48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference
 15th
23 - 26 April 2007, Honolulu, Hawaii

AIAA 2007-1878

Advanced Visualization Techniques for Trade Space Exploration

John J. O'Hara[1], Gary M. Stump[2], Michael A. Yukish[3]
Penn State Applied Research Laboratory, State College, PA, 16801

E. Nathan Harris[4], Gregory J. Hanowski[5]
Lockheed Martin Space Systems, Denver, CO, 80201

and

Atherton Carty[6]
Lockheed Martin Aeronautics, Palmdale, CA, 93599

Recent advances in modeling and simulation technology have made it feasible to generate large datasets of design alternatives and their attributes in a relatively short amount of time. However, tools to understand and explore these datasets are limited. To this end, the Applied Research Laboratory at Penn State University has been developing a tool, entitled the ARL Trade Space Visualizer (ATSV) to support multi-dimensional trade space exploration. The ARL, in conjunction with the Lockheed Martin Corporation, has extended the tool to tackle several real world design challenges. In response to the needs of the engineering teams at Lockheed Martin, several key enhancements to the ATSV have been designed and implemented. These enhancements include contour plotting in two dimensions; isosurface generation in three dimensions; multiple independent brushing controls; and k-means cluster analysis. This paper will describe the full capabilities of the tool, as well as give an example of the types of design optimization performed by Lockheed Martin. The paper will focus on using the advanced visualization techniques to discover relationships within the dataset that would otherwise prove difficult to extract using traditional analysis techniques.

I. Introduction

The design by shopping paradigm, originally proposed by Balling[1], treats the design process like shopping for a car. Rather than using black box optimization, Balling proposes providing the designer with "a 'car lot' of competitive engineering designs." Once the user sees these designs, they can begin to formulate an exact preference, and find the design that is truly optimal. In his paper, Balling noted the need for research in two key areas. The first research area was in efficient methods for calculating the Pareto frontier. The second research goal was to design interactive graphical computer tools to aid the user in the decision making process. In this paper, we will show how the ATSV software solves the second issue, namely by providing advanced visualization techniques to support trade-space exploration and multi-disciplinary design optimization. Specifically, we will show how the inclusion of such features as contour plotting, cluster analysis and multiple independent data brushes allow for a more robust understanding of the underlying trade space. For more information about calculating the Pareto Frontier more efficiently, please see the work of Dr. Yukish[2].

---

[1] Associate Research Engineer, Information Systems, PO Box 30, State College, PA 16801
[2] Research Assistant, Product & Process Design, PO Box 30, State College, PA 16801
[3] Department Head, Product & Process Design, PO Box 30, State College, PA 16801
[4] Staff Systems Engineer, Technology Planning & Analysis, PO Box 179, Denver CO, 80201
[5] Staff Mechanical Engineer, Technology Planning & Analysis, PO Box 179, Denver CO, 80201
[6] Manager, ADP Conceptual Design Tools and Processes, Palmdale, CA 93599

II. Motivation

The ever-increasing sophistication of modeling and simulation techniques available today makes a design by shopping approach feasible. Currently there are a number of capabilities available to facilitate integrated design analysis and optimization, resulting in a reduction in the cycle time associated with performing Multi-Disciplinary Design (MDD), Analysis (MDA), and Optimization (MDO). At this point, many candidate frameworks already exist which allow the integration of existing analysis and application codes in a modular format, providing an efficient user environment for generating rich trades that are excellent candidates for the design by shopping paradigm. Typical trade spaces that can be generated from such software often contain thousands of designs, with each design supporting dozens of individual attributes.

While these rich datasets have the potential to provide the user with a large amount of insight, we must first find ways to interpret and understand the underlying data. When attempting to understand such large datasets, one quickly realizes that traditional data representations (such as scatter plots, histograms, and carpet plots) fail to visualize all the complexities inherent in the trade space. While traditional two-dimensional visualization techniques remain an excellent presentation formats for the attributes of a small number of candidate designs, they fail to capture the inherent multi-dimensionality of the unfiltered data. Even the addition of artifacts to the two-dimensional representations, such as color or shading, fails to adequately capture the underlying data. For this, we need to extend our visual representations into three dimensions, in order to fully appreciate the interconnectedness of the data.

Consider, for example, a simple trade space of only three variables (x, y, z) and a single relationship among them ($z=x^3-3xy^2$) This will create a mathematical curve known as a "monkey saddle" (Fig 1) This plot is named because it looks like a typical saddle curve, but contains one additional depression (for the monkey's tail). Such a curve, when represented in 3D, can be easily understood and studied. However, it is very difficult to construct a 3D mental model of the shape in question by simply looking at the 2D projections of the data, even if with the addition of color in each 2D plot to represent the variable not present. As you can see, an extremely simple tradespace, with only 3 variables and a single constraint function, benefits greatly from the addition of 3D representation. This type of 3D visualization quickly becomes required when attempting to decipher a trade space that consists of dozens of variables generated from hundreds of competing constraint functions.

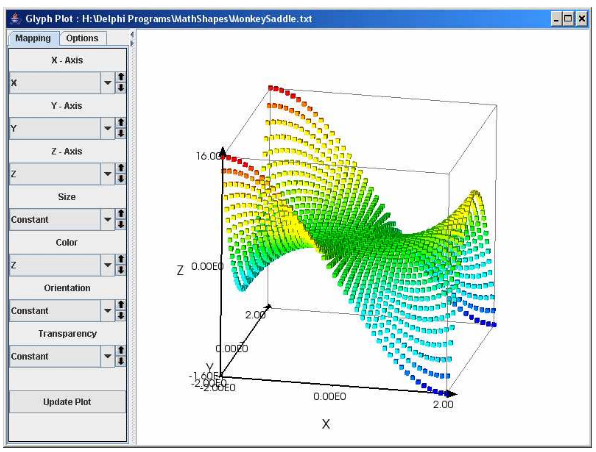

Figure 1. Three dimensional representation of the "monkey saddle" graph.

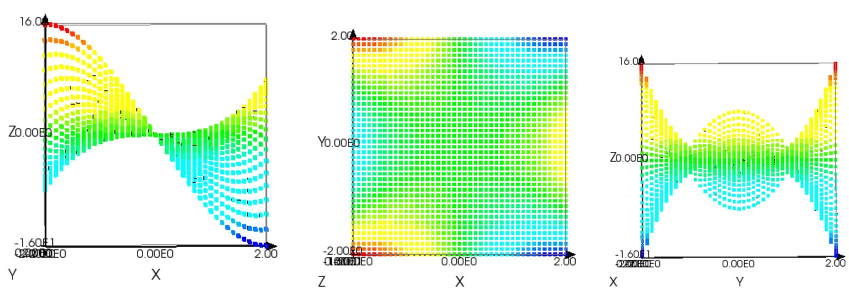

Figure 2(a-c). Two dimensional cross sections of the "monkey saddle" with color used to represent the third dimension

III. Previous Work

Previous publications[3,4] discuss the architecture, design, and capabilities of the original ATSV tool, especially the functionality and capabilities of the 3D plotting capabilities. For a detailed understanding of the underlying software, please refer to these papers. For this work, it is sufficient to understand the basics of the application's functionality.

The ATSV provides a visual interface for the exploration of trade space data. Input is provided in the form of a large matrix, each row of the matrix representing a particular design, and each column representing an attribute of that design. This matrix is used to create various plots in real time at the request of the user. In addition to the aforementioned 3D plotting, ATSV offers scatter plots, scatter matrices, three types of histogram plots, and parallel coordinate plots. Each plot can be customized to display all the available data, or only a particular subset from the trade space. Figure 3 illustrates some of the plotting techniques available in the current version of the ATSV. In this paper, we will focus mostly on the extensions to the 3D glyph plots.

Another feature of note, in addition to the plotting, is the brush/preference control panel, show at the bottom of Figure 3. This control panel allows the user to dynamically alter his preference structure, as well as dynamically constrain the data currently being viewed. Any change in the brush/preference control panel, such as restricting the maximum and minimum ranges of a particular variable, will dynamically update all the currently visible plots in response to such constraints. Previous versions of the ATSV software relied on a single control panel for the entire application, but this paper will discuss some enhancements to overcome this limitation.

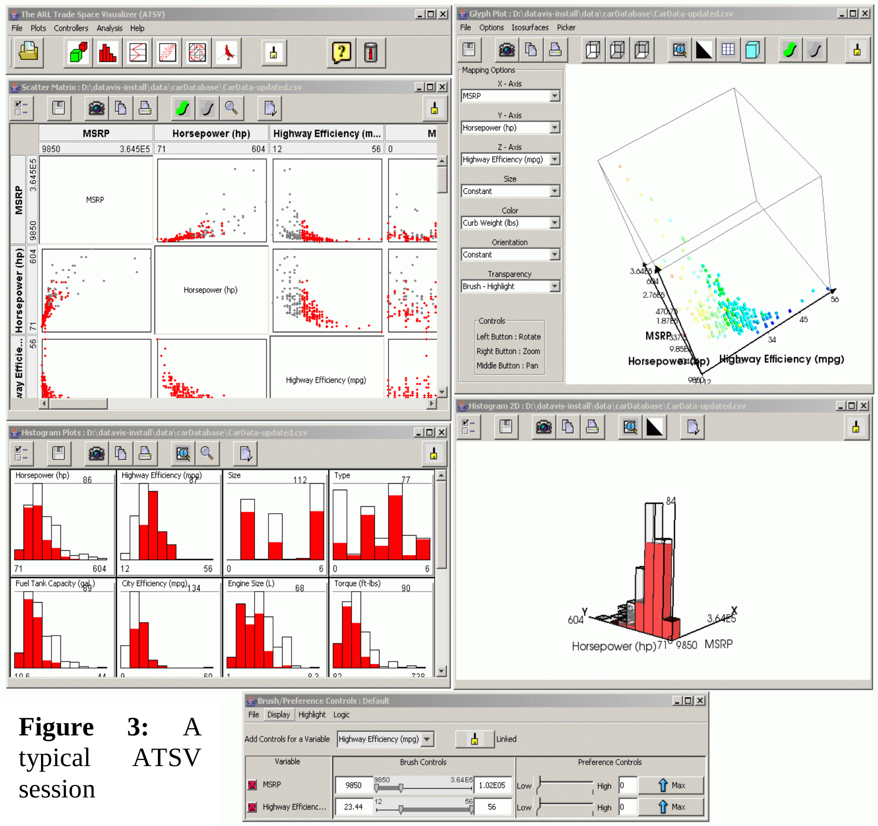

Figure 3: A typical ATSV session

IV. Recent Advancements

Recent development efforts on the ATSV software have focused on three distinct areas. These focus areas are designed to better support trade space exploration at Lockheed Martin.

A. Contour Plotting in Three Dimensions

Historically the process of exploring an aircraft performance design space is deeply rooted in traditional data representations of the most relevant vehicle and system level performance. As mentioned above, chief among these data representation formats is the traditional performance carpet plot, which is simply one way to display data dependant on two variables. Many generations of aircraft have been designed with these tools and still today it serves as the 'tool of choice' of designers and analysts for the dissemination of aircraft performance data. The elegance of this technique lies in its simplicity…not of the data itself, but of the manner in which the inherently inter-related nature of the data is captured. That said, these tools clearly have their limitations, many of which become painfully obvious as they are employed to investigate complex 5, 6, or 7 dimensional datasets, 3 dimensions at a time. These traditional methods are certainly still applicable, and will continue to be used as they have been for decades. As the size of the these datasets continue to increases concurrently with the capability of tools, frameworks and computational resources which were used to generate them, more advanced visualization techniques are required.

In order to more effectively support the interactive process of design space investigation an ATSV may be employed to provide an introspective capability not typically afforded by conventional techniques. The existing ATSV environment brings a host of features and functionality detailed earlier in this paper as well as in previous publications. Several recent key improvements in functionality of these advanced visualization tools and techniques have greatly increased their overall effectiveness and applicability to complex system level design problems such as that of classical aircraft design. These essentially boil down to two new core capabilities which have enhanced the existing capability 1) The ability to add multiple constraint planes for multiple data subsets on a single plot and 2) The ability to 'turn-off' the data that falls outside of these constraints

A significant improvement to the utility of the existing ATSV environment is the ability to display multiple constraint planes within the 3D glyph plot. This is accomplished through the tessellation of the data sets to view 3D surfaces using the ATSV environment and by displaying these dynamic planes to represent complex constraint surfaces through this multidimensional data. In addition, the creation of clipping planes which may then be “pushed” through the 2D planar projection results in the display of points lying within that plane. Building off of this core capability several complimentary capabilities also become viable such as the ability to selectively remove points falling above, below or on these 3D constraint surfaces, and resulting in the definition of a new design sub-space. In this way a user may begin to ‘carve’ away at the design space as one might cut up an apple with a knife. Just as in the case of the apple, there will be portions of a given data space that will be undesirable and which the user will want to remove, such as the apple core and seeds. The user essentially requires a tool which allows them to ‘core’ the design space without unnecessarily disposing of the ‘fruit’ in the process. The ability of the ATSV to dissect a dataset by cutting along constraint lines, gives the user that power to carve away the parts of the design trade space which do not satisfy constraints which have been imposed, and leaves only the viable potential solutions. They can then effectively utilize existing ATSV functionally such as the brush preference controls to further peel away the remaining data to identify more optimal families of solutions.

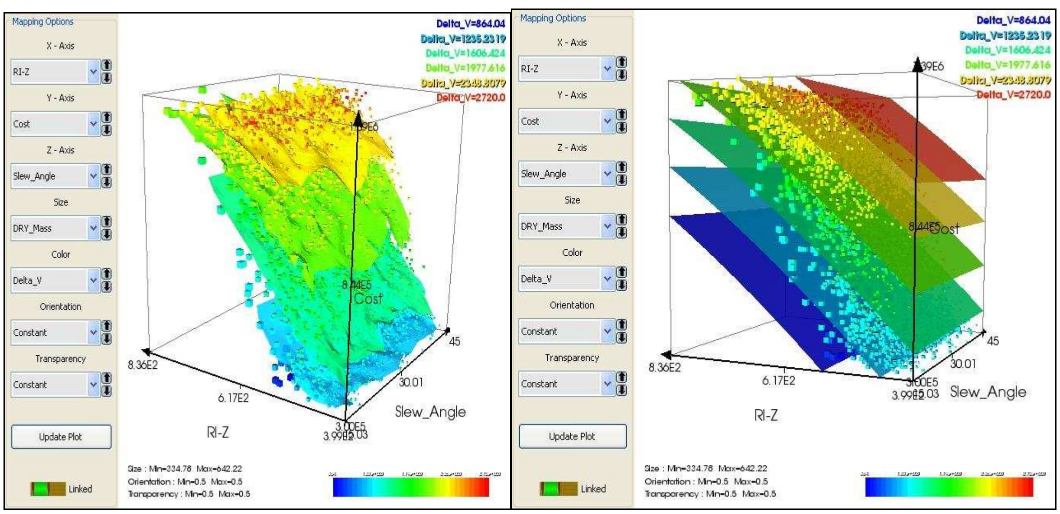

Figure 6: Constraint surfaces in the design space using an exact fit (left) and a linear fit (right)

In Figure 4, we show two of the potential uses of the contour curves. On the left, we show a design space that has been tessellated to show which points fall along a particular constraint. On the right, the same points have been used as the basis for a 2D linear fit. Once these constraint curves have been generated, the user then has the ability to dynamically down sample the dataset to see only those designs which lie above, below, or along that constraint. By combining several of such restrictions, we can down sample our trade space to only those points which match all of our requirement criteria. In this way, we can begin to transition our focus from thousands of potential designs, to a few desirable candidates.

B. Cluster Analysis in Three Dimensions

Another useful tool in trade space exploration is the ability to automatically detect and manipulate clusters. In multidimensional datasets, it is very common to find clusters, especially when considering that several of the variables may be interrelated to each other. For this type of analysis, we have implemented a K-means clustering algorithm[5]. The user can specify which particular columns they want to have included in the analysis. In this way, they are not limited to the space which is currently being used. Once the n-dimensional analysis (one dimension for each attribute of interest) is complete, then the results are appended to the data matrix as simply another attribute. In this way, the cluster to which each design belongs simply becomes another attribute of the dataset. This allows the user to visualize that information in the same way they would visualize any other attribute in the design space. When taken with the surface generation abilities discussed above, it is possible to create 3D shells that represent the bounds of the clusters in three dimensional space. A visual representation of the cluster analysis result can be seen in figure 5.

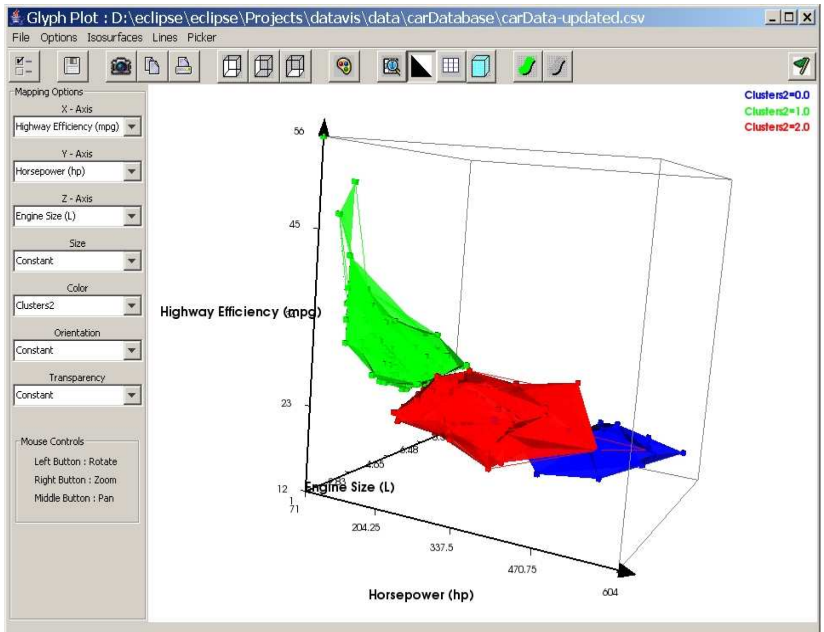

Figure 5: The results of a k-Means cluster analysis being visualized with one surrounding isosurface for each cluster.

C. Independent Preference Structures in Three Dimensions

Finally, the ability to display multiple n-dimensional visualization mappings of same data set with a preference structure independent to each view allows the user to visualize changes in separate 3D views. In this way individual views may be tailored specifically to a given disciplinary perspective. A key point to note is that although the datasets are represented in separate windows, all of these representations of the data are linked such that constraint planes may be imposed in one perspective view, and subsequently be represented in all views as applicable. This includes lower dimensionality data representations such as the 2D planar cuts and histograms. This attribute of independent data views, representing interlinked constraints is a key functionality particularly as the user considers more complex higher dimensionality trade spaces which are no longer effectively captured with planer graphical representation and in which interdependent relationships are often extremely complex and non-intuitive. In these cases traditional design intuition is frequently challenged, not because it is any less applicable, but because the data is not represented in such a way as to lend itself to straightforward introspection. It is these types of applications which most critically require tools such as the ATSV - not to draw conclusions…that is the user's job…but rather to represent the data in such a way that the user has a conclusion to draw.

As an example, in designing an aircraft, multiple disciplinary system 'performance' constraints might be collected up in one view, perhaps all 'cost' related constraints are gathered up into another, and perhaps 'operational' constraints are in still a third. The design dataset may then be simultaneously interrogated from a number of disciplinary perspectives simultaneously. In addition, by pairing this capability with the capability to 'carve' up the design described above, the implications of constraints or design decisions in one 'discipline' may now clearly be captured in another, seemingly unrelated discipline area. As a result, the user has at their disposal a powerful new toolset with which to begin to understand system level implications of the design decisions which are being made, and to engage in a truly multi-dimension selection process. Ultimately this capability can allow use of Response Surface Methodologies[6] (RSMs) to drive these high power visualizations, resulting in a dynamic trade space environment. With this, a user can consider design trade-offs, as well as assess the impacts of changing constraints and problem assumptions, the infusion of new technology or capabilities to the system.

V. Future Research Goals

Development work is still continuing on the ATSV. We are focusing on improving the user interface to make the system easier to learn and use. We are constantly adding new plotting mechanisms based on the latest research in information visualization. We are also working on data management techniques in order to expand the tools ability to visualize data sets with over 100,000 unique designs. This research involves adding some multithreading capabilities to both to calculate the Pareto frontier, and also to increase the render speed of all the plots.

Additionally, we are working on extending the ATSV software to work in a networked setting. We are exploring methods for users to share not only preference structures and brush settings, but the ability to share plots across the network. Ideally, many designers should be able to sit down and cooperatively iterate over the design space to determine the best design. The interface could take into account the differing preferences of the various designers. We are especially interested in the case where multiple designers have competing objectives. For example, an vehicle designer is trying to minimize the total mass of a aircraft, while a trying to maximize the amount of weight that can be carried. Resolving these types of design roadblocks will require some form of dispute resolution theory, or perhaps some applications of game theory to find the optimal design.

VI. Conclusion

In conclusion, we have developed an advanced visualization interface that makes it easier to understand the complex design space inherent in the modern design process. We have demonstrated how this technology is useful in real world applications. We believe that this tool will ultimately provide users with an efficient tool for analyzing and understanding complex, real world design spaces.

References

[1]Balling, R. *Design by Shopping: A New Paradigm*. in *Proceedings of the Third World Congress of Structural and Multidisciplinary Optimization (WCMSO-3)*. 1999. Buffalo, NY.

[2]Yukish, M. A. and Simpson, T. W. *Analysis of an Algorithm for Identifying Pareto Points in Multi-Dimensional Data Sets*. In *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. 2004. Albany, NY.

[3]Stump, G.M., Yukish, M., Simpson, T. W., Harris, E. N. *Design Space Visualization and Its Application to a Design by Shopping Paradigm*. in *Proceedings of DETC 2003, ASME*. 2003. Chicago, IL, Paper No. DETC2003/DAC-48785.

[4]Stump, G.M., Yukish, M., Simpson, T. W., and O'Hara, J. J. *Trade Space Exploration of Satellite Datasets Using a Design By Shopping Paradigm*. in *IEEE Aerospace Conference*. 2004. Big Sky, MT.

[5]J. B. MacQueen (1967): "Some Methods for classification and Analysis of Multivariate Observations", Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability, Berkeley, University of California Press, 1:281-297

[6]Carty, A., "An Approach to Multidisciplinary Design, Analysis and Optimization for Rapid Conceptual Design", 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Atlanta, Georgia, September 4-6, 2002, AIAA 2002-5438