11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference
6 - 8 September 2006, Portsmouth, Virginia

AIAA 2006-7056

Exploring Mass Trade-Offs In Preliminary Vehicle Design Using Pareto Sets

Joseph Donndelinger[1]
General Motors Research & Development Center, Warren, Michigan 48090

Scott Ferguson[2] and Kemper Lewis[3]
University at Buffalo – State University of New York, Buffalo, New York 14260

Our goal in this work is to develop analytical tools to support the definition of balanced and compatible sets of vehicle specifications in the early stages of vehicle development. In this paper, we discuss the development and application of a Technical Feasibility Model (TFM) that may be used in preliminary design to assess the technical feasibility and optimality of specified combinations of vehicle performance targets. For this paper, we have exercised the TFM specifically to explore the relationships between vehicle mass, vehicle performance measures, (such as acceleration, fuel efficiency, and interior roominess), and high-level vehicle design parameters (such as overall exterior dimensions, occupant positions, and selection of a powertrain). The TFM is developed by first applying a Multi-Objective Genetic Algorithm to a multidisciplinary design framework to generate a set of Pareto-optimal design solutions, then applying response surface methods to generate a smooth mathematical representation of the Pareto set, and finally using geometric construction to analyze the position of a test point relative to the representation of the Pareto set. Results of this analysis include an assessment of the feasibility and optimality of the test point as well as a variety of projections from the test point to the representation of the Pareto set that may be used to identify opportunities for refining, relaxing, improving, or prioritizing performance specifications. The mapping between performance space and design space has been preserved, allowing for investigation of relationships between performance specifications and design variable settings.

In this paper we broadly demonstrate the application of the TFM, beginning with its basic capabilities of testing the feasibility of a specified combination of performance measures, quantifying the available amount of design freedom for a specified combination of performance measures, and quantifying the change in each performance measure required to attain a Pareto-optimal solution. In addition, we will demonstrate how the capabilities of the TFM may be leveraged specifically for exploring relationships between vehicle mass, vehicle performance measures, and vehicle design parameters by generating response surfaces to identify compatible sets of vehicle performance targets at specified levels of vehicle mass and quantifying the sensitivity of performance measures to changes in vehicle mass. Collectively, these capabilities make the TFM a powerful tool for managing vehicle mass and ensuring vehicle design feasibility in the earliest stages of the vehicle development process.

[1] Staff Researcher, Vehicle Development Research Laboratory; joe.donndelinger@gm.com, Member AIAA.
[2] Graduate Research Assistant, Department of Mechanical and Aerospace Engineering; smf7@eng.buffalo.edu. Student Member AIAA.
[3] Professor, Department of Mechanical and Aerospace Engineering; kelewis@eng.buffalo.edu. Associate Fellow AIAA.

I. Introduction

ONE of the fundamental challenges in preliminary vehicle design is developing a design that is both highly desirable and technically feasible. This challenge is compounded by competitive pressures to develop designs rapidly, as this both lowers development cost and reduces time to market. Thus vehicle manufacturers have invested considerable effort into developing and applying technologies for rapidly assessing the feasibility of preliminary designs. Virtual design environments[1,2] have substantially improved the efficiency of design feasibility assessment. Application of approximate models in these environments[3,4] has improved efficiency even further. Application of design optimization techniques[5] has proven helpful in balancing feasibility and desirability of designs modeled and explored in virtual environments. Some practitioners have found it helpful to generate sets of Pareto-optimal solutions[6] for both technical and non-technical reasons. This approach relieves difficulties in formulating objective functions. It also separates the exploration of design efficiency from assessment of the merit of solutions, which often facilitates decision-making in the product development process[7].

The goal of this research is to explore the culmination of the technical trends in this area; specifically, to generate an approximate representation of a set of Pareto-optimal design solutions and to explore its application to feasibility assessment and decision-making in preliminary vehicle design. In this paper, we discuss the development and application of a Technical Feasibility Model[8] (TFM) that may be used in preliminary design to assess the technical feasibility and optimality of a specified combination of vehicle performance targets. The TFM presented in this paper is tailored specifically for exploring the relationships between vehicle mass, measures of vehicle performance, and high-level vehicle design parameters.

II. Model Development

A TFM is developed by applying a Multi-Objective Genetic Algorithm[9] (MOGA) to a multidisciplinary design system to generate a set of Pareto-optimal design solutions[10,11], fitting a smooth, continuous mathematical representation to the performance measures of the Pareto-optimal design solutions, and implementing a mathematical algorithm for assessing the feasibility and optimality of a test point relative to the representation of the Pareto set. It is also very desirable, but not necessarily required, to provide means for preserving the correspondence between variables in performance space and design space for the Pareto-optimal solutions. The techniques used to develop the TFM presented in this paper are discussed below. Alternative techniques for TFM development have also been explored[12].

A. Model Scope

The TFM developed in this work spanned five measures of vehicle performance (including acceleration, fuel economy, and measures of interior accommodation) and eleven high-level vehicle design parameters (including overall exterior dimensions, occupant positions, and specification of the vehicle's powertrain). Representation of vehicle mass in a TFM poses unique and interesting challenges. Strictly speaking, mass is neither a design variable nor a performance objective; rather, it is a function of many design variables and it strongly influences many performance objectives. Regardless, vehicle mass is an important consideration in preliminary vehicle design and it is therefore desirable to represent mass in the TFM. When mass complements other performance measures included in the TFM, representing vehicle mass as a performance measure may seriously compromise the fidelity of the approximation of the Pareto set through confounding of model coefficients. This issue manifested itself during the development of this TFM; attempts at including mass as a performance response along with acceleration and fuel economy yielded unstable regression models. For this reason, vehicle mass is represented as a design variable in this TFM; thus the mass for any Pareto-optimal combination of performance measures may be queried through performance-to-design mapping.

B. Pareto Frontier Generation

The generation of the Pareto frontier used in this TFM has been discussed extensively elsewhere[8,13] and will be reviewed only briefly here. The performance and design spaces were each discretized prior to generation of the Pareto set. This was necessary not only for application of the MOGA, but also to provide a schema for mapping between the performance and design spaces. Both MOGA efficiency[14] and Pareto set quality[15-17] were monitored throughout generation of the Pareto set. The most densely populated Pareto set was generated with 20,000 evaluations of the multidisciplinary design system, although it has been shown that satisfactory results were achieved with far fewer evaluations[13]. Once generated, the Pareto set was represented with a pure quadratic response surface.

C. Feasibility Assessment

Given this representation of the Pareto frontier, the next step in developing the TFM is to define an algorithm for determining whether a given test point is feasible, Pareto optimal, or infeasible. In this work, the feasibility of a point in multiobjective performance space is assessed based on its location relative to the utopia point and to the Pareto frontier, as shown in Figure 1.

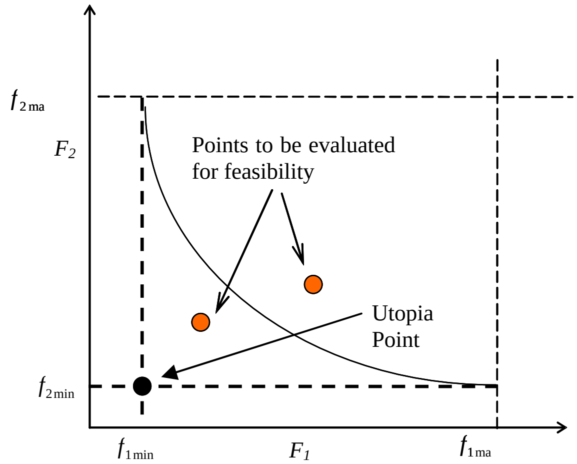

Figure 1. Geometrical References for Feasibility Testing

The utopia point is defined as the point in performance space that is best in every objective; this point is necessarily infeasible whenever the objectives compete. The feasibility of a test point may be assessed geometrically based on its location relative to the Pareto frontier. First a ray is constructed from the Utopia point through the test point. Next, the distances from the Utopia point to the Pareto frontier and to the test point are evaluated. There are three possible outcomes:

- If the distance from the Utopia point to the test point is greater than the distance from the Utopia point to the Pareto frontier intersection point, then the test point is infeasible.
- If the distance from the Utopia point to the test point is equal to the distance from the Utopia point to the Pareto frontier, then the test point is feasible and non-dominated.
- If the distance from the Utopia point to the test point is less than the distance from the Utopia point to the Pareto frontier intersection point, then the test point is feasible and dominated.

A measure of distance from the test point to the Pareto frontier is a by-product of the feasibility test. However, the Surface Intersection Point is not necessarily the closest point to the test point on the Pareto frontier. Depending on the location of the test point relative to the Utopia point and on the curvature of the Pareto frontier in the vicinity of the test point, the Surface Intersection Point may be much farther away from the test point than the closest point on the Pareto frontier, as shown in Figure 2.

The distance to the closest point on the Pareto frontier might be a more reliable measure of the degree of feasibility (or infeasibility) of a design solution than the distance to the Surface Intersection Point. However, it is difficult to interpret either of these measures because both are measured in a non-uniform coordinate system. This measure is only meaningful when rates of exchange between the performance objectives are known throughout the entire range of performance space, which is a highly unlikely circumstance.

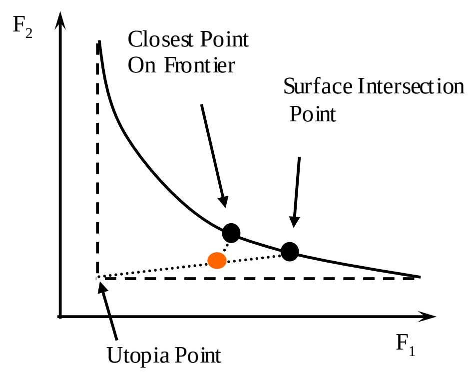

Figure 2. Surface Intersection vs. Closest Point on the Pareto Frontier

For this reason, the distance from the test point to the Pareto frontier along each individual axis of performance is also measured and reported. These measures of feasibility (or infeasibility) may be interpreted much more easily as they represent the change in one performance objective required to reach Pareto optimality when all other performance objectives are held constant. This procedure is illustrated in Figure 3. A test point with performance (F1,F2) is tested for feasibility in a two-objective performance space and is found to be infeasible. The distances to the Pareto frontier are measured along each axis and reported. To reach the boundary of feasibility, a tradeoff of either α units of F1 or β units of F2 is required.

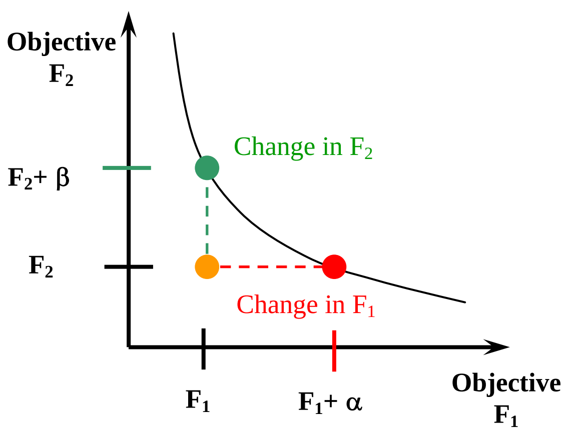

Figure 3. A Tradeoff in One Objective to Reach Pareto Optimality

Note that while quantification of the change in one performance objective required to reach Pareto optimality with all other performance objectives held constant is desirable, it may not always be possible. Depending on the location of the test point and on the local geometry of the Pareto frontier, an orthogonal projection from the test point may not intersect the representation of the Pareto frontier within the bounds of performance space. When this occurs, the TFM returns a message that the required tradeoff in said objective to achieve Pareto optimality is not available.

D. Performance-to-Design Mapping

Another issue addressed in developing the TFM is establishing correspondence between the Pareto-optimal solutions in the performance space and the corresponding variables and configurations in the design space. Mapping between performance space and design space is not a new problem in engineering design and has been recognized as a challenging task because the mapping can be one-to-many, with one objective function point mapping to multiple design points. If the mapping can be established, however, the knowledge of this correspondence provides valuable insight into the robustness of a given combination of performance targets by quantifying the amount of design freedom available to the designer for achieving a desired performance.

In this work we have approached the mapping of performance space to design space through the following steps:

- Discretize the performance and design spaces using indifference thresholds specified by the designer
- Represent the Pareto frontier as a collection of elements in the performance space
- Identify the corresponding elements in design space for each element in the Pareto frontier

When applying this procedure for performance-to-design mapping, the nature of the relationships between a performance space element and its corresponding design space elements can be one of three fundamental types:

Type 1: One performance space element maps to one design space . In this case, the centroid of the design space element is the design variable vector and the design freedom is defined to be half the size of the element. This is illustrated in Figure 4.

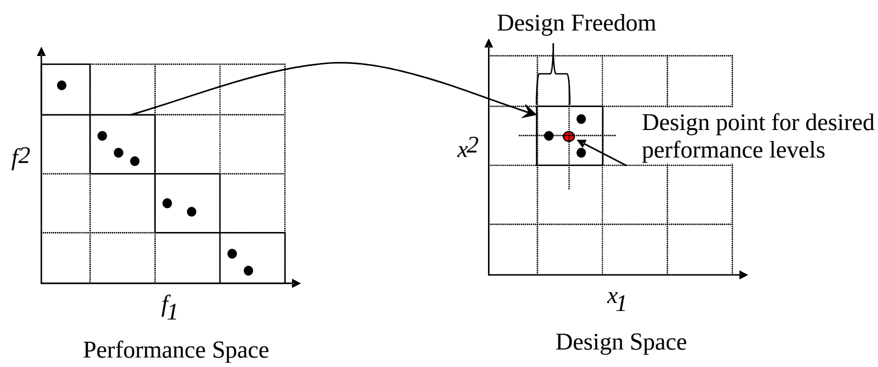

Figure 4. One-to-One Element Mapping (Type 1)

Type 2: One performance space element maps to multiple, adjacent design space elements. In this case, one super-element may be defined in design space to encompass all the adjacent elements mapped to a single element in performance space. The design variable values for this element are defined as the centroid of the super-element and the design freedom is defined to be half the size of the super-element. This is shown in Figure 5.

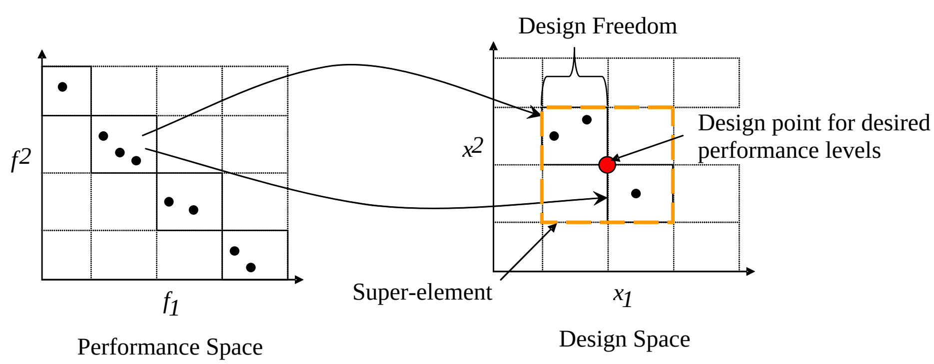

Figure 5. One-to-Many Adjacent Element Mapping (Type 2)

Type 3: One performance space element maps to multiple, non-adjacent design space elements, as shown in Figure 6. In our previous work[13], we found it generally acceptable to define one super-element encompassing the entire mapped region in design space. The validity of this procedure may be tested as follows: For a set of non adjacent, mapped design space elements, a larger super-element is defined to envelop these elements. Design points from each element are then evaluated; if the objective function values for every element within the design space super-element all fall within the same element in performance space, then the design variable and design freedom values may be defined in the same manner as for Type 2 mapping.

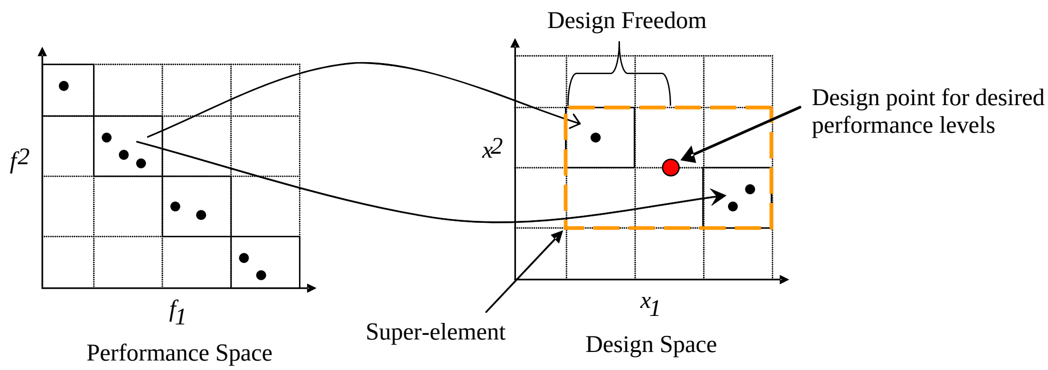

Figure 6. One-to-Many Non-Adjacent Element Mapping (Type 3)

Collectively, these capabilities make the TFM a powerful tool for managing vehicle mass and ensuring vehicle design feasibility in the earliest stages of the vehicle development process. This will be illustrated in the next section through an example of TFM application.

III. Application

In our previous work[8], we have demonstrated the application of the TFM by testing the feasibility of performance specification for 78 late-model sedans. Encouragingly, all vehicles with specifications within the domain of the TFM were found to be feasible, with newer vehicles falling closer to the Pareto frontier and older vehicles falling farther from the Pareto frontier. This exercise helped to establish the credibility of the TFM; however, it demonstrated only a small portion of its full capability. In this section we broadly demonstrate the application of the TFM, beginning with its basic capabilities of testing the feasibility of a specified combination of performance measures, quantifying the available amount of design freedom for a specified combination of performance measures, and quantifying the change in each performance measure required to attain a Pareto-optimal solution. In addition, we will demonstrate how the capabilities of the TFM may be leveraged specifically for exploring relationships between vehicle mass, vehicle performance measures, and vehicle design parameters by generating response surfaces to identify compatible sets of vehicle performance targets at specified levels of vehicle mass and quantifying the sensitivity of performance measures to changes in vehicle mass.

Suppose that a vehicle manufacturer is planning its next-generation entry into a hypothetical midsize car market segment. After careful analysis of marketplace conditions, the vehicle program manager has developed an initial set of targets for measures of the vehicle's interior accommodation, acceleration, and fuel economy. These targets, scaled relative to the performance of the current market entry, are shown in Table 1.

	Objective 1	Objective 2	Objective 3	Objective 4	Objective 5
Current Vehicle	1.000	1.000	1.000	1.000	1.000
Target for New Vehicle	1.000	1.054	1.011	1.026	1.122

Table 1. Performance Targets for Hypothetical New Product

A vehicle integration engineer (VIE) is tasked with applying the TFM to assess the technical feasibility of this point in performance space and finds that, unfortunately, it lies in the infeasible region. Detailed outputs from the TFM are shown in Table 2. There are two noteworthy features in this output table. First, there are appreciable differences in some performance measures between the surface intersection point and the closest point on the Pareto frontier, attributable to the test point's position relative to the Utopia point and to the curvature of the Pareto frontier in this region. Second, there is no tradeoff available within the domain of the TFM for reaching Pareto optimality by changing only Objective 4 or Objective 5.

	Objective 1	Objective 2	Objective 3	Objective 4	Objective 5
Test Point	1.000	1.054	1.011	1.026	1.122
Surface Intersection Point	**0.952**	**1.034**	**1.009**	**1.014**	**1.115**
Closest Point on Frontier	**0.949**	**1.036**	**1.008**	**1.025**	**1.126**
Tradeoff in Objective 1	**0.892**	1.054	1.011	1.026	1.122
Tradeoff in Objective 2	1.000	**1.011**	1.011	1.026	1.122
Tradeoff in Objective 3	1.000	1.054	**0.969**	1.026	1.122
Tradeoff in Objective 4	Not Available				
Tradeoff in Objective 5	Not Available				

Table 2. Tradeoffs Generated by the TFM

Next, the VIE evaluates each alternative using the TFM to estimate its mass and design variable settings. The masses of the alternatives are shown in Figure 7, scaled to the minimum mass observed among the alternatives. Interestingly, mass varies by more than 12% across five non-dominated solutions within a relatively small region of performance space. The VIE observes another interesting trend in design variable settings between alternatives. Four of the super-elements in design space overlap almost completely, whereas the fifth (the Tradeoff in Objective 1) is significantly different in both in nominal design variable settings and in available design freedom. The differences in design variables, scaled relative to the centroid of the four overlapping super-elements, are shown in Figure 8. The differences in design freedom for this alternative are shown in Figure 9. The differences in design freedom are reported in numbers of indifference thresholds to help distinguish the more robust design variables (such as E and F) from the more sensitive ones (such as A and G). The VIE then presents the results in the program team's Town Hall meeting[18] for discussion and definition of next steps.

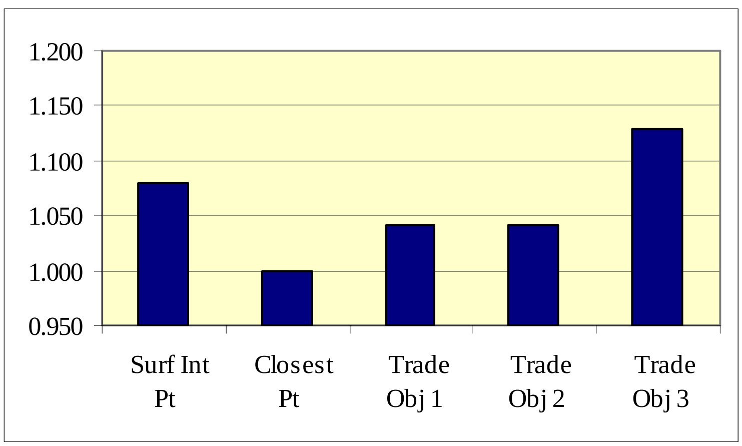

Figure 7: Relative Mass of Tradeoffs Generated by the TFM

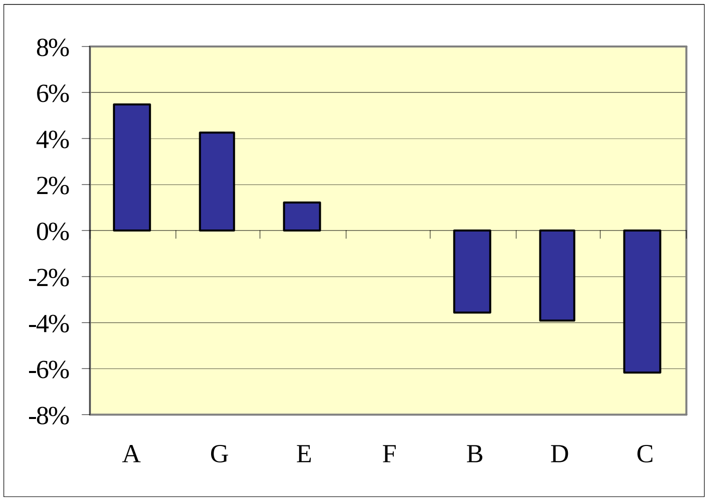

Figure 8: Differences in Design Variables Between Two Solutions

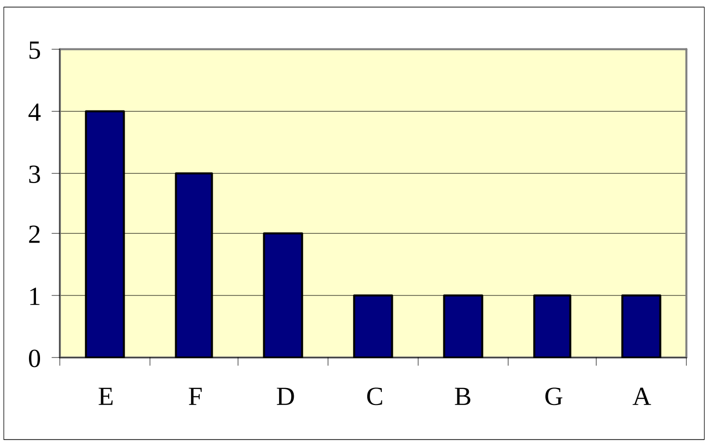

Figure 9: Differences in Design Freedom Between Two Solutions

After some discussion of the relative merits of each proposed tradeoff, the program team reaches consensus on their direction. The team is not prepared to accept any of the tradeoffs presented by the VIE without further investigation. The team does, however, express a preference for the tradeoffs at the closest point on the Pareto frontier, on Objective 1, and on Objective 2, partially due to the low estimated masses for these solutions. The VIE agrees to generate a local representation of the Pareto frontier representing all of the potential trade-offs between Objectives 1, 2, and 3 within the Equivalent Test Weight Classes of the team's preferred solutions. The VIE collects the required data by sampling the TFM using a standard experimental design in a software integration system. The VIE then applies linear regression to represent all of the feasible combinations of performance objectives within each Equivalent Test Weight Class with smooth response surfaces. The VIE is pleased to find that pure quadratic surfaces with Objective 3 as the dependent variable fit the data extremely well, yielding $R^2 > 0.995$ in both cases. One of these surfaces is shown in Figure 10. Finally, the VIE evaluates the slopes of both surfaces at each of the team's preferred tradeoff points. Using these slopes, the VIE constructs linear approximation models to estimate the change in each performance objective for a movement of one Equivalent Test Weight Class. The relative changes in Objectives 1, 2, and 3 for a movement of one Equivalent Test Weight Class at one of the team's preferred tradeoff points is shown in Figure 11.

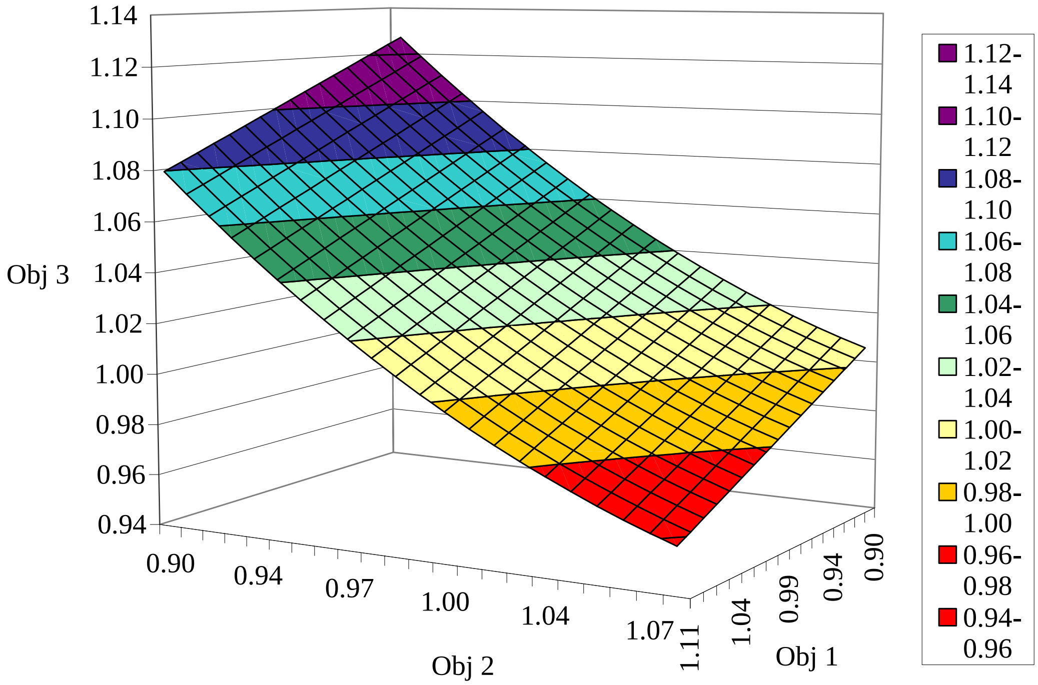

Figure 10. Compatible Performance Specifications Within an Engineering Test Weight Class

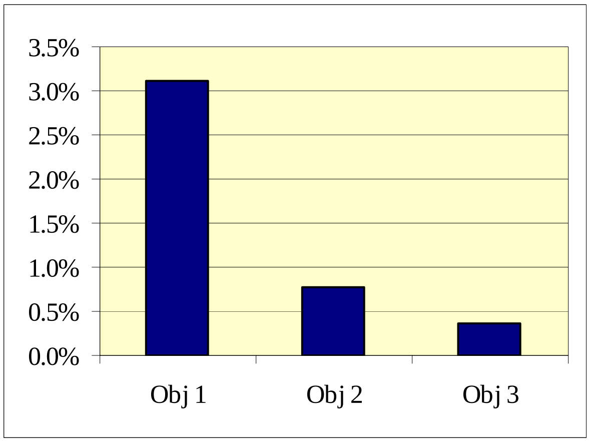

Figure 11. Relative Changes in Objectives For a Movement of One Engineering Test Weight Class

The vehicle program manager is provided with a rich representation of the feasible and efficient combinations of performance objectives within the region bounded by the initial performance targets. A revised set of vehicle performance targets is selected from the set defined by the TFM based on other criteria outside the technical feasibility domain. The program team proceeds with the vehicle development process, confident that the vehicle they will design and produce will remain within its targeted test weight class and will realize the targeted performance.

IV. Conclusion

The TFM shows potential for becoming a powerful tool for establishing compatible vehicle performance targets along with a feasible preliminary vehicle design during the early stages of the vehicle development process.

Although the TFM was originally designed only to model tradeoffs between vehicle performance objectives, it has been demonstrated that important design variables or intermediate quantities such as vehicle mass may be tracked or included as decision variables by sampling and selective post-processing of solutions from the TFM.

Perhaps the greatest challenge in future TFM development lies in managing the dimensionality of the performance and design spaces. The challenge in generating high-quality sets of non-dominated solutions scales geometrically. Likewise visualization of results becomes challenging whenever the number of interrelated quantities to be viewed simultaneously exceeds three. Fundamentally there are two approaches for addressing these scale-related issues. The first is to apply technical enablers, such as optimally-tuned genetic algorithms or systems specifically designed for visualization of high-dimensional data, to allow the scale to increase. The second is to actively reduce the scale, perhaps through analytics or by applying formal decomposition methods. For effective application of TFMs in the design of complex products, it is likely that a combination of both approaches will be necessary.

Acknowledgments

The authors extend their thanks to Dr. Srinivasan Rajagopalan of GM North American Product Development and Dr. Jeffrey Robinson of the GM Research and Development Center for the continuous support and valuable feedback and insight they have given throughout the course of this project.

References

[1]Crabb, H., 1998, *The Virtual Engineer: 21st Century Product Development*, Society of Manufacturing Engineers, Dearborn, MI.

[2]Fenyes, P. A., Donndelinger, J. A., and Bourassa, J.-F., 2002, “A New System For Multidisciplinary Design and Optimization of Vehicle Architectures”, 9th AIAA Symposium on Multidisciplinary Analysis and Optimization, Atlanta, GA, AIAA-2002-5509.

[3]Barthelemy, J. and Haftka, R., 1993, “Approximation Concepts for Optimum Structural Design - a Review,” *Structural Optimization*, 5 (3), pp. 129-144.

[4]Gu, X., Fenyes, P., and Donndelinger, J., 2006, “Comparison of Performance Trade-off Evaluations for Conceptual Vehicle Design Using Different Levels of Approximation in IFAD (Integration Framework for Architecture Development)”, 11th AIAA Symposium on Multidisciplinary Analysis and Optimization, Portsmouth, VA, AIAA-2006-6922.

[5]Papalambros, P. and Wilde, D., 1988, *Principles of Optimal Design*, Cambridge University Press, New York, NY.

[6]Pareto, V., 1906, *Manuale di Econòmica Polìttica*, Società Editrice Libràia, Milan, Italy; translated into English by A. S. Schwier, as Manual of Political Economy, Macmillan, New York, 1971.

[7]Cafeo, J., Donndelinger, J., Lust, R., and Mourelatos, Z., 2005, “The Need for Nondeterministic Approaches in Automotive Design: A Business Perspective,” *Engineering Design Reliability Handbook*, Ch. 5, CRC Press.

[8]Ferguson S., Gurnani A., Lewis K., and Donndelinger J., 2006, “A Constraint-Based Approach to Feasibility Assessment in Preliminary Design,” *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, in print.

[9]Goldberg, D.E., 1989, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, MA.

[10]Viennet, R., Fontiex, C., and Marc, I., 1996, “Multiobjective Optimization Using a Genetic Algorithm for Determining a Pareto Set”, *Journal of Systems Science*, Vol. 27(2), pp. 255-260.

[11]Eddy, J. and Lewis, K., 2001, "Effective Generation of Pareto Sets Using Genetic Programming”, ASME Design Technical Conferences, Design Automation Conference, DETC2001/DAC-21094.

[12]Padmanabhan, D., 2005, "Techniques for Construction of Pareto Frontier Approximations," 46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, Austin, Texas, AIAA-2005-1817.

[13]Ferguson, S., Gurnani, A., Lewis, K., and Donndelinger, J., 2005, “A Study of Convergence and Mapping in Preliminary Vehicle Design,” *International Journal of Vehicle Systems Modeling and Testing*, Vol. 1, No. 1/2/3, pp. 192-215.

[14]Deb, K., and Jain, S., 2002, “Running Performance Metrics for Evolutionary Multiobjective Optimization”, KanGAL Report No. 2002004, Kanpur Genetic Algorithms Laboratory, Indian Institute of Technology Kanpur, India.

[15]Zitzler, E., and Thiele. L., 1998, “Multiobjective Optimization using Evolutionary Algorithms – A Comparative Case Study”, In Eiben, A. E., et al., Proceedings of the 5th International Conference: Parallel Problem Solving from Nature – PPSNV, Amsterdam, The Netherlands, Springer, pp. 292-302.

[16]Azarm S., Reynolds, B. J., and Narayanan, S., 1999, “Comparison of Two Multi-objective Optimization Techniques with and within Genetic Algorithms”, ASME Design Engineering Technical Conferences, DETC99/DAC-8584.

[17]Wu, J., and Azarm, S., “Metrics for Quality Assessment of a Multiobjective Design Optimization Solution Set”, *Journal of Mechanical Design*, 123(1), pp. 18-25.

[18]Ulrich, K. and Eppinger, S., 1995, Product Design and Development, McGraw-Hill, New York, NY.