SPACE SHUTTLE EXTERNAL FUEL TANK DESIGN OPTIMIZATION†

Massimo Usan*
Massachusetts Institute of Technology
Engineering System Division
Sloan School of Management, School of Engineering

ABSTRACT

A design optimization study of the Space Shuttle External Fuel Tank (SSEFT) is performed using a model that, although simplified, captures some of the important aspects of the system's attributes and behavior. The goal of the optimization is to determine the values of the geometric characteristics of the system that maximize the ROI of the project and the payload that the Shuttle orbiter can carry. The process is articulated in several steps. First, a preliminary design exploration is performed using a DOE technique. Then an optimization for a single objective (ROI) follows. Both gradient-based methods and heuristic techniques are applied, the results of the different methods compared and pros and cons for each of the techniques, in the context of the specific application, highlighted. Finally a full multi-objective analysis is conducted.

Within the limitations of the model, the problem is shown to exhibit a modest coupling. ROI is particularly sensitive to the geometry of the cylindrical portion, namely length, radius and thickness as well as the stresses that insist on this part.

Evidence is that a sequential quadratic programming algorithm outperforms a genetic algorithm both in execution time and results. The optimization end values are shown to have a good agreement with actual data. Finally the payload is demonstrated not to be a separate objective from ROI.

Optimization is viewed not only as a method to determine the best design, but, equally importantly, as a process through which understanding the system under investigation and unveiling its intrinsic trade-offs.

INTRODUCTION

The Space Shuttle External Fuel Tank (SSEFT) is the largest single element and the only major non-reusable component of the Shuttle system. The SSEFT is 46.88 meters long, 7.8 meters in diameter and its empty weight is 29,937 kg. It carries more than 2 million liters of cryogenic propellants that are fed to the orbiter's three main engines during powered flight (see Fig. 1).

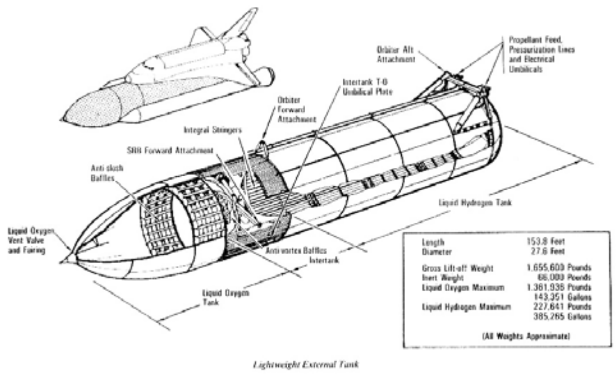

Fig. 1: The Space Shuttle External Fuel Tank

When the main boosters are shut down, the SSEFT is jettisoned, enters the Earth's atmosphere, breaks up, and impacts in a remote ocean area. It is not recovered.

The SSEFT has three major components: the forward liquid oxygen tank, an unpressurized intertank that contains most of the electrical components, and the aft liquid hydrogen tank [1] [2] .

A simulation model is built that even though rather simplified, is able to provide quantitative responses to complex aspects of the system, such as the stress levels in the structure, its resonance frequency, its aerodynamic drag and system cost. The model is then used to perform a design optimization with the goal of determining the tank's geometric characteristics that maximize the ROI of the tank and the payload of the Space Shuttle for a given amount of carried fuel. Requirements are also put to maximum allowed stress levels in the structure and to the minimum natural frequency of the system.

The optimization process is carried on in three main steps: 1) design space exploration, 2) single-objective optimization, 3) multi-objective optimization. The

*Aeronautical Engineer, System Design and Management Fellow optimization is eased by the use of a commercial optimizer package, iSIGHT v.07 Academic from Engineous [3] . The step-by-step approach allows not only the otherwise uncertain success of the optimization process, but also an insightful investigation of the existing trade-offs.

In the description, a managerial perspective is taken, so the emphasis of the exposition is neither on the details of the model, nor on the underlying mathematical / numerical implementation. The accent is instead rather on the design process as a whole and on the valuable insights that a rigorous and scrupulous optimization analysis can bring and which could be key to a "better" product. For this reason, the depth is sometimes sacrificed in favor of the focus on the big picture.

MODEL DESCRIPTION

Overview

The model, although with several simplifying assumptions, is conceived so as to provide a quantification of the main important system characteristics. It consists of more than sixty equations with six design variables and more than thirty parameters (considered fixed for the analysis in hand). In what follows each of the different modules is described succinctly.

Model Mechanics

The tank (see Fig. 2 for a close look) is modeled as composed by three main sections: a cone, a cylinder and a hemisphere (Fig. 3). The design variables used to describe the system are:

- length of the cylinder: L
- cylinder radius: R
- cone height to radius ratio: h/R
- skin thicknesses: of the cylinder (t_{cyl}), cone (t_{cone}) and hemisphere ($t_{h\text{-}sphere}$).

Fig. 2: Close view of the SSEFT

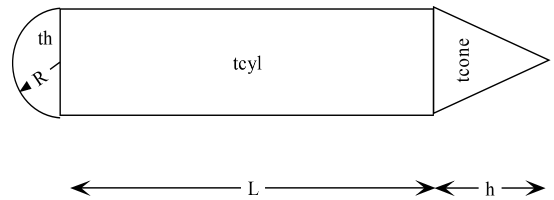

Fig. 3: schematic of the SSEFT

The identified variables are used to calculate several characteristics of the system and of the project.

Geometry and Weight

Using basic geometry formulae, the volume enclosed by the three elements is derived as well as the volume of structure material. Weight of the structure is then inferred applying material density (material assumed for the calculation is Aluminum).

Structural behavior

Stresses are calculated as resulting from pressure of the gas contained in the tank, within a thin membrane vessel approximation. Formulas available in the literature are used (see, for example, [4]). Stress factors are then obtained by dividing the resulting stress by the maximum allowable stress of the material.

When the first natural frequency of the tank is lower than a certain value, the structure can be subjected to dangerous instability phenomena. For that reason, this aspect has been included in the structural model. The frequency is calculated assimilating the tank to a beam and using simple structure formulae available in the literature [4] .

Cost

In order to keep the model manageable in size, the cost of the task is calculated simplistically as the sum of material cost and the cost of the welds that connect the tank's three main components. Empirical relationships are used that link material cost to material weight and skin thickness. The cost of the seams is calculated as the seam unit cost (known empirically) multiplied by seam length calculated using linear geometry.

Aerodynamics and Payload

The tank aerodynamic drag represents an additional force that adds to the Shuttle weight and contrasts the thrust provided by the engines. All else being equal, the lower the drag, the higher the payload (weight of the equipment for scientific research) the Space Shuttle is able to bring to orbit. As such, it is an important performance metrics of the project. A semi-empirical formula is used to compute the aerodynamic drag of the tank that takes into account its wetted area, the cross sectional area of the cylindrical section and the bluntness of the cone. The drag is actually calculated as difference from the value of a reference tank and added or subtracted to the related nominal payload yielding the actual payload.

Project Financials

Return On Investment (ROI) is used to evaluate the "health" of the project and as a main objective for the optimization. ROI is calculated as:

$$\frac{P - C_{launch}}{C_{launch}}$$

Where P is the price the Customer is willing to pay and is linearly dependent on the allowable payload (the linear coefficient is assumed known), whilst C_{launch} is the total cost to launch the Space Shuttle, calculated as $C_{tank}+C_{other\ than\ tank}$; $C_{other\ than\ tank}$ is based on existing data and considered a parameter for the analysis.

THE OPTIMIZATION PROCESS

Step0: Problem Formulation

The optimization problem is formulated as follows:

Select $x = [L, R, h/R, t_{cyl}, t_{cone}, t_{h\text{-}sphere}]$

In order to Maximize *ROI(x)& Payload(x)*

subject to the following constraints:

$$g_v = \frac{V}{V_{req}} = 1$$

$$g_{cyl} = \frac{\sigma_{cyl}}{\sigma_{\max}} \leq 0.8$$

$$g_{cone} = \frac{\sigma_{cone}}{\sigma_{\max}} \leq 0.8$$

$$g_{h-sphere} = \frac{\sigma_{h-sphere}}{\sigma_{\max}} \leq 0.8$$

$$g_{vib} = \frac{f}{f_{\min,req}} \geq 1$$

In plain English, this reads as follows. Select the values of the design variables in order to maximize the Return of Investment of the Project and the Payload, while ensuring that the volume of liquid that the tank is able to carry equals the one required (g_v=1), that the stresses are 20% lower than the maximum stress allowable by the material (g_{cyl}, g_{cone}, $g_{h\text{-}sphere}$ <0.8), and that the first natural bending frequency exceeds a certain minimum required value (g_{vib}>1).

Step 1: Design Space Exploration

A preliminary exploration of the design space is performed applying a Design of Experiment (DOE) technique to the model described. The goal of this phase is manifold: 1) to understand the degree of coupling of the system, 2) to evaluate how and to which design variables the objectives and the constraints are sensitive; 2) to be able to set a starting point for the subsequent optimization analysis.
For each of the six factors (i.e. the design variables), three levels are selected to capture the curvature of the effect. Detail of the settings are in **Tab. 1**.

Design Variable	L [cm]	R [cm]	t_{cyl} [cm]	t_{cone} [cm]	$t_{h\text{-}sphere}$ [cm]	h/R
Min	2000	250	0.1	0.1	0.1	1
Med	4000	500	0.2	0.2	0.2	2
Max	6000	750	0.3	0.3	0.3	3

Tab. 1: Variables Arrangement for DOE

Several arrays with different strength levels are tried from a L27 to a L243. The system, however, is found to be relatively uncoupled and the L27 array is found to offer the best accuracy vs. computational costs.

In fig Fig. 4-Fig. 7 a selected sample of the results is shown. It can be noted that some of the dependencies are quadratic and some of them are linear. If the curvature were big the linearization resulting from the use of two levels only could lead to significant misinterpretations, that's why 3 levels

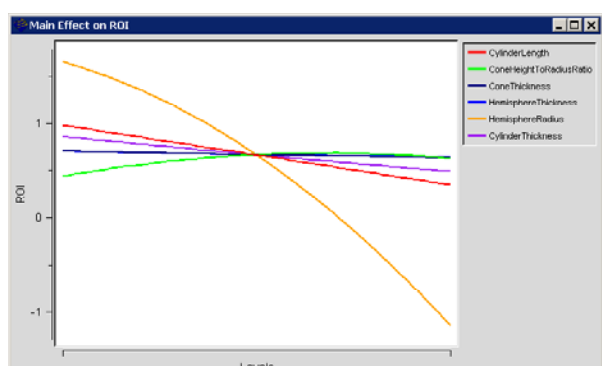

Fig. 4: Main Effect on ROI

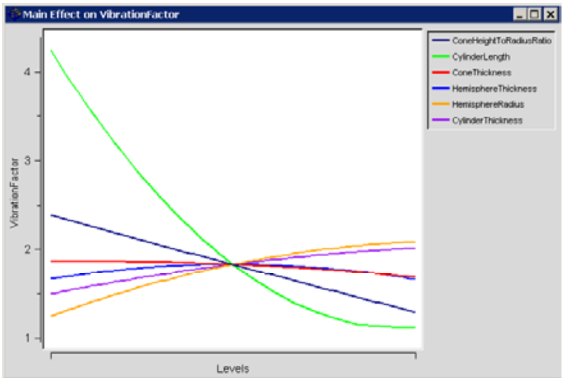

Fig. 5: Main Effect on Vibration Factor

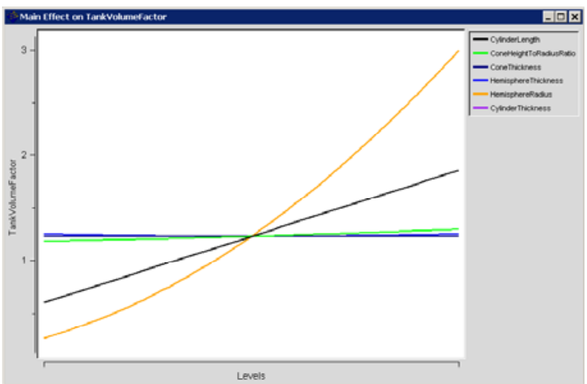

Fig. 6: Main Effect on Tank Volume Factor

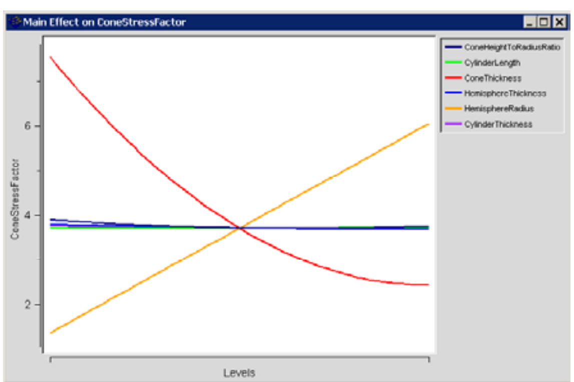

Fig. 7: Main Effect on Cone Stress Factor

The main findings of the analysis are as follows. The top three factors that influence the different objectives / constraints are (in order of importance): 1) tank radius, 2) cylinder length and 3) cylinder thickness. As confirmed by the sensitivity analysis (not reported for brevity's sake) the second order effects are negligible, so the results obtained with the orthogonal arrays have high significance.

The results of the DOE are used to identify a suitable initial design vector for the following optimization phase. The final selection is shown in Tab. 2.

Design Variable	L [cm]	R [cm]	t_{cyl} [cm]	t_{cone} [cm]	$t_{h\text{-}sphere}$ [cm]	h/R
Value	5000	350	1	1	1	2

Object. Constr.	g_v	g_{cyl}	$g_{h\text{-}sphere}$	g_{cone}	g_{vib}	ROI [%]	PL [kg]
Value	0.72	0.53	0.61	0.53	0.93	10.3	31,506

Tab. 2: Design Variables Setting and Objective / Constraints after DOE phase

Step 2: Single Objective Optimization

Multi-objective optimization is often delicate and complex, and implies that careful trade-off decisions be made along the process. Before adventuring in such an intricate domain, a process where the design is optimized for a single objective function is pursued. In addition to the obvious simplification that arises, this approach has also got the advantage to allow identifying the maximum possible value for the single most important objective, and therefore being able to crisply evaluate the penalty associated to the introduction of other objectives we want the design to be optimized for. As the single most important objective to be maximized in the design process ROI is selected, as it is a major business indicator.

In solving the optimization problem, algorithms belonging to different categories are applied. In particular, an algorithm is chosen among the ones in the category of the so-called gradient-based methods, and one in the pool of the heuristic techniques. The confidence to have really identified a global optimum is gained through a cross comparison of the results from different algorithms and with different algorithm settings.

During this phase the shortcomings and the advantages of each algorithm for the specific problem are also identified.

Gradient Based Algorithm

Gradient-based algorithms are designed to search the minimum or the maximum of an objective function *J(x)* using some information about its gradient. Various techniques and implementations are possible, with different level of complexity and varying computational effort. The selected algorithm is a particular variant of Sequential Quadratic Programming (SQP), known as NLQPL. NLPQL solves general nonlinear mathematical programming problems with equality and inequality constraints. It assumes that all problem functions are continuously differentiable. The internal algorithm is a sequential quadratic programming (SQP) method. Proceeding from a quadratic approximation of the Lagrangian function and a linearization of the constraints, a quadratic sub-problem is formulated and solved to get a search direction. Subsequently a line search is performed with respect to two alternative merit functions and the Hessian approximation is updated by the BFGS-formula [5] .

Gradient-based methods require an initial design vector to start the search with. The vector identified after the exploratory phase is used as a starting point for the NLPQL. Convergence is achieved very rapidly, after 99 iterations and 72s of execution time on a Pentium 4 1.8 GHz laptop. Results are shown in **Tab. 3**.

Design Var.	L [cm]	R [cm]	t_{cyl} [cm]	t_{cone} [cm]	$t_{h\text{-}sphere}$ [cm]	h/R
Value	5000.8	393.35	0.745	0.860	0.750	1.593

Obj. Constr.	g_v	g_{cyl}	$g_{h\text{-}sphere}$	g_{cone}	g_{vib}	ROI [%]	PL [kg]
Value	0.91	0.8	0.8	0.8	1.04	16.37	33,251

Tab. 3: Design Variables Setting and Objective / Constraints after NLPQL ROI Optimization (no scaling)

Convergence histories show that all variables but one converged very early, coming close to the final value after 20 iterations (an example is shown in Fig. 8); the

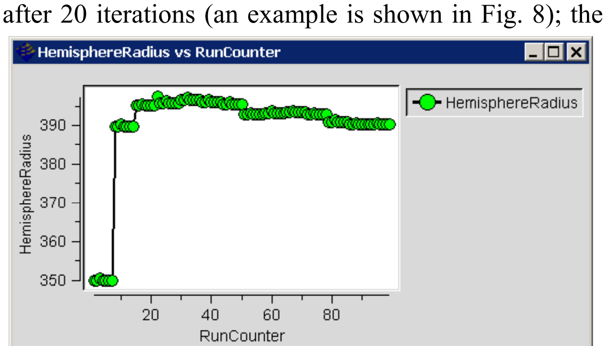

Fig. 8: Convergence histories of R during the optimization process

one that converged last was h/H (Fig. 9). This gives

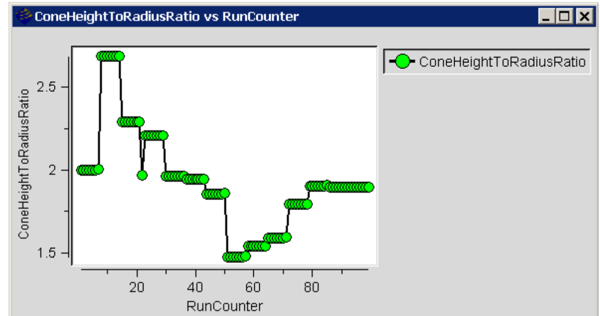

Fig. 9: Convergence history during the optimization process

some evidence that ROI exhibits low sensitiveness with respect to h/R.

Gradient-based methods are, however, prone to numerical errors when approximating the gradient and the Hessian using finite differencing. To mitigate this issue, scaling is introduced. The purpose of scaling is essentially to make the rate of change in the objective function, in region of interest, similar and independent of which design variable is considered, so that each in performing computation, cancellation errors can be avoided. In its simplest, but also most effective practical implementation, scaling of the design variables is obtained by calculating the diagonal entries in the Hessian matrix and, if they are greater than 100 or lower than 0.01, dividing each design variable by the square root of the Hessian diagonal (for a more rigorous and detailed description, see [6]). The values of the second derivative of the ROI with respect to the design variables at the optimum point are summarized in Tab. 4. Scaling of cylinder length and cylinder radius is consequently performed, by defining two new variables as: L'=L/1000, R'=R/100 and using those, instead of the original ones, as design variables handled by the optimizer.

Results of optimization after scaling are shown in Tab.

$\frac{\partial^2 J}{\partial L^2}$	$\frac{\partial^2 J}{\partial R^2}$	$\frac{\partial^2 J}{\partial (h/R)^2}$	$\frac{\partial^2 J}{\partial t_{cyl}^2}$	$\frac{\partial^2 J}{\partial t_{cone}^2}$	$\frac{\partial^2 J}{\partial t_{h-sphere}^2}$
$1.88\ 10^{-7}$	$2.33\ 10^{-5}$	8.49	6.37	8.37	1.72

Tab. 4: Diagonal Entries of The Hessian Matrix at the Optimum Point

5. As can be noted, values have improved significantly, thereby demonstrating the effectiveness of the scaling technique.

Design Var.	L [cm]	R [cm]	t_{cyl} [cm]	t_{cone} [cm]	$t_{h\text{-}sphere}$ [cm]	h/R
Value	4853.6	393.05	0.744	0.860	0.750	2.373

Obj. Constr.	g_v	g_{cyl}	$g_{h\text{-}sphere}$	g_{cone}	g_{vib}	ROI	PL [kg]
Value	0.90	0.8	0.8	0.8	1.00	20.51%	34,436

Tab. 5: Design Variables Setting and Objective / Constraints after NLPQL ROI Optimization after scaling

Heuristic Techniques

Several heuristic tools have evolved in the last decade that facilitate solving optimization problems that were previously difficult or impossible to solve. These tools include evolutionary computation, simulated annealing, tabu search, particle swarm, etc. In particular, evolutionary algorithms are stochastic search methods that mimic the metaphor of natural biological evolution. Evolutionary algorithms operate on a population of potential solutions applying the principle of survival of the fittest to produce better and better approximations to a solution. At each generation, a new set of approximations is created by the process of selecting individuals according to their level of fitness in the problem domain and breeding them together using operators borrowed from natural genetics. This process leads to the evolution of populations of individuals that are better suited to their environment than the individuals that they were created from, just as in natural adaptation.

Evolutionary algorithms model natural processes, such as selection, recombination, mutation, migration, locality and neighborhood. They work on populations of individuals instead of single solutions. In this way the search is performed in a parallel manner.

Evolutionary algorithms differ substantially from more traditional search and optimization methods, as gradient-based. The most significant differences are: 1) evolutionary algorithms search a population of points in parallel, not a single point; 2) they do not require derivative information or other auxiliary knowledge; only the objective function and corresponding fitness levels influence the directions of search; 3) they use probabilistic transition rules, not deterministic ones; 4) they can provide a number of potential solutions to a given problem and the final choice is left to the user.

Starting with the same initial condition supplied by the previous DOE phase, optimization is attempted with one variant of evolutionary algorithms known as Multi-island Genetic Algorithm (GA).

Convergence proved to be very difficult to achieve in the particular problem under examination. A great deal of work has been done in selecting properly upper and lower bound of the design variables, population size and mutation rate. The algorithm is terminated not because of satisfaction of convergence criteria, but because the maximum number of iterations was reached: runs totaled 10,000 with an execution time of 6180s on the same laptop used for the gradient-based solution, so hundredfold the NLPQL case. The final results are shown in Tab. 6.

Design Var.	L [cm]	R [cm]	t_{cyl} [cm]	t_{cone} [cm]	$t_{h\text{-}sphere}$ [cm]	h/R
Value	4089.1	411.62	0.780	0.780	0.900	4.259

Obj. Constr.	g_v	g_{cyl}	$g_{h\text{-}sphere}$	g_{cone}	g_{vib}	ROI	Payload [kg]
Value	0.90	0.8	0.8	0.8	2.06	15.28%	32,941

Tab. 6: Design Variables Setting and Objective / Constraints after ROI Optimization - GA

Analysis of convergence histories shows that the solution oscillates a lot between the extremes and struggles to converge. For typical examples, see Fig. 10.

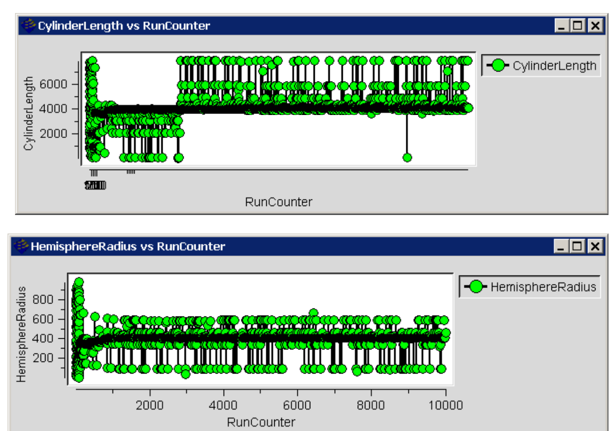

Fig. 10: GA Optimization – Convergence history, cylinder length and cylinder radius

Comparison of results and conclusions of the single objective optimization phase

Fig. 11 summarizes the results obtained in the various optimization steps. ROI is plotted for the different solutions as well as a graphical representation of the resulting SSEFT geometry.

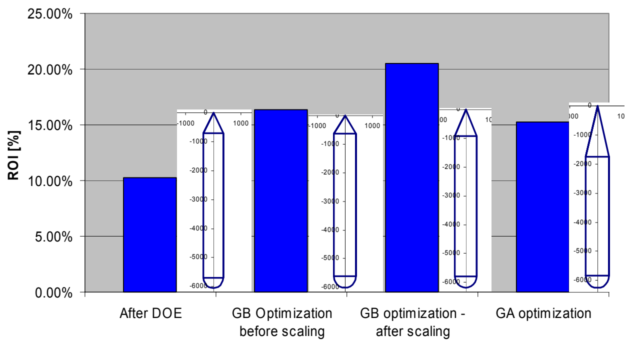

Fig. 11: Single Objective Optimization Results Summary

We note that the "best" result was achieved using a gradient-based method with scaling. We also note that *h/R* is the parameter that varies most from one solution to the other. This provides further evidence that ROI is very little sensitive to *h/R*. Since ROI is known to be very sensitive to R, it can be concluded that ROI is insensitive to the value of *h*, the height of the cone and that's the main reason why solutions with similar ROI can be significantly different in shape.

The best design is compared with the actual system to assess its plausibility. Dimensions such as overall length, diameter and weight are chosen because data are readily available (Fig. 13). The predicted length is higher than the actual by more than 25%, but diameter is very close (less than 10%) and weight is nearly exact.

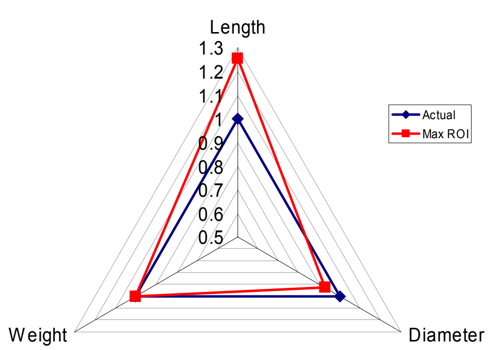

Fig. 13; Comparison of obtained results to values of the actual system

While the model may need some improvement, the results are judged to be good, considering the extreme simplified assumptions that are at its basis.

Step 3: Multi-Objective Optimization

With all the knowledge about the design space and the problem in hand gathered during the previous exploratory and single objective optimization phases, we can now head to the last step of analyzing the problem in its full complexity. As stated in the very beginning, the ultimate goal is to maximize the ROI of the SSEFT and, at the same time, the payload of Space Shuttle.

Relationship between ROI and payload

In order to understand the relationship between the two objectives, a Monte Carlo analysis is performed where values of the different design variables are selected randomly in a defined interval and ROI and payload computed. Constraint violation is allowed, but it's relevance is considered minor with respect to the focus

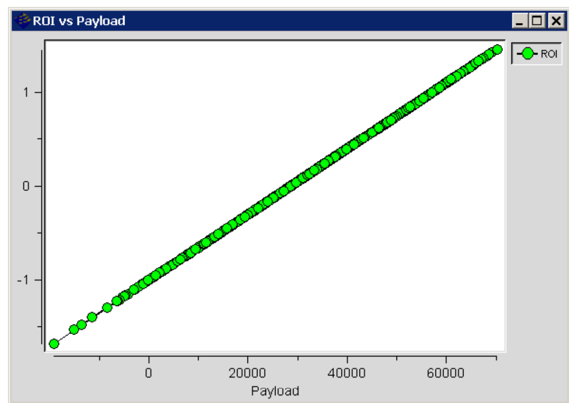

Fig. 12: ROI – Payload relationship

of our investigation. Fig. 12 shows the related plot.

It can be noted that the two objectives are supportive. Moreover, there is a linear relationship between the two; therefore maximization of one objective actually leads to the maximization of the other and the planned multi-objective analysis degenerates to the single objective optimization performed so far.

A closer look at the constraints

Sometimes it might be arguable whether a quantity should be considered as an objective or a constraint. In the analysis conducted up to now, for example, we have set the stress factors to be less or equal then 0.8, in order to have stress levels with a 20% safety margin with respect to the maximum allowable stress. This margin was set on the basis of experience and conservative best practice. However, it's not sharply clear whether a 20% margin is too conservative, too risky or correct. More importantly, we don't know how sensitive is ROI with respect to the stress level that we

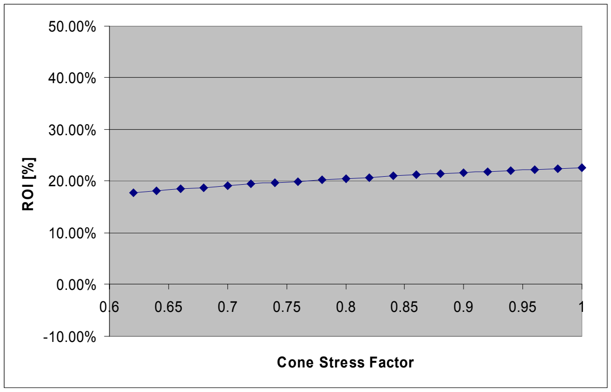

Fig. 14: ROI - Cone Stress Factor Pareto Front

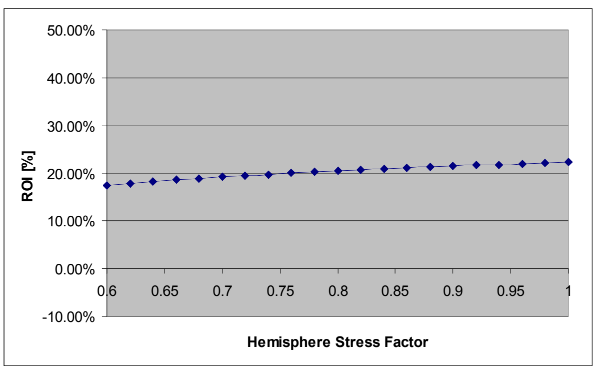

Fig. 15: ROI - Hemisphere Pareto Front

allow and therefore what penalty do we pay by pursuing a higher safety margin. It may therefore seem well grounded to consider stress factors no longer as constraints but as objectives and perform a multi-objective analysis.

Traditional multi-objective analysis relies on the weighted sum approach, where a weight is given to each of the objectives, proportional to the relative "importance" of that objective with respect to the other and optimization of the weighted sum of the two is pursued. The formulation is certainly valuable in some cases, but one shortcoming is certainly the fact that the end result is very much sensitive to the weights that are selected *a priori*.

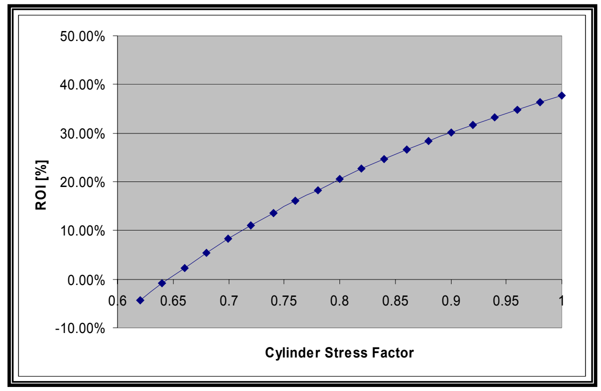

Fig. 16: ROI - Cylinder Stress Factor Pareto Front

In this particular case it's deemed far more important to understand the trade-off between two counteracting objectives. By exploring the ROI-Stress level Pareto front, it's in fact possible to get an understanding of how one objective trades against the other.

The scaled variable version of the model is used since it enables a better accuracy. The constraint stress limits are varied in steps from 0.6 to 1 one at a time and the optimum ROI solution is determined by the optimizer. Results are presented in Fig. 14, Fig. 15, and Fig. 16.

Within the assumptions of the model, while ROI is rather insensitive to the limits we put to the stresses on the cone and hemisphere, it is however fairly sensitive to any safety margin we want to take on the cylinder stress level. This insight offers some leverage for further optimization. In essence, what it implies is that the ROI penalty which is paid when the cone and hemisphere structure are dimensioned generously is negligibly small. This can allow devising a policy that directs to increase conservatively the thickness of those two elements, without the need of long hours of engineering effort in tweaking the structure for the minimum thickness. The resources thus saved, can, vice versa, be exploited to tailor the dimensioning of the cylinder structure that allows the tank to sustain a higher level of stress without compromising safety.

As verification of the advantages of the above described line of action, a run with cone and hemisphere stress factors reduced to 0.7 and cylinder stress factor raised to 0.9 is run. Results are presented in Tab. 7.

Design Var.	L [cm]	R [cm]	t_{cyl} [cm]	t_{cone} [cm]	$t_{h\text{-}sphere}$ [cm]	h/R
Base	4089.1	411.62	0.780	0.780	0.900	4.259
Revised constr.	4865.2	393.67	0.663	0.854	0.984	2.149

Obj. Constr.	g_v	g_{cyl}	$g_{h\text{-}sphere}$	g_{cone}	g_{vib}	ROI	Payload [kg]
Base	0.90	0.8	0.8	0.8	1.00	20.51%	34,436
Revised Constr.	0.90	0.9	0.7	0.7	1.00	**27.37%**	**36,396**

Tab. 7: Design Variables Setting and Objective / Constraints after Constraint optimization

The resulting benefit is huge: from 20.5% to 27.4% ROI, a 33% increase.

This example shows that the optimization process may be extended further than a mere function maximization and that it can be utilized as a tool for resource allocation or, more generally, for the maximization of design value.

CONCLUSIONS AND FURTHER WORK

Acknowledging the limitations of the simple model of the SSEFT, the present work has allowed unveiling several aspects of this system and its characteristics that were not evident at the outset.

First, ROI and payload are tightly linked and maximization of one objective leads to maximization of the other.

Each of the objective or constraint is sensitive to only two or three design variables among the six, the effect of the others being negligible. In particular, the identified main objective, ROI, is dependent on three main design variables, all related to the cylindrical portion of the tank: cylinder radius, cylinder length and cylinder thickness. In addition it is very sensitive to the cylinder stress constraints as well.

The problem is relatively uncoupled, meaning that the dependency of the objective from the second, or third order effects between the design variables is of a minor importance. This makes the use of orthogonal arrays particularly effective in determining a proper set of values for the design variables that bring the resulting system very close to the ROI optimum.

For the particular problem in hand and within the limitations of the specific algorithm implementation in iSIGHT, gradient based methods, and in particular SQP, vastly outperform a Heuristic technique such as Genetic Algorithm, both in the execution time (100 times lower) and in the results (ROI about 30% more). SQP, however, proved to be particularly sensitive to the starting point vector and to the scaling of the design variables. If not properly set, the algorithm might converge to local minima or terminate prematurely. In both cases a suboptimal solution results. Genetic algorithm, on the other hand, need tuning of upper and lower bounds of the design variables, population size and mutation rate. This is essential in order to ensure convergence and a relatively smooth convergence history.

The model built for the optimization analysis is relatively simple to allow the computational viability of the optimization process. It's certainly amenable to further improvements to take into account of different other phenomena equally important in the design of the system or to improve the accuracy of current prediction levels. Envisaged further areas of work that can e enhance the quality of the results at the expense of a modest increase in complexity and that can be considered for further work are:

- Improvement of the accuracy in the calculation of the tank volume. The top element is not a simple cone but has a sensibly different lofted shape that encases a higher volume, while the bottom part is actually more similar to a squeezed hemisphere with a lower contained volume.
- Inclusion the stresses not generated by the liquid pressure, such as those arising from the fuel and equipment mass. At the same time, some model that takes into account the buckling of parts of the system may also be added. Literature is available on this topic [7] .
- Improvement of the tank drag calculation, by introducing some elements that take into account the interaction with the orbiter and the boosters
- Modeling assembly costs in addition to material and welding costs.

REFERENCES

[1] http://science.ksc.nasa.gov/shuttle/technology/sts-newsref/et.html#et

[3] http://www.engineous.com/index.htm

[4] Warren C. Young , Richard Budynas “Roark's Formulas for Stress and Strain” *McGraw Hill*.

[5] http://www.uni-bayreuth.de/departments/math/~kschittkowski/nlpql.htm

[6] Panos Y. Papalambros and Douglass J. Wilde, “Principles of Optimal Design – Modeling and Computation”, 2nd edition, ISBN 0 521 62727 3, (paperback), Cambridge University Press, 2000

[7] R. D. Young, M. P. Nemeth, T. J. Collins, and J. H. Starnes, Jr., NASA Langley Research Center - Hampton, Virginia “Nonlinear Analysis of the Space Shuttle Superlightweight External Fuel Tank”, NASA Technical Paper 3616, December 1996