model. Other factors such as machine precision, code precision, relaxation parameters and type of time integrator used in the thermal solver, etc. also affect the result, but the most direct control is usually afforded through specification of allowable per-time-step errors. The driving issue is not necessarily that a numerically converged result has not been obtained; the maximum deviation of the unsmooth curve in Figure 6 from the "smooth" curve is acceptably small. Rather, the fundamental issue is the stochastic nature of the deviations, which can create local extrema that the optimizer can converge to, or which can prevent the optimizer from converging at all. As it turned out, none of the nonlinear programming algorithms tried could successfully navigate the design space at the less strict set of error tolerances listed in Figure 6, but much better success was obtained at the tighter set of tolerances.

The stricter error tolerances that smoothness demands can increase tremendously the compute time per forward evaluation. In our problem the additional CPU time varied considerably over the parameter space, but typically factors of 2 to 10 times more compute time was required to produce the data points of the

smooth curve in Figure 6 than of the unsmooth curve.

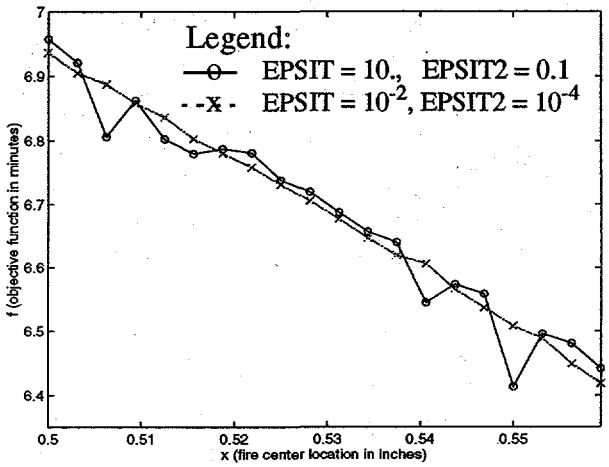

Figure 6 Detail of Figure 4 showing effect of QTRAN convergence tolerances on objective fcn. nonsmoothness

4. THERMAL MODEL-to-OPTIMIZER INTERFACE

The thermal model for solving the forward problem was coupled with various software packages for solving the optimization problem via an object-oriented software environment described in reference [4]. The basic mode of operation is the following. The optimizer generates a parameter set (radius and location) describing a heating configuration for which the value of the objective function must be determined. A software script translates these parameters into heat flux boundary conditions for all of the affected elements on the roof of the safing device. The thermal model is then run until the objective function can be calculated. This involves monitoring the temperatures of computational nodes on the capacitor winding and underside of the SSA plate, sorting to find peak temperatures, suspending the computation after both weaklink and stronglink have reached their failure temperatures, and calculating the value of the objective function from the respective failure times. Finally, the value of the objective function is passed to the optimizer which, with the benefit of this and previous information, comes up with a new set of heating parameters to try.

5. OPTIMIZATION RUNS

The optimizations were performed over the parameter ranges $0 \le x \le 2.9$ and $0.5 \le r \le 5.8$. Runs were executed on a single processor of an IBM SP2 multi-processor computer. Several widely available optimization routines were tested, as described below. The function evaluations with the thermal model consumed the overwhelming proportion of computation time in the optimizations. An average evaluation took about 20 CPU minutes on the SP2 machine.

This application was challenging from an optimization perspective due to the nonsmoothness and nonconvexity of the design space. Several nonlinear programming optimization packages were employed for the solution of this problem,

including DOT[5], NPSOL[6], and OPT++[7]. In general, the Newton-based optimizers (NPSOL's sequential quadratic programming algorithm, DOT's BFGS quasi-Newton method, and OPT++'s quasi-Newton methods) performed poorly due to inaccuracy of the Hessian approximation caused by the nonconvexity of the design space (*i.e.* negative curvature in the viscinity of the minimum). Conjugate gradient methods (DOT's Fletcher-Reeves conjugate gradient and OPT++'s Polak-Ribiere conjugate gradient) were much more successful. Furthermore, choice of finite difference step size (FDSS) for computation of gradients proved to be important. Table 1 shows the results for conjugate gradient optimizers with varying FDSS for EPSIT=10^{-2} and EPSIT2=10^{-4} (see Figure 6 for effect of EPSIT tolerances). An asterisk (*) in the function evaluations column indicates that the optimization did not run to convergence (premature termination occurred due to search directions that failed a descent-direction test).

Optimizer	FDSS	Initial Values (r, x)	Initial Obj. Fcn.	# Fcn Evals	Final Values (r, x)	Final Obj. Fn.
OPT++ P.-R. CG	4%	(1.4, 0.5)	7.2045	34*	(1.5812, 0.75016)	2.8293
OPT++ P.-R. CG	1%	(1.4, 0.5)	7.2045	100*	(1.6038, 0.76547)	2.5956
OPT++ P.-R. CG	0.1%	(1.4, 0.5)	7.2045	73	(1.6086, 0.76895)	2.5546
OPT++ P.-R. CG	0.01%	(1.4, 0.5)	7.2045	28*	(1.6016, 0.57370)	4.9477
OPT++ P.-R. CG	1%	(1.0, 1.0)	86.954	31*	(1.9321, 0.62026)	5.9543
OPT++ P.-R. CG	1%	(1.2, 0.9)	34.831	106	(2.1044, 0.50665)	6.9526
DOT F.-R. CG	1%	(1.4, 0.5)	7.2045	34	(1.6435, 0.77498)	2.5973
DOT F.-R. CG	1%	(1.0, 1.0)	86.954	40	(1.6374, 0.77591)	2.5362
DOT F.-R. CG	1%	(1.2, 0.9)	34.831	38	(1.6475, 0.77393)	2.6217

Table 1 Optimization results with conjugate gradient (CG) optimizers for EPSIT = 10^{-2}, EPSIT2 = 10^{-4}

The first four rows of Table 1 illustrate the effect of FDSS on the optimization: FDSS should be as small as possible to allow for effective convergence to a minimum (0.1% is better than 1% which is better than 4% since the gradients are less accurate locally for larger FDSS), but still large enough that small-scale nonsmoothness does not cause erroneous gradients (0.01% is too small; the optimizer cannot successfully navigate the design space since the FDSS is on the order of the design space noise). The last five rows show that DOT's conjugate gradient (CG) optimizer was more robust and efficient than OPT++'s CG optimizer, through the fact that DOT was successful from 3 different starting points and OPT++ from only one, and through the lower number of function evaluations that were required. The chief cause for these differences was not the version of conjugate gradient being used (in fact, Polak-Ribiere is generally regarded to be superior to Fletcher-Reeves [8]), but rather was DOT's superior line search routine. The line-search algorithm

in the OPT++ package is currently being improved.

To achieve the best answer possible, the error convergence tolerances were tightened 2 additional orders of magnitude (EPSIT=10^{-4}, EPSIT2=10^{-6}) and the FDSS was set to 0.1%. DOT's Fletcher-Reeves conjugate gradient algorithm was used to obtain the lowest objective function value of 2.5309 minutes at (r=1.6204, x=0.78205). *This safety margin is an order of magnitude lower than the result estimated with adhoc trial and iteration methods traditionally applied to these types of problems in our work!*

6. CONCLUSIONS AND RECOMMENDATIONS

Formal optimization was successfully applied to the difficult and computationally intensive inverse problem of identifying the worst-case heating configuration for a safing device. Our efforts to date in this and other applications involving expensive function evaluations with large complex nonlinear mechanics codes have been largely successful. Along the way we've learned many things and identified difficulties that still need to be addressed in more general and efficient ways. We summarize our findings:

- Finite computing precision, finite discretizations in space and time, and approximate methods that resolve nonlinearities inexactly (out of necessity to make the problem tractable) are unavoidable factors that lead to stochastic numerical nonsmoothness of the predicted objective function. If the scale of the nonsmoothness is large relative to the optimization step size and/or to the local undulations of the "exact" objective function, then gradient-based optimizers will have difficulties converging to true extrema. After the model has been built, the analysis code has been picked, and the hardware has been identified on which to run the optimization, solution-algorithm parameters such as iterative convergence criteria, per-time-step error tolerances, and relaxation parameters are the most convenient parameters to adjust to alleviate numerical nonsmoothness. However, such adjustments may be extremely expensive in terms of added computational cost. It is generally true that the extra requirement of smoothness that optimization brings to numerical modeling of complex nonlinear phenomena drives the acceptability or nonacceptability of finite errors much more than numerical convergence. Therefore, "smarter" gradient-based techniques must be devised to look at the scale of the objective function and its gradients, at the history of the optimization, and at other cues to enable it to efficiently navigate response surfaces with small numerical nonsmoothness.
- When using finite difference gradients in applications with nonsmoothness, an effective finite difference step size is not easily determined. It must be small enough to allow convergence to the actual minimum, but large enough to not be adversely affected by acceptable levels of nonsmoothness.
- Hessian inaccuracy can be a problem in nonconvex design spaces, causing poor performance in many Newton-based methods [Trust region methods (to be evaluated in future research) have the potential to overcome this difficulty while maintaining the theoretical strength of second-order optimizers].
- Since conjugate gradient optimizers do not rely upon an approximate Hessian,

they remain effective in the presence of nonconvexity and reasonable nonsmoothness. While conjugate gradient optimizers are not as efficient as Newton-based optimizers in smooth, convex problems, approximate curvature information can prove to be a liability in some nonsmooth applications.

- Since our application involved very expensive function evaluations and events of unknown duration, we tried stopping the simulations at a preselected time and extrapolating to find the value of the objective function if necessary. However, we found this approach to be inaccurate in some cases and wasteful in others. We then implemented an adaptive time-stop strategy to terminate the simulation just after the information vital to calculating the objective function had been determined. This minimized CPU usage while maintaining solution accuracy.

References

[1] PATRAN Plus version 2.5 User Manuals, January 1990, PDA Engineering, Costa Mesa, CA.

[2] P/THERMAL Analysis Package User Manuals, Release 2.6, March 1993, PDA Engineering, Costa Mesa, CA.

[3] Clark, R.K., Foley, J.T., Hartmann, W.F., and Larson, D.W., "Severities of Transportation Accidents," Sandia National Laboratories report SLA-74-0001 (U), printed September 1976.

[4] Eldred, M.S., Outka, D.E., Fulcher, C.W., and Bohnhoff, W.J., "Optimization of complex mechanics simulations with object-oriented software design," proceedings of the 36th AIAA Structures, Structural Dynamics, and Materials Conf., April 10-13 1995, New Orleans, LA., part IV, pp. 2406-2415.

[5] DOT Users Manual, Version 4.10, VMA Engineering, Colorado Springs, CO, 1994.

[6] Gill, P.E., Murray, W., Saunders, M.A., and Wright, M.H., "User's Guide for NPSOL (Version 4.0): A Fortran Package for Nonlinear Programming," System Optimization Laboratory, TR SOL-86-2, Stanford University, Stanford, CA, Jan. 1986.

[7] Meza, J.C., "OPT++: An Object-Oriented Class Library for Nonlinear Optimization," Sandia National Laboratories report SAND94-8225, Livermore, CA, printed March 1994.

[8] Coleman, T.F. and Li, Y., Eds., *Large-Scale Numerical Optimization*, Society for Industrial and Applied Mathematics, Philadelphia, 1990.