

ADAPTIVE SELECTION OF PARETO OPTIMAL ENGINE TECHNOLOGY SOLUTION SETS

Dr. Bryce Roth, Mr. Mathew Graham, Prof. Dimitri Mavris
Georgia Institute of Technology, Atlanta, GA 30332-0150
and
Mr. Noel Macsotai
GE Aircraft Engines, Cincinnati, OH 45215

Keywords: *genetic, algorithm, engine, technology, selection, Pareto*

Abstract

Successful selection of propulsion system technologies for development and incorporation into new engine designs requires careful balance among many competing design objectives (i.e. performance, cost, risk, etc.). One seldom has sufficient development resources available to fully explore all promising concepts and must therefore choose a few technologies that show the greatest promise to meet program objectives. This paper describes a method of selecting optimal combinations of engine technologies. This method employs a technology impact forecasting environment in conjunction with genetic algorithms to find Pareto-optimal technology solution sets. These results are illustrated using Technology State Transition Diagrams to show how technologies move into and out of the Pareto-optimal sets. An edge search procedure is introduced as a means to efficiently characterize the objective space, the results of which are presented in the form of ternary plots. These plots show how technologies benefit multiple (often-conflicting) objectives and help find robust or compromise technology combinations. Finally, these methods are applied to select engine technology combinations for a commercial engine system of current interest.

1. Introduction

As tomorrow's aircraft engines become more complex and pressure for reduced design cycle times grows, it becomes increasingly difficult to evaluate and select engine technologies. These issues in the aircraft engine industry are driving the need to develop advanced analysis methods to assist decision-makers in product development and resource allocation. The objective is to select sets of engine technologies that will deliver the most "bang for the buck" so as to meet performance goals within

program budget and schedule constraints. This need is particularly important in light of the extremely high capital investment that usually accompanies a technology development decision.

Engine technology selection can be thought of as a constrained combinatorial optimization problem. The objective is to select an optimal set of technologies from a list of discrete technology choices. Optimal in this context means those technologies that represent the best fit to a given set of conflicting requirements and program objectives. Program objectives typically encompass competitive advantage, time and budget constraints, and minimization of development risk. Some technologies may be incompatible with program objectives, incompatible with other technologies, or be dependent upon others in complex ways. Additionally, the number of permissible technology combinations grows geometrically with the number of technology options available such that it is usually unfeasible to investigate every possibility. It is therefore imperative to *efficiently* search for technology combinations that best fit the given requirements, account for incompatibilities, enforce enabling relationships, and do so with a high level of accuracy and confidence in the analysis.

Roth et al. [1] and Kirby & Mavris [2] have shown that genetic algorithms used within the Technology Identification, Evaluation, and Selection (TIES) method are an extremely effective means of solving this constrained combinatorial optimization problem described above. The method works by using a response surface representation for the impact of any given technology in terms of system-level figures of merit (FoMs). This technology impact model is then interrogated using a genetic algorithm (GA). The GA works by creating a pool of technology combinations and evaluating them in the technology impact model to yield estimates of how each technology combination performs in the system. These combinations are then compared to one another and the best combinations are kept in

the pool while the worst are discarded. The surviving combinations are then used as “parents” to create a new generation of combinations that replace the discarded combinations in the pool. This process is repeated over many generations until the population has converged to an optimal set of technologies. The surviving technology combinations are taken to be the best solutions for the given objective function.

This approach to technology selection is useful for several reasons. First, it allows one to create a generic technology model that can easily be extended to include new technology options as they emerge. Moreover, this model can be created at minimal expense and incorporates a combination of expert opinion and analytical data. Third, the genetic algorithm is an analytically repeatable means of obtaining an optimal technology solution set for a given technology model. Finally, it is very easy to incorporate many types of data into the genetic algorithm objective function, including subjective data, analytical data, non-numerical data, probabilistic data, etc.

This technique has already been demonstrated in prior work, using a single objective function weighting. In reality, explicit objective function weightings are seldom given, and one is usually interested in finding those technologies that are the most robust compromise to conflicting objectives.

The objective of this study is to show how the TIES method can be used to find optimal technology solutions over the spectrum of without the need to explicitly define objective weights. The GA-enabled TIES technique is used to investigate parametric variations in objective function weightings. The result is a Pareto front, which represents the *technology frontier* (or range of optimal solutions) achievable with a given list of technology candidates. The technology Pareto front goes beyond a simple technology ranking by showing how the set of optimal technologies changes with shifting objectives, a key to understanding compromise and robust designs. Finally, this technique is applied on a technology selection problem for commercial turbofan engines based on that previously described by Roth et al.¹ The ideas developed and demonstrated herein are motivated by aircraft engine technologies and systems. However, they are broadly applicable to any complex system where the problem is to meet objectives by selecting a subset of optimal technologies.

2. Pareto Fronts and Technology Frontiers

A sample technology Pareto front for a two-objective optimization problem is illustrated in Fig. 1., where the axes represent mission fuel burn and technology risk. We may postulate that the minimum risk and fuel burn scenarios occur when there is no consideration of the other objective (i.e. sub-optimization). As the relative weighting of a combined objective function is parametrically varied between each extreme, the technology mix will gradually evolve from a minimum risk set to a performance-optimal set. The locus of fuel burn-cost points formed by these optimized solutions is the Pareto front. It represents the bound of approach to the ideal solution and gives a clear visual indication as to how closely one may approach it. The payoff comes as one uses the Pareto front to tailor the performance-cost technology mix, and find maximum desirability.

An interesting feature of technology frontiers is that they are discontinuous and are formed by an essentially infinitesimal set of technology combinations taken from a large but finite solution space. To understand this, consider a typical technology optimization problem consisting of 40 technology candidates. If one presumes that each technology can either be selected or rejected* and ignoring any compatibility constraints, then there are 2^{40} or ~ 1.1 trillion possible technology combinations. Since there are a finite number of technology combinations, it follows that the Pareto front must consist of several hundred technology combinations—the *Pareto set*. It is these technology combinations that represent the locus of optimal technology solutions (the proverbial needle). The remaining trillion+ possible combinations are non-optimal solutions (the haystack).




Fig. 1 Pareto Front of Fuel Consumption Versus Manufacturing Cost.

* This assumption of binary discrete states is usually adequate for most technology problems. For those cases where there may be more than one technology option (e.g. – material selection) it may be useful to model the technology as having more than two possible states.

Statistically speaking, one can think of the points in the Pareto set as being the outliers in a group of 1.1 trillion points. Fig. 2 illustrates this idea by showing a grouping of 10,000 randomly selected technology combinations plotted on the technology risk versus fuel burn axes (grouped at the upper right). The string of points at the lower-left of the figure defines the Pareto set, and is found by using a combined GA-TIES model. It is clear from this figure that the best solutions found using 10,000 random trials do not even approach the optimal solution set. This is because the 10,000 randomly selected technology combinations constitute one millionth of one percent of the total technology space. It can be shown statistically that the only way to have high confidence in finding the outliers in a population by taking random samples is to make the sample size so large as to constitute essentially the entire population.[3] This quickly becomes infeasible as the number of technologies increases, so it is necessary to use an “intelligent” search algorithm to find the outliers defining the Pareto set.

The use of TIES in conjunction with a GA is a powerful means of selecting technologies because it allows one to visualize the technology frontier and intuitively make tradeoffs among the objectives without the need to specify objective weights. For instance, the technology frontier shown in Fig. 2 shows the entire spectrum of trade-offs amongst two objectives. One can easily select the point along this frontier that is the best fit with overall objectives and interrogate the model to determine what technologies correspond to that point. Moreover, this figure gives a good indication about the absolute limits for the technologies under consideration. For instance, Fig. 2 clearly shows that no combination of technologies considered in this example can yield more than a 7% improvement in mission fuel burn. Although Fig. 2 shows trade-offs for only two objectives, the basic method can be extended to more dimensions.

3. Method

The Pareto-optimal technology selection method consists of four basic steps: 1) create a TIES technology model, 2) implement a GA Pareto search, 3) an edge search, and 4) visualize the results. The basic analysis flow is depicted in Fig. 3. Note that each step builds upon the previous, and steps 2 and 3 in particular involve recursive application of previous steps. However, it should be noted that each step provides useful results in its

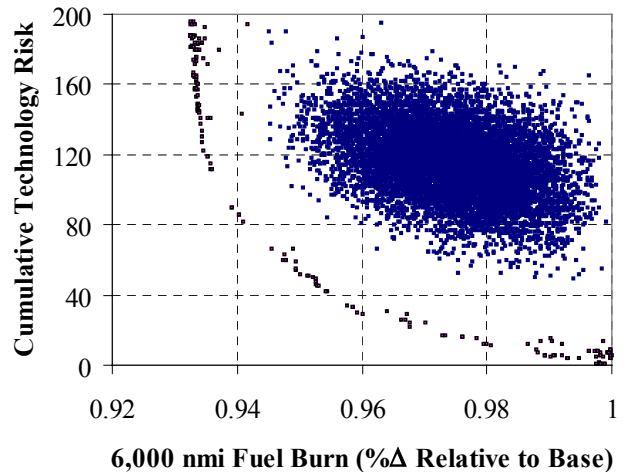


Fig. 2 10,000 Random Technology Combinations Plus Technology Pareto Front.

own right, so one can obtain useful information during each stage of the analysis.

3.1. Step 1: TIES Technology Model

The first step in the analysis process is to create a TIES technology model. The TIES modeling process has been described extensively in Refs. 1 and 2. For the purposes of this paper, it is sufficient to regard the TIES model as a function that takes a binary technology vector of the form:

$$\vec{T} = \{t_1 \quad t_2 \quad \dots \quad t_m\} \tag{1}$$

where $t_i = \begin{cases} 1 \Rightarrow \text{(Technology 'i' is used)} \\ 0 \Rightarrow \text{(Technology 'i' is not used)} \end{cases}$

and maps it into a response vector of the form:

$$\vec{R} = \{r_1 \quad \dots \quad r_n\} \tag{2}$$

where n is the number of objectives of interest for a given problem. A properly constructed TIES model is a compact and accurate representation of technology impact for arbitrary technology combinations. The Pareto-GA relies on the ability to use this model to perform quick function calls, as the GA requires many technology evaluations.

3.2. Step 2: GA-Pareto Front Procedure

The second step is to “wrap” a GA routine around the TIES model constructed in the previous step. As described in Ref. 1, the resulting GA-TIES model can be regarded as a function whose input is an objective function weighting of the form:

$$\vec{O} = \{o_1 \quad \dots \quad o_n\} \tag{3}$$

where o_i is the objective function weighting on objective ‘i’ with $\|\vec{O}\|=1$. The outputs from this

function are in the form of a binary vector containing the optimal technology solution set:

$$\vec{T}_{opt} = \{t_1 \quad t_2 \quad \dots \quad t_m\} \tag{4}$$




Fig. 3 Pareto-Optimal Technology Selection Method.

and a pair of objective values that form the x-y coordinates of a point on the Pareto front.

The purpose of this GA-TIES model is to efficiently solve the combinatorial selection problem to yield an optimal technology solution set for any given objective function weighting. The objective function is implemented in the GA evaluation routine by using the objective weights as probabilities governing how frequently the various objectives are selected by the GA during pair-wise tournaments amongst population members. The result is a robust, compact, and efficient means of incorporating competing evolutionary pressures in the genetic algorithm.

Once this GA-TIES model is in place, one can create a technology Pareto front by parametrically varying objective weights over a series of GA optimizations. This iterative Pareto front technique is used instead of more efficient techniques (such as the niched-Pareto method) because it yields explicit information relating optimal technology sets to specific objective weightings. The result is a Pareto front describing the frontier between two objectives, as shown in Fig. 2, and a series of optimal technology solution sets, one for each Pareto case.

The evolution of the optimal technology solution set can be described in terms of a *technology state transition diagram*. This can be thought of as a grid plot having technology index on the abscissa versus objective function on the ordinate. Boxes in the grid represent the technology state. A shaded box means the technology is part of the optimal set, while an empty box means it is not. This plot illustrates how the optimal technology set evolves as objective weights change.

3.3. Step 3: Edge Search Procedure

Up to this point, a method has been defined which allows one to find and examine the outlying (Pareto) points in an objective space. This procedure is valid for any number of technologies and any number of competing objectives. If there are three objectives, the Pareto front shown in Fig. 2 becomes a *Pareto surface*. For four or more objectives, one must think in terms of a *Pareto hypersurface*. The GA-TIES function described

previously finds optimal technology sets that correspond to points on this hypersurface. The Pareto front method described in the previous section allows one to find a *line* on the hypersurface. The next logical step is to use this tool to build up a comprehensive characterization of the entire Pareto hyperspace and the underlying technology combinations that give rise to it.

We are typically interested in those technologies that occupy the largest area of the Pareto hypersurface (best satisfy all objectives). It therefore follows that one must have a means to estimate the hypersurface area over which each technology is a member of the optimal solution set. One possible means of doing this is to use a grid search procedure wherein each dimension of the objective function is discretized and a GA technology optimization is run for all combinations of objective weights. However, the drawback to this approach is that the number of points to be run increases as d^n , where d is the number of discrete points per objective and n is the number of objectives. For problems with three or fewer objectives, this is relatively easy. As the number of objectives increases, the dimensionality of the problem makes an exhaustive search impractical.

One way to circumvent this problem is to employ an edge search procedure. The idea behind an edge search procedure is to examine the periphery of an area and use that information to interpolate across the area. Since edges are an essentially infinitesimal portion of the entire area, one can evaluate them at minimal cost, yet still get a very good idea of what lies between the edges. So for a technology selection problem involving n objectives, the number of edges to be examined is:

$$Edges = \sum_{n=2}^n n-1 = {}^n C_2 \quad (5)$$

where each edge is essentially a Pareto front between two objectives. Thus, one can recursively apply the Pareto procedure described in step 2 to obtain a good estimate of how much “Pareto hypersurface area” each technology covers and do so at minimal computational cost.

3.4. Step 4: Results Visualization

The final step in the technology analysis process is to present the results in a form that is useful for making technology decisions. As mentioned previously, a Pareto front and a technology state transition diagram are sufficient to completely characterize the optimal technology set along one edge. If only two objectives are being optimized, then these two plots characterize the entire objective space.

For problems involving three objectives, the objective space has three edges (Pareto fronts), and can be visualized in the form of a ternary plot, as shown in Fig. 4. The ternary plot has three vertices, each representing a pure-optimal solution for one of the objectives. For instance, the vertex labeled objective 1 in Fig. 4 represents an optimal technology solution having 100% weighting on objective 1. Contours of constant objective weighting plot as straight lines parallel to one of the sides, as shown in Fig. 4. *The area inside the triangle represents every possible combination of weightings between three objectives.* Therefore, at every point inside the triangle, a given technology will either be part of the optimal set or it will not. The best “all around” technologies are those that cover the most area inside this triangle.

This technique will be demonstrated later in this paper for a commercial engine technology selection problem involving three conflicting objectives. Although the ternary plots become cumbersome for problems involving more than three objectives, the basic method still applies. However, the resultant hypersurface area of n-objective problems cannot be visualized as easily and it quickly becomes more practical to express technology desirability in terms of numerical hypersurface area.

4. Implementation

This section details an application of the ideas described previously to a technology selection problem for a commercial turbofan engine. This study is typical of those conducted during the preliminary design stages of engine development and involves 40 engine technologies as possible candidates for implementation into a next-generation commercial aircraft engine. It is not feasible to incorporate all of these technologies into a single design, so the objective is to determine the subset of technologies representing the optimal compromise between increased performance (reduced fuel burn), reduced manufacturing cost, and lowest possible development risk.

The specific example cited in this paper is based on a study described in detail in Ref. 1. It is important to note that GEAE technologists evaluated the technologies considered in this study and the study was conducted in close cooperation with GEAE designers. This study therefore represents an actual industrial application of the techniques described herein, and not merely an academic study.

4.1. Problem Description

The baseline engine used for this study is a current technology large commercial turbofan engine. This engine was modeled using GEAE analysis tools and incorporates thermodynamic cycle and flowpath analyses. This model also includes provisions for modeling engine technologies such that their impact can be estimated. The primary outputs of interest from this analysis are engine performance data (an engine deck), engine weight, and fan diameter. The latter impacts engine installation, so the base aircraft configuration is adjusted as the engine weight and nacelle drag vary. The aircraft model characterizes a typical mission, which is used to assess the impact of engine technologies on the engine/airframe system.

The baseline aircraft used in this study is a 300 passenger, twin engine, long-range transport. The aircraft empty weight (less engine weight) is fixed for the purposes of this study and standard commercial mission rules and assumptions are applied. The aircraft mission was modeled using a standard mission analysis code and reflects current technology for this class of aircraft. The primary aircraft performance figure of merit used in this study is fuel burn for a 6,000 nmi (6K) mission.

Manufacturing cost and technology risk were also important considerations in this study. Since




Fig. 4 Ternary Plot for Visualization of 3-Objective Technology Selection Problems.

expediency was a key requirement, a very simple technology readiness model was devised. This involved first estimating technology risk using NASA's technology readiness level (TRL) rating system. [2] Next, technology risk is calculated by taking the compliment of technology readiness (i.e. risk = 10-TRL). Finally, cumulative technology risk for an arbitrary group of technologies is assumed to be the sum of the individual risk scores.

In like fashion, manufacturing cost was rated based on expert opinion using a "relative manufacturing cost index." In this scheme, the manufacturing cost of today's state-of-the-art turbofan engines is taken as the baseline and assigned a score of zero. Each technology is then rated relative to the baseline, with numbers less than 0 indicating decreased cost of manufacture, and vice-versa. The relative cost of an arbitrary group of technologies is then taken to be the sum of the individual cost scores. This provided a suitable mechanism to capture technology impact on manufacturing cost while avoiding the need to create complicated engine cost models. It is important to note that the ability to simultaneously optimize on analytical data as well as qualitative scores is one of the chief strengths of the GA approach to technology selection.

4.2. Analysis Method

The analysis method used to evaluate these 40 engine technologies is that described in Ref. 1, with some minor additions. The first step was to create a TIES model for the technology set. This was accomplished using GEAE analysis tools in conjunction with GEAE technologists. For this example, the model was validated and found to be reasonably accurate in predicting technology impact on system (i.e. aircraft) performance.

This model was implemented in the MATLAB® environment such that it could be used as part of a larger "technology toolbox". In addition, a tournament-style genetic algorithm was implemented in MATLAB® to find optimal technology solution sets using the TIES model as its evaluation function. A Pareto front generation routine was developed to calculate technology skylines and find Pareto fronts. This Pareto routine works by recursively calling the GA-TIES routine for a variety of objective weights. As a practical matter, it usually requires 30-50 objective weighting cases along a Pareto front to get a good definition of the technology skyline. Therefore, one should

regard the 'd' parameter mentioned previously in the method section as being on the order of 40 for most technology studies of practical interest.

5. Results

Results for the 40-technology commercial turbofan technology optimization problem are shown in Fig. 2. Recall that this figure depicts a technology frontier for cumulative technology risk versus 6,000 nmi mission fuel burn for the 40 technologies. Mission fuel burn is given as a percentage, relative to the baseline (no technology) case. Fig. 2 clearly shows that the best 6K fuel burn achievable using these 40 technologies is roughly a 7% reduction relative to the baseline. This reduction comes at the expense of considerable technology risk. This plot also shows that if risk is a concern, one could elect to use a technology set that yields a 6% reduction in fuel burn while having 60% less technology risk than the optimal fuel burn technology set. If risk is a major concern, a 4% reduction in fuel burn with minimal technology risk might be an acceptable solution. Fig. 2 shows that such a solution is possible.

Fig. 2 was created by parametrically varying the relative weighting between reduction in fuel burn and technology risk. Therefore, the points in the upper left corner of this figure represent solutions for the pure fuel burn case while those in the lower right corner represent the minimum technology risk solutions. Each point on this figure therefore represents one possible combination of technologies. It is useful to know how this optimal technology solution set changes from one extreme to the other.

The change in optimal technology solution set as a function of relative objective weightings can be described in terms of a *technology state transition diagram*, as shown in Fig. 5. The abscissa of this figure is a listing of technologies, numbered 1 through 40. The ordinate shows the objective function weighting. Technology state is depicted as a grid of squares in the figure, with a black square indicating that the technology is part of the optimal solution set. It is evident from this figure that no technologies are optimal for the 100% risk-weighted case (which corresponds to the lower right corner of Fig.2). As the relative weighting is parametrically shifted towards a pure fuel burn-weighted solution, most of the technologies are eventually subsumed into the optimal solution set (corresponding to the upper left in Fig. 2).




Fig. 5 Optimal Technology State Transition—Fuel Burn Versus Technology Risk.

Fig. 5 shows that almost all technologies considered in this study provide a beneficial impact on fuel burn. Note that the state transition point is roughly centered around 50% weighting, but the exact transition point for any given technology may be quite far from the mean. This point is punctuated by the fact that some technologies never reach a transition point, an indication that they are not beneficial to either reduced risk or fuel burn (technologies 5, 18, 24, 26, and 30).

If the same Pareto analysis is run multiple times, the exact transition point of each technology may fluctuate slightly due to the stochastic nature of the GA, but these fluctuations are on the order of a 1-2% shift in objective weightings. This suggests that the “technology skyline” shown in Fig. 5 is a function of the risk to benefit ratio provided by each technology. Low risk/benefit technologies are among the first to enter the optimal set, while high risk/benefit technologies are slow to be included in the optimal set. Consequently, *the height of each dark band in the technology skyline can be used as an ordinal measure of technology desirability*, and is even useful as a quantitative measure of desirability.

A similar exercise can be used to examine the Pareto optimal set describing the frontier of 6K fuel burn versus relative manufacturing cost, shown in

Fig. 6. This figure looks qualitatively very similar to Fig. 2 in that the best fuel burn reduction is still on the order of 7% better than the baseline. However, note that even the manufacturing-optimal solutions exhibit a 1% reduction in fuel burn *in addition to being less expensive* than the baseline. This is because several of the technologies considered in this study have a beneficial impact on both fuel burn and manufacturing cost, and are therefore always part of the Pareto optimal set regardless of objective weighting. Fig. 6 is an explicit illustration of how new technologies can drive the frontier of propulsion

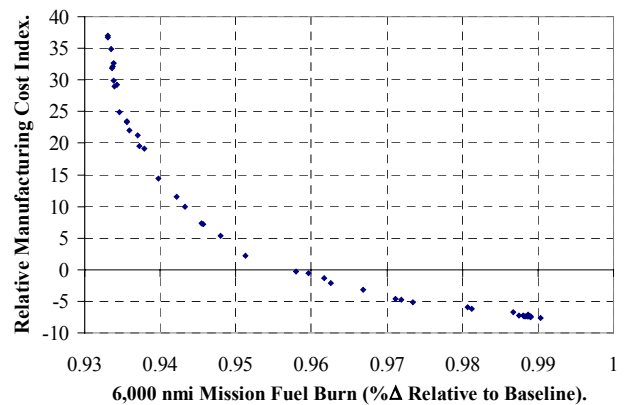


Fig. 6 Pareto Front of 6,000 nmi Mission Fuel Burn Versus Relative Manufacturing Cost.




Fig. 7 Optimal Technology State Transition—Fuel Burn Versus Manufacturing Cost.

capabilities into regions previously unreachable using older technology.

The technology state transition diagram for fuel burn versus manufacturing cost is shown in Fig. 7. This figure shows that technologies 5, 22, 24, 26, 27, 35, and 38 are desirable for reducing both fuel burn and engine manufacturing cost. Technologies 5 and 35 are desirable for reducing manufacturing cost, but do so at the expense of increased fuel consumption, and so exhibit state transition opposite to that exhibited by most other technologies. Technologies 18 and 30 are not beneficial to either fuel consumption or manufacturing cost, and are therefore never part of the Pareto optimal set.

A third Pareto front of interest is risk versus manufacturing cost. This can be constructed in like fashion to the previous two, and is shown in Fig. 8. This figure shows that it is possible to reduce relative manufacturing cost considerably below the baseline, but only by incurring considerable technology risk. Conversely, one can reduce

technology risk to zero, but only for the trivial solution. The state transition diagram corresponding to this is shown in Fig. 9. Once again, there are no technologies present in the optimal risk solution set, as it is the trivial solution. The optimal manufacturing cost solution set contains technologies 5, 22, 24, 26, 27, 35, and 38, just as it




Fig. 8 Pareto Front of Cumulative Technology Risk Versus Relative Manufacturing Cost.




Fig. 9 Optimal Technology State Transition—Technology Risk Versus Manufacturing Cost.

did before in Fig. 7. Not surprisingly, most technology transitions occur near the 50% weighting case. The skyline of Fig. 9 shows that technologies 26 and 38 are the most desirable from a combined cost and risk perspective.

It is possible to categorize each technology’s set membership based on the results of the previous three Pareto analyses. Specifically, the 40 technologies considered herein can be divided into three overlapping optimal sets: fuel optimal, risk-optimal, and cost-optimal sets, as shown in Fig. 10. The technologies present in all three sets are the most desirable, followed by technologies contained in two sets and lastly, those contained in only one set. Those technologies not contained in any set are not desirable under any set of objective weightings.

The state transition diagram shown in Fig. 9 indicates that no technologies are desirable from a risk point of view. This solution is trivial and relatively obvious: any new technology, even if it is a mature technology, presents some degree of risk from an implementation point of view. Therefore, if risk is the only concern, the best thing to do is nothing. However, the trivial solution of no technologies is not particularly interesting. One could alternatively define technologies with a TRL of 9 to be risk optimal, for after all, a TRL of 9 is the best possible risk for a technology.[†] If this is the case, then there are three technologies that are part of the risk-optimal set: 3, 8, and 34. Examination of Fig. 5 shows that these technologies are also part of the fuel burn-optimal set, as reflected in Fig. 10.

It is evident from Fig. 10 that most of the technologies considered in this study are primarily

[†] This definition of risk optimality has an implied assumption that the collective TRL of a technology group is equal to the sum of its part and there are no interactions between TRLs. This is not always a valid assumption, but it is satisfactory for the purposes of this study.

beneficial only to fuel burn, while two technologies are beneficial only to manufacturing cost. Only a handful of technologies belong to overlapping sets. Technologies 18 and 30 are undesirable for *any* combination of fuel burn, risk, and cost weighting.

This “technology set representation” is useful for classifying each technology according to the type of benefit it yields. Non-contributing technologies are clearly identified, and contributing technologies are classified according to their usefulness in improving the three objectives. However, this set membership representation does not show any information as to the *degree of membership* to each set. Rather, it is valid only at the extreme weightings: pure fuel burn-optimal, pure risk-optimal, etc. Yet, *the weighting scenarios having near-equal balance are the ones that are of the most practical interest* because the design objective in selecting technologies is to find the set yielding a balance amongst many conflicting objectives.

Since this is a three-objective problem, the global behavior of the optimal technology solution set can be depicted using ternary plots as discussed

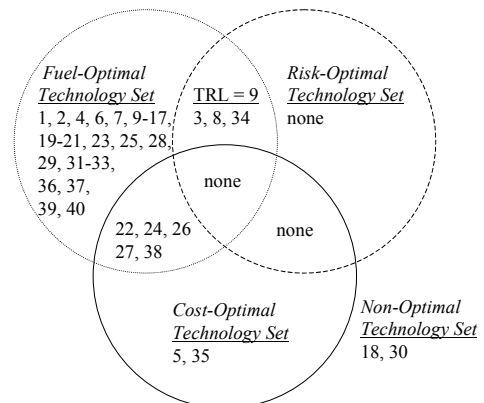


Fig. 10 Set Membership of Fuel, Cost, and Risk-Optimal Technology Solution Sets.




Fig. 11 Ternary Plots for 40 Technology Candidates Mapped Against 3 Objectives.

previously. If an edge search procedure is used, one need only create Pareto fronts along three edges (as done previously in this section). The technology state transition data in Figs. 5, 7, and 9 is all that is necessary to characterize the technology space.

The result is a set of 40 ternary plots, one for each technology, as shown in Fig. 11.[‡] These figures each have a shaded region, and this region represents the portion of the objective space over which that technology is part of the optimal solution set. For example, if the entire triangle were solid black, it would mean that the technologies were desirable for any possible combination of weights. Conversely, a white triangle indicates that the technology is not optimal for any weighting scenario. This figure visually shows which are the best “all around” technologies and which are not.

This same set of plots can be superimposed on one another to give a better feel for the relative area covered by each technology, as shown in Fig. 12. This figure shows the 40 technologies stacked in order of increasing desirability, with the best technologies being on bottom and the least desirable technologies being on top. Note that two technologies (18 and 30) do not cover any space in the ternary plot, as one might expect based on the result from Fig. 10.

Figs. 11 and 12 clearly show the order of desirability for all 40 technologies. The dominant technologies are 22, 38, 27, 26, 5, 40, and 8. These technologies were selected without need to specify objective weightings, but are rather the technologies that exhibit the broadest and most robust utility over the range of all possible objectives. One could therefore think of these as the “robust” optimal technology set. This set should be the point of

[‡] Note that in addition to three edges, and fourth Pareto front was run to get data along the perpendicular bisector between fuel and cost. This was done to test how much curvature is actually in the boundaries. The results of Fig. 11 show that the curvature is small, thus giving some measure of confidence that the results obtained with an edge search are accurate.




Fig. 12 Technology Desirability Expressed as Superimposed Ternary Plots.

departure for detailed technology studies prior to a final down-select on engine technologies.

6. Conclusions

This paper illustrated a GA-TIES approach to quickly analyze combinatorial optimization problems typically encountered in technology trade studies. It was shown that this technique: 1) simultaneously optimizes quantitative and qualitative technologies, 2) can handle any number of technologies, and 3) can evaluate against any number of objectives (though visualization becomes more complicated). This technique allows visualization of technology optimality over wide ranges of objective weightings, and this was demonstrated using ternary plots for the 40-technology commercial engine problem. If used properly, this approach circumvents the need for explicit objective weightings, shows the limits of a given technology set, and provides a metric to measure robustness.

7. Acknowledgements

The authors would like to thank GE Aircraft Engines and the Office of Naval Research for supporting this research.

8. References

- [1] Roth, B.A., German, B.J., Mavris, D.N., Macsotai, N.I., “Adaptive Selection of Engine Technology Solution Sets from a Large Combinatorial Space,” Presented at the 37th Joint Propulsion Conference, Salt Lake City, UT, July 2001, AIAA2001-3208.
- [2] Kirby, M.R., Mavris, D.N., “A Method for Technology Selection Based on Benefit, Available Schedule and Budget Resources,” Presented at the 5th World Aviation Congress and Exposition, San Diego, CA, October 10-12, 2000, SAE 2000-01-5563.
- [3] Hayter, A.J., *Probability and Statistics for Engineers and Scientists*, PWS Publishing, Boston, 1996, pp 331.