10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference
30 August - 1 September 2004, Albany, New York

AIAA 2004-4314

A Framework for Aircraft Conceptual Design and Environmental Performance Studies

Nicolas Antoine* and Ilan Kroo†
Stanford University, Stanford, California, 94305

Karen Willcox‡ and Garret Barter§
Massachusetts Institute of Technology, Cambridge, Massachusetts, 02139

Although aircraft environmental impact has been a concern since the beginning of commercial aviation, continuous growth in passenger traffic and increasing public awareness make aircraft noise and emissions critical considerations in the design of future aircraft. This research explores the feasibility of including environmental performance as an optimization objective at the aircraft conceptual design stage, allowing a quantitative analysis of the trade-offs between environmental performance and operating cost. A program for aircraft design and optimization was developed, using a multiobjective genetic algorithm to determine optimal aircraft configurations and to estimate the sensitivities between the conflicting objectives of low noise, low emissions, and operating costs. The design tool is based on a new application integration framework incorporating a detailed noise prediction code, engine simulator, and aircraft analysis and optimization modules. This paper describes the framework and design approach, including initial results that illustrate environmental performance trades and explore the feasibility of very low-noise and low-emissions designs that could dramatically decrease the environmental impact of commercial aviation.

I. Introduction

Continuing growth in air traffic and increasing concern over aircraft noise and emissions have made environmental considerations one of the most critical aspects of commercial aviation today. It is generally accepted that significant improvements to the environmental performance of aircraft will be needed if the long-term growth of air transport is to be sustained. The IPCC has projected that, under an expected 5% annual increase in passenger traffic, the growth in aviation-related nuisances will outpace improvements that can be expected through evolutionary changes in engine and airframe design.[1]

While considerable progress has been made to reduce the noise signature of airliners, the public's perception of noise continues to grow, as illustrated by the ever-increasing number of public complaints. This can be attributed to increasing air traffic as well as further encroachment by airport-neighboring communities. As a result, noise has become a major constraint to air traffic, with 60% of all airports considering it a major problem and the nation's fifty largest airports viewing it as their biggest issue.[2]

*Doctoral Student, Department of Aero/Astro, Stanford University, AIAA Student Member
†Professor, Department of Aero/Astro, Stanford University, AIAA Fellow
‡Assistant Professor, Department of Aero/Astro, M.I.T., AIAA Member
§Graduate Student, Department of Aero/Astro, M.I.T., AIAA Member

The release of exhaust gasses in the atmosphere is the second major environmental issue associated with commercial airliners. The world fleet releases approximately 13% of CO_2 emissions from all transportation sources, or 2% of all anthropogenic sources.[3] The expected doubling of the fleet in the next twenty years[4] will highlight the issue: the fractional contribution of aviation is expected to increase by a factor of 1.6 to 10, depending on the fuel use scenario.

Although dramatic simultaneous improvements in aircraft noise, emissions, and operating cost have been achieved since the 1960s,[5] gains since the mid-eighties have not been as significant and the point seems to have been reached where future improvements through technological advances will be possible only by trading off operating costs for environmental performance (Figure 1). Quantifying the terms of this trade-off — critical for the design of efficient future aircraft — is one of the main topics addressed by this research.

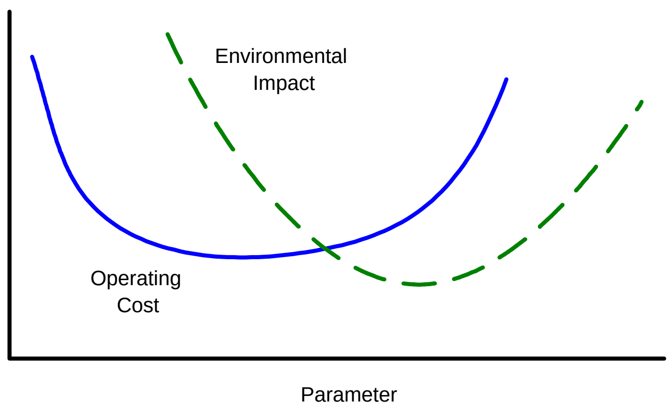

Figure 1. Cost vs. environmental impact.

In addition to the compromises between cost and environmental impact, tradeoffs between different aspects of noise and emissions arise. The Airbus A380 design was reportedly modified late in the design process to meet nighttime restrictions at Heathrow airport, using an engine fan substantially larger than required for lowest fuel consumption, necessitating a redesign of the engine, nacelle, pylon and wing. These modifications resulted in a 1-2% increase in fuel burn for a 1-2 dB noise reduction,[6] considered a very expensive trade-off. In a related example, the ICAO's Committee on Aviation Environmental Protection, at its 6th meeting in early 2004, concluded that it could not demand, for new aircraft entering service in 2008, a reduction of aircraft NO_x of more than 12% relative to today's aircraft.[7] The issue was not related to technology risk, but rather a lack of information regarding the impact of further NO_x reductions on noise and other emissions. It was agreed that demanding a reduction in one type of emissions only to obtain an increase in another — by an unknown quantity — was not a viable solution. Clearly, there is a need for integrating environmental considerations early in the aircraft design process, and for more systematic investigation and quantification of the tradeoffs involved in meeting specific noise/emissions constraints.

This research intends to contribute by proposing a conceptual design tool structured to generate optimized preliminary aircraft designs based on specified mission parameters. Existing aircraft design codes were extensively modified to incorporate the parameters required to model environmental performance. Various optimizers were also created to explore the design space, while high-fidelity noise prediction codes and an engine simulator were integrated into the automated design process.

The design tool is composed of a library of routines used to compute many aspects of aircraft design and performance. NASA Langley's Aircraft Noise Prediction Program (ANOPP) is used for jet, fan, and airframe noise modeling. Engine performance is estimated using NASA Glenn's Engine Performance code (NEPP). The tool integrates these programs with other disciplinary analyses ranging from component weights to stability and control and mission performance. This was accomplished using a framework that facilitates the coupling of multidisciplinary analyses and optimization. In the present application, approximately twenty different analysis modules were combined with nonlinear optimization and a database management system to allow rapid reconfiguration of the design variables, objectives, and constraints. These approximate methods are particularly well-suited for optimization due to their rapid execution and robustness.

In addition to traditional performance constraints such as range and field performance, maximum cumulative certification noise (the sum of the noise at each certification location) is included. This approach allows the increase in environmental acceptability to be explicitly specified: from slight improvements in noise to silent aircraft.[8] NO_x emissions are constrained through the emission index of the engine and CO_2 emissions are modeled via fuel burn. Design variables include parameters from the aircraft configuration, propulsion, and flight profile.

Although other tools are available to assess the environmental impact of existing aircraft, the present approach seeks to evaluate the potential for new aircraft to satisfy the conflicting demands for low noise, emissions, and cost.

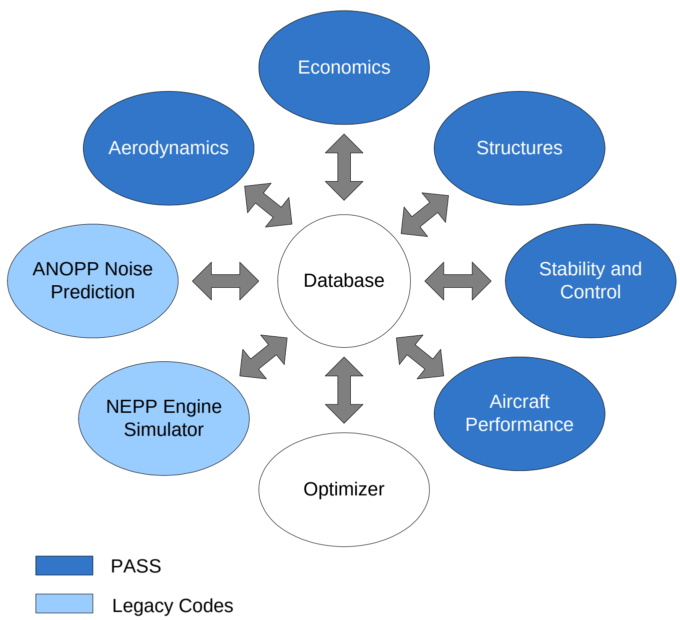

Figure 2. Elements of the Aircraft Conceptual Design Framework

II. System Summary

The system is built from a set of fundamental blocks shown in figure 3. These include a set of custom analyses specific to this design task, a group of aircraft performance modules comprising PASS, A Program for Aircraft Synthesis Studies, and legacy codes for propulsion and noise from NASA, all integrated within a generic framework for engineering design and optimization.

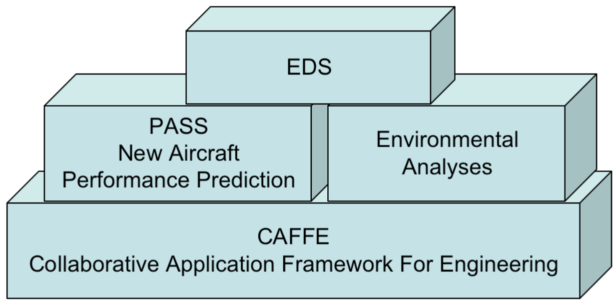

Figure 3. The system for assessing new aircraft environmental performance consists of a variety of analyses, including custom and legacy codes, integrated with the CAFFE framework.

A. CAFFE Framework

1. Overview

Underlying the aircraft design system is a software integration framework, developed by some of the authors, that simplifies the development of complex, modular design applications. This collaborative application framework for engineering (CAFFE) is intended to be more than just an application integration tool, but rather a true framework for design applications, handling data management and optimization, but also providing a platform and API for customizable user interfaces. Some of its features include:

- Platform independence with modern Java / XML implementation
- Dynamic linking of analysis modules
- Legacy code support without source
- Hierarchical database API
- Remote execution
- Integrated optimization suite

Developed initially to facilitate distributed, multi-level optimization,[9] the framework is used here as a single-level analysis integration and optimization tool.

2. Data Managment

Two approaches to data management are common among application integration frameworks: those in which applications communicate with files or a centralized persistent database, and those that keep data in memory. CAFFE permits both forms of data management with XML-based data storage and an in-memory database accessible to native applications with reduced overhead.

The database provides support for various data types, including multi-dimensional arrays and cased parameters. The latter provides convenient access to arrays, permitting code development that deals with generic parameters rather than knowing about specific instances. For example, a routine that computes drag coefficient need not know that drag coefficient is an array with individual elements corresponding to different flight conditions, rather it *gets* altitude, weight, Mach number and other scalars from the database, does the relevant computations, and *puts* the scalar C_D into the database. The framework utilities *get* and *put* recognize that scalars such as altitude and C_D are actually elements of arrays and place them in the appropriate positions based on the flight condition.

Some data is stored in memory and accessed through a hash table, again transparent to a user or developer. A program can assign a local variable to an entry in the database through a call such as: CL = db.getd("cltotal"), which returns a double precision copy of the airplane CL at the current flight condition.

Data may also be read or written to disk in XML format. CAFFE provides simple tools for reading and writing data to and from XML, reducing the complexity of developing new components for the application. XML is used here because of the many commercially-available programs for manipulating and viewing data in this form and because of the ease in which it may be transformed (Figure 4).

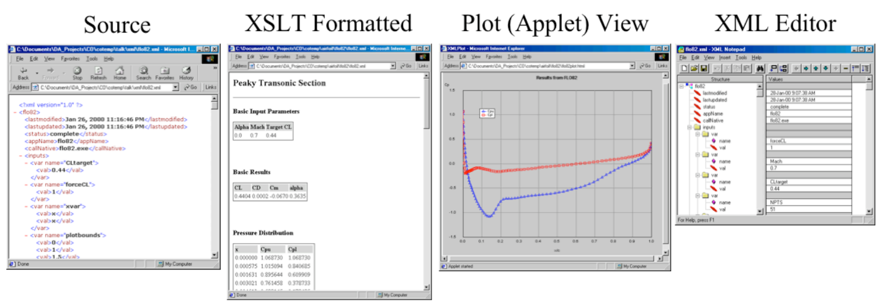

Figure 4. Data is represented in CAFFE in an internal data structure and with a persistent XML database, easily presented to users in many forms.

3. *Code Structure*

CAFFE may be used to integrate existing codes without changes to their structure – in fact, even without access to their source – by wrapping them in a simple native class. Communication with the database is done in the wrapper and data may be transmitted to a legacy application through its native input and output file formats. Although this allows the use of applications not written to be a part of the existing system, it is slow and restricts the architecture of such components to the simplest input-compute-output model. To permit more general architectures and to allow complete flexibility in program control, CAFFE defines an API (applications programming interface) for new classes using the Java programming language. The use of Java provides access to a wide range of classes, enables platform independence and dynamic linking of modules, and does not require knowledge of a framework-specific scripting language. This allows developers to create more complex analysis modules, such as adaptive quadrature methods and optimization tools that must interact with the database and other modules in a more complex manner than that available with many other integration frameworks. The following code illustrates, in a very simple case, how a Java code module may interact with other parts of the framework.

```
public class drag extends CaffeClass {

    // Define other classes to be called:
    pdrag parasite = new pdrag();
    idrag induced = new idrag();
    cdrag compress = new cdrag();
    trim trm = new trim();

    public void compute() {
        double CL, CD, loverd, doverl;
        double cdp, cdi, cdc;

        // Get required inputs from database:
        CL = db.getd("cltotal");

        // Calculations:
        trm.compute();
        parasite.compute();
        induced.compute();
        compress.compute();

        cdp = db.getd("cdp");
        cdi = db.getd("cdi");
        cdc = db.getd("cdc");

        CD = cdp + cdi + cdc;
        loverd = 0.;
        if (Math.abs(CD) > 0) loverd = CL/CD;

        // Repack database:
        db.putd("cd", CD);
        db.putd("l/d", loverd);
    }
}
```

4. Optimization and Trades

Several optimization algorithms are available from CAFFE including gradient methods, nonlinear simplex, and genetic algorithms. The framework also permits custom optimization methods to be incorporated directly and a new multi-objective optimizer was developed for use in the present application.

Although gradient-based nonlinear optimization methods provide efficient design tools, they require smooth objective and constraint functions. Especially when one is employing file-based legacy codes that may not have been written with this requirement in mind, discontinuous and noisy function values may make the use of gradient methods difficult. This is the case for the design problems considered here. Figure 5 shows how the predicted noise (from the NASA ANOPP code) changes with engine size and bypass ratio. Note that the scale is expanded greatly to illustrate the small numerical noise, but the large local gradients call for unconventional design tools.

Population-based methods tend to be computationally expensive, but offer a level of robustness not found in sequential methods: in particular, convergence to local extrema can be avoided and large local gradients can be tolerated. In this research a multi-objective, floating point genetic algorithm was used to allow tradeoffs between competing objectives. Elements of this algorithm, developed at Stanford, but based on methods described in Deb,[10] include a penalty function approach to constraints, dominance-based ranking, and solution niching

At each iteration, non-dominated solutions in the population of designs are assigned a rank of 1, and form the Pareto set of that generation. The Pareto set obtained from the final population constitutes the set of optimal solutions to the problem.

Through the generations, the genetic algorithm drives the population toward better solutions. Eventually,

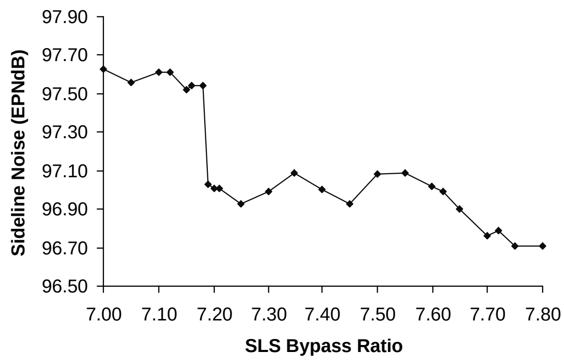

Figure 5. ANOPP numerical noise.

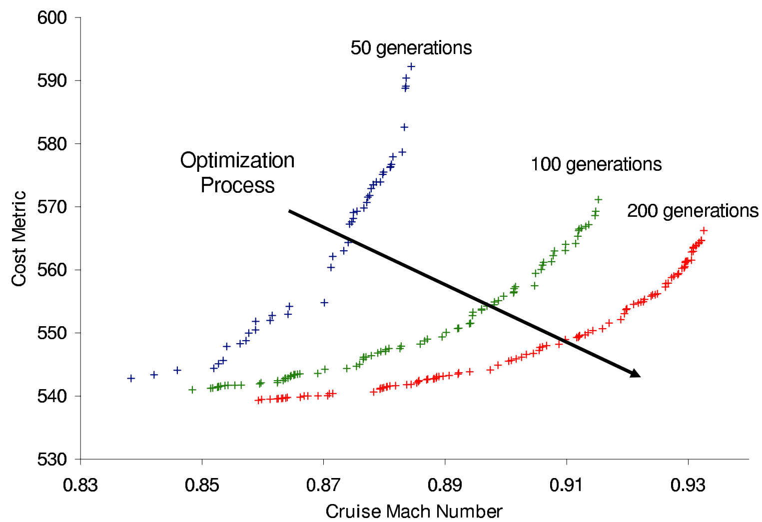

Figure 6. The Pareto front indicates the set of non-dominated solutions in a given generation. The optimization process drives the population towards their optimal values.

the quality of the population stops improving and the resulting Pareto set contains the optimal solutions. This is illustrated in Figure 6. In this example, a 200-seat, 6,000nm range aircraft is optimized simultaneously for both minimum cost and maximum cruise Mach number. With each generation, the Pareto front is pushed towards higher Mach numbers and lower costs. Eventually, the front no longer progresses and the set of optimal trade-offs between Mach number and cost is obtained.

5. Interface API

The framework was designed with the idea that engineering design is much more than application integration and optimization. Designers should be able to interact deeply with the data in ways that a numerical analyst or optimization expert might not expect. Thus, rather than just including the standard set of response surface, sensitivity, and visualization tools common to multidisciplinary optimization frameworks, CAFFE also provides an API for custom interface modules, allowing designers to manipulate geometry, view CFD solutions, and graphically explore sensitivities in ways that cannot be accomplished with a generic interface. Since the framework is Java-based, these interface modules are standard Java applets that may be embedded in html pages and viewed over the web, or may be integrated directly in the design application. In either case, the framework allows access to the database and to all of the low level tools available to the analysis routines. Such interface modules are completely optional and the complete code may be accessed entirely via command-line interactions. This is how many of the results described here were generated, but graphical versions using many of the same analysis modules have been developed (Figure 7).

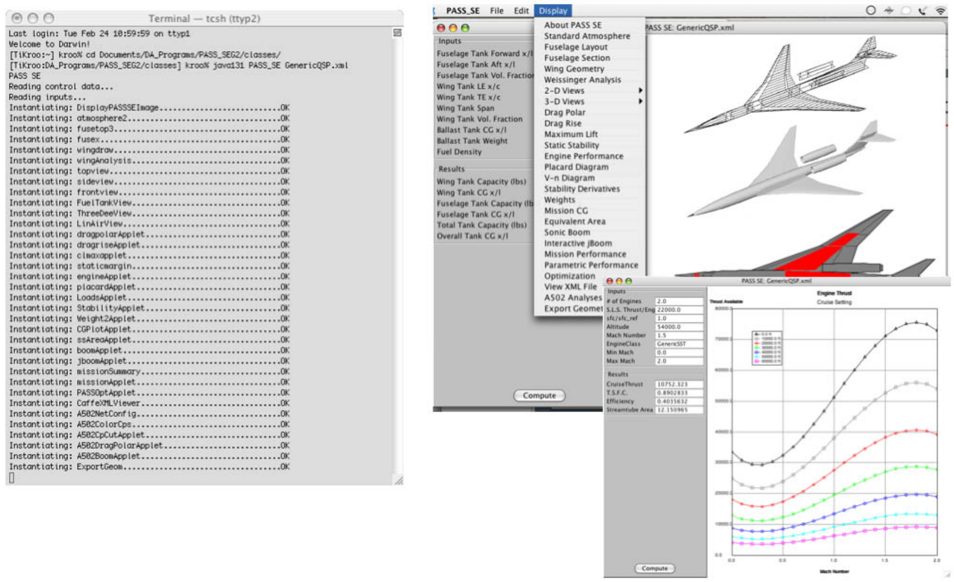

Figure 7. Interfaces may be customized – from simple command-line interaction to custom user interfaces, CAFFE defines an API, not just for analyses, but for user display and interaction.

B. Analyses

1. Airframe Performance

The Program for Aircraft Synthesis Studies (PASS) is an aircraft conceptual design tool based on a collection of McDonnell-Douglas methods, DATCOM correlations, and new analyses developed specifically for conceptual design. PASS allows the rapid generation of a design and contains modules to compute many aspects of aircraft design and performance: from vehicle geometry, to drag and weight build-ups, stability and trim calculations, and aircraft performance. In addition, these codes have been specifically designed to be integrated with an optimizer and utilize the CAFFE framework for data management.

The system has evolved over more than 15 years[11] and now forms the basis of a two-quarter graduate-level aircraft design course at Stanford University taught, over the years, by Professors Richard Shevell, Ilan Kroo, and Juan Alonso. A detailed description of these methods may be found on the AA241 Aircraft Design: Synthesis and Analysis website[12] and are very similar to those described by Schaufele .[13] These methods were extended by the authors to improve low speed aerodynamics modeling and incorporate sensitivities of drag and weight to some of the additional engine parameters introduced in this study. Further details on the methods including comparisons with existing aircraft are given in Antoine[14].[15]

2. Propulsion Modeling

The propulsion decks in PASS are rubberized to allow engine sizing but do not allow investigation of certain engine parameters that may be particularly important in the design of aircraft for which environmental characteristics are paramount. For this reason, and to demonstrate the integration of legacy codes in the CAFFE framework, a NASA propulsion simulator code was used to compute engine characteristics.

Developed at NASA Glenn, NEPP is a 1-D steady thermodynamics analysis program. At the design point, NEPP[16] automatically ensures continuity of mass, speed, and energy by varying the scale factors on the performance maps for the compressor and turbine components. Off-design operation is handled through the use of component performance tables and minimization of work, flow, and energy errors. The engine is then balanced by altering free variables of available components.

Variable controls can also be used to obtain a specified performance. For example, airflow or combustion temperature can be varied to reach a desired thrust level. Controls are also used to limit and optimize engine parameters. For the purpose of the present design tool, the range of variables has been selected to accommodate technology that would be available by the end of the decade, including increased combustion temperatures and higher turbomachinery efficiencies — for instance, bypass ratios ranging from 4 to 15 are acceptable (in this study, bypass ratio is calculated at sea-level static thrust conditions).

3. Noise and emissions

Because a 1-2 dB difference in aircraft noise is significant, high-fidelity noise prediction is essential, even during conceptual design. The Aircraft Noise Prediction Program (ANOPP) is a semi-empirical code that incorporates publicly available noise prediction schemes and is continuously updated by NASA Langley.[17] As progress is made in the field of aeroacoustics, ANOPP is enhanced with the latest prediction methods. Hence, using ANOPP involves accepting a certain technology level – all designs considered feature the same noise prediction methodology: a "state-of-the-art" is assumed. As part of this research, three noise sources are considered: fan turbomachinery, jet, and airframe. Other noise sources, such as combustion, turbines, and compressors were not considered because of their relatively minor contribution to total aircraft noise.

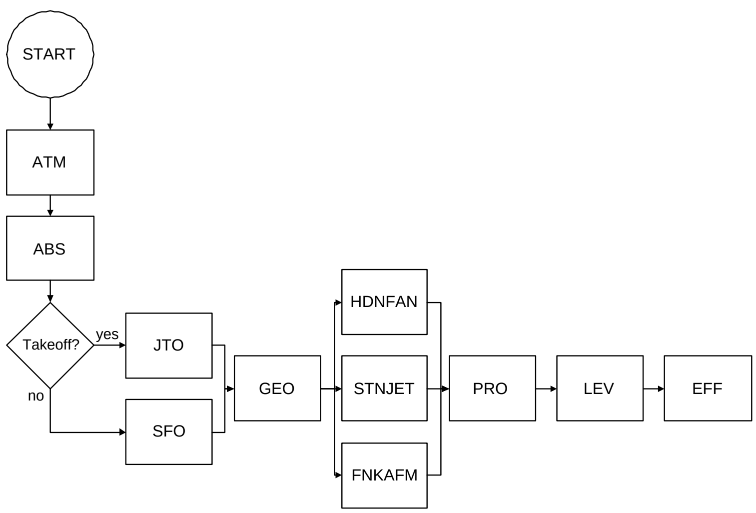

Figure 8. Flow chart of ANOPP program modules.

A flow chart of the ANOPP system is shown in Figure 8. The procedure begins by defining an atmosphere using the Atmosphere Module (ATM), followed by the atmospheric absorption module (ABS). The steady flyover module (SFO) is used for the approach measurement point, and the jet takeoff module (JTO)for sideline and takeoff measurement points. The geometry module (GEO) computes the range and directivity angles from the observer to the noise source. At this point, the various noise sources modules are run: Heidmann's[18] for fan noise (HDNFAN), Stone's[19] for coaxial jet noise (STNJET) and Fink's[20] for airframe noise (FNKAFM).

Once data has been generated by the noise source modules, the propagation module (PRO) applies corrections to the noise data in the source frame of reference to transfer it to the observer frame of reference.

Atmospheric absorption effects are applied at this point. The noise levels module (LEV) computes the tone-corrected Perceived Noise (PNLdB),and the effective noise level module (EFF)is run next to compute the EPNdB levels used as noise metrics in this research.

Both particulate and gaseous pollutants are produced through the combustion of jet kerosene (products in italic stem from non-ideal combustion):

REACTANTS	Air $N_2 + O_2$
	Fuel $C_nH_m + S$
PRODUCTS	$CO_2 + H_2O + N_2 + O_2 + NO_x + UHC + CO + C_{soot} + SO_x$

The greenhouse gases carbon dioxide CO_2 and water H_2O are the major products. Minor emissions formed during combustion include nitrous oxides (NO_x), unburned hydrocarbons (UHC), carbon monoxide (CO), and soot (C_{soot}).

ICAO regulations for the landing-takeoff (LTO) cycle cover NO_x, CO, unburned hydrocarbons, and smoke emissions.[21]

NO_x emissions are computed based on engine fuel flow (expressed in kg/s) and the combustor emission index (EI, expressed in g of NO_x formed per kg of jet fuel used), a strong function of power setting, during a take-off and landing cycle involving four different throttle modes: 100% (take off), 85% (climb), 30% (approach) and 7% (idle). Time in mode is simulated as follows: 0.7 minutes for take off, 2.2 minutes for climb, 4 minutes for approach, and 26 minutes for taxi/ground idle. The sum of the emissions at these four conditions (expressed in kg) is used to determine the amount of NO_x emitted per LTO cycle:

$$\text{LTO NO}_x = \sum \text{Fuel Flow} \times \text{EI}_{\text{NO}_x} \times \text{Time in Mode} \tag{1}$$

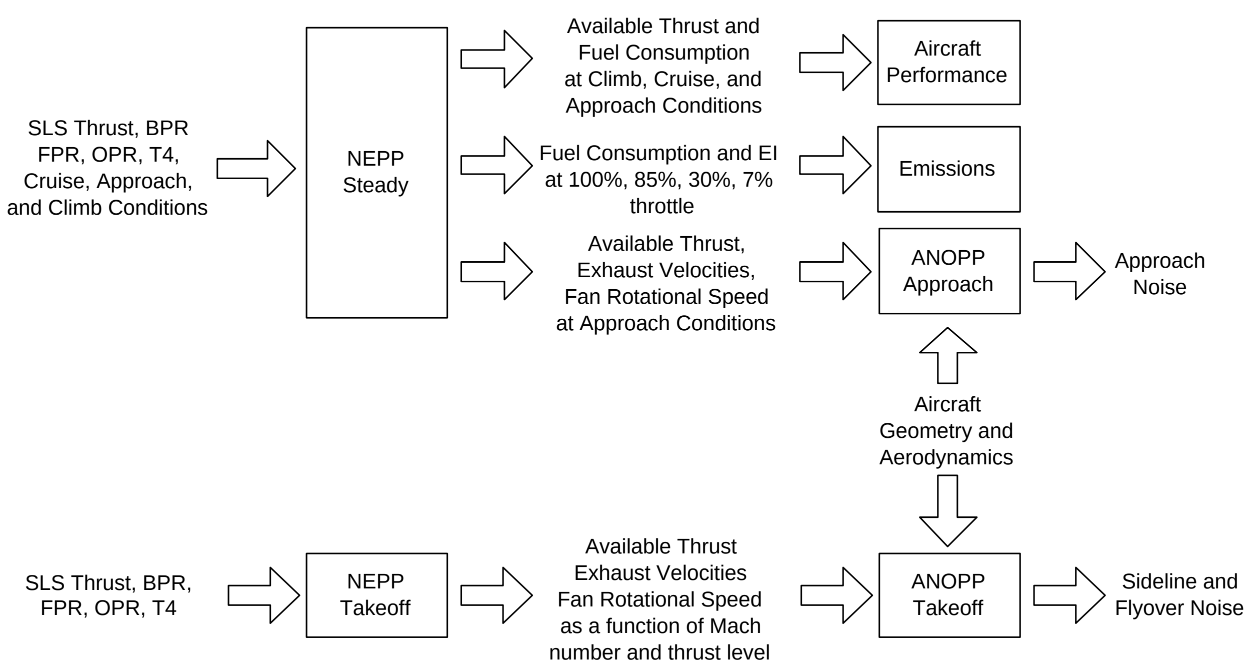

Figure 9. ANOPP and NEPP integration in the framework.

The engine design point is determined by running NEPP at sea-level static (SLS) condition, given combustor exit temperature, overall pressure ratio, desired sea-level static thrust, bypass ratio, and fan pressure ratio. The engine is run off-design for a variety of conditions, as required for noise prediction, emissions, and overall aircraft performance (Figure 9).

At off-design conditions, for example part-power operation, the engine must be balanced using a free variable: burner exit temperature (T_4) is decreased to obtain the desired fraction of maximum thrust.

PASS also requires available thrust and fuel consumption at various conditions to compute overall aircraft performance. Engine out performance is required to meet emergency climb requirements.

III. Example Results

The design tool allows trade-offs (using the multi-objective genetic algorithm) and case studies to be quickly completed in order to visualize the interrelationships between noise, emissions, and other aircraft performance metrics. A selection of example results is included below.

1. Studies regarding the effects of cruise altitude on the global atmosphere have focused solely on existing aircraft operating off-design. Using the present tool, we optimized aircraft to operate at lower cruise altitudes, minimizing the operating cost impact. Preliminary results are shown in Figure 3 for a 250-seat, 6,000nm range aircraft.

2. Employing multi-objective optimization the tradeoff between conflicting objectives may be visualized. Figure 4 illustrates this capability: selecting an aircraft with a cumulative certification noise reduction of 6 EPNdB results in a 45% increase in LTO NO_x. Each point of the Pareto set represents a design.

3. To understand the penalty associated with such design trades, an operating cost metric was minimized together with trip CO_2 emissions and cumulative certification noise (Figure 5).

A. Aircraft Mission, Variables, and Constraints

This section illustrates the optimization process performed by the design tool in the case of a 280-passenger, twin-engine airliner with a 6,000 nm range, and takeoff, cruise, and landing performances in line with industry standards for similarly-sized aircraft. The 15 design variables are listed in Table 1, split in three groups: aircraft geometry, engine parameters, and performance. Constraints are shown in Table 2.

Variable	**Units**	**Min**	**Max**
Maximum Take-Off Weight	lbs	280,000	550,000
Wing Reference Area	ft^2	1,500	4,000
Wing Thickness-over-Chord	%	0.07	0.20
Wing Location along Fuselage	%	0.2	0.6
Wing Aspect Ratio	—	4.0	15.0
Wing Taper Ratio	—	0.1	0.7
Wing Sweep	deg	0.0	40.0
Horizontal Tail Area	ft^2	225	600
Sea-Level Static Thrust	lbs	40,000	100,000
Turbine Inlet Temperature	°F	3,000	3,300
Bypass Ratio	—	4.0	15.0
Engine Pressure Ratio	—	40.0	60.0
Initial Cruise Altitude	ft	20,000	40,000
Final Cruise Altitude	ft	20,000	50,000
Cruise Mach Number	—	0.65	0.95

Table 1. Variable names, units, and minimum and maximum allowable values for the optimization problems.

Constraint	Units	Value
Cruise Range	n.miles	$\leq$ 6,000
Takeoff Field Length	ft	$\leq$ 9,000
Landing Field Length	ft	$\leq$ 8,000
Engine Out Climb Gradient	—	$\geq$ 0.024
Drag-to-Thrust Ratio	—	$\leq$ 0.88
Stability Margin	—	$\geq$ 0.18
Wing Cruise Lift Coeff. Margin	—	$\geq$ 0.01
Tail Rotation Lift Coeff. Margin	—	$\geq$ 0.01
Tail Cruise Lift Coeff. Margin	—	$\geq$ 0.01
Tail Landing Lift Coeff. Margin	—	$\geq$ 0.01
Wing Span	ft	$\leq$ 260.0

Table 2. Constraints for the optimization problems.

B. Operating Cost vs. Cruise Emissions, LTO NO_x Emissions, and Noise

The process of obtaining a low-rank Pareto front can be significantly accelerated by first computing the extreme points of the fronts. This is done by running a single-objective version of the genetic algorithm. These optimal designs are subsequently inserted into the initial population of the multiobjective problems. The resulting Pareto fronts of fuel carried, NO_x emissions, and cumulative noise margin vs. cost are shown in Figure 10. Key parameters for the optimized extreme designs are summarized in Table 3.

The configuration leading to minimum operating cost (Design A) was computed first by running the design tool without specifying any noise or emissions constraints. This aircraft is considered as the baseline and is representative of existing aircraft. Reflecting the impact of block time on the cost function, the cruise Mach number is higher than would be required for minimal fuel burn (Design B). Fuel plays a dominant role in the cost calculation, as illustrated by the similarities in the designs for minimum cost and minimum fuel carried (and therefore, minimum cruise CO_2, SO_2, and H_2O). This tight coupling is also reflected in the relatively small fuel-cost trade space (notice the fuel-cost Pareto front is narrow). At the engine level, noticeably, both designs attain high fuel efficiency via large pressure ratios and high turbine inlet temperatures.

Optimizing the aircraft for lowest fuel yields a 10% decrease in fuel carried (propagating through the design to yield a 5.4% decrease in maximum takeoff weight) but with a cost increase of 2% relative to low-cost Design A.

Since the generation of NO_x emissions is a strong function of combustor exhaust temperature and compression ratio, Design C compromises fuel efficiency for low NO_x emissions by reducing the engine overall pressure ratio and combustor temperature. The resulting 12% reduction in sea-level static thrust relative to the low-cost design mandates that the aircraft fly slower (Mach 0.67, close to the lower allowable limit, versus Mach 0.84) and at lower altitudes (the initial cruise altitude is reduced from approximately 33,000 ft to 28,000 ft). The result is a 53.5% drop in LTO NO_x emissions for an 9% and 13% increase in operating cost and fuel consumption, respectively.

These divergent requirements for low NO_x and low fuel designs are well illustrated by a wide, and very smooth, Pareto front. As a result of the lower cruise Mach number, wing sweep is reduced from 34 to 11 degrees. With significantly reduced available thrust, the wing taper ratio is increased from 0.10 to 0.39 to increase the maximum lift coefficient during takeoff and initial climb. For similar reasons, the wing area is enlarged by 12%, contributing to an increase in maximum takeoff weight of 9%. Combined with lower available thrust, the climb performance of the low-NO_x design is significantly deteriorated: thrust-to-weight ratio drops to 0.296 — resulting in the highest cumulative certification noise of any design, over 3.5 dB louder than the baseline Design A.

	Units	Design A Min Cost	Design B Min Fuel	Design C Min NO_x	Design D Min Noise
Objectives					
Relative Cost	—	1.0	1.02	1.09	1.26
Fuel Carried	lbs	119,018	106,707	134,796	138,840
LTO NO_x	kg	30.88	29.68	14.36	41.09
Relative Noise	EPNdB	0.0	-5.13	3.66	-14.98
Variables					
Max. Take-Off Weight	lbs	372,539	352,515	407,516	473,532
Wing Reference Area	ft^2	3,461	2,942	3,887	3,578
Wing t/c	%	11.7	13.5	12.8	11.5
Wing Location	%	39.2	41.2	48.1	48.2
Wing Aspect Ratio	—	7.38	9.99	8.94	14.43
Wing Taper Ratio	—	0.10	0.10	0.39	0.1
Wing Sweep	deg	33.70	26.17	11.22	14.25
Horizontal Tail Area	ft^2	929	766	953	1,431
SLS Thrust (per engine)	lbs	68,404	67,311	60,264	100,000
Thrust-to-Weight Ratio	—	0.367	0.382	0.296	0.422
Turbine Inlet Temp	°F	3,203	3,215	3,147	3,300
Bypass Ratio	—	9.59	10.35	10.32	14.87
Engine Pressure Ratio	—	59.91	59.63	40.27	59.78
Init. Cruise Altitude	ft	32,937	30,746	28,381	31,674
Final Cruise Altitude	ft	40,790	38,734	33,288	35,486
Cruise Mach Number	—	0.844	0.739	0.669	0.664

Table 3. Data for the optimal extreme designs obtained with the single-objective genetic algorithm.

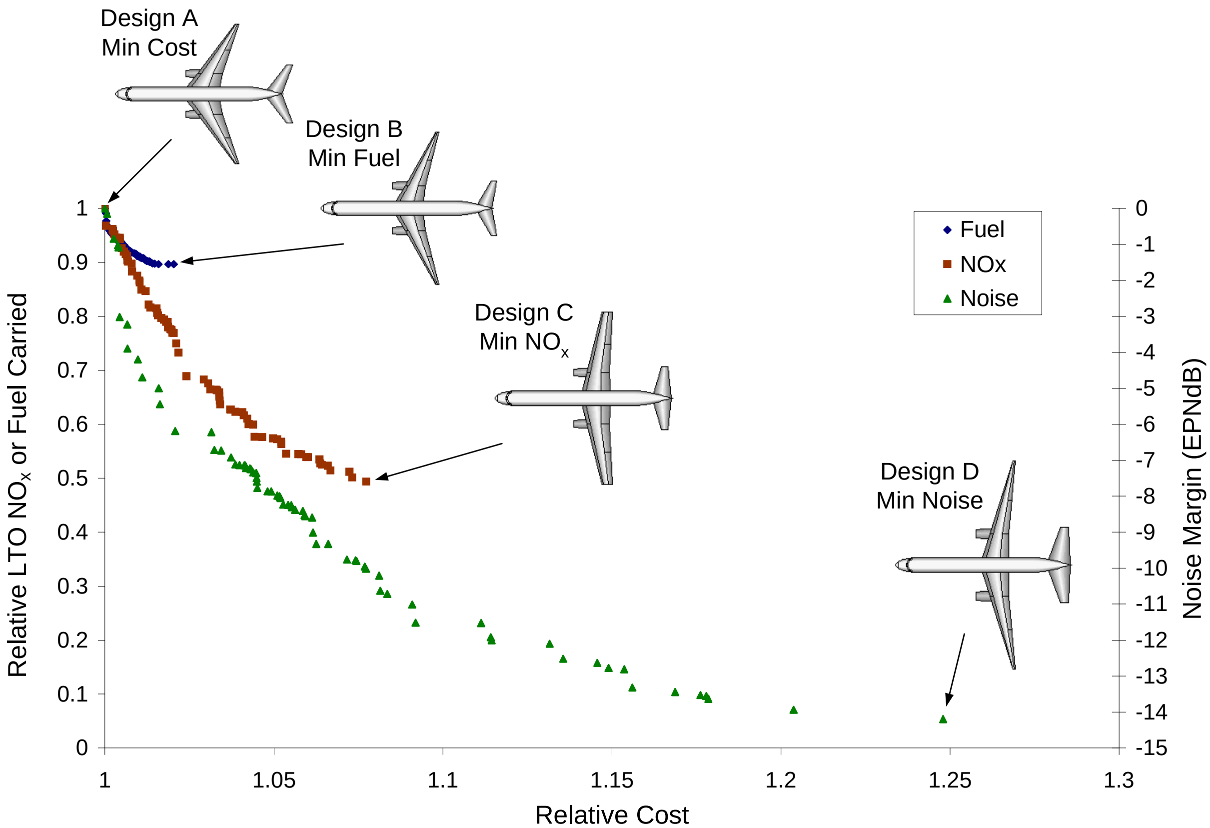

Figure 10. Pareto fronts of fuel carried, LTO NO$_x$, and cumulative certification noise vs. operating cost. Only rank 1 designs are shown. Average rank for all fronts is under 4.

The large fan necessary to reduce noise to the minimum (Design D, with a bypass ratio very close to the maximum allowable value of 15) requires more power, resulting in the selection of the highest allowable combustion temperatures and overall engine pressure ratio. The result is a 15 cumulative EPNdB reduction in noise relative to the low cost design, equivalent to a 25-fold reduction in noise energy. The penalty is a 26% increase in operating cost and 16% in fuel carried, along with NO$_x$ emissions that are 33% higher due to the increased combustion temperature. The higher thrust levels required by this high bypass-ratio engine at altitude are significant: sea level static thrust is raised from 68,404 to 100,000 lbs (the maximum allowable), a 46% increase. Higher thrust enables the aircraft to climb faster, increasing the distance to the flyover certification point and decreasing measured noise. These enormous, and therefore very heavy, engines cause a 27% maximum takeoff weight increase relative to Design A. The large frontal area, and therefore increased drag of the design, leads the aircraft to fly slower than the low cost candidate (Mach 0.66 vs. Mach 0.83). Similarly to the low NO$_x$ aircraft, the reduced cruise Mach number results in a reduced sweep of 14 degrees. A summary quantifying the trades between designs is shown in Table 4.

C. Cruise Emissions vs. LTO NO$_x$ Emissions

To explore the interrelationship between the conflicting requirements of reducing NO$_x$ and fuel-based emissions (CO_2, H_2O, and SO_2), the multiobjective optimizer was applied to the min-NO$_x$/min-fuel problem. The resulting Pareto front is shown in Figure 11, with Design B (low-fuel) and Design C (low-NO$_x$) the

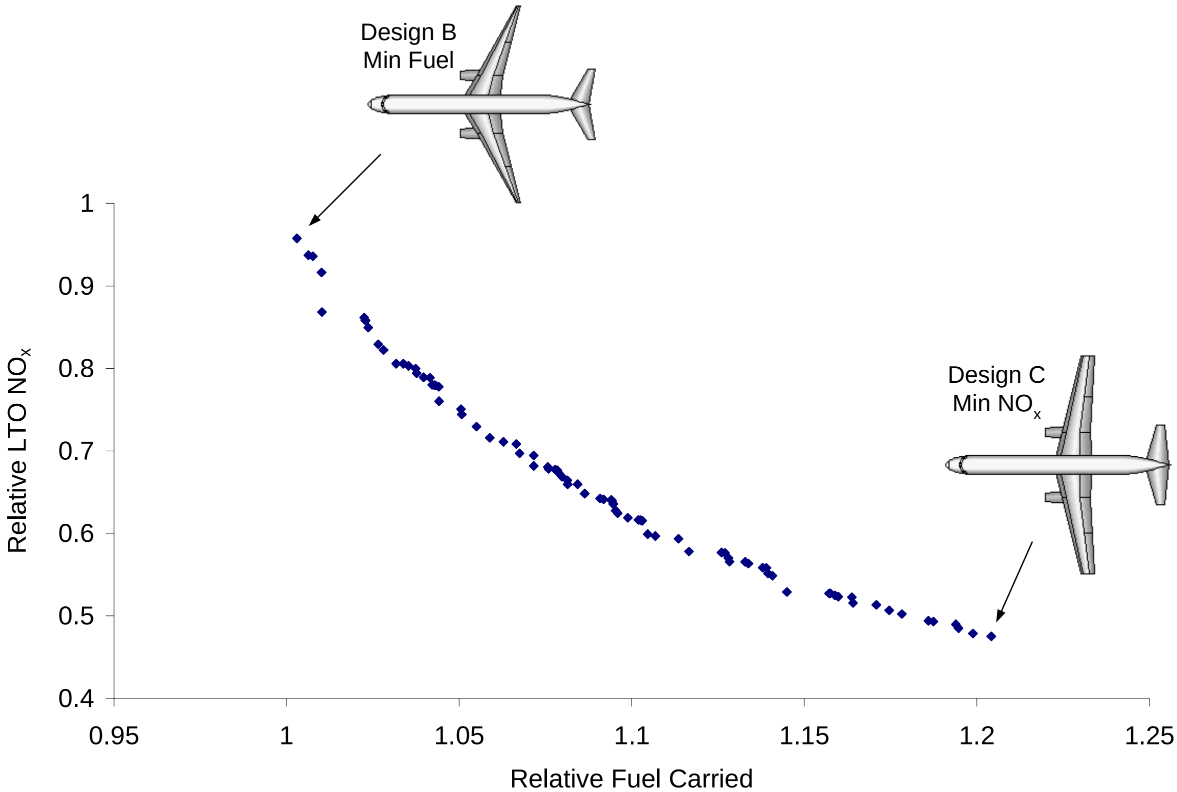

Figure 11. Pareto front of LTO NO$_x$ vs. Fuel Carried. Only rank 1 designs are shown (Average rank = 3.47).

extreme points discussed previously. According to these results, a decrease in LTO NO$_x$ of 12% (as recommended by ICAO for new aircraft after 2008 under CAEP/6) would require an increase of approximately 2% in fuel consumption and related emissions. As the demand for reductions in NO$_x$ increase, this penalty grows: the next 12% require a further 4% increase in fuel. These results illustrate the delicate trade-off that must be resolved as new regulations come into play: what is the "value" of trading one type of emissions for another?

D. Noise vs. Cruise vs. LTO NO$_x$ Emissions

This trade-off approach is expanded to include a third objective, cumulative noise. The surface that is obtained, as well as the location of the three extreme points, are shown in Figure 12). The conflicting design requirements for the min noise (Design D) and min NO$_x$ (Design C) aircraft are well illustrated here: the low-noise aircraft is also the design with highest NO$_x$, and conversely, the aircraft with lowest min NO$_x$ is the noisiest. Indeed, Designs D and C are costly to obtain and require almost complete deterioration of the other two objectives. The minimum fuel design (Design B), however, is obtained without entirely forgoing gains in noise or NO$_x$ emissions.

The usefulness of this Pareto surface is not limited to the extreme designs. Every design on the surface it is optimized for a combination of noise, fuel, and NO$_x$ performance; the impact of reducing one objective on the two others can be estimated directly from the surface.

Displaying three objectives also allows the selection of the objective to forego in order to improve the

For this increase in cost	Can reduce one of these by		
	Fuel Carried	LTO NO_x	Cumulative Noise
1%	7%	10%	3 EPNdB
2%	10%	20%	6 EPNdB
9%	10%	51%	10 EPNdB
25%	10%	51%	15 EPNdB

Table 4. Fuel carried, LTO NO_x, or cumulative noise can be traded with operating cost.

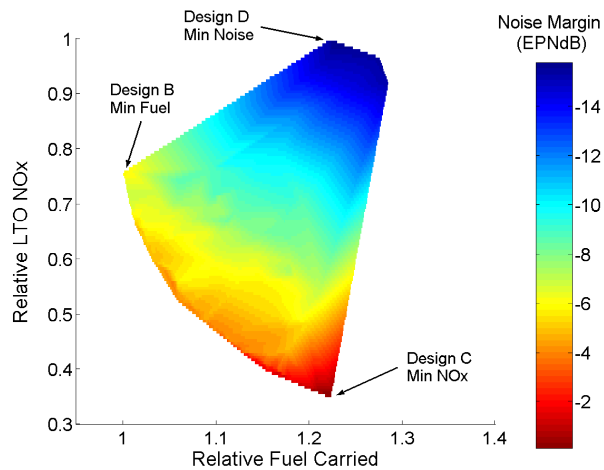

Figure 12. Pareto surface of LTO NO_x vs. Fuel Carried vs. Cumulative Noise. Only rank 1 designs are shown.

design. If the goal is to trade noise and fuel efficiency for a 20% decrease in NO_x, for example, a whole family of designs is applicable. Each aircraft features a different fuel and noise trade in order to attain the desired reduction in NO_x. The final decision for selecting the appropriate design lies with the user: higher-level information, such as certification or operational requirements, is required.

E. Cruise Altitude Study

Reducing the cruise altitude of commercial aircraft would reduce contrail formation and potentially reduce the net impact of aircraft emissions. Today's aircraft operate at 30-40,000 ft, the optimal altitudes considering range and cruise speed. In order to minimize the impact on fuel economy, the aircraft needs to be designed to operate at these altitudes.

The single-objective version of the genetic algorithm was run with the initial cruise altitude fixed to 24,000 ft and 28,000 ft. A maximum cruise climb of 4,000 ft is allowed.

As expected, Designs E and F fly slower (Mach 0.762 and 0.728) in order to negotiate the increased drag inherent to lower altitude cruise. The corresponding increase in operating cost is 4% for Design E and 7% for Design F. The amount of fuel carried to complete the mission is decreased by 5% if the aircraft is designed to fly at 28,000 ft (a similar altitude to the aircraft optimized for lowest-fuel, Design B) and increased by 7% for a design altitude of 24,000 ft. Changes in NO_x production are minimal: with all three designs optimized for minimal operating cost, the optimizer is driven to select high pressure ratios and combustion temperatures, regardless of cruise altitude.

The purpose of operating aircraft at lower altitudes would be to reduce the net impact of the emissions on the atmosphere. Accurately estimating these net effects requires more information than the amount of

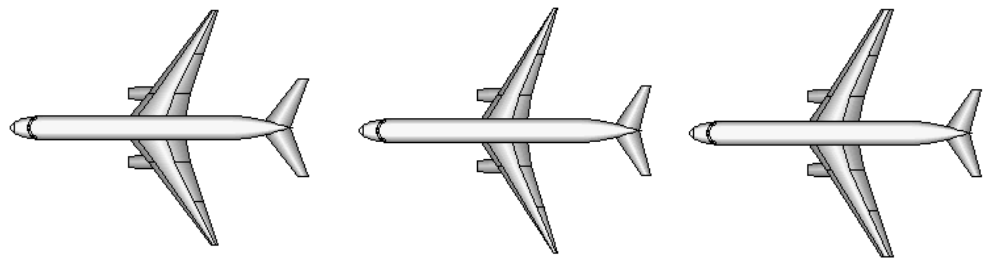

Figure 13. Top view of aircraft optimized for cruise at altitudes of 28,000ft (Design E) and 24,000 ft (Design F) with minimum cost Design A.

emissions generated by the engines. Indeed, to truly understand the effects of the combustion products on the atmosphere, a detailed study of the propagation and absorption characteristics of the upper troposphere would be required.

F. Impact of Future Airframe Technologies

Although much of the progress on noise and emissions reduction has been achieved through improvements to aircraft engines, the design tool was applied here to investigate three examples of potential future airframe technologies: laminar flow, advanced structures, and induced drag reduction. In particular the following three assumptions were explored:

- Increased laminar flow: for wing sweeps less than 20 degrees – it is assumed that natural laminar flow is achieved over an average of 60% chord.
- Advanced structures – a factor of 0.8 is applied to the aircraft structural weight to reflect use of new materials or other innovative structural concepts.
- A 10% reduction in induced drag – enabled, perhaps with span extensions and active load alleviation, for example.

These technologies, meant to improve aerodynamic efficiency or reduce structural weight, are of interest because they also have significant impact on the environmental performance of the aircraft. Figure 14 illustrates the changes to the aircraft fuel-NO_x performance. The advanced technology lowest-NO_x aircraft (Design K) produces 34% less NO_x per LTO cycle than the conventional low-NO_x aircraft (Design C), at 9% lower operating cost. Similarly, the advanced low-fuel candidate (Design J) requires 27% less fuel (and therefore produces 27% fewer fuel-proportional emissions) to complete the mission while releasing 32% fewer NO_x emissions than the conventional lowest-fuel aircraft (Design B). Taking advantage of the drag benefits associated with increased laminar flow, the two advanced designs feature wing sweep under 20 degrees — without any impact on cruise Mach number, thanks to the thinner wing afforded by the reduced fuel capacity requirement.

Incorporating these advanced technologies negates some of the adverse effects of optimizing the aircraft for low-noise or low-emissions. Indeed, the advanced low-NO_x aircraft (Design K) generates 30% of the NO_x emissions generated by the conventional low-cost design (Design A), at the same operating cost. From the reduced structural weight and fuel load (and therefore maximum takeoff weight), the two advanced designs require approximately 30% less installed thrust, resulting in a cumulative noise margin of 5.5 EPNdB for Design J and 0.8 EPNdB for Design K.

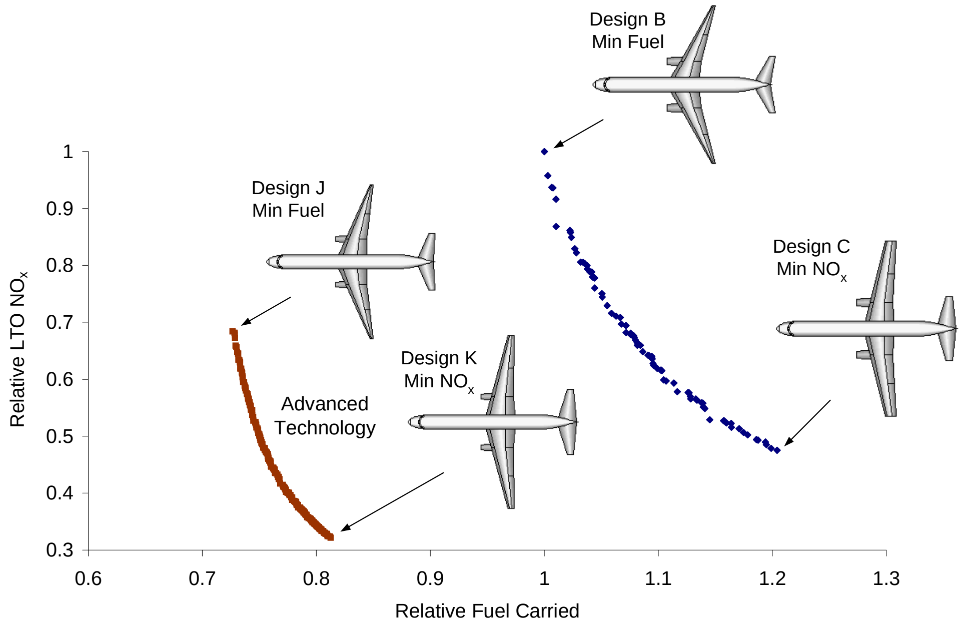

Figure 14. The benefits of increased laminar flow, reduced induced drag, and lower structural weight are illustrated on the fuel-cost Pareto front.

IV. Conclusions and Future Work

The objective of this research was to assess the feasibility of including environmental performance in the initial phase of aircraft design. The underlying design framework permitted integration of a wide variety of analysis and optimization tools including high fidelity engine and noise models. Multidisciplinary, multiobjective optimization proved a useful approach to quantifying trade-offs between noise, emissions, and operating cost.

Results suggest that large reductions in aircraft environmental impacts are possible even without dramatic changes to propulsion systems, but at significant cost. The study established a tradeoff between noise, emissions, and cost performance. The resolution of these diverging requirements will largely depend on the environmental regulations applying in the markets served by the aircraft. Significant reductions in emissions and perceived noise were found to be possible for aircraft specifically optimized with these objectives in mind. For an increase in operating cost of 9%, NO_x emissions could be reduced by as much as 50%, while cumulative certification noise could be lowered by up to 15 EPNdB for a cost increase of 26%. Optimized aircraft are able to satisfy stringent environmental constraints at far lower cost than aircraft that are design principally for cost, but operated to improve environmental characteristics.

Future work may include the following elements:

- Incorporation of modeling uncertainties and input parameter probability distributions directly in the framework.
- Analysis and design of fleet-level impacts, including system-of-systems optimization to provide more realistic assessments of environmental effects (cf. Antoine[22]).
- Development of analysis methods to accommodate unconventional designs such as strut-braced wings and blended-wing-body configurations.
- Enhanced user interface modules to enhance understanding, especially with probabilistic and fleet-level design optiona.

Acknowledgments

This work was made possible by a grant from NASA Glenn Research Center. The CAFFE framework was developed and provided to the project by Desktop Aeronautics, Inc.. The authors gratefully acknowledge the contributions to this work by our colleagues at Stanford, M.I.T., and NASA's Langley and Glenn Research Centers who provided, data, software, and many stimulating technical discussions.

References

[1]Penner, J.E., *Aviation and the Global Atmosphere*, Cambridge University Press, 1999, pp. 1-15.

[2]Sietzen Jr., F., "New Blueprint for NASA Aeronautics," Aerospace America, August 2002, p.25.

[3]Penner, J.E., *Aviation and the Global Atmosphere*, Cambridge University Press, 1999, p. 6.

[4]Boeing Current Market Outlook 2002 [Online], URL: http://www.boeing.com/commercial/cmo, The Boeing Company, 2002 [Cited April 8 2003].

[5]Smith, M.J.T., *Aircraft Noise*, Cambridge University Press, 1989.

[6]Pacull, M., "Transport Aircraft Noise Technologies," *Proceedings of the International Symposium: Which Technologies for Future Aircraft Noise Reduction?*, Arcachon, France, October 2002.

[7]FAA's Preliminary Briefing on CAEP/6 [Online], URL: http://www.ueet.nasa.gov/toi/viewtoi.php?id=119, NASA Glenn Research Center, February 2004 [Cited May 17 2004].

[8]The Cambridge-MIT Silent Aircraft Initiative [Online], URL: http://sai.eng.cam.ac.uk/, Silent Aircraft Initiative, 2004 [Cited June 18 2004].

[9]Kroo, I., Manning V., "Collaborative Optimization: Status and Directions," 8th AIAA / NASA /USAF / ISSMO Symposium on Multidisciplinary Analysis and Optimization, Long Beach, CA, Sept. 6-8, 2000, AIAA Paper 2000-4721.

[10]Deb, K., *Multi-objective Optimization Using Evolutionary Algorithms*, Wiley, 2001, pp. 190-196.

[11]Kroo, I.M., "An Interactive System for Aircraft Design and Optimization," AIAA Aerospace Design Conference, Irvine, CA, Feb 3-6 1992, AIAA Paper 92-1190.

[12]Kroo, I.M., "AA241 Aircraft Design: Synthesis and Analysis," Course Notes [Online], URL: http://adg.stanford.edu/aa241/AircraftDesign.html, January 2003 [Cited August 1 2004].

[13]Schaufele, R., *The Elements of Aircraft Preliminary Design*, Aries Publixations, Santa Ana, CA, 2000, pp. 311-323.

[14]Antoine, N.E., *Aircraft Optimization for Minimal Environmental Impact*, Ph.D. Thesis, Stanford University, August 2004.

[15]Antoine, N.E., and Kroo, I.M., "Aircraft Optimization for Minimal Environmental Impact," AIAA Journal of Aircraft, Volume 41, Number 4, pp. 790-797, July-August 2004.

[16]Klann, J.L. and Snyder, C.A., *NEPP User's Manual*, NASA Glenn Research Center, March 1997.

[17]Gillian, R.E., *Aircraft Noise Prediction Program User's Manual*, NASA Langley Research Center, 1982.

[18]Heidmann, M.F., "Interim Prediction Method for Fan and Compressor Source Noise," NASA Technical Memorandum X-71763, 1979.

[19]Stone, J.R., Groesbeck, D.E., and Zola, C.L., "An Improved Prediction Method for Noise Generated by Conventional Profile Coaxial Jets," 7th AIAA Aeroacoustics Conference, Palo Alto, CA, Oct 5-7 1981, AIAA Paper 81-1991.

[20]Fink, M.R., *Airframe Noise Prediction Method*, Federal Aviation Administration FAA-RD-77-29, 1977.

[21]Anon., *ICAO Annex 16 Volume II — Aircraft Engine Emissions, Second Edition*, International Civil Aviation Organization, 1993.

[22]Antoine, N.E. et al., "Fleet Assignment Using Collective Intelligence," 42nd AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, January 5-8 2004, AIAA Paper 2004-0622.