A COLLABORATIVE OPTIMIZATION ENVIRONMENT FOR TURBINE ENGINE DEVELOPMENT

Peter J. Röhl*, Beichang He*, Peter M. Finnigan†
General Electric Corporate Research and Development
Schenectady, NY 12301

Abstract

A MDO scenario for the design and manufacturing process of gas turbine engine disks is developed. High-fidelity engineering analysis and process simulation tools are integrated into an optimization environment. While different formal MDO approaches are discussed, a sequential optimization approach seems to be best suited or this specific problem. The forging optimization results in a minimum-weight forgeable disk that meets all constraints in terms of process parameters. The optimization of the heat treatment process reduces residual stresses while maintaining required cooling rates through the modification of surface heat transfer coefficients. Optimization of both the forging and the heat treatment process individually has been successful, but the complete MDO scenario still faces a number of obstacles. Parametric CAD tools are not as robust for complicated geometry as it would be necessary in an automatic optimization environment. The same applies to the interface between CAD and CAE tools. Computational resources constitute another bottleneck - formal MDO algorithms tend to be slow in their convergence behavior, which makes them less well suited for problems requiring high-fidelity analysis codes with their long execution times. Despite all these obstacles, though, progress towards a comprehensive disk MDO environment is apparent.

Introduction

The design and manufacturing of gas turbine engines is a highly coupled multidisciplinary process involving design of the thermodynamic cycle, flow path and airfoil design, rotordynamics, and thermo-mechanical design for life prediction. An important aspect is the design and optimization of the manufacturing process of the mechanical components, requiring detailed simulation of forging, heat treatment, and machining processes. With the economic pressures which exist today, the need to develop affordable, high-performance defense systems, with shorter product development cycle times has never been greater. Propulsion systems are no exception considering their intrinsic complexity and strong system coupling with their associated aircraft or launch vehicles. The successful development of integrated propulsion systems is critically linked to our ability to perform system, subsystem, and component-level simulations of the design and manufacturing processes. Today, the problems are compounded because of the geographically distributed intra- and inter-company partnerships, including second and third tier suppliers, which are formed out of economic, technical, and product necessity. The ability for industry to develop, and cost-effectively deploy these systems, is predicated on its ability to rapidly simulate both products and processes to achieve globally optimized designs. To that end, there are a number of key technologies which are being developed and demonstrated under the DARPA-funded RaDEO (Rapid Design Exploration and Optimization) program[1] as part of the propulsion scenario. Under the RaDEO contract, the GE Research and Development Center is teamed with Engineous Software, Inc. (ESI) to develop a collaborative optimization environment based on iSIGHT[2], Engineous Software's optimzation framework. One focal point is the development of an optimization toolkit which enables the user to easily formulate an MDO problem and cast is into the form of one of the "formal" MDO algorithms supported by this toolkit. Another is the extension of the iSIGHT environment to facilitate the integration of CAD and CAE systems with the help of two toolkits, the Product Modeling Toolkit (PMTK) and the Discrete Analysis Modeling Toolkit (DMTK). The engine disk design problem is one of the application demonstrations to be addressed in this project.

The Engine Disk MDO Problem

The individual steps of the disk design process, broken down into the mechanical design and process/manufacturing aspect, are shown in fig. 1. Each of the five steps in the process can be further subdivided into a number of individual sub-steps with analyses at varying levels of complexity. The thermo-mechanical design, for example, starts with a simple 1-d analysis to obtain a rough thickness distribution of the disk. As knowledge about the design increases, more complex analysis models are created up to a full 3-d finite

---

* Staff Engineer, Member AIAA
† Manager, Mechanical Design Methods and Processes Program

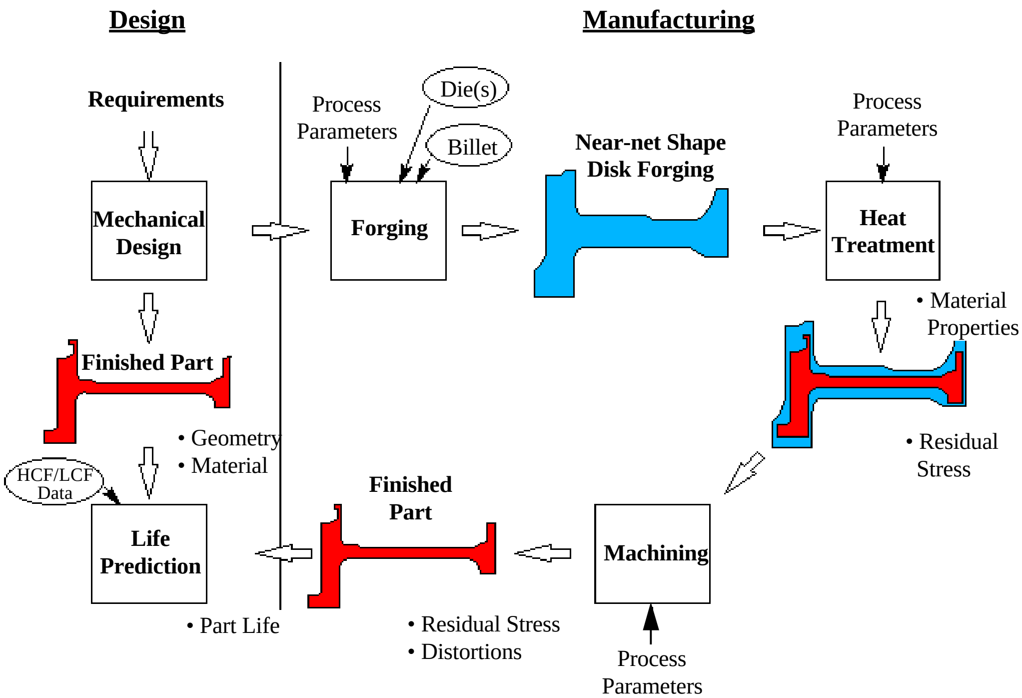

Fig. 1: Engine Disk Design and Manufacturing Process

element analysis with tens of thousands of elements. A thermal transient analysis is performed on the disk to supply the mechanical design group with thermal loads for the different points in the design missions. These thermal loads are iteratively adjusted as the design progresses. Objective during the mechanical design phase is first and foremost the determination of the final disk shape, as early as possible in the design timeline in order to be able to release the forgings which tend to require a long lead time. A shape is to be determined that meets mission requirements at minimum weight and/or minimum cost. A detailed simulation of the manufacturing process is necessary in order to determine both residual stresses and final distortions of the finished part after machining operations. These residual stresses, in turn, are used in the subsequent life prediction of the part. Objective during the simulation of the forging process is the determination of the die shape on the one hand and of an optimum forging process on the other that ensures proper die filling without compromising mechanical properties of the disk through the violation of stress, strain, strain rate, or temperature limits. The subsequent heat treatment process is designed to generate acceptable mechanical properties in the forged disk. A simulation of the machining process results in the final disk shape with accurate residual stresses and distortions. If the distortions are within acceptable limits, an accurate life prediction of the part will be performed. Otherwise, the heat treatment or forging process need to be improved in order to achieve acceptable distortions. If that is not possible, the finished disk shape needs to be changed and the mechanical design - at least in parts - be repeated. The same applies in the case that the design does not meet life requirements.

As this description demonstrates, an integrated procedure that addresses both mechanical design and manufacturing processes is absolutely necessary because of the iterative nature of the process and the prohibitive costs involved if changes become necessary once actual parts are being produced. Simulation tools for each individual stage are available and widely used. But opportunities for mathematical optimization of the individual process steps are currently not fully utilized, and an integrated procedure which is the ultimate goal of this research is missing altogether.

If we try to recast the disk design problem in the form of a formal MDO problem, weight can be considered as the overall system objective, and the different objectives of some of the individual subsystems can be formulated as constraints. Weight here would be the billet weight of the forging, which, of course, also includes the weight of the final part, both of which need to be minimized. Since the forging billet weight is inherently much larger that the final part weight, a linear combination of the two in the following form could be considered as the system objective:

$$F = a\ W_{final} + (1 - a)\, W_{billet} \qquad (1)$$

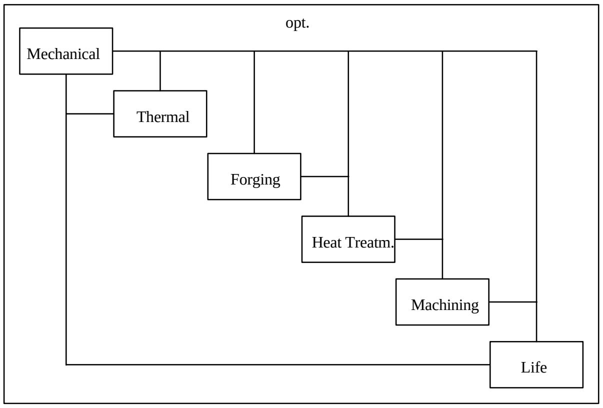

Fig 2: Multi-Discipline Feasible Formulation

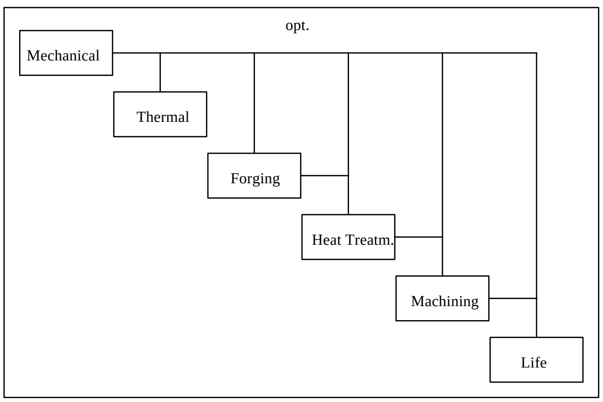

Fig 3: Individual Discipline Feasible Formulation

In a multidiscipline-feasible type scenario (fig. 2), each of the disciplinary analyses would contribute a number of constraints to the system level optimization problem. In an individual discipline-feasible type scenario (fig. 3), the feedback loops from life analysis to mechanical design and from thermal analysis to mechanical design are severed at the cost of additional system level constraints accounting for the interdisciplinary discrepancies introduced. Both of these standard formulations are not very satisfactory for this type of problem out of several reasons. First, a large number of system level constraints would be introduced which are purely disciplinary in nature. It makes no sense for the system level optimizer to be bothered with all the intricacies of the forging optimization problem, for example. Additionally, the heat treatment problem is an optimization problem in itself, but it does not directly contribute to the overall objective, weight, but rather addresses producibility and the satisfaction of constraints for distortion and material properties. Therefore, the disk design problem calls for an approach where the optimization itself is distributed, and where each disciplinary optimization problem does not necessarily contribute directly to the overall objective. Both the Concurrent Subspace Optimization (CSSO)[3] and the Collaborative Optimization (CO)[4] methods have been looked at as possible solutions, but it seems that neither one of them really captures the salient features of the disk design problem. CSSO assumes a common objective that each discipline is somehow contributing to, and requires an approximation of the non-local states in each discipline. This means one would have to create an approximation of the forging problem inside the heat treatment problem and so on, which is not very practical. CO introduces artificial non-physical objectives for the disciplinary optimization problems so that for the designer it is somewhat difficult to follow the progress of the optimization from a disciplinary point of view. Besides, slow convergence rates in conjunction with long analysis times (in the order of several hours per analysis for the forging problem, for example) render this approach impractical. Therefore, it seems that in this case a sequential optimization within an integrated framework seems most promising (fig. 4), where we start with the mechanical design problem and simple 1-D and 2-D axisymmetric tools to obtain an initial disk shape, use this to design a near net shape forging process, then optimize the heat treatment process, and finally perform the machining and life analyses.

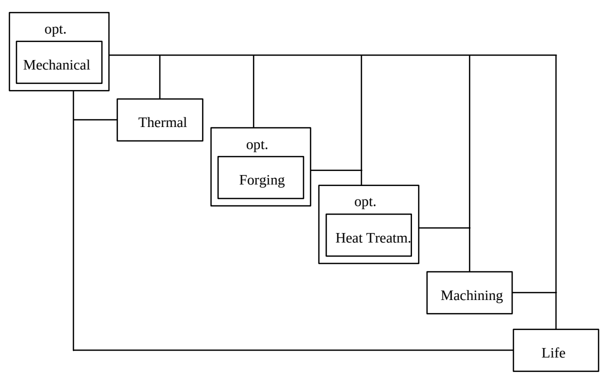

Fig 4: Sequential Optimization Approach

In subsequent loops, full 3-D analysis tools are used in the mechanical design phase. This sequential approach is possible in this specific case because the near-net shape forging optimization will not compromise the thermo-mechanical minimum weight design, and the optimum heat treatment process has no influence on either one of the two.

Simulation and Optimization Tools

Optimization Framework

The basic framework for the optimization environment under development is iSIGHT, a software product developed and marketed by ESI. iSIGHT is conceptually a follow-up product to Engineous[5,6], the optimization framework developed at GE CR&D during the 1980s. Both products facilitate easy integration of both commercial and company proprietary software into an overall optimization environment which makes use of the concept of interdigitation[7] where the user has a suite of optimization tools available, including gradient based and heuristic search techniques, genetic algorithms, and simulated annealing, which can be used in any combination during the optimization process. Experience over the years has shown that one optimization strategy alone is often unable to solve a problem, but that a combination of different strategies leads to improved results. iSIGHT enables the user to formulate a sequence of different optimizers and then apply this sequence to the optimization problem. Another strong point of iSIGHT is the ease with which analysis programs can be integrated into the framework, including large-scale engineering applications like finite element codes. These codes can reside on their respective platforms, irrespective of where iSIGHT is installed, an important point with respect to software leasing and maintenance cost for software which may be licensed only on a certain workstation. A number of toolkits are under development in conjunction with the RaDEO project, among them the Product Modeling Toolkit (PMTK) to support product data modeling and the interaction with commercial CAD software, and the Discrete Analysis Modeling Toolkit (DMTK), which facilitates the interaction of analysis models of different disciplines and levels of fidelity. Both of these toolkits are heavily leveraged in the engine disk design scenario.

Process Simulation

Nowadays advanced process simulation tools are becoming more and more available for all stages of the disk design and manufacturing process. Simulation tools such as DEFORM[8] and ABAQUS can accurately predict the mechanical behavior and properties during the manufacturing process. Therefore, these tools have become the state-of-the-art and are widely used. In combination with numerical optimization techniques, these tools offer the opportunity to improve individual steps in the overall process[9]. In this application, DEFORM was chosen as the tool to be applied in the forging and heat treatment optimization procedure.

CAD Tool

The CAD tool of choice is Unigraphics[10], developed by EDS, which is the adopted CAD software at GE Aircraft Engines. Parametric master models control the geometry and engineering analysis “views” which support analyses at different levels of fidelity. These analysis-“views” - or context models - are de-featured models capturing the essential geometry for the respective analysis. They can also contain additional information necessary to generate the analysis models like boundary conditions and load and mesh information. PMTK will allow the user to graphically pick geometric design variables from the CAD model and automatically link them with the optimizer for topology optimization.

Mechanical Analysis

Finite element analyses are performed using ANSYS, with model preparation done partly in ANSYS and partly in PATRAN. PATRAN’s P/THERMAL module supplies the required heat transfer data and temperature boundary conditions for the stress analyses. Different approaches are being evaluated for automatic analysis model generation from the CAD representation. One is the use of “tags” in the CAD model, where the CAD model would house all the information necessary to generate the analysis model. Another is the use of scripts that are reusable. This approach relies on a constant topology of the geometric model and entity consistency of the geometry import into the CAE tool.

Current Status

MDO Algorithms

Implementations of both the CSSO and CO algorithms inside iSIGHT have been developed and tested[11,12]. Since convergence of the CO algorithm tends to be very slow, its usefulness in detailed design applications requiring high-fidelity engineering analysis remains doubtful. One promising possiblity is the combination of the CO algorithm with response surfaces in order to reduce the number of analyses at the subsystem level. An initial implementation of the CSSO algorithm has been validated both with standard textbook-type example problems and with two more realistic problems representing a welding design and an idealized turbine blade. This CSSO implementation is currently being evaluated in connection with the disk design problem.

Forging Shape Opitmization

A procedure has been developed to address the forging shape optimization problem, integrating Unigraphics and DEFORM with iSIGHT, leveraging functionality of the product modeling toolkit. Reference 13 describes the system in greater detail than is warranted here. The objective of the forging shape optimization problem is the design of a minimum weight forged shape that satisfies constraints on both forging press capacity, strain and strain rates, die filling, and minimum coverage of the final part shape.

In the present study, forging is modeled as a time-dependent, plastic-deforming, either isothermal or non-isothermal process. Since the forging simulation is conducted in an optimization environment, some of the process and geometry parameters are modified in each DEFORM run. Therefore, it is necessary to regenerate the mesh and redefine the boundary conditions. Furthermore, it is necessary to post-process the analysis results and extract information on optimization objective and constraint functions. Several modules have been developed that drive DEFORM to accomplish following tasks:

- import geometry and regenerate die and billet meshes,
- create appropriate boundary conditions,
- start DEFORM simulation in batch mode,
- monitor DEFORM runs, and
- postprocess simulation results to extract maximum press load, strain, temperature, etc.

Each of these modules acts like a separate executable, or "simcode" in iSIGHT terminology. iSIGHT executes these "simcodes" in a pre-defined sequence, including potential looping and branching.

Consider the forging shape optimization of a generic turbine disk. A cylindrical billet is forged into a disk of the shape shown in figure 5. The die geometry is captured in a Unigraphics parametric model. Several fillet radii R_1-R_6 have been chosen as design variables. Both invalid geometry and intrusion into the minimum coverage over the so-called "shipped" shape, the intermediate shape in which the forging vendor supplies the part, and which is used for testing purposes, can be prevented by putting simple bounds on the design variables. It should be noted that simple bounds may not be sufficient to guarantee geometric validity in a more general situation. They work in this case because there is no coupling among the selected design variables.

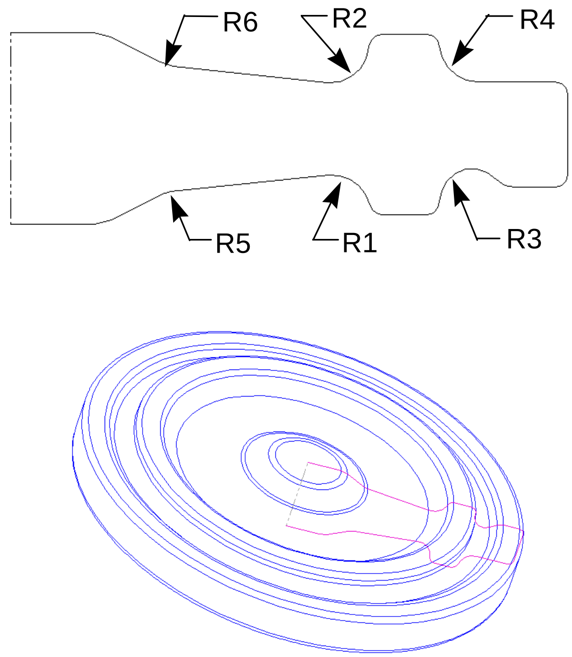

Figure 5: Turbine disk and its cross-section

Thus, the optimization problem is formulated as

$$\begin{aligned} &\min \quad V, \\ &s.t. \quad R_{ilb} \le R_i \le R_{iub} \quad (i = 1, \ldots, 6), \\ &\qquad\; P \le P_{ub}, \end{aligned}$$

where V is the volume of the work-piece, P and P_{ub} are the maximum press load and its upper bound, respectively, and R_{ilb} and R_{iub} are given lower and upper bounds of the fillet radii, respectively. The most aggressive shape, which corresponds to the lowest volume V, has been chosen as the initial design. It is relatively easy to get this shape from the specified disk design by adding a minimum cover. However, the press load constraint is usually violated for this design, and thus the fillet radii R_i have to be increased which results in a larger volume. Subject to the press load constraint P_{ub}, the optimizer will choose the optimal values of design variables R_i . As it is pointed out in the previous section, iSIGHT provides a suite of optimization algorithms. The modified method of feasible directions from ADS[14] is employed in this study. Since analytical design sensitivities are not available, the gradient information has to be obtained through finite differencing.

In the example we consider a time-dependent, plastic deforming, isothermal, closed-die forging process. The top and bottom dies are assumed to be rigid. The maximum load *P* normally occurs at the end of the forging stroke as the dies fill out and the material starts to move into the flash region. The load changes rapidly with the stroke at this stage of the process. Therefore, it is difficult to accurately compare the loads at the end of the stroke from different die designs due to the inherent noise in the load predictions. For this reason we artificially set *P* to be the stroke-averaged load in between 98-99% of the final stroke. A good estimate on the real maximum press load may be obtained by multiplying *P* with a correction factor.

There are about 1600 quadrilateral elements on the workpiece, and automatic mesh regeneration is enabled to accommodate the large deformation that is inevitable in the forging process. Four design variables R_1 - R_4 are used in this application, and the time step is taken to be $\Delta t = 0.1$s. Due to repeated remeshing during the forging simulation, non-smoothness is introduced in the finite element solution. Therefore, we chose a 10% perturbation on the design variables during sensitivity analysis using finite differencing to smooth out the design space. Although the design sensitivities so calculated may not be very accurate locally, they provide the optimizer with the right search directions in a global perspective.

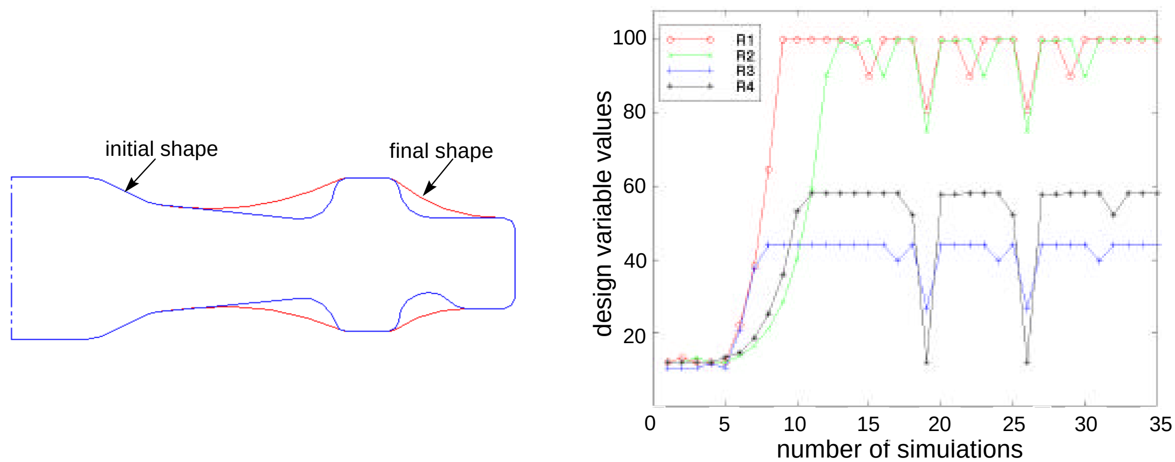

Figure 6: Initial and final shapes (left), and design variable values versus the number of simulations (right) for isothermal the forging process

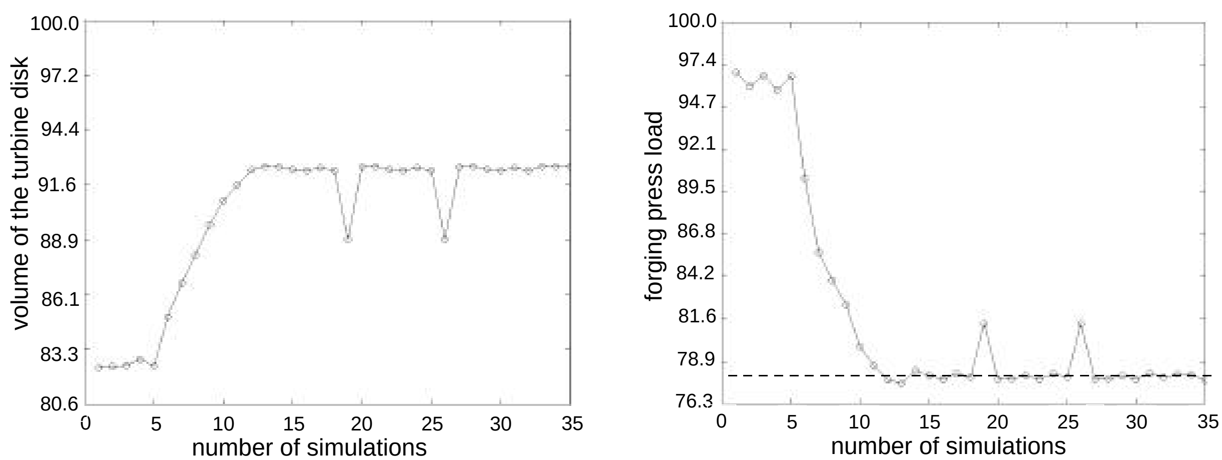

Figure 7: Disk volume and press load vs. the number of forging simulations for the isothermal forging process

The initial and final shapes of the disk are shown in Figure 6 (left). The history of design variable values against forging simulation runs is given in Figure 6 (right). The objective (volume) and constraint (press load) function values versus simulation runs are shown in Figure 7. All numbers have been normalized. The results suggest that the optimization is close to convergence after 15 simulation runs. There are some downward spikes in the figures. The smaller ones are the result of finite difference perturbation, while the larger ones are due to line search of the optimizer. Since the abscissa shows the number of simulation runs as opposed to the number of optimization iterations, the results of both finite differencing and line search have been included. The upper bound of the press load $P_{ub} = 77.8$ is shown as a dashed line in Figure 7 (right). It is apparent that the forging press load far exceeds this limit initially. As a result of the optimization, the press load drops from 96.9 of the initial design to 77.8, which is the upper bound, a 19.7% reduction. The volume, however, has been increased by 12.4% from 82.4 of the initial minimum-weight shape to 92.6 of the final optimized shape. In addition to the single step process described here, a multi-step forging process is presented in ref. 13.

This work can be extended in several aspects: first, new interprocess communication mechanisms may be introduced to improve data passing between processes; second, a more comprehensive forging simulation should be conducted that includes the effect of heat loss during transport of the billet and positioning of the tools; third, a larger design space may be explored by incorporating more geometric parameters as design variables; finally, additional constraints, such as those on strain rate and temperature, should be considered to model more realistic situations.

Heat Treatment Optimization

The purpose of the heat treatment process is to develop the necessary mechanical properties in the forged part. This is achieved by heating the part to solution temperature and then cooling it rapidly. During the cooling phase residual stresses are introduced. In the case of Ni-based superalloys that are considered here, a certain minimum cooling rate has to be maintained to generate the needed creep and tensile properties. On the other hand, the faster the cooling process is, the higher are the resulting residual stresses which can lead to excessive part distortions after machining to the final shape.

Traditionally, an oil quenching process has been employed which ensures fast cooling and thus a high cooling rate, but the oil quenching process introduces high residual stresses, and, from a process optimization point of view, offers very little room for improvement as there are very few parameters which can be controlled. Therefore, fan cooling is gaining larger acceptance where it is possible to control the airflow on individual sections of the part and thus influence the local surface heat transfer coefficients. Obviously, the heat transfer coefficients that can be achieved with fan cooling are lower than those for oil quenching, so that for thick parts it may not be possible to satisfy cooling rate requirements, but for moderately thick parts fan cooling offers clear advantages. For very thin parts like engine seals, where machining distortions due to residual stresses are especially critical, fan cooling may be the only process that produces acceptable parts at all.

The challenge here is to formulate a fan cooling optimization problem without actually having to execute a combined heat transfer-stress analysis each time the optimizer needs a new design point. An accurate heat transfer analysis requires small time steps in the simulation, and a stress analysis, in turn, requires a fine finite element grid, therefore the combination of both is the most computationally expensive analysis possible. In general, though, the stress analysis is much more time consuming than the heat transfer analysis alone. Since it is known that spatially uniform cooling reduces residual stresses, the idea is to formulate an objective function that penalizes non-uniform cooling and at the same time ensures fast cooling at or above the target cooling rate. These are obviously two conflicting objectives since fast cooling always means uneven cooling as the heat can only be extracted at the surface of the part. Therefore, the objective function for the heat treatment optimization problem is formulated in a quadratic form that penalizes the deviation from the cooling rate target:

$$obj = \sum_{nodes} \begin{cases} w \cdot \left(\bar{\dot{t}}_{\text{target}} - \bar{\dot{t}}\right)^2 & if \;\; \bar{\dot{t}} < \bar{\dot{t}}_{\text{target}} \\ (1-w) \cdot \left(\bar{\dot{t}}_{\text{target}} - \bar{\dot{t}}\right)^2 & if \;\; \bar{\dot{t}} \geq \bar{\dot{t}}_{\text{target}} \end{cases} \quad (2)$$

W is a user-defined weighting factor between 0 and 1 that penalizes under- and over-achievement of the target cooling rate differently. A value close to one (but less than one, of course) seems to give the best results. The target cooling rate is also a material-dependent value. Design variables are the surface heat transfer coefficients, h_i, which can be related back to a certain airflow produced by the fan cooling apparatus.

A total number of up to ten or twelve design variables seems to be in the range of what can be controlled by current fan cooling fixtures. The procedure developed here gives the user a choice in terms of optimization constraints. He can impose a hard constraint on the cooling rate:

$$c_{cr} = \frac{1}{nodes} \sum_{nodes} \begin{cases} \bar{\dot{t}}_{\text{target}} - \bar{\dot{t}} & if \;\; \bar{\dot{t}} < \bar{\dot{t}}_{\text{target}} \\ 0 & if \;\; \bar{\dot{t}} \geq \bar{\dot{t}}_{\text{target}} \end{cases} \quad (3)$$

This constraint has a discontinuity at zero, exactly where it is active, and will never assume a value less than zero, that is satisfied and not active, caused by the "if" in the constraint formulation. This discontinuity leads to problems with gradient-based optimizers, which will always see a zero constraint gradient for a satisfied or active constraint, therefore in the case of constraint satisfaction the constraint value of zero is replaced with the difference of the target cooling rate and the minimum of all nodal cooling rates:

$$c'_{cr} = \bar{\dot{t}}_{\text{target}} - \bar{\dot{t}}_{\text{min}} \quad (4)$$

In this fashion at least the sign of the constraint gradient that the optimizer sees above and below a constraint value of zero will be equal. An additional constraint can be placed on the nodal fraction that fulfills the cooling rate target which has to be equal to 1.0 if the target is met everywhere. The two constraints may seem somewhat redundant, but depending on the optimization strategy used, one or the other or a combination of both lead to the best convergence.

The heat treatment optimization procedure described above was applied to a generic turbine disk. Figure 8 shows the heat treatment geometry and the distribution of the nine design variables employed.

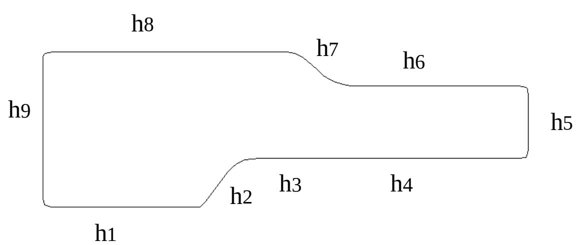

Fig. 8: Turbine Disk Geometry and Fan Cooling Variables

In order to cut down on analysis time, the optimization was started with all heat transfer coefficients linked to only one design variable. This problem was executed for six iterations, using the sequential linear programming technique from ADS inside iSIGHT, until both constraints were active. The full convergence history is depicted in figure 9, and the constraint history is shown in figure 10. Negative constraint values indicate a satisfied constraint.

At this point, all nine design variables were activated, and the new optimization problem converged within seven more iterations, that is 13 total. For this segment, the modified method of feasible directions, also from ADS, was chosen as the optimization technique. The deviation function was initially reduced from a value of 1.4 to about 0.6 and then further down to under 0.2. These numbers as such have no physical meaning, but the significance can be seen in a comparison of the initial and final cooling rate distribution (fig. 11 and 12), normalized with respect to the target value, indicating a much more uniform cooling than at the starting point.

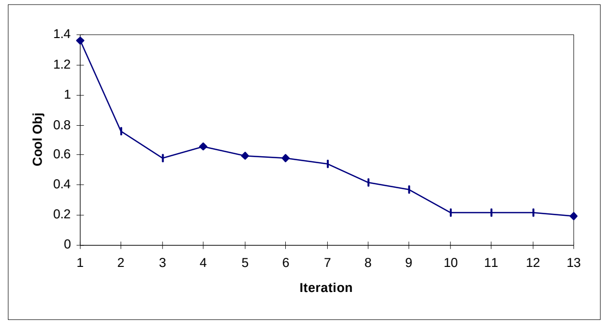

Fig. 9: Objective History

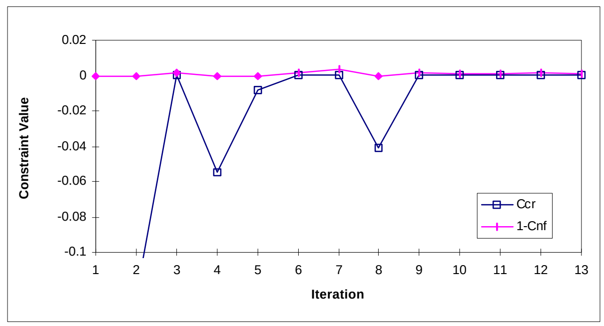

Fig 10: Constraint History

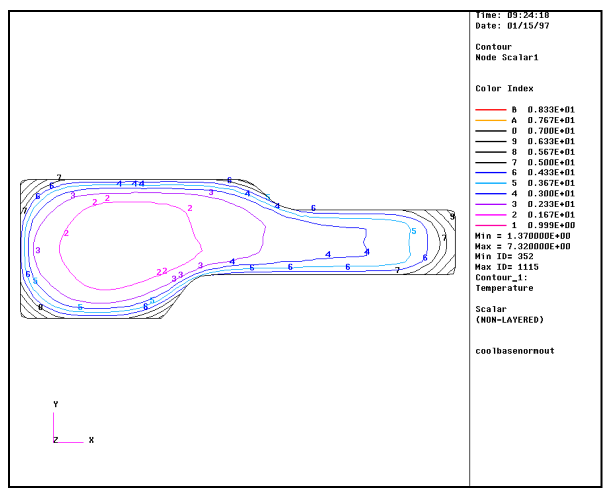

Fig. 11: Initial Cooling Rate Distribution

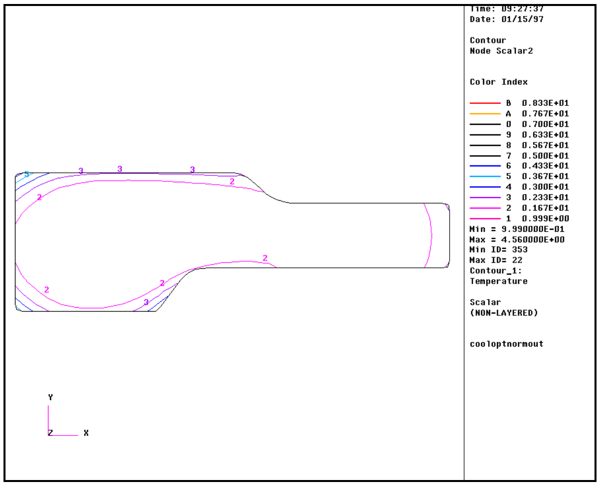

Fig. 12: Final Cooling Rate Distribution

The question still to be answered is what effect this optimization procedure, which is based on heat transfer analysis only, has on the residual stresses of the part which is what we are ultimately interested in. Therefore, a combined heat transfer-stress analysis was performed on both the starting configuration and on a disk with the final heat transfer coefficient distribution. For comparison purposes, an analysis of a typical oil-quenching process was also performed.

Fig. 13 through 15 show the resulting hoop stresses for the three cases, all normalized with respect to the maximum tensile stress of the oil-quenched part. The stresses are highest for the oil-quenched disk, closely followed by the non-optimized fan-cooled disk with uniform high fan blowing all around. The residual stresses for the optimized disk, in turn, are considerably lower, almost by one order of magnitude compared to the oil-quenched part in terms of tensile stresses. The reductions in compressive stresses are not quite that large, but still by a factor of between six and seven. These results clearly show the advantage of a numerically optimized fan cooling process compared to the traditional oil-quenching. Ref. 15 describes the heat treatment optimization process in greater detail. These findings were confirmed during multiple runs with different starting points on actual geometries which are of proprietary nature and cannot be shown here. The formulation of the objective function as a quadratic clearly aids in this behavior.

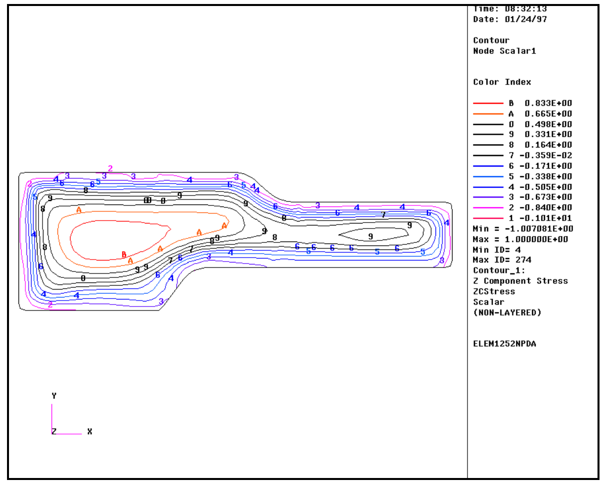

Fig. 13: Hoop Stress, Oil-Quench Process

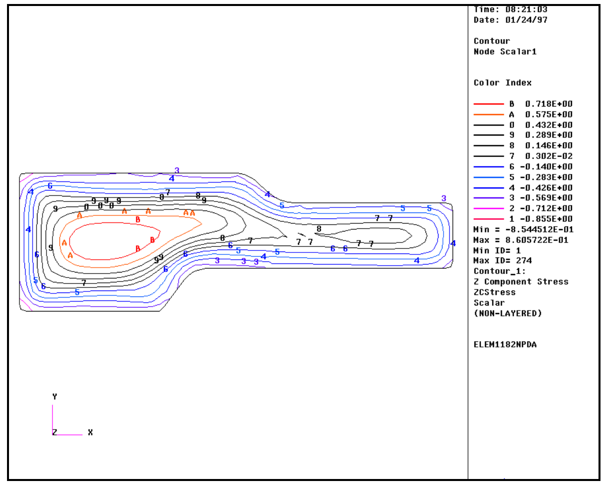

Fig. 14: Hoop Stress, Starting Point

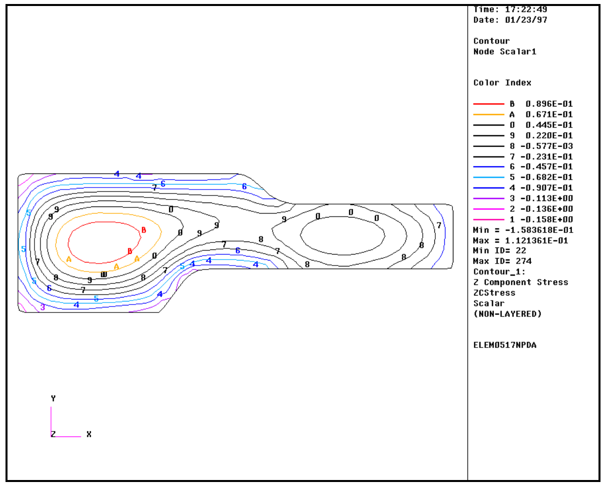

Fig. 15: Hoop Stress, Final Configuration

Thermo-Mechanical Design

The mechanical design and engineering analysis portion of the integrated process is currently lagging behind the efforts on the processing side. This has several reasons, one being that the development of fully parametric master models proved to be more time-consuming than anticipated. But a major bottleneck is the automatic generation of high-fidelity finite element analysis models complete with loads and boundary conditions which update with parametric changes of the model. Several pilot projects have been ongoing since the last year, evaluating different concepts of relating analysis-related information to the geometry. One approach is the "tagging" of the geometry, applying basically CAE-type information on the geometry on the CAD side. A major hurdle here is reluctance from the side of analysis engineers who rather want to work within their CAE tool of choice instead of the CAD system. Also, the processing of "tags" inside the CAE tool has proven not to be very robust. Another approach is the use of scripts for the CAE tool, where the engineer prepares the model once manually and then saves the session log file for subsequent reuse. This approach demands entity-consistent import of the geometry from the CAD tool into the CAE system, which again is not robust at the moment. This approach certainly breaks down in the presence of topological changes. Before the issue of reliable, repeatable automatic generation of analysis models for complex 3-D-geometries is resolved, any effort to use optimization on the mechanical design side beyond conceptual studies is premature.

Outlook

The plan is to complete one full manufacturing process exercise by the end of the year. How fast the developments on the mechanical design side will be able to catch up remains to be seen and depends largely on external factors beyond GE CRD's control. In order to reduce analysis times for the forging optimization, the use of approximate models and response surfaces will be investigated. The machining simulation will be integrated with the heat treatment optimization package, so that the final machining distortions will be available automatically without manual intervention. Once the system is in place for the complete process simulation and process optimization, the question of the applicability of formal MDO algorithms will be revisited.

In parallel, various strategies will be further investigated on how to capture analysis model information and make it reusable in a robust fashion so that analysis models for complex geometries will finally automatically update without human intervention. Once this obstacle has been cleared, 3-D-shape optimization during the mechanical design phase can be addressed, probably initially limited to relatively simple features comparable in complexity to the 2-D-forging shape optimization discussed earlier.

Computer resources continue to be a problem in conjunction with the long analysis times required for the solution of industrial size problems. A forging simulation as it is considered here may take 6 to 7 hours on an SGI workstation. Finite differencing could potentially be done in parallel, but there are the problems of software licensing and maximum number of processes one user is allowed to run at any given time. It seems clear that the computer resource issues will remain a major bottleneck for the application of MDO to industry problems.

One of the highlights so far in this project has been the optimization and integration framework itself, iSIGHT, which has performed very well, although still under development. It fits with GE's paradigm shift away from proprietary software development to the use of commercially available CAD and CAE tools, which require a loose and non-intrusive coupling of the individual analysis modules.

References

[1] A Collaborative Optimization Environment (COE) for MADE, Technical Proposal, General Electric Corporate Research and Development and Optimum Technologies, Inc., August 1995

[2] iSIGHT Version 3.0 User Manual, Engineous Software, Inc., Raleigh, NC, 1997

[3] Wujek, B.; Renaud, J.E.: *Design Driven Concurrent Optimzation in System Design Problems Using Second Order Sensitivities*, 5th AIAA/NASA/USAF/ISSMO Symposium on Multidisciplinary Optimization, Panama City, FL, September 1994

[4] Braun, R.D.; Kroo, I.M.: *Development and Application of the Collaborative Optimization Architecture in a Multidisciplinary Design Environment*, Aug. 1995

[5] Engineous User Manual, General Electric Corporate Research and Development, Schenectady, NY, 1995

[6] Lee, H.; Goel, S.; et al.: *Toward Modeling the Concurrent Design of Aircraft Engine Turbines*, Presented at the International Gas Turbine and

Aeroengine Congress and Exposition, Cincinnati, OH, May 1993

[7] Powell, D.J.: *Inter-GEN, A Hybrid Approach to Engineering Optimization*, General Electric Technical Report, Feb. 1991

[8] Design Environment for Forming (DEFORM), Version 5.0, Online Users Manual, Scientific Forming Technologies Corporation, Columbus, OH, August 1997

[9] Kumar, V.; German, M.D.; Srivatsa, S.: Design Optimization of Thermomechanical Processes with Application to Heat Treatment for Turbine Disks, Presented at the Manufacturing International Conference, Atlanta, GA, 1990

[10] Unigraphics V13 Online User Documentation, EDS, Cypress, CA, Oct 1997

[11] Tappeta, R.; Nagendra, S. et al.: *Concurrent Sub-Space Optimization (CSSO) MDO Algorithms in iSIGHT*; GE CRD Technical Report, January 1998

[12] Conversations with S. Kodiyalam, Engineous Software, Raleigh, NC

[13] He, B.; Röhl, P.J. et al.: *CAD and CAE Integration with Application to the Forging Shape Optimization of Turbine Disks*. To be Presented at the 39th AIAA/ASME/ASCE/AHS/ASC Structural, Dynamics, and Materials Conference, Long Beach, CA, April 1998

[14] Vanderplaats, G.N.: ADS - A FORTRAN Program for Automated Design Synthesis, Version 3.00, VMA Engineering, March 1988

[15] Röhl, P.J.; Srivatsa, S.K.: *A Comprehensive Approach to Engine Disk IPPD*. Proceedings of 38th AIAA/ASME/ASCE/AHS/ASC Structural, Dynamics, and Materials Conference, pages 1250-1257, Kissimmee, Florida, April 7-10, 1997