GE Aviation

FAA CLEEN II

Flight Management System Final Report – Public Version

28 January 2020

FAA CLEEN II

Flight Management System Final Report – Public Version

28 January 2020

Publication No. YV4129
Deliverable In Response to Government OTA No. DTFAWA-14-R-73573

Levent Ileri, Program Manager
Federal Aviation Administration
800 Independence Avenue, S.W.,
Washington, DC 20591

Submitted by: GE Aviation
1 Neumann Way
Cincinnati, OH 45215
CAGE Code: 07482

Technical Contact: Nicholas Visser
616-224-6613
nicholas.visser@ge.com

Program Manager: Theresa Zeug
513-914-2208
theresa.zeug@ge.com

GE Aviation Systems LLC
3290 Patterson Ave. SE
Grand Rapids, Michigan 49512
geaviation.com

Contents

Figures

1. Executive Summary

GE Aviation developed a TRL6 Connected Flight Management System with real-time optimization that achieves over 1% fuel burn reduction compared to the Legacy product

In the Continuous Lower Energy, Emissions, and Noise II (CLEEN II) program, the GE Aviation Flight Management System (FMS) development group studied fixed-wing aircraft flight dynamics with the purpose of implementing new vertical control policies for Part 25 commercial aircraft to achieve the Federal Aviation Administration (FAA) goals of reduced fuel usage and emissions in the National Air Space (NAS). New vertical control policies are implemented in a unified manner that use high-fidelity weather forecast data, as compared to legacy systems. The new controls include:

- Variable-speed and variable-thrust climb.
- Variable-speed cruise.
- Cruise step-climbs and step-descents informed by high-fidelity weather forecast.

See Figure 1 and Figure 2 for an example control profile and weather data collection.

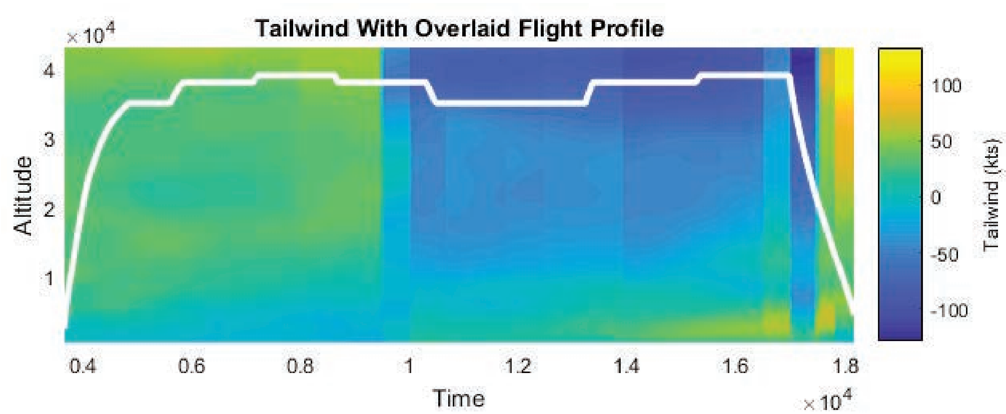

Figure 1. CLEEN optimized vertical flight.

Optimized vertical flight has a fuel-efficient climb profile and cruise step-climbs and step-descents that track favorable weather.

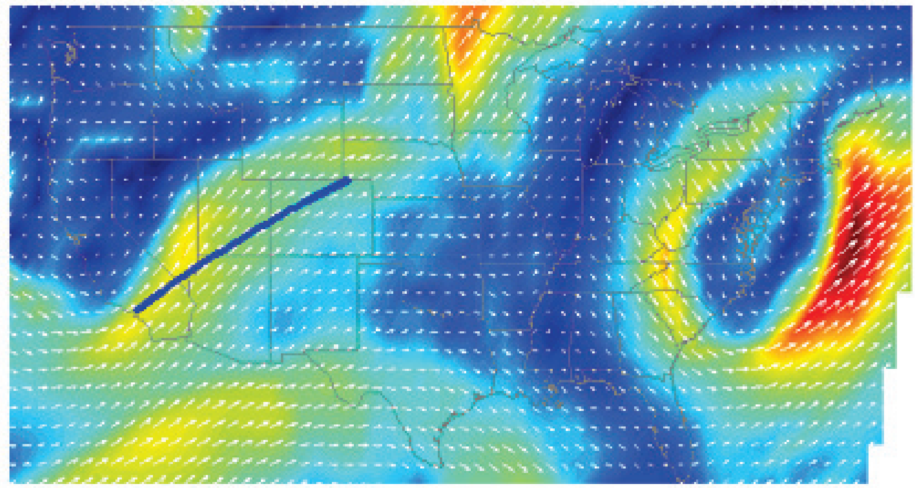

Figure 2. Weather data for use in vertical path optimization.
Weather data is automatically extracted from 4D high-fidelity weather sources.

Optimization techniques are employed to find the maximum possible cost reduction for a full flight cycle. The methods developed are implemented in a TRL 6 real-time system and therefore constrained to follow the instrument flight rules. The software is implemented as an Electronic Flight Bag (EFB) application with bi-directional communication with the Flight Management System via an Aircraft Interface Device (AID); see Figure 3. To maintain safety, security, and support certification, GE's approach uses the Connected Flight Management System (CFMS) Software Development Kit (SDK). This approach ensures that the FMS maintains full control authority, and the higher Design Assurance Level (DAL) maintains the ultimate responsibility for validating the control inputs prior to data exchange into the aircraft control domain. This approach provides a practical implementation with minimal operational impact for pilots and controllers and allows the technology developed to readily move into production. The TRL 6 implementation is constructed with modular optimization capabilities allowing reuse in other products (such as Air Traffic Management applications) and other use cases (such as community noise reduction and lateral planning).

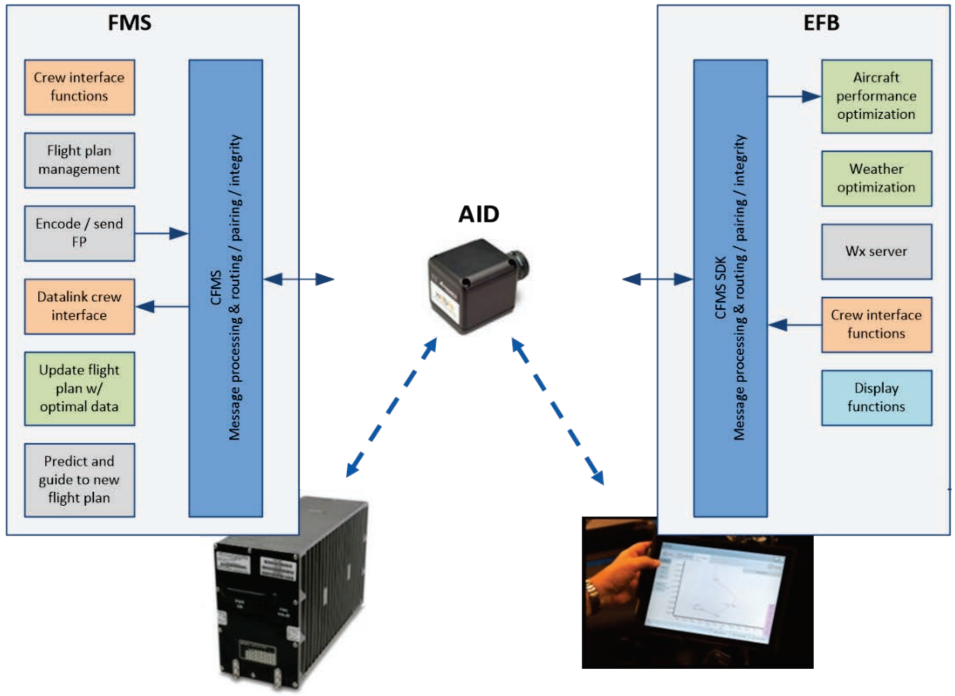

Figure 3. System architecture.

The System architecture enables communication between FMS and EFB-hosted CLEEN optimizer via AID

Fast-time computer simulation is used to quantify a statistically significant benefit across a wide range of aircraft types, routes, weather, and passenger loading in a way that is infeasible without many months to years of revenue-service flight testing; see Figure 4. The goal of this assessment is to compare the CLEEN technologies with the best-in-class baselines without the impact of difficult-to-quantify effects such as Air Traffic Control actions, sensor errors, and aircraft modelling errors. These effects are assumed to impact both the CLEEN and legacy baselines by the same amount and thus have no impact on the average benefit. In total, technologies developed under the CLEEN program reduce fuel burn by **1.02%** against a typical in-service FMS (termed “Legacy” in this report), and **0.40%** against an FMS with decision-aiding ground tools representative of today’s best-in-class technology (termed “Legacy++” in this report); see Figure 5. Note that fuel burn reduction is equivalent to Direct Operating Cost (DOC) reduction when the cost of time is considered irrelevant; see paragraph 3.1 for details on the distinction.

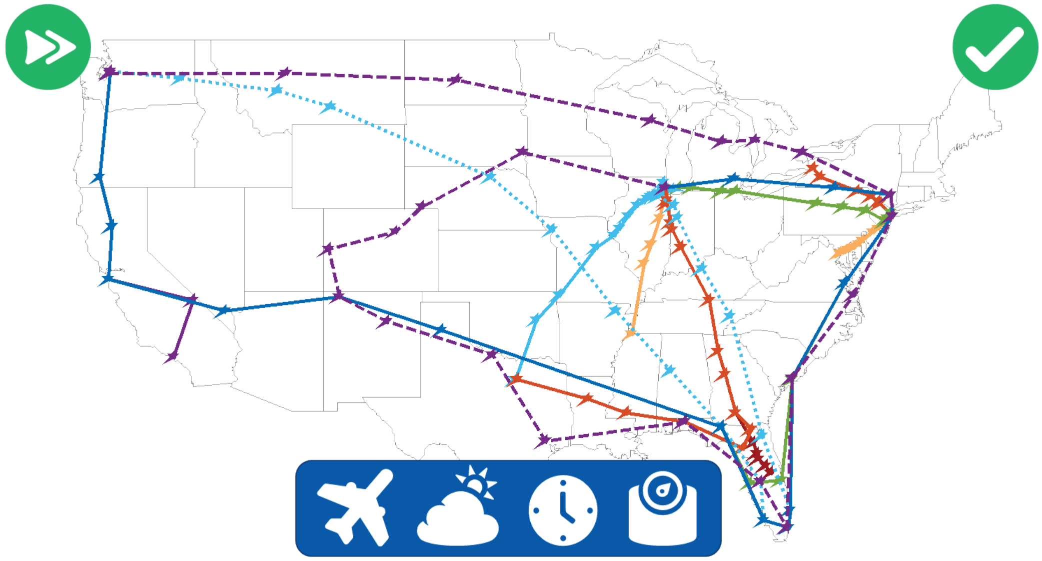

Figure 4. Computer simulation.
20 different routes and 5 different aircraft models are assessed in fast-time simulation for a variety of conditions to produce an assessment of fleetwide benefit.

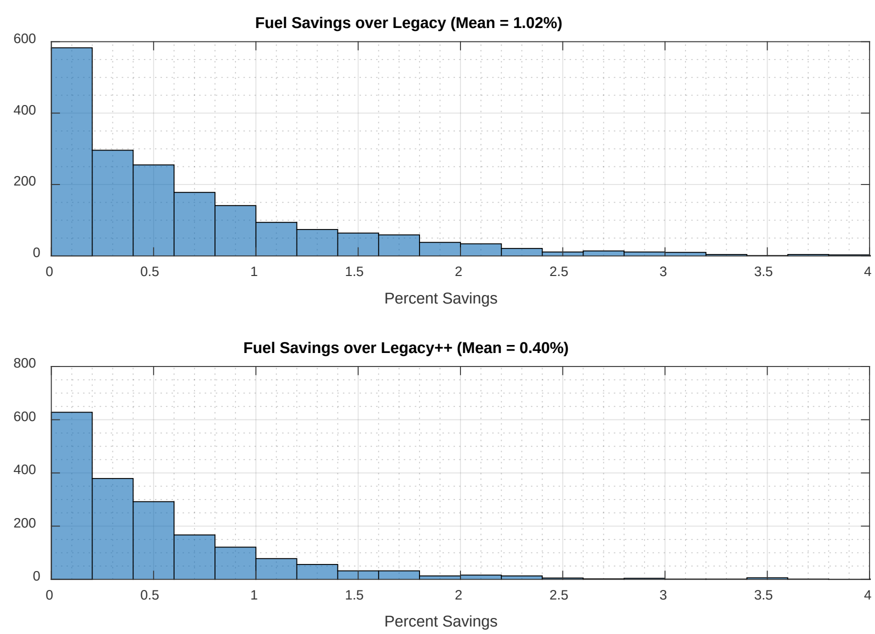

Figure 5. Fuel burn comparison.
Histogram of fuel burn savings for the CLEEN technologies shows up to 4% savings over legacy technologies.

While most of today's FMSs can construct a flight path to meet a Required Time of Arrival (RTA) constraint, these types of constraints generally lead to increased fuel consumption. The CLEEN algorithms calculate optimal control for minimal fuel consumption in the presence of RTA constraints at any location along the flight path. This feature enables FAA NextGEN Air Traffic Management policies without a large associated individual flight penalty. A large-scale assessment using Monte Carlo fast-time computer simulation shows fuel burn reduction by **0.81%** against a typical in-service FMS (Legacy), and **0.47%** against an FMS with decision-aiding ground tools representative of today's best-in-class technology (Legacy++); see Figure 6.

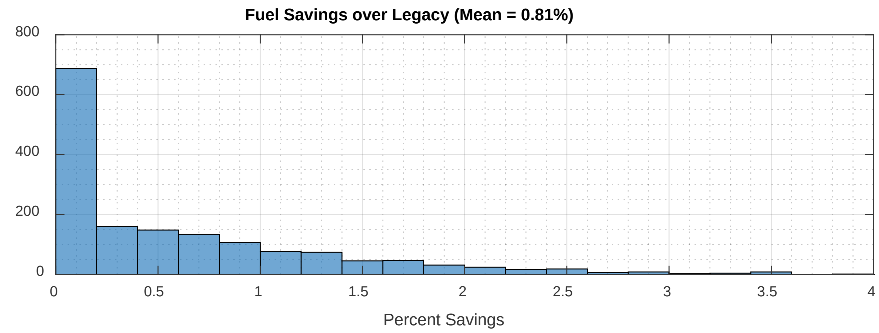

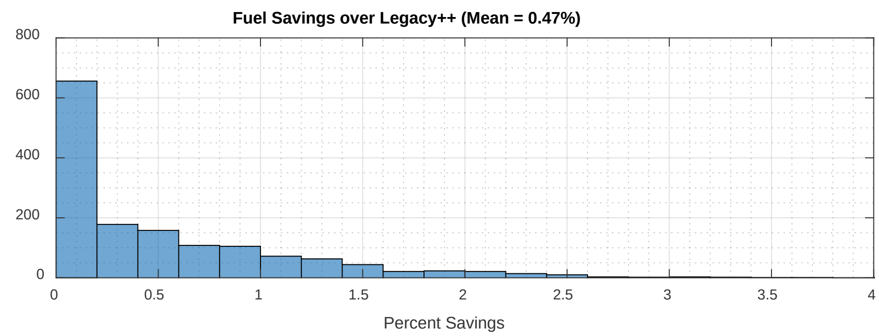

Figure 6. Monte Carlo computer simulation.

Histogram of fuel burn savings for the CLEEN technologies with Required Time-of-Arrival shows benefit can be achieved while meeting flight path constraints.

The benefits of this technology provide a commercially viable introduction into service, and GE will continue to further develop this technology beyond conclusion of the FAA CLEEN II program with the goal of deployment for both retro-fit and new aircraft.

2. Introduction

In 2015, the FAA awarded GE Aviation a contract to research and develop FMS technology. The work was jointly funded by the FAA and GE Aviation as part of the CLEEN II program. These concepts expand upon the CLEEN Climb Optimization technology to improve the vertical flight control function of an FMS to operate an air transport more efficiently and thereby achieve the goals of the CLEEN program. Specifically, this project focuses on optimal control algorithms and their benefit across a fleet of various aircraft types. The technology is implemented as a TRL 6 prototype that optimizes the entire climb, cruise, and descent flight cycle.

2.1 CLEEN Goals

According to the FAA's CLEEN fact sheet available online[*]:

> "CLEEN is the FAA's principal environmental effort to accelerate development of new aircraft and engine technologies and advance alternative jet fuels. The program is a key element of the NextGen strategy to achieve environmental protection that allows for sustained aviation growth. The FAA launched the initial CLEEN I in 2010. Over the five-year course of the program several technologies have been tested and are in use today. Based on that success, the FAA is currently in a second phase, CLEEN II, which runs from 2015 through 2020."

Aircraft performance is largely tied to the route an aircraft takes to reach a destination way point. The GE FMS CLEEN II project has a focus on uncovering new technologies that will construct routes with lower direct operating cost (DOC). GE's Connected FMS applications will be updated with the final version of these algorithms to enable decision support for end users (airlines) resulting in savings of 1% direct operating cost reduction over their fleet.

A reduction in fuel burn means a reduction in the average thrust produced. A reduction in thrust yields a proportionate reduction in NOx emissions and noise. Thus, an adjacent benefit to the fuel savings is reduced emissions and noise. GE has statistically characterized the resulting flight profiles. Engineers at the Georgia Institute of Technology (on behalf of the FAA Ascent Project 037) will estimate the reduced emissions and noise for the fleet of commercial transports in the United States.

2.2 Development Method

The GE FMS CLEEN II development is performed in four generations (termed, "Generation A" through "Generation D") as described below. Each generation has an up-front TRL 4 analysis phase to determine the expected savings from a given algorithm. At the follow-on phases, the concept is matured to TRL 6. The final demonstration uses existing test equipment, when possible, to demonstrate how a pilot or user would improve the flight and complete a new flight plan update. The flight plan update, calculated on an EFB, is sent to the FMS via the Connected FMS. This approach requires minimal modification to the existing FMS.

The development method is broken down into two phases such that the savings can first be quantified with enough certainty to decide if the benefit justifies maturation to a higher TRL. In the first phase, called the *Concept Evaluation,* the technical approach is researched and developed in a TRL 4 environment using mathematical models and a simplified simulation. At this decision gate, a determination is made on whether the incremental benefit warrants full-scale development. If so, then the second phase, called the *Prototype Development and Demonstration,* is performed. In the second phase, TRL 5 is achieved using a full-scale fleet-wide benefit assessment of a detailed fast-time computer simulation. After this assessment, TRL 6 is achieved by developing a prototype FMS and integrating it with a real-time simulation of the vehicle and the engines.

The concept evaluations are preceded by the development and validation of models of the vehicle and engines called the *Advanced Technology Test Bed.*

The following paragraphs describe the development process in terms of the work elements in the Statement of Work and the development phases.

Work Element – Design

The design process begins by defining and specifying requirements and performance objectives for the subject technology. A literature survey is performed to learn the current professional and academic design methods and technology developments. The best methods are applied or developed to design the algorithms and logic that are then modeled mathematically. The models are coded for computer simulation and integrated with the FMS model. The Simulink[†] modeling language is used for this analysis. Since performance improvements are defined relative to the performance of the legacy FMS, simulation trials are performed first using the legacy design and then using the modified design. A comparative analysis is performed to quantify the fuel savings. GE then performs a Preliminary Design Review (PDR) to demonstrate the high-level design met the system

[*] "Fact Sheet – Continuous Lower Energy, Emissions, and Noise (CLEEN) Program", FAA, 4 June 2019. [Online]. Available: https://www.faa.gov/news/fact_sheets/news_story.cfm?newsId=22534.

[†] Simulink is a registered trademark of MathWorks

requirements with acceptable risk and within cost and schedule constraints – to establish the basis for proceeding with the detailed design activity.

If the performance benefit is sufficient to justify the development of a prototype design, high-level requirements are specified and allocated to implementation as software. The review shows that satisfactory design options had been selected, interfaces have been identified, and verification methods are defined. The software is then designed according to the requirements. When the design phase completes, GE performs a Detailed Design Review (DDR) to demonstrate the design is sufficiently mature to proceed with implementation, integration, and test. The DDR determines that the design and development are on track to complete the test-article hardware and software development in accordance with the performance requirements.

The design and development is performed in accordance with GE Aviation's model-based methods and best practices for an experimental system. GE Aviation uses these engineering and integration methods and tools to verify that the process and prototype system design comply with the schedule, cost, and technical requirements and applicable design assurance, safety, and airworthiness regulations that apply to TRL 6.

Work Element – Fabrication/Implementation

The Advanced Technology Test Bed consists of a computer simulation of the air vehicle, the engines, and the other vehicle systems that the prototype system interacts with to perform the vehicle-level function. The FMS and EFB hardware are integrated with the improved software to implement the test article, test equipment, and special tools required to install and test the prototype system in the laboratory.

The implementation and integration process are compliant with GE Aviation's model-based methods and best practices for an experimental system. Said methods and tools are applied to plan and perform the reviews, analyses, audits, and tests necessary to verify the prototype system complied with the function, performance, design assurance, safety, and airworthiness requirements and regulations that apply to TRL 6.

Work Element – Performance Demonstrations

A fleet-wide benefit assessment is perfomed using the Advanced Technology Test Bed, and the assessment is validated using the real-time Flight Management Workstation (FMWorkstation) production simulation. The tests required to demonstrate compliance with key high-level requirements are performed in the laboratory using the FMWorkstaiton. In this environment, the test article (FMS + EFB) is integrated with a real-time computer simulation of the air vehicle, the engines, and the other vehicle systems that the test article interacts with functionally to perform the vehicle-level function or functions. The performance of the modified system is measured and compared to the performance of the legacy system. The test article is avionics hardware that GE Aviation produces for the Boeing 737.

The methods and tools used for the demonstration comply with GE Aviation's processes and best practices to ensure the work was performed in accordance with schedule, cost, and technical requirements.

Work Element – Assessment and Reporting

Finally, the system design and quantified benefits are summarized in monthly meetings with the FAA, bi-annual consortium presentations, design reviews, and within this final report.

2.3 Purpose and Scope

This document, prepared by GE Aviation Systems LLC (GE Aviation), provides a full report of work performed under DTFAWA-15-A-80013 for Flight Management System Technologies. This document is organized with an overview of the existing flight path optimization technology prevalent in FMS today, followed by several incremental generations of improvements and their associated benefit assessments.

3. Technology Overview

3.1 Problem Formulation

The primary objective of the CLEEN II program is to minimize fuel costs incurred by flights following an FMS-determined flight profile. The pursuit of an FMS design that produces flight profiles that consume less energy and produce less emissions while still permitting airlines to meet the operational mandates and incentives of Air Traffic Control (ATC) demands a solution that considers both time and fuel.

In the commercial aviation industry, Direct Operating Cost (DOC) is defined as:

$$DOC\ [\$] = Fuel\ Cost\ [\$] + Time\ Cost\ [\$]$$

To define DOC as a function of service time, it is expressed in integral form:

$$\int_{t_0}^{t_f} (c_t + \dot{c}_f) dt$$

where c_t is the time-related cost rate to operate the airplane, $\dot{c}_f$ is the cost rate of fuel, t_0 is the departure time from the origin, and t_f is the arrival time at the destination. The cost rate of fuel, $\dot{c}_f$, in units of dollars per hour may be expressed as the product of fuel flow rate, w_f, and the cost of fuel, c_f. Substituting gives:

$$DOC = \int_{t_0}^{t_f} (c_t + w_f c_f) dt$$

Historically, speed and altitude are how the pilot controls the longitudinal motion of the airplane in cruise. To formulate the problem using these control variables, the independent variable is changed from time to distance:

$$DOC = \int_{R_0}^{R_f} \left(\frac{c_t + w_f c_f}{V_g} \right) dr$$

where r is along-track position, R_0 is the position at t_0, R_f is the position at t_f, and V_g is ground speed. On most commercial transports today, the operator specifies the cost of time and the cost of fuel as a single parameter called the Cost Index, CI, that is defined as:

$$CI\ [100\ lbhr] = \frac{Time\ Cost\ \left[\frac{\$}{hr}\right]}{Fuel\ Cost\ \left[\frac{¢}{lb}\right]}$$

$$100 \cdot CI = \frac{C_t}{C_f}$$

Re-arranging and substituting,

$$DOC = \int_{R_0}^{R_f} \frac{100 \cdot CI \cdot c_f + W_f \cdot c_f}{V_g} dr$$

Dividing both sides by c_f and defining a new cost function, ψ_c,

$$\psi_c = \frac{DOC}{c_f} = \int_{R_0}^{R_f} \frac{100 \cdot CI + W_f}{V_g} dr$$

Thus, the problem is to find the speed and altitude that minimizes ψ_c. Note that the cost of fuel burned to traverse from R_0 to R_f depends on weight, altitude, Mach number, engine deterioration, atmospheric conditions, and aircraft trim (control surface deflections). For pure fuel minimization, the Cost Index is set to 0 (cost of fuel is infinitely more than the cost of time).

3.2 Technology Iterations

Throughout the CLEEN II program, vertical flight plan optimization technology to minimize direct operating cost proceeds in incremental development (see Figure 7). In each phase, three major steps are performed:

- Algorithm development and initial testing – Optimization Algorithm deployed as standalone MATLAB module (TRL 4)
- Integration with real FMS software in a fast time environment for benefit assessment – Flight Plan Predictions Driver (FPPD) in MATLAB (TRL 5)
- Integration with real-time FMS and HMI in a laboratory prototype – FMS + EFB (TRL 6).

This iterative development in each generation allows for technology tollgates rapid assessment of optimization features with increasing fidelity. For example, as described in subsequent sections, GE determined that Gen A and UCCD technology provide a solid foundation for further development, but do not provide enough benefit alone to be considered a stand-alone product. These features are paused at the TRL 4 phase and assessed at higher TRL with the later technology generations.

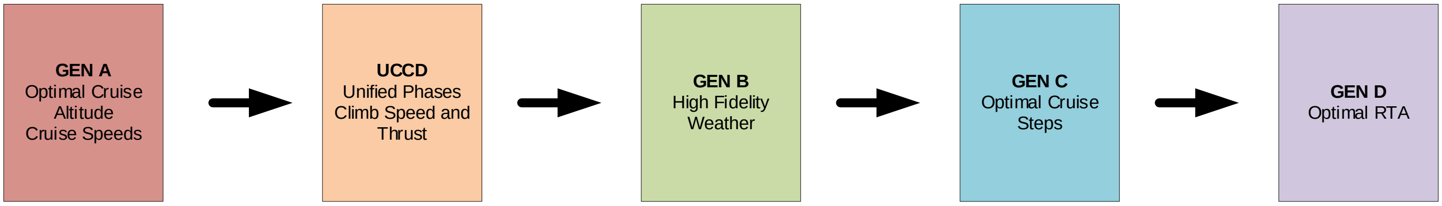

Figure 7. Progression of CLEEN flight path optimization generational development.

3.3 Legacy Methods

It is well known in industry that the optimum cruising altitude for a transport aircraft generally increases as the aircraft weight decreases, due to varying aircraft performance at different altitudes. In an unrestricted airspace, the optimum cruise profile is a “drift-up” trajectory, where the cruise altitude steadily increases as fuel is burned throughout the cruise phase. However, this flight profile is not practical in crowded and controlled airspace where vertical spacing between aircraft is necessary for safety concerns. As a result, Air Traffic Control (ATC) generally requires that aircraft fly at intervals of either 1,000 or 2,000 feet (depending on flight direction or one-way status of jetways). Furthermore, the pilots must request from ATC the clearance to change cruising altitudes in advance of doing so.

These ATC considerations, and the lack of weather data at multiple cruise altitudes drive many legacy FMS solutions to consider only a constant cruise altitude for the entire flight (referred to as the Legacy system below). Some FMSs with multiple weather altitudes can perform a better constant cruise altitude selection (referred to as Legacy+), and finally, some aircraft perform “step climbs” at specified points in the cruise phase (Legacy++). Regardless of how the location of the step climb is selected (i.e., via manual entry from the pilot, a native function of the FMS, or an advisement from a ground operator), the process for the step climb is as follows: the pilot must request the climb and receive clearance from ATC, then the pilot manually sets the new cruise altitude and the FMS performs the climb at maximum thrust to capture the new cruise altitude. This process is repeated as necessary; on longer flights, the cruising altitude can change upwards of four to five times. Further steps may occur as ATC routes aircraft around particularly crowded airspace or to avoid dangerous weather patterns. However, these unique scenarios are unpredictable and not the subject of this technology.

Two industry-standard methods exist for computing the location to initiate step climbs. For Flight Management Systems that do not provide a native computation of the step location, a ground tool may be used. The ground-calculated step points are either communicated to the pilot, who manually enters them into the FMS and executes it or are uplinked via datalink. The former has obvious concerns of pilot workload and each potentially sacrifices optimality – the aircraft and weather conditions may be different than those assumed in the ground performance tool. Some Flight Management Systems feature a simple computation of the optimal step location that includes weather data entered into the FMS; however, this function generally includes only the capability to compute one optimal step, which must be a step climb. These shortcomings in the industry standard solutions prompt the development of the Gen C technology.

In order to determine the merits of the UCCD Gen C technology, two baseline legacy methods are used for comparison, titled Legacy and Legacy++ (with capitalization to distinguish them from the generic “legacy” term). The latter is titled to distinguish it from the Legacy+ method that was the basis of comparison for the Gen B technology. These methods are described in detail in paragraph 3.3.

- **The Legacy method** is representative of a simple baseline FMS. This FMS does not compute step locations natively, and thus, no step climbs are performed in this method. The low fidelity cruise weather model contains tailwind data at only one altitude. This is meant to represent a less sophisticated method currently available in the market, and the benefit of Gen C relative to this method is expected to be higher.
- **The Legacy+ method** (compared against in earlier generations of CLEEN software) is identical to the Legacy baseline, but with an extension of the weather data model to include more data points. This baseline is modeled upon a more advanced FMS optimization capability than the Legacy.
- **The Legacy++ method** is representative of a more advanced FMS in conjunction with a nominal ground performance tool. The FMS computes one optimal step climb that accounts for weather. The weather model has slightly higher fidelity than the Legacy method, in that there is tailwind stored at four different altitudes in cruise. The ground tool computes the remaining three step climbs assuming standard atmosphere conditions and is based on GE’s best understanding of the capabilities of such a tool without having direct access to one. This combination is meant to represent the best method currently available in the market, and the benefit of Gen C relative to this method is expected to be lower.

All methods use a maximum thrust climb profile and simple table lookups for optimum cruise altitude and speed. Neither method can produce a step descent during cruise.

3.3.1 Legacy Approach for RTAs

In general, the FMS supports the selection of a parameter called a Cost Index to determine the operational balance between fuel and time costs. This input is usually determined by the airline operator and requires comparing the price of fuel to costs of time such as crew wages. This Cost Index is then used by the FMS function7s to determine the flight speed that minimizes the Direct Operating Cost. In the lack of RTA constraints, the Cost Index essentially determines flight time as balanced with fuel cost.

Now, a Required Time of Arrival constraint applied to a waypoint in the flight path specifies that the aircraft must arrive at that waypoint within a certain tolerance of the specified time. In general, RTAs are considered “at” constraints, as opposed to an “at or before” or an “at or after” constraint. The constraint time is usually specified by ATC in order to properly manage traffic flow through a congested airspace, such as the approach of an airport. However, the constraints may occur at any location in the flight.

In the presence of an RTA constraint, the time of arrival at the waypoint of concern is specified. This means that the flight time for the RTA segment (from the aircraft’s present position to the RTA waypoint) will be a constant, specified value. The Cost Index that has been specified by the operator is irrelevant for this segment, as the cost of time will be fixed as well. Thus, the FMS uses the Cost Index as a mechanization to vary aircraft speed until the RTA has been achieved. Since the flight controls are optimal for any given Cost Index, it will be optimal for this selected *RTA Cost Index* as well; since time is fixed, this corresponds to a minimum fuel case. After the RTA constraint distance has been passed, the aircraft may return to normal operations, such as economy speed, using the *True Cost Index* as specified by the operator.

Figure 8 provides a generic depiction of this process. Given different control speeds, the flight profile will arrive at the RTA location with varying time-of-arrival (on the x-axis) and fuel usage (on the y-axis). Varying the Cost Index and then looking up the tabulated economy flight data (cruise altitude and speeds) for that Cost Index will produce flights on the lower portion of this cloud, known as the Pareto Frontier. In the general case where this solutions cloud is convex, the eventual Cost Index that generates a flight that intersects the vertical RTA time line will be found.

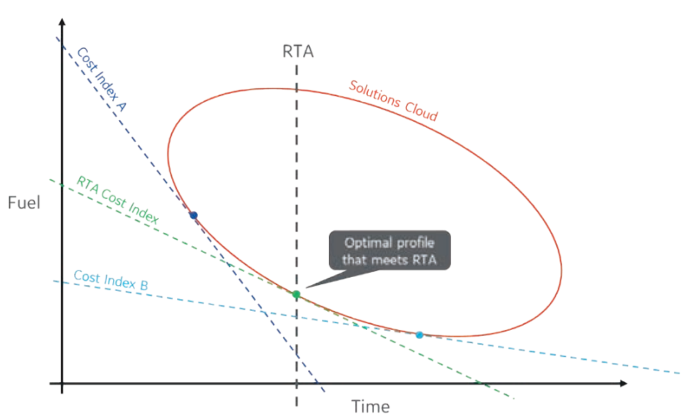

Figure 8. Legacy Cost Index loop.

Various Cost Indices are trialed until the one that generates a flight that meets the RTA time (constant vertical dashed line) is found.

There are a few shortcomings with this technique. The first is that the method can only modify aircraft speed, and not cruising altitude, to meet the RTA. The altitude profile is already set – whether it is a constant cruise altitude or contains step climbs. This eliminates the possibility to determine a more optimal altitude that allows cruising at more efficient speed to meet the RTA. This also means that the presence of different winds at different altitudes are not accounted for. Another major shortcoming arises when the solutions cloud in Figure 8 is non-convex, which could be caused by weather patterns. These shortcomings in the industry standard solutions prompt the development of the Gen D technology.

3.4 Benefit Assessment Methodology

Measuring the benefit of the CLEEN technology over the Legacy baseline methods is accomplished through a comparison of the total Direct Operating Cost (DOC) throughout the flight, which includes the fuel burned and time burned. See paragraph 3.1 for a detailed description of DOC calculations.

Vertical Flight Plans are constructed using both the CLEEN technology and legacy methods, and these plans are flown through fast time aircraft simulations to compute fuel burn. A single metric benchmark for “fleet-wide savings” for comparison gives the general formulation of percent fuel savings for a given set of test cases as:

$$\%\,Savings = \frac{\sum_i (LegacyCost_i - CLEENCost_i)}{\sum_i LegacyCost_i} * 100$$

where the summation happens over each case i in the test set, and a positive value indicates saved fuel. Intuitively, this number indicates the sum of fuel saved over the sum of fuel burned, which is slightly different than simply taking the average of the percent cost saved for each case.

The general philosophy of comparison throughout this program is to compare a candidate technology to the best possible profile the baseline FMS software can produce, and *not* to make comparisons to flights as they were flown using real recorded flight data. In practice, several unpredictable events can occur that reduce the optimality of the flight, including extra airspace constraints, ATC demands of cruise altitude, missed approaches, and so on. If a comparison occurs between an actual flight that experienced these events and an optimization technique that is not likewise constrained by them, much larger savings are reported. Several other studies performed by other parties compare their optimization methods to as-flown trajectories, artificially inflating the reported benefit. To make a fair comparison, these types of unpredictable suboptimalities are removed from the equation[‡].

For a detailed description of the benefit assessment process for each TRL, including case setup and selection and omitting outliers, refer to Appendix B.

3.5 Terminology

Below is some standard terminology used throughout this document, and the corresponding definitions.

- **UCCD:** Unified Climb, Cruise, and Descent, an iteration of CLEEN software that merges flight phase optimization (see Section 5)
- **Generation A (Gen A):** Iteration of CLEEN software, including optimal constant cruise altitude for standard day (see Section 4)
- **Generation B (Gen B):** Iteration of CLEEN software, including optimal constant cruise altitude with weather (see Section 6)
- **Generation C (Gen C):** Iteration of CLEEN software, including optimal cruise altitude steps with weather (see Section 7)
- **Generation D (Gen D):** Iteration of CLEEN software, including optimal cruise altitude steps for Required Time of Arrival (see Section 8)
- **Legacy:** A baseline for comparison including legacy FMS constant cruise altitude technology
- **Legacy+:** A baseline for comparison including legacy FMS constant cruise altitude technology with a more sophisticated weather model
- **Legacy++:** A baseline for comparison including optimal steps calculated from today's sophisticated ground tools
- **Standard Day:** The set of atmospheric conditions with no winds and a temperature that varies as a function of altitude according to the 1975 International Standard Atmosphere model
- **Cost Index:** A parameter used to determine the relative cost of fuel and cost of time to an airline operator for a single flight (see paragraph 3.1)
- **DOC:** Direct Operating Cost, the sum of the cost of fuel and the cost of time (see paragraph 3.1)
- **TOC (T/C):** Top of climb
- **TOD (T/D):** Top of descent
- **EOD (E/D):** End of descent
- **ETA:** Estimated time-of-arrival
- **ATA:** Actual time-of-arrival
- **ETG:** Estimated time-to-go
- **RTA:** Required time-of-arrival constraint on a waypoint in the flight plan
- **Altitude Quantization (Separation):** The Air Traffic Control mandated vertical separation between aircraft during the cruise phase, usually in increments of 1,000 or 2,000 ft depending on the route

[‡] D. Lax, M. Darnell, O. O'Keefe, B. Rhone, N. Visser, R. Ghaemi and E. R. Westervelt, "Quantifying Operating Cost Reduction from Aircraft Performance Optimization", in *Integrated Communications, Navigation, Surveillance Conference (ICNS)*, Herndon, VA, 2018.

4. Generation A: Cruise-Only Optimization

4.1 Executive Summary

The primary objective of the CLEEN II program is to minimize fuel costs incurred by flights that follow an FMS-determined flight profile. To meet this objective, the Generation A Cruise Optimization technology is developed to generate a constant-altitude, variable speed cruise profile that minimizes direct operating cost for the cruise flight phase by treating the mass of the aircraft as a state variable during the cruise phase calculation. This technology represents a direct improvement over the Legacy method using more accurate calculations with the available data and no new control policy or additional information.

Testing conducted with a TRL 4 implementation on several variants of the optimization demonstrated cruise fuel savings on the order of 0.1% as compared to the Legacy methods, averaged over an array of flight distances and for different altitude separation constraints; see Figure 9. Implementation into TRL 5 Predictions software was not pursued for this generation due to the small savings. The Generation A optimizer serves as a basis for the further generations of cruise optimization technologies.

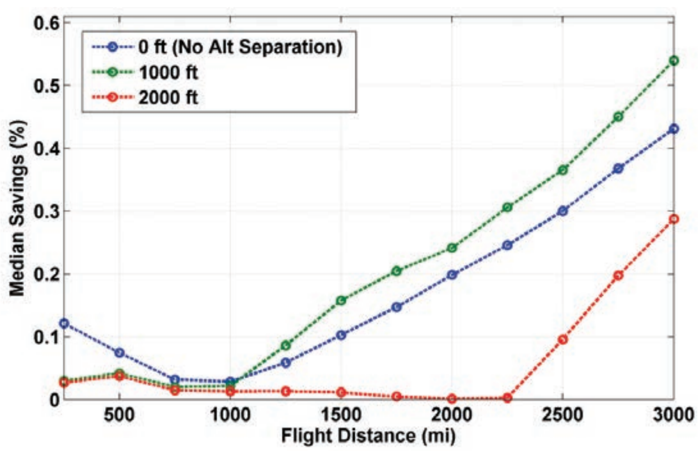

Figure 9. Generation A optimization.

Optimization results show that optimal altitude selection including effect of distance (and mass change throughout flight) result in small overall cost savings depending on altitude separation constraint.

4.2 Development

4.2.1 Optimal Constant Altitude, Variable Speed

The goal of the Generation A Cruise Optimizer is to select constant cruise altitude and variable cruise speeds resulting in the minimum cost given the distance travelled in cruise phase, and an approximation of the cost to climb and descend from the cruise altitude. This method removes assumptions of trip distance (see Figure 10) and inaccuracies in tabulated Legacy economy data, replacing Legacy lookups of optimal control with live calculation using mass as a state variable. This framework provides the basis for all future generations of CLEEN software, with increasingly more available controls and higher fidelity input data.

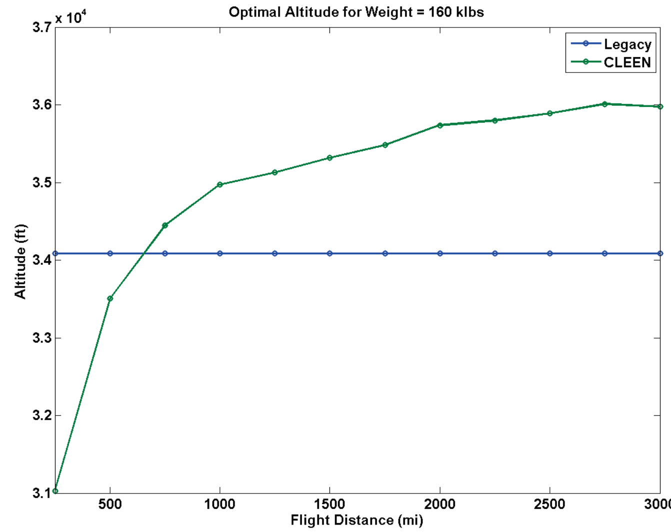

Figure 10. Cruise optimizer.

Optimal cruise altitude (CLEEN) depends on cruising distance; versus Legacy approach based only on aircraft weight.

As described in paragraph 3.1, this requires finding the altitude and velocity that minimize the cost function over the range $[R_0, R_f]$ in feet of the flight:

$$\psi_c = \int_{R_0}^{R_f} \frac{CI + W_f}{V_g} dr$$

where W_f is the aircraft fuel flow in pounds per second and is primarily a function of altitude and velocity in cruise, V_g is the ground speed in feet per second, and CI is the Cost Index converted from the standard units of 100s pounds per hour to pounds per second by multiplying by $\frac{100}{3600}$. Thus, the problem is to find the speed and altitude that minimizes ψ_c.

Minimization of cruise DOC necessarily involves mathematical optimization methods. Several different optimization strategies were considered, and a qualitative analysis of available optimization software was performed. In one set of studies, speed was varied continuously through cruise at a constant Legacy recommended trip altitude. In another set of studies, constant altitude and variable speed are selected through live optimization.

Using the cost function to optimize both altitude and speed is the most optimal, as expected, but also that altitude optimization provides a much greater share of the benefit over Legacy or constant selection strategies than speed does. However, the benefit that speed optimization achieves is non-zero, and thus the optimization of both values is pursued for the Gen A technology and beyond.

The optimal cruise profile has been shown to not be sensitive to initial cruise weight; it is highly sensitive to cruise range. Figure 10 shows the effect of cruise range on the optimal altitude for the B737-800 airframe, operating at 160,000 pounds gross weight. As is apparent from the graph, shorter flight profiles resulted in Gen A choosing optimal altitudes lower than that of the Legacy optimization and longer profiles lead to higher altitudes. This plot also provides evidence that the Legacy method is intentionally optimized for short to medium range flight profiles. For this scenario, a flight range of approximately 750 nmi appears to have been the design case.

4.2.2 Optimization Algorithm

In its final form, the Gen A optimizer is a trajectory optimizing utility that generates optimal control speed targets in a series of discretized steps at an optimum altitude, determined via an outer golden-section search loop.

Inherent in the current design of the Gen A optimizer are several assumptions that limit the ability to determine the optimal solution. The current implementation of the optimizer does not include wind speed in the calculation of the aircraft true airspeed, which is equivalent to presuming a windless atmosphere. This shortcoming is addressed in Generation B – see Section 6. Also present in the current design are gross simplifications of the climb and descent distances and fuel burn, as well as an assumption that a single cruise altitude must be used for the entire cruise profile. This latter assumption is only the case for short range flights. Medium to long range flights often involve successful clearance to implement step climbs, thereby allowing multiple altitudes in a single cruise profile. This shortcoming is addressed in Generation C – see Section 7.

4.2.3 Comparing with NLP Method

Validation of the optimal profile produced by the Gen A optimizer is accomplished through the comparison with an independently generated non-linear programming method that uses commercially available optimization algorithms. More generic Non-Linear Programming (NLP) methods simultaneously optimize both speed and altitude, providing for an analysis of the coupled effect of adjusting each. This coupled analysis inevitably results in a more accurate optimal solution; however, their prohibitively long response times make them infeasible for implementing on an embedded FMS.

See Figure 11 for an example comparison. Note the start and end regions of these plots are not representative of the cruise flight phase, and exhibit some transition effects between climb, cruise, and descent. A rudimentary analysis of this cruise region reveals that the largest difference in calculated optimal speed is 4.5 KTAS, or approximately 2.4 KCAS at altitude, and ultimately resulted in 0.036% more fuel consumed over the course of the cruise trajectory. This data implies that the Gen A optimizer is sufficient to extract all meaningful fuel savings from the cruise profile.

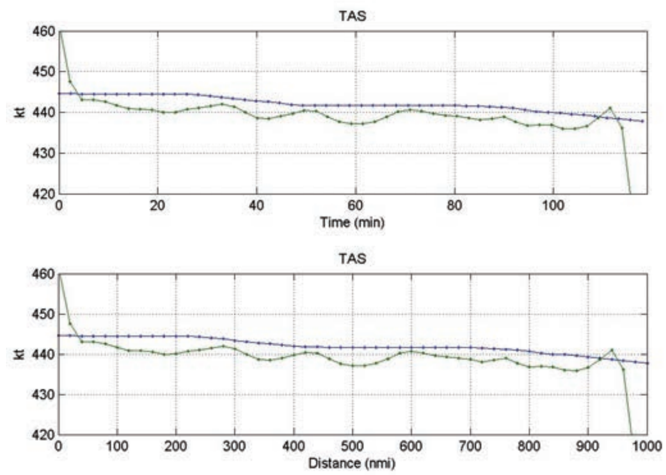

Figure 11. Comparison of NLP and Gen A optimizer cruise profiles.
The comparison shows zero distance as the top of climb for 1000-mile cruise.

The initial version of Gen A involved dividing the cruise distance into the maximum number of discretized steps (200); however, longer than acceptable response times prevented such a small distance step. Further study demonstrated that the step distance has a small effect on the optimal solution and that 10 discretized steps – as implemented in the final version of Gen A – is enough to derive the available fuel savings from flying the optimal solution.

4.3 Benefit Assessment

A MATLAB-based Cost Calculator Analysis (CCA) is conducted on a set of 468 flights of the B737-800 airframe to determine fuel savings. The test set sweeps through a representative range of aircraft initial gross weights, flight ranges, and mandated altitude separations (i.e., odd/even, flight level, or none). The Legacy and CLEEN II optimized cruise profiles are calculated and direct operating cost is roughly integrated to determine the cost of each flight.

The cost calculator approach uses simplified flight dynamics to determine the fuel consumed by each method's optimal solution. More specifically, the calculations simplify the model of fuel flow in that speed changes are assumed to be instantaneous. This simplification also assumes that slow vehicle dynamics dominate the behavior of the aircraft and the resultant forces and fuel flow. Given this formulation, the Legacy and CLEEN II optimized cruise profiles are compared in Figure 12, resulting in about 0.1% benefit (see Figure 13).

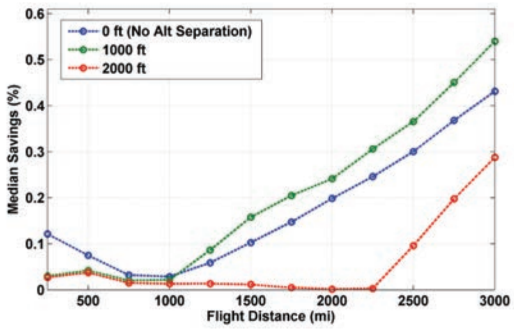

Figure 12. Median fuel savings of Gen A Cruise over Legacy.

A brief survey of data on FlightAware[§] reveals that 75% of flights traversing more than 2700 nmi involved step climbs, as did 50% of flights traversing around 1400 nmi. The same data indicated that no flights on the order of 900 nmi involved step climbs. The effect of step climb calculation is neglected in this Gen A analysis but will be handled in Gen C in Section 7.

Figure 13 contains the calculated fuel savings associated with runs with different altitude rounding implementations that mimic the typical air traffic constraints for vertical separation.

Altitude Rounding	Median Savings
0 ft (no rounding)	0.14% (20 lbs.)
1000 ft	0.16% (20 lbs.)
2000 ft	0.019% (1.7 lbs.)

Figure 13. Effect of altitude rounding on fuel savings.

This data, combined with that shown in Figure 12, demonstrate the sensitivity of fuel savings to altitude rounding and cruise distance. When a 2000-ft altitude separation is imposed, the fuel savings become insignificant. As can be seen in Figure 12, this imposed separation prevents fuel savings until flight ranges increase beyond 2,250 nmi, but savings then increase dramatically thereafter.

The testing made apparent the relationship between initial weight and cost savings. Where lighter aircraft can fly relatively efficiently at a wider range of speeds, heavier aircraft require larger deflections in flight surfaces to trim the aircraft at different speeds and this can have a pronounced effect on the forces on the aircraft and, consequently, its efficiency (see Figure 14). The more accurate optimum speed of Gen A drives higher fuel savings as compared to the more generically formulated Legacy optimizations.

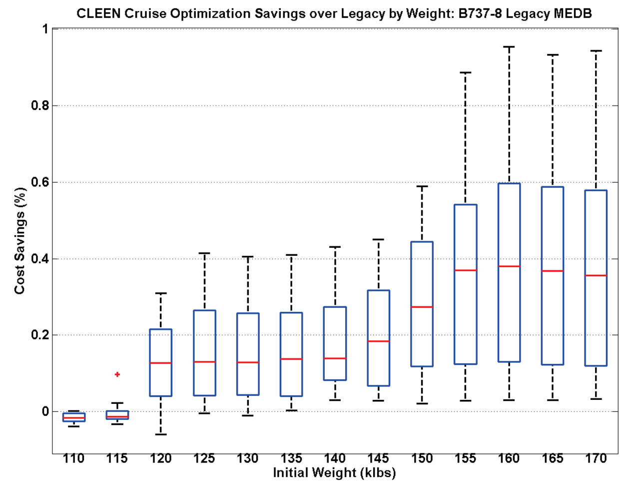

Figure 14. Relationship between weight and savings observed in testing.

Figure 15 provides evidence that Legacy method is intentionally optimized for short to medium range flight profiles. For this test set, a flight range of approximately 750 nmi appears to have been the design case, as demonstrated by the minimal fuel savings with little to no variation at the 750 nmi flight distance. For every other flight distance, there is a strong observable direct correlation between the magnitude of fuel savings and the difference of flight distance and 750 nmi. This holds true even for shorter flights, where the cost of climb and descent have a larger impact on the ability of Gen A technologies to choose a profile optimized specifically for cruise. It is likely that these more restricted scenarios are what drive the higher variability in the observed savings, shown as a larger blue box. It also holds true for longer flights, where there appears to be a linear correlation between the observed savings and total flight distance. Given the generic formulation of the Legacy methods with a 750 nmi design case, there is reason to believe even longer flights (like those of transatlantic trips) will follow this same trend.

[§] "FlightAware", FlightAware, [Online]. Available: https://flightaware.com/.

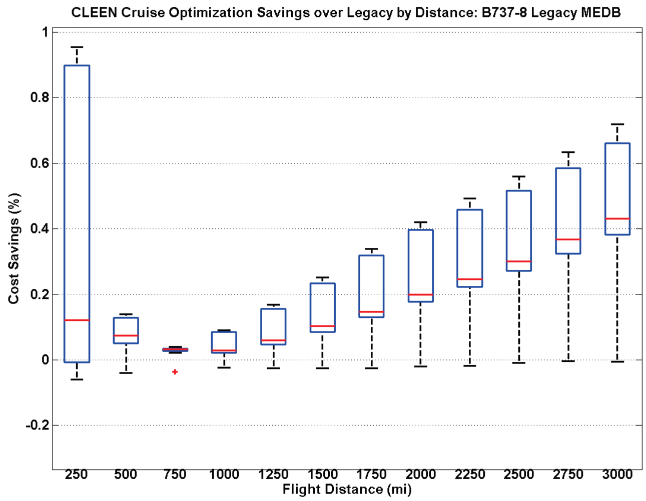

Figure 15. Relationship between flight distance and savings observed in CCA.

5. Unified Climb, Cruise, and Descent (UCCD)

5.1 Executive Summary

Previous technologies focused solely on the optimization of individual flight phases: CLEEN I on the climb profile, and CLEEN II Generation A technology on the cruise profile; see Figure 16. The work of the technology described in this section is to improve the climb phase optimization, investigate optimizations of the descent profile, and to unite their optimization of the different phases to produce a flight profile that minimizes the cost of the entire flight. Of note is the finding that maximum climb rated thrust does not provide a minimal fuel profile; that is, a reduced throttle improves life of the engine and reduces cost to operate.

In its final form, the UCCD implementation is a trajectory optimizing utility that computes the best overall flight profile and provides the full-flight optimization framework for future technology generations. An impact assessment of the resulting vertical flight profile indicates that the new control calculations have some operational impact on air traffic management and pilot training. These new policies account for typical IFR and ride quality, and typically producing a longer time to cruise altitude than Legacy controls. The operational impact of this control should be further vetted with air traffic management authorities.

TRL 5 testing described in paragraph 5.3.2 demonstrates an average **0.47%** fuel savings for Cost Index 0 as compared to the Legacy method (the average is 0.33% across all Cost Indices evaluated). The effect of engine de-rate is studied and does not affect the resulting savings. Operationally, when larger vertical separation limits are used (2000 ft), the average savings is reduced by approximately 0.05%.

While this savings is significant, it does not fully account for tailoring for specific weather patterns or a variable cruise phase control policy (for example, one that allows for altitude steps) achieved in Generation B, C, D. The initial UCCD algorithms and software provide a framework to further improve the vertical flight profile and take advantage of additional tailoring.

5.2 Development

5.2.1 Variable-Thrust Climb Optimization

Unconstrained Legacy climb profiles are generated assuming a constant CAS/Mach pair flown at maximum climb thrust to a tabulated optimum cruise altitude. These Maximum-Thrust Constant-Speed (MTCS) profiles are a rough approximation of optimal control and provide for consistent and predictable aircraft performance that can be leveraged by air traffic control and other airspace governing bodies to more easily control aircraft interactions.

The intuitive sub-optimality of the maximum-thrust, constant-speed climb profile motivates the investigation of how best to control a flight profile to better approximate optimal performance. Historical research** focuses on variable-speed profiles flown at maximum thrust and generally thought them to be the only significant source of increasing optimality. The work of the CLEEN I program sought to utilize this knowledge with modern processing power to compute a more optimal variable-speed (but still maximum-thrust, titled MTVS) climb profile that provided savings over Legacy constant CAS/Mach pair approach. See paragraph C.1 for a description of this work.

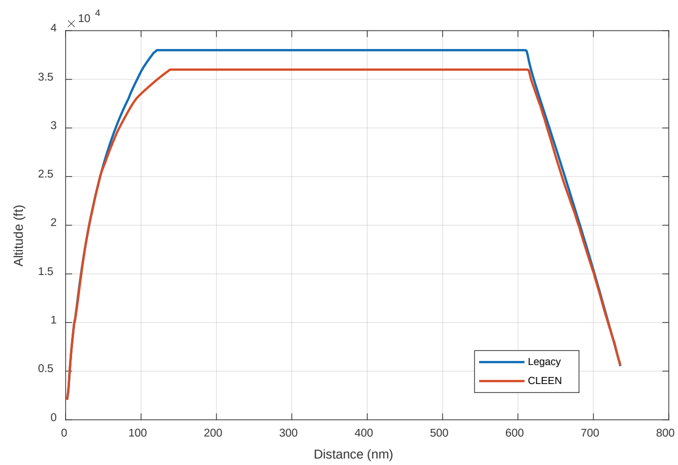

Figure 16. Legacy and CLEEN climbing profiles.

UCCD optimization technology automatically balances between climbing higher and resulting performance gains at cruise altitude

However, the work of investigating unifying the climb and cruise phases lead to the conclusion that the maximum thrust profile leaves fuel savings to be extracted. The Legacy algorithm assumes engine operation at maximum climb thrust, which is substantially higher than the cruise thrust setting. This thrust transition from a high climb setting to a lower cruise setting is executed and controlled by the auto-throttle system during transition to the cruise flight phase. During this time, closed loop feedback within some auto-throttle systems targets the selected (cruise) altitude and the throttles dial back as necessary to achieve its vertical speed targets while still flying the constant Mach. This results in a throttle transition that is sub-optimal and motivates consideration of a variable thrust (or throttle) and speed profile. The three graphs in Figure

** H. Erzberger and H. Lee, "Constrained Optimum Trajectories with Specific Range", *Journal of Guidance and Control,* vol. 3, pp. 78-85, 1979.

17 respectively depict the distance-based altitude, thrust, and fuel savings profiles for one flight.

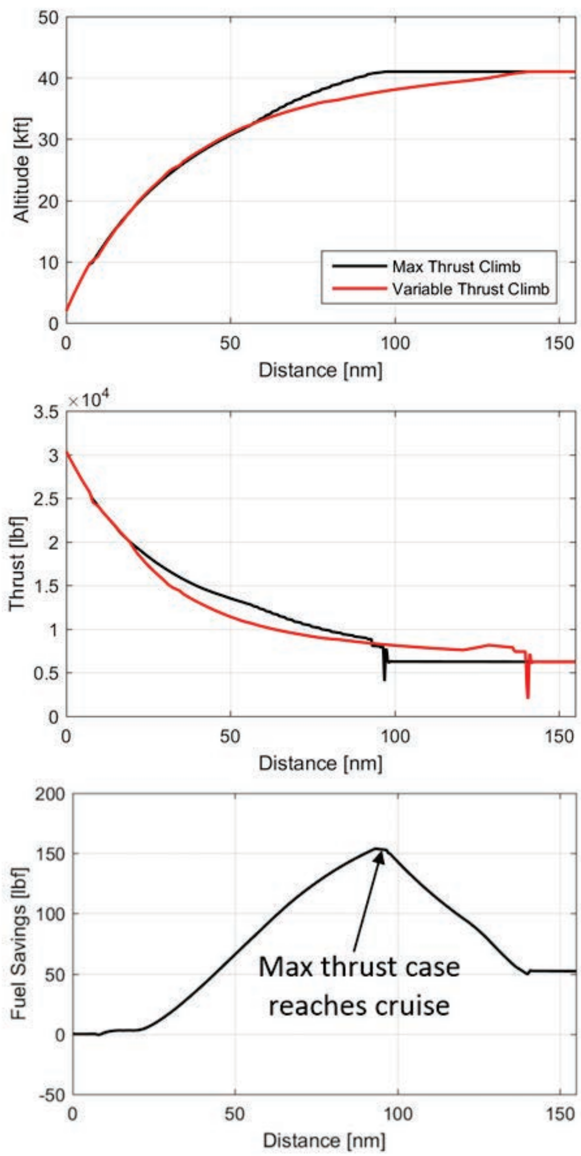

Figure 17. Variable thrust profile and savings impact.

In each of the plots, the red line represents a variable thrust climb profile and the black line the Legacy maximum climb thrust profile. Approximately 25 nmi into this flight, the optimum variable thrust control (red line) deviates from the maximum thrust (black line). At this same point the bottom graph begins to show fuel savings, as the reduction in thrust inevitably leads to a reduction in fuel flow. Although variable, this comparative reduction in thrust (and fuel consumption) persists until the Legacy profile achieves the cruise altitude (around 90 nmi along-path) and the thrust is dialed back to its cruise setting (where thrust is equal to drag). From this point until the point at which the variable thrust profile reaches the cruise altitude, the Legacy profile is at a lower thrust setting and is consuming less fuel. This is reflected in the bottom plot that shows decreasing but still positive savings from 90 nmi until about 140 nmi when both profiles are flying at cruise thrust at the cruise altitude.

A closer look at the altitude-based cost of different thrust settings further confirms the ability of a variable thrust profile to save fuel. Figure 18 is a contour map of the climb cost function evaluated as a function of altitude and thrust value. At each altitude, an ascending aircraft can achieve a positive rate of climb at a wide range of thrust settings. The cost of operation at each of the altitude and thrust combinations is represented by the contour lines that are shown on a temperature color scale with blue being lower costs and red/yellow higher.

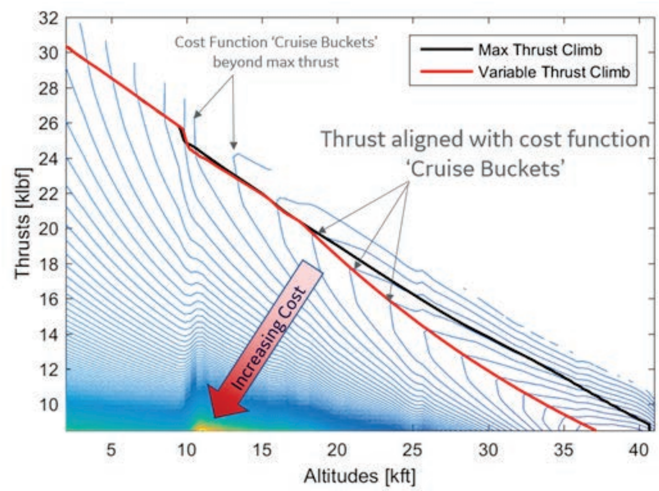

Figure 18. Contour map of climb cost function.

Again, the black line represents a maximum thrust profile and the red the variable thrust climb profile. Though difficult to discern from the contour colors, the cost of each thrust and altitude combination tends to increase as each value is reduced. Similarly, at each altitude the instantaneous cost tends to decrease as the thrust increases to its maximum climb value. This holds true at all altitudes below approximately 20,000 ft. Above 20,000 ft there is a minimum in the cost profile as a function of thrust that lies within the viable thrust region. These minima create the referenced "cruise buckets" whose minimums (at the top left apex of each curve) lie below the maximum thrust value, which makes them viable for use as optimal control. This indicates that an increase in thrust beyond this minimum results in an increase in instantaneous cost. This drives the optimal variable thrust profile to deviate from the maximum thrust value, which is exactly what is observed. As the profiles approach 20,000 ft the variable thrust profile (in red) begins to depart from the maximum thrust profile (in black) and does so to follow the apex of the cruise buckets as expected. This deviation near 20,000 ft was widely observed and correlates strongly with the thrust deviation in the middle graph of Figure 17.

Beginning with the previously formulated CLEEN I equations of motion in climb (paragraph C.1), the simple difference to create a variable-thrust climb profile is that thrust is now varied in the search for optimality rather than assumed constant at its maximum climb value. Thus, the direct operating cost function is minimized over values of speed and throttle (or equivalently thrust). In order to produce such a profile, an optimizer that can handle the optimization of more than one variable simultaneously is required. This new method of climb optimization is Variable-Thrust, Variable-Speed (VTVS) and unifies the climb thrust to the cruise thrust in an optimal transition.

Each of the methods (CLEEN I MTVS and the new VTVS) are tested and compared with Legacy functionality (MTCS) to evaluate their efficacy for inclusion in the UCCD optimizations. Also compared is a Variable-Thrust Constant-Speed (VTCS) method, for sake of completeness.

Each climb optimization method (including that of the Legacy software) is formulated in the TRL 5 environment and compared using simplified flight dynamics over a set of 156 test cases for each optimization method (12 flight distances from 250 nmi to 3,000 nmi at an interval of 250 nmi and 13 takeoff gross weights from 110 klbs to 170 klbs at an interval of 5 klbs). For each test case, vertical flight plan states are generated (predicted) for both Legacy and CLEEN profiles, assuming a start in climb at 2,000 ft, flight up to the optimization method determined cruise altitude along a straight-line route, and an estimation of the Legacy descent. For all non-legacy methods, the cruise altitude was chosen via golden-section search based on the direct operating cost of each altitude (including the cost to climb to that altitude using the optimization method being evaluated). This approach differs from the final implementation that conducted a brute-force search across a viable altitude range. Each flight prediction subject to physical and operational restrictions, including minimum and maximum vertical speed requirements, speed and thrust limits, and acceleration limits.

The resulting profiles are evaluated using the CLEEN II Advanced Technology Testbed (ATT) to simulate how the actual aircraft would fly the predicted flight. The ATT calculates and records fuel burn, flight time, speed, altitude, thrust, and observed drag. Figure 19 presents the aggregate fuel and time savings for each method across the entire 156 test flights.

Only three optimzation methods are shown on each graph – Maximum Thrust With Variable Speed (MTVS), Variable Thrust With Variable Speed (VTVS), and Variable Thrust With Constant Speed (VTCS). The fourth method, the Legacy Maximum Thrust Constant Speed (MTCS), is not shown because it is used as the control for comparisons, so each improvement or diminishment is made relative to the MTCS profile.

Both multi-variable climb control optimizations made use of MATLAB's fminseaerchcon function to produce the optimal control used in the profile generation. This function produces equivalent results to the Active Set method used in the final version of the solvers.

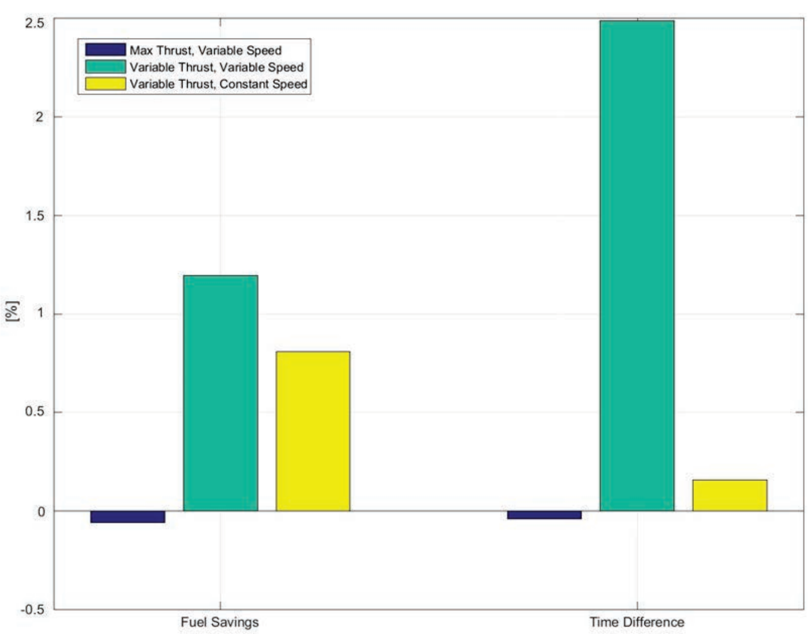

Figure 19. Climb optimization method.
This shows the VTVS method provides significant benefit.

5.2.2 Variable-Thrust Descent Optimization

Analogous to the variable-thrust variable-speed climb, a study of the benefit of descent thrust optimization is performed. Descent optimization concludes that savings can be improved from 0.47% savings from the UCCD assessment in paragraph 5.3.2 to 0.78%. However, implementation of the descent optimization in the TRL 6 software posed several challenging problems related to the handling of flight constraints and arrival procedures during descent and was determined to be out of scope for the CLEEN II project.

5.2.3 Phase Unification

Intuitively, it can be reasoned that the cruise altitude selection affects the cost of the climb and descent flight phases. Climbing to a higher altitude inevitably includes a longer climb phase that takes place at a higher thrust setting than that of cruise and demands more fuel. A similar effect occurs for descent due to the longer duration required to descend from a higher altitude. Does the cost to ascend and accelerate to the best cruise altitude and speed yield the lowest total cost or would a lower-cost climb to a lower, less-efficient cruise altitude be cheaper? The answer depends on range — how long the airplane operates at the more efficient cruise altitude. Keep in mind the adjacent effect of weight. A lower cost climb means a heavier airplane at cruise altitude and, as range increases, the required amount of fuel increases and the airplane gets heavier.

The problem formulated here must either be solved iteratively or simultaneously to find the best combination of climb, cruise, and descent controls. To minimize the total cost, then, the cost of the climb, cruise, and descent must be solved simultaneously — either analytically, iteratively using numerical methods, or a hybrid mix of both approaches.

The UCCD software design solves this problem through exchange of planning information and optimal control results between individual optimizers for each phase. Convergence loops and brute force search over cruise altitude is used to find the overall optimal solution.

The profiles generated by the UCCD technologies are traditional flight trajectories; that is, they include the three standard phases of flight – a climb followed by a cruise and then a descent. The cruise and descent phases are essentially identical to that of the Legacy profile, only differing in the altitude and speeds chosen to fly. The climb profile used in the UCCD technologies has a deviation from the Legacy climb profile in that both speed and thrust were optimized as variable aircraft controls.

Figure 20 depicts the impact of the UCCD technology on the flight profile of two example flights. In both graphs, the black line is the trajectory produced via Legacy profile prediction methods and the red is that of the UCCD technologies. The graph on the top shows that the UCCD methods choose a lower cruise altitude for shorter duration flights while the graph on the bottom shows its tendency to pick higher altitudes on longer flights. Intuitively, this makes sense. While turbofan engines tend to operate better at higher altitudes, the cost to climb to that higher altitude may outweigh the efficiency savings, as is the case in the left graph. When the aircraft can maintain that higher altitude for a substantial duration, the efficiency savings from a higher altitude accumulate and provide for a lower total flight cost, including the cost to climb to that higher altitude, and drive the selection of that higher altitude as the optimum cruise altitude.

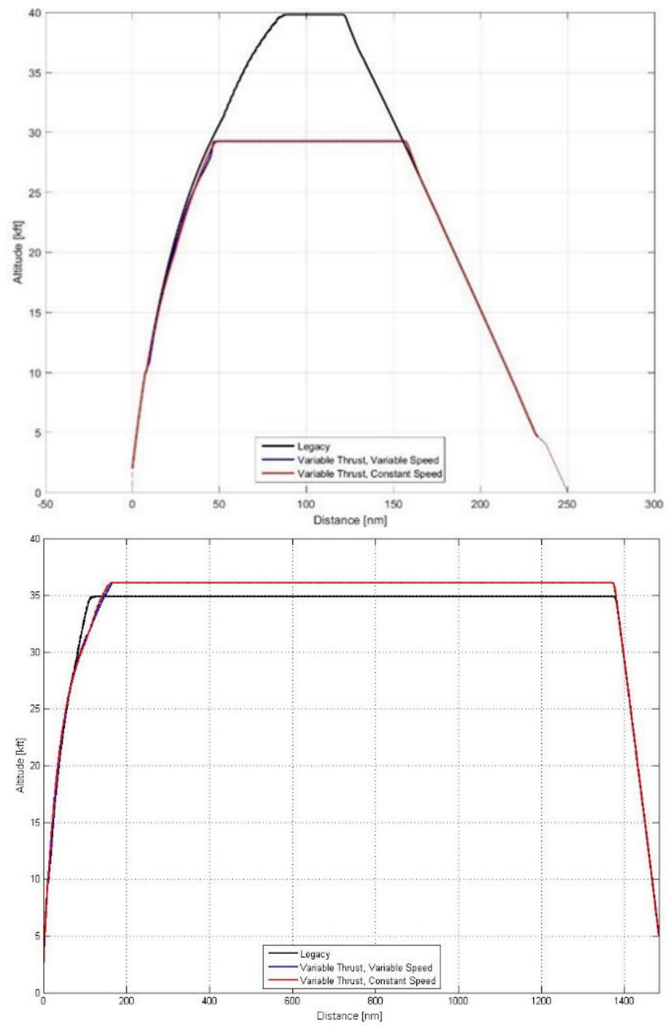

Figure 20. Examples of UCCD profile impact.

Although these modifications to the flight trajectory are driven by aircraft control that drives a mathematically optimal profile, there is concern about their ability to be adopted into commercial airspace. Operators in FAA control centers use well established dead reckoning techniques to coordinate the flow and order of airspace traffic and any changes in aircraft behavior must gain the acceptance of all stakeholder communities to avoid resistance to the purchase and use of the UCCD optimizations. In addition, the variable-thrust climb profiles extend the aircraft's time and distance in climb. This can lead to a climb profile occurring through more than one ATC sector, which drives the need for cruise entry coordination across control centers. This is a change to the standard operating procedure. However, there is a large variation today in aircraft type, weight, and wind profiles that can drive substantially different climb profiles, possibly on the order of the changes driven by the variable thrust profiles. It is recommended that further research is conducted, in coordination with ATC representatives, to investigate the difficulty in adopting these extended climb UCCD profiles.

5.2.4 Optimization Algorithm

In the final design, the UCCD problem is solved by iterating through potential cruise altitudes, solving the climb problem for each, generating a constant-altitude cruise profile, and appending the descent trajectory. Upon completion of this process, the full flight profile that yields the overall lowest cost is selected. In this formulation, the fuel flow rate and consequently the fuel cost is calculated via solving the system equations for the most optimal climb thrust and velocity and cruise altitude and velocity.

After the optimal data is determined valid, the produced optimal control is incorporated into the flight plan along with all applicable constraints via the execution of an FMS prediction event. At this point the entire flight profile has been optimized in a unified fashion and the means of controlling the aircraft to the optimal trajectory is stored and available for access through aircraft model functions.

5.3 Benefit Assessment

5.3.1 Test Methods

The UCCD trajectory optimization technologies are validated through execution of an extensive test set of aircraft flights. Aerodynamic, thrust, drag, and fuel flow data for the Boeing 737-800 airframe with CFM International LEAP-1B engines is used to construct straight-line flight profiles beginning at 12 different distances ranging from 250 nmi to 3,000 nmi at an interval of 250 nmi. Eleven different test sets each consisting of combinations of these 12 distances and 13 gross weights (ranging from 110 klbs to 170 klbs at 5 klbs intervals) are predicted, simulated, and analyzed to assess the fuel savings available through the UCCD optimizations.

Cost Index, engine de-rate, and altitude quantization were varied across the nine test sets to allow for the assessment of their impact on the ability of the optimization methods to extract fuel savings. Figure 21 contains the mapping of each of the inputs varied across test sets. The same set of distance and gross weight conditions are used for each test set.

Each test flight is predicted using point mass integrations of both the UCCD and Legacy flight optimization methods. Both prediction methods begin with the aircraft in climb at 2,000 ft and 250 KCAS, fly up to their respective optimum altitudes, and descend to destination altitude. These predicted profiles are then input into the CLEEN II ATT to simulate the flight of a physical aircraft along the predicted flight profile, tracking control outputs of the predictions and performing integrations of “real-time” aircraft states.

Note that this assessment process differs from TRL 5 assessments performed for subsequent technology enhancements. In later generations, ss the technology gets more sophisticated, considering more specific flight plan tailoring (and thus extracting more benefit) it is impractical to enumerate and run all possible scenario variations. For this initial UCCD assessment, all full sweep of all cases is practical, and is performed.

Test Set	1	2	3	4	5	6	7	8	9
Cost Index	0	0	0	100	100	0	100	0	25
Derate 1	0	0	0	0	0	1	1	0	0
Derate 2	0	0	0	0	0	0	0	1	0
Altitude Quantization (ft)	0	1000	2000	0	2000	0	0	0	0

Figure 21. Varied inputs in UCCD test sets.

5.3.2 Test Results

The UCCD and Legacy simulated flights are compared to assess the achieved fuel savings, using the direct operating cost formulations described in paragraph 3.1. Figure 22 depicts the aggregate cost of each test set relative to the Legacy profile, with positive percentages indicating savings over the Legacy profile prediction method. The aggregate savings across all test sets is approximately 0.33% of the fuel used without UCCD optimizations.

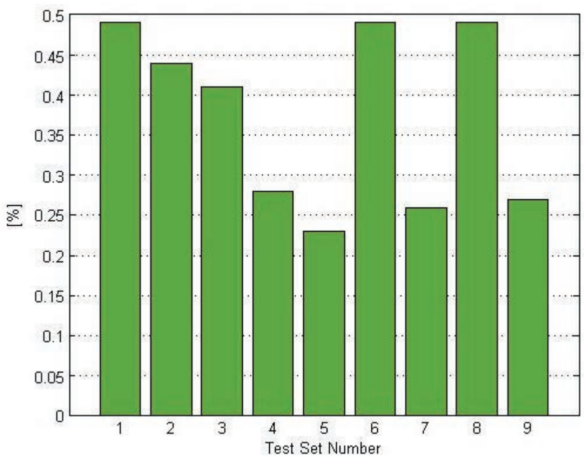

Figure 22. Cost savings of UCCD method over Legacy method for each test set.

The available flexibility of control in the optimization directly effects the achievable benefit. This correlation can be seen in test sets whose only uncommon input is Cost Index (for instance, test sets 1, 9, & 4 or 6 & 7). By definition, when the Cost Index is zero the time spent in flight has no impact on the total flight cost. A non-zero Cost Index serves as a conversion factor between time and fuel, providing a way to factor the time cost in the total direct operating cost. Consequently, the aircraft must fly faster to reduce the cost (time and fuel) of the flight, which forces unfavorably high speeds that drive engine operation in inefficient regimes. Cases with a zero-valued Cost Index demonstrated approximately 0.46% fuel savings that corresponded to an average of about 56 lbs. of fuel, or 8.3 gallons of Jet-A1 Fuel.

Commercial airline flights cruise at discretized flight levels, with east-bound flights utilizing odd flight levels and evens for west-bound flights. Test sets 2 and 3 address the effect of this limit. The effect of altitude quantization is small, but significant, reducing the savings by approximately 10% for each additional 1000 ft.

Engine de-rates are a common method of commercial airlines to reduce takeoff and climb thrust in order to extend the life of the engines and potentially reduce the fuel consumed. Many airlines implement two engine de-rate settings, so test sets 6, 7, and 8 are run to assess their impact on savings in typical commercial operation. In these sets, the Legacy profile prediction is generated assuming the de-rate reduced engine thrust. Test sets 1, 6, and 8 result in the same percentage of fuel savings, demonstrating that the engine de-rate setting has no effect on the ability of the UCCD algorithms to extract fuel savings.

Figure 23 provides insight into the source of the cost savings and the additional benefit of the UCCD technologies not reflected in the Cost Index 0 sets – time. With the exception of test set 3, the zero-valued Cost Index test sets not only saved fuel, but also reduced the time required to complete the test flight as compared to the Legacy method. Operationally the reduced flight time from origin to destination can lead to significant monetary benefit and a reputational boost. Many business customers of the commercial airline industry simply require the quickest way to get from point A to point B and actively seek to book flights with the shortest flight time. Using the UCCD technologies serves to satisfy these customers without compromising fuel savings in this test set. In fact, with few exceptions, all test sets showed both time and fuel savings. Figure 24 collects the time, fuel, and cost savings, grouping the test flights by their input Cost Index value. It serves to highlight the unaccounted-for benefit.

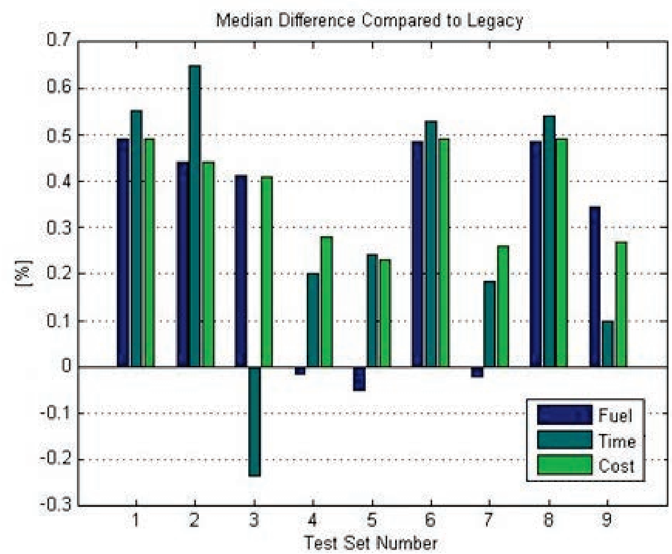

Figure 23. Fuel, time, and cost savings of UCCD method over Legacy method for each test set.

	Cost Savings	Fuel Savings	Time Savings
CI = 0	0.47%	0.47%	0.47%
CI = 25	0.27%	0.34%	0.09%
CI = 100	0.27%	-0.03%	0.26%

Figure 24. Cumulative UCC savings.

There is also substantial variation across flight groupings within each test set. The test set numbers are indexed in such a way that each group of 12 consecutive test cases corresponds to a single gross weight and an increasing flight distance. Thus, as can be seen in Figure 25, at lower weights, the shorter distances save more fuel than the longer flights, but the opposite is true at the higher weights. This is indicative of the optimizer ability to account for both off-average aircraft weight and cruise distance in its calculation of the optimum cruise altitude. As the aircraft increasingly deviates from the nominal condition documented in the database tabulated optimal control, its deviation from the true optimum increases. This leads to a higher savings margin achieved by the UCCD technologies.

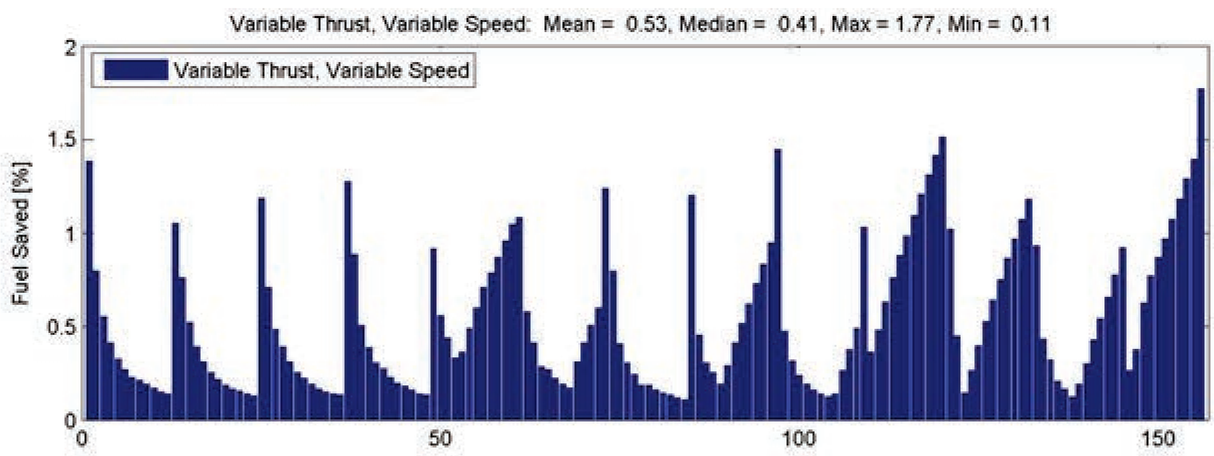

Figure 25. Test Set 1 fuel savings.

6. Generation B: Weather-Optimized UCCD

6.1 Summary

The work of the Generation A and the initial UCCD developments focused on using engine and aerodynamic models to produce a unified optimal flight profile when operating at standard day conditions. The work of the Generation B technology described in this section is to incorporate non-standard day weather data and produce an optimal flight profile that harnesses favorable weather conditions.

Current methods used by Flight Management Systems to account for weather effects on flight profiles rely upon the assumption that the known weather data poorly reflects what will be experienced throughout the flight. As such, many airlines pursue overly simplified wind and temperature models for use in profile predictions. This “bad data in, bad data out” approach drives performance predictions that deviate substantially from the observed performance. These imprecise predictions have a cascading effect on Required Time of Arrival performance and descent profile predictions, driving frequent updates to each that confuse operating crews and increase pilot workload. The Generation B technology described in this section incorporates high-fidelity data for both wind and temperature (see Figure 26) into the existing Unified Climb, Cruise, and Descent (UCCD) framework.

In its final form, the Generation B technology performs the following steps:

1. Collect weather data along a 2D grid (altitude and distance) that follows an approximation of the lateral path associated with the entered flight plan
2. Conducts a brute force cruise altitude search – with an embedded multi-variable climb path optimization and Legacy descent path estimation at each altitude – to determine the optimum altitude and control commands that produce a profile with the lowest flight cost for any origin and destination combination.
3. Predict the flight plan using the optimal control.

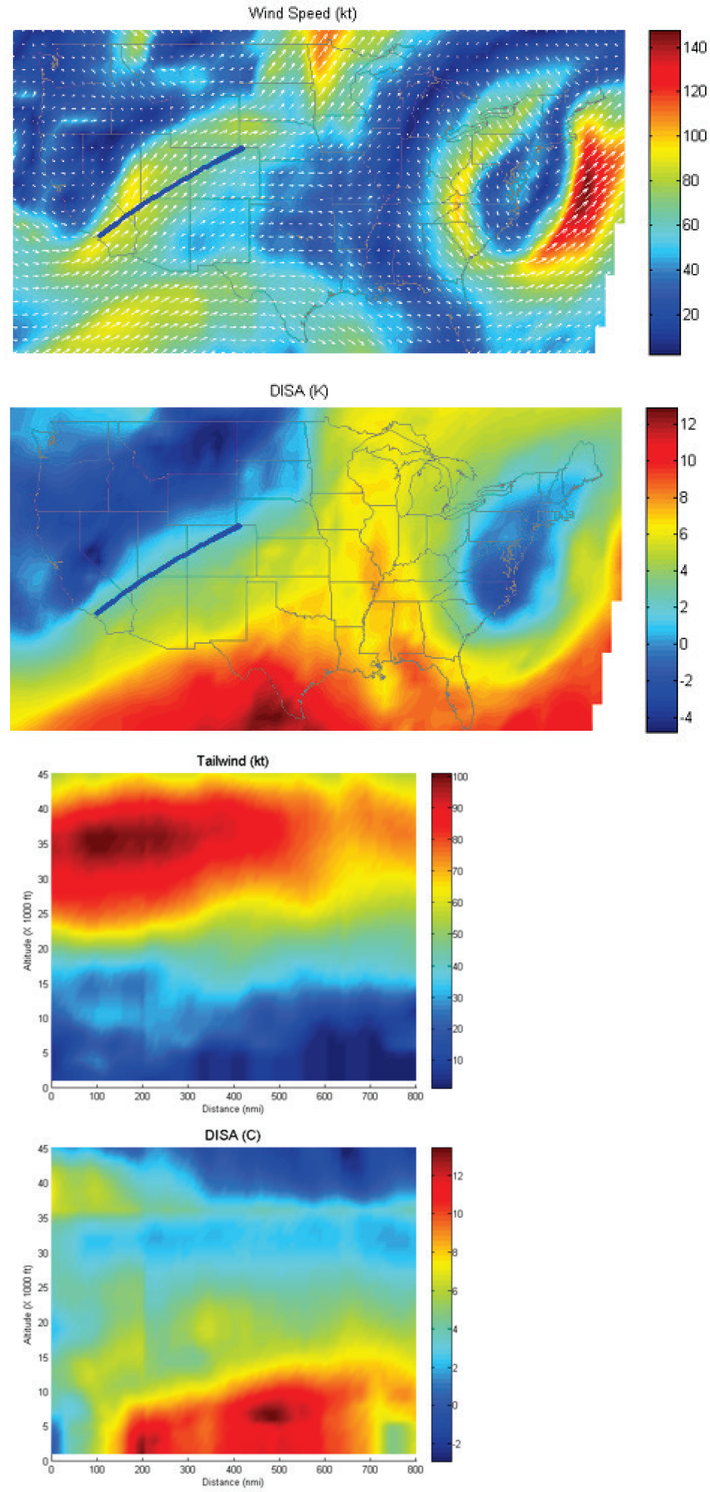

Figure 26. High-fidelity wind and temperature example.

High fidelity wind and temperature along the path enables more accurate cost calculations.

In TRL 5 benefit assessment testing (described in paragraph 6.4), Generation B demonstrates that an operator can achieve **0.54%** fuel savings as compared to the Legacy method and **0.44%** as compared to Legacy+ method (Figure 27). The Generation B technology produces an optimal flight profile that accounts for weather and minimize direct operating cost in compliance with all requirements, providing an ability to extract substantial cost savings. As with the Generation A and UCCD developments, the Generation B technology serves as a development step towards future generations of the optimizer.

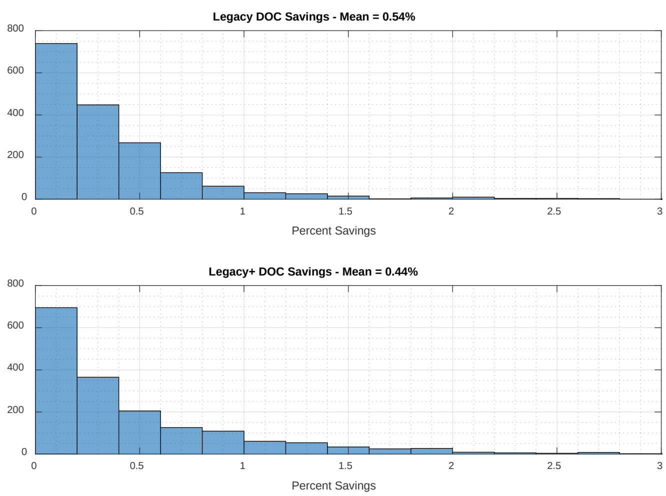

Figure 27. TRL 5 benefit assessment testing.

Monte Carlo study in TRL 5 simulation testbed shows significant savings from Generation B technology.

6.2 Development

6.2.1 Weather Data Source

Several publicly available weather data sources exist, each providing a 4-dimensional model (latitude, longitude, altitude/pressure, and time) of observed and forecasted weather. Data is available in various resolutions depending on the data source and region, with higher fidelity typically available in higher traffic regions.

The NOAA National Operational Model Archive and Distribution System (NOMADS) is a network of data that accesses and integrates models and other data stored in geographically distributed repositories into heterogeneous formats. Initially, the Rapid Refresh system (RAP) forecast from NOAA is selected for use with the CLEEN program; later replaced by NOAA with the High-Resolution Rapid Refresh (HRRR) forecast.

Both RAP and HRRR are regional weather forecast models of North America, with separate sub-grids (with different horizontal resolutions) within the overall North America domain. These models produce forecasts hourly with time resolution of one hour out to 18 hours in the future. Refer to NOAA publications for more information.

RAP data is stored as a 4-dimensional grid (latitude, longitude, pressure altitude, and time) and can be retrieved from an online archival system for each of the 365 days that precede the retrieval date. Once downloaded, this data can be stored on a local hard drive and queried by a weather server and a client user. This historical cache of weather data is used for CLEEN analysis purposes, where live data would be used in operational practice.

From the weather subscriber's perspective (in this case the optimization algorithm's perspective), the weather service provides weather information extracted from a specified day's weather record/forecast on a point-by-point basis. The service provides this information properly transformed to the requested local point's latitude, longitude, and altitude, and doesn't require any additional manipulation or interpolations (other than units conversion depending on your requirements).

6.2.2 Weather Data Processing

Figure 28 shows a map of wind speed data across the United States at an arbitrary selected altitude for display purposes. The bold line represents a sample flight path.

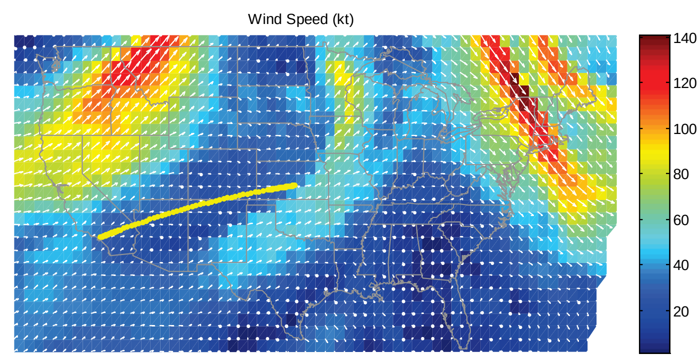

Figure 28. Weather data after extraction example.

An initial estimate of the flight profile is used to extract weather data resolved along two-dimensions: along-path distance and altitude. These data include wind speed and temperatures (color mapping), and wind direction (shown as vectors in Figure 28).

Along this two-dimensional grid of distance and altitude, the weather data is retrieved for equally spaced points at 10 nmi distance increments and 50 equally spaced altitude levels from 1000 ft to 42000 ft. The grid method was chosen as opposed to using some "tube" strategy around the initial predicted profile as it demands only a slightly higher amount of data and guarantees coverage of all possible flight paths. Figure 29 depicts an example of this weather grid for one flight profile, as is used within the Gen B optimization. Compared to the Legacy methods of weather data retrieval and entry (see paragraph 3.3), this method produces a much higher-fidelity representation of the weather conditions.

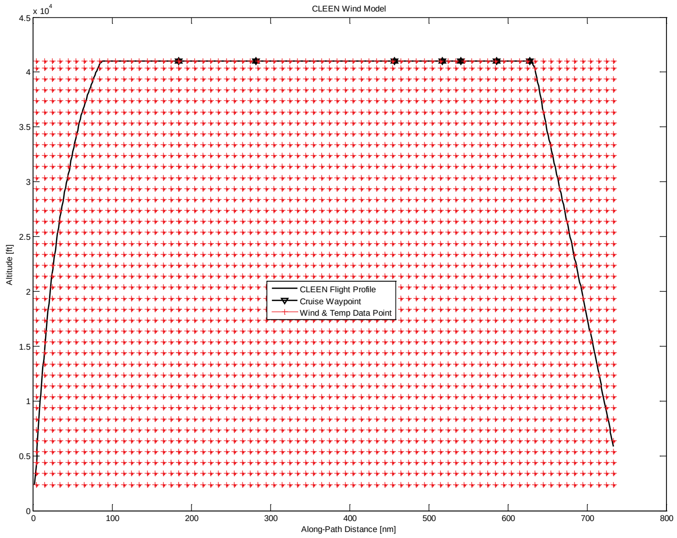

Figure 29. CLEEN weather model - equally spaced 2-dimensional grid.

Figure 30 shows the wind and temperature models used within the Gen B optimizer (top-most graphs), as extracted along the route depicted in the two lower figures. The tailwind shown in the lower graph becomes substantially larger at higher altitudes. This data agrees with expectation for a westward flight that has the possibility of exploiting the jet stream.

6.2.3 Optimization Algorithm

The optimization algorithm of Gen B is nearly identical to the previous UCCD algorithm. The only changes made were inclusion of the weather data in all aircraft model calculations. Following collection and extraction of this weather information, the data is available for use throughout the optimization and predictions calculations. Unlike Gen A that assumed standard day conditions, the fuel flow and time of flight parameters in the direct operating cost function use calculated groundspeed and temperature deviation. The 4D state predictions similarly use the weather information. Most of the changes for the Gen B optimization software is the inclusion of weather data gathering in predictions, which is a part of the TRL 5 software design.

6.3 Algorithm Validation Testing

This section describes analysis of direct operating cost savings and robustness of the Gen B optimizer using contrived weather scenarios designed to stress the system using the TRL 4 MATLAB implementation of the optimization algorithms.

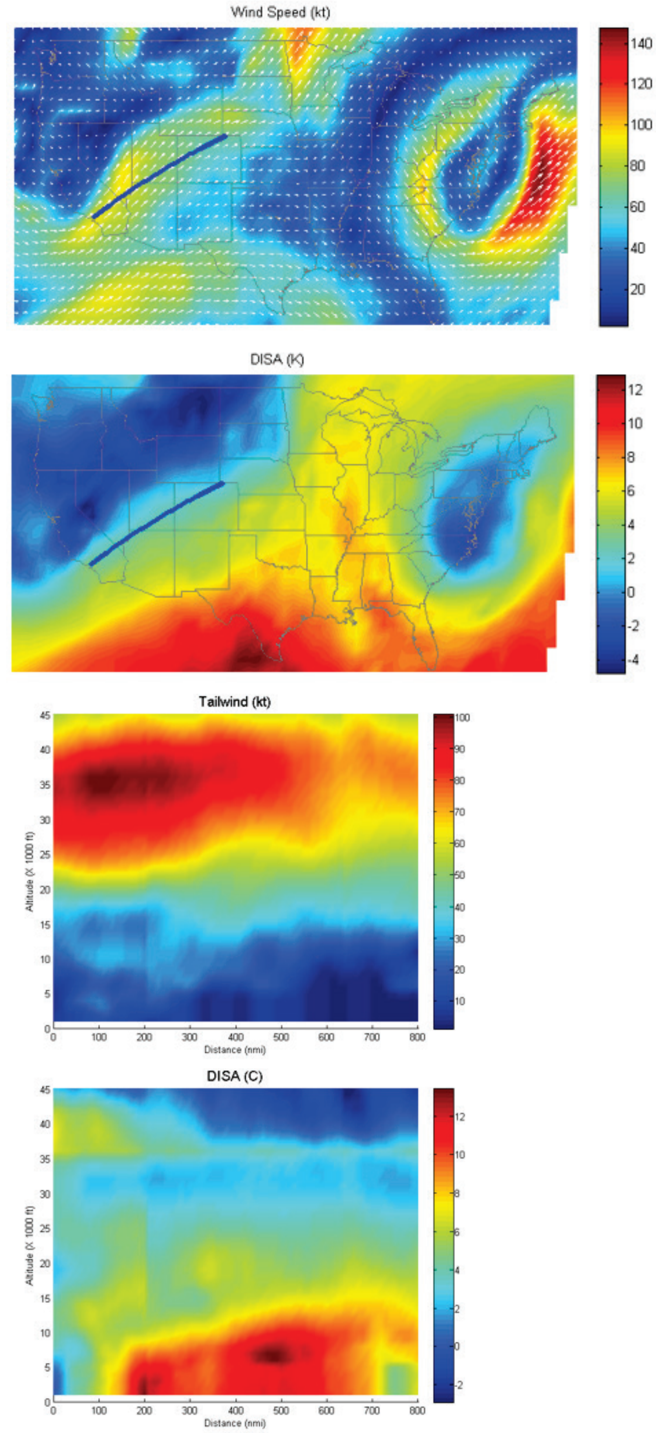

Figure 30. Raw weather data (top) and extracted along-path weather data (bottom).

6.3.1 Test Methods

Both the legacy systems (Legacy and Legacy+) and Gen B algorithms are implemented in a TRL 4 MATLAB test suite using simplified point-mass integration of the trajectory. Once a planned control trajectory is calculated, separate point-mass simulation is performed to compute the cost of each profile.

See paragraph B.1 for a detailed description of the TRL 4 testing method.

The eight test sets shown in Figure 31, totaling 13,584 flights, are developed to evaluate the technologies. In all, these flights drive 4,528 comparisons between the three vertical flight planning methods (Legacy, Legacy+, and CLEEN).

#	Description	Cases
1	Constant Wind	840
2	Constant DISA	1008
3	Ramp Wind	1512
4	Wind “Jet Stream”	2880
5	DISA “Jet Stream”	2880
6	Wind Sine Wave	2304
7	Weather Days (No Error)	1296
8	Weather Days (Real Error)	864

Figure 31. TRL4 test set flight breakdown.

Test sets 1 through 6 represent the contrived weather scenarios. These scenarios use relatively simple mathematical formulas to drive weather patterns that have clear effects on the optimal solution. For instance, test set 4 simulates large tailwinds at arbitrary altitudes, like the effect of the Jet Stream. Test set 5 tests an analogous scenario to test the effect of large temperature deviations from standard day atmosphere. These scenarios were generated with the understanding that the CLEEN and Legacy+ optimization implementations should choose cruise altitudes close to the most favorable tailwind.

Test sets 7 and 8 use actual weather data for several specific dates. The weather source described in paragraph 6.2.1 is used for these test sets. Test set 7 demonstrates the entitled benefit when the forecast weather matches exactly with the experienced/simulated weather. Test set 8 is designed to be more representative of real-use scenarios, in that it drives discrepancies between weather used to generate the optimal profile (forecast) and what is experienced during the simulation of the flight. This real-use scenario is intended to reveal the impact of forecast error on the benefit of the Gen B technologies as compared to existing optimization methods.

6.3.2 Test Results

The relative benefit achieved by the Gen B technology for each test set (1 - 7) is presented relative to the Legacy implementation (on bottom) and the Legacy+ (on top) in Figure 32. The savings achieved compared to the Legacy are significantly greater than that compared to the Legacy+ method. Intuitively, this makes sense. If correctly implemented, a method of optimization that accounts for an average tailwind throughout cruise should provide a relatively substantial benefit as compared to one that does not. Accordingly, the Legacy+ method uses substantially less fuel than the Legacy method and the relative Gen B benefit is reduced.

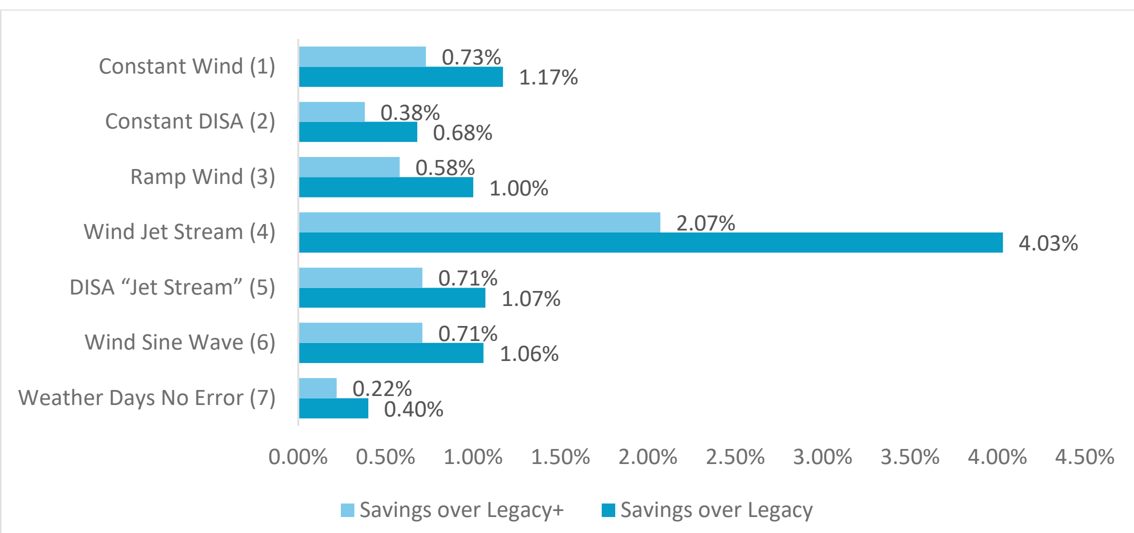

Figure 32. Gen B fuel savings for test sets 1 – 7.

Comparing the relative benefit of the CLEEN technology in test set 2 versus that of 1 demonstrates the impact of harnessing tailwinds as compared to temperature deviations. The increased benefit achieved through the harnessing of tailwind implies that wind is a larger factor in optimal profile calculations than temperature deviation, however these parameters are different units so they may not be the same relative scale of deviation.

Also interesting in Figure 32 is the magnitude of savings achieved in test set 4, especially as compared to set 7. This may be partially due to the large magnitude of wind speed used for analysis. Chosen discretely at -100 and 100 knots, these values and their invariance with distance drive large penalties to the Legacy profile when the headwind “jet stream” aligns with the Legacy optimal altitude (that does not consider wind). While 100 knot winds are not unreasonable for flights that encounter the jet stream, their presence is known and accounted for by dispatching operations that select cruising altitudes that favor tailwinds.

Test set 8 is not shown in Figure 32 due to its fundamentally different nature. Breaking the data out in greater detail does well to illustrate the relevance of the test set. Test set 8 explores the effect of weather error and how it erodes the savings that the Gen B method achieves over the Legacy and Legacy+ methods. In this set, all three methods (Gen C, Legacy, and Legacy+) are given the same weather data, which is designed to be inconsistent with the actual weather experienced in the simulation of the flights. In this section, “predicted weather” refers to the data passed to the optimization techniques and “observed weather” refers to that used for simulation. To measure the effect of producing a flight profile with predicted weather that is a measurable amount different than the observed weather, the *Mean Absolute Error* (MAE) statistic is used.

For a given set of predicted weather and observed weather, the MAE between the two is calculated along the flight path. It is computed from absolute differences between the predicted and observed values of both tailwind and DISA at points along the flight path.

In order to generate differing predicted and observed weather for this test, “forecast lead time” is used. For example, given a flight that is scheduled to depart at 6 pm, a 4-hour lead time weather forecast means that the predicted weather used to generate the flight profile comes from the weather from 2 pm that same day. Note that a forecast lead time of 0 hours indicates that the predicted and observed weather are identical.

In Figure 33, the boxplots are colored to represent the scatterplot data from the right. The left subplot shows how the savings erode with an increasing staleness. The right subplots show the same data but expanded as plots against the tailwind and DISA MAE. As expected, a more stale weather data corresponds to a larger MAE. Observe, for instance, the violet scatter points on the right edge of the upper right subplot. In these cases, a 12 hour staleness can result in tailwind MAE of about 15 kts and loses up to 1.5% relative to the 0 hour staleness case. This implies that if the Gen B optimizer saved 2% over the Legacy+ method in this case, the weather error has the potential to erode those savings to 0.5%. In some cases, this results in the Legacy+ algorithm performing better than Gen B.

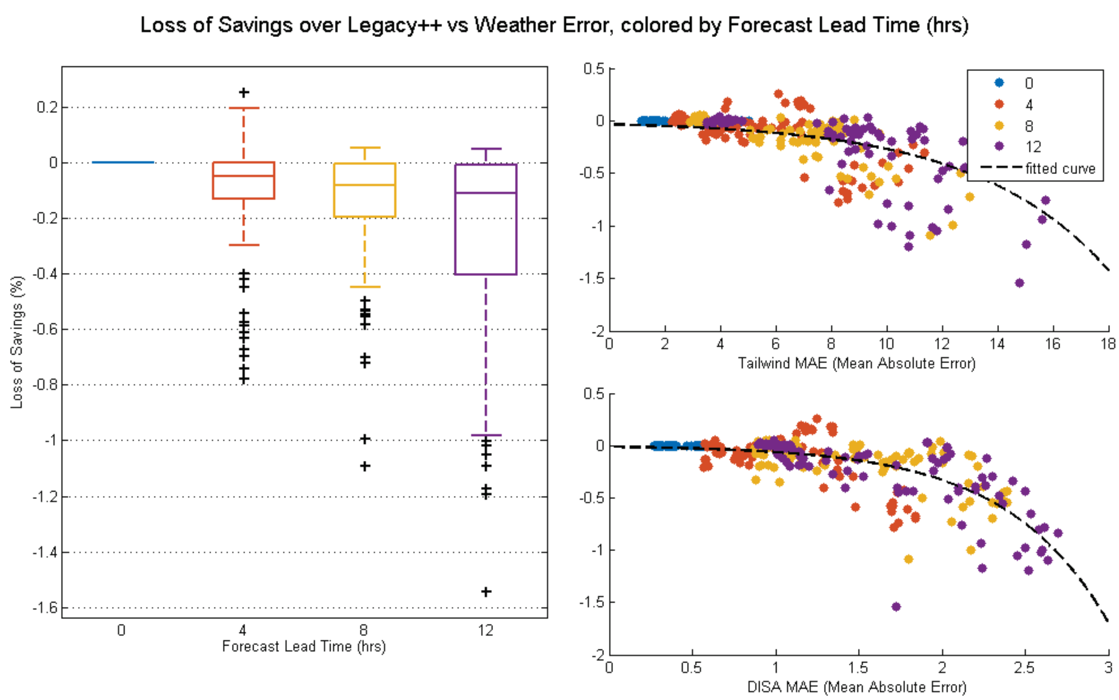

Figure 33. Test set 5.

Test set 5 shows how error in weather can erode savings.

To highlight data trends, an exponential curve is fitted to the data in the scatter subplots. The boxplots also help identify trends. For instance, if the weather staleness is eight hours or less, then the yellow boxplot shows that the 25th to 75th percentiles of savings erosion are 0 to 0.2%. If the forecast lead time is restricted to four hours (red box), the savings erosion is almost always less than 0.3% except for in extreme outlier cases, and the median is less than 0.1%. The scatter plots also generally convey this trend. In addition to being more tightly grouped, data for the 4 hour forecast lead time have lower MAEs (appear closer to the y-axis) and savings losses lower than that for the 8 hour forecast lead time.

Figure 34 shows how the MAE changes as a function of forecast lead time for a given example day. Generally, as the lead time increases, so does the MAE. For this day, when the forecast lead time is less than 6 hours, the MAE of tailwind is less than 3 kts and the DISA MAE is less than 0.7 deg C. Combining this data with the exponential trends shown in Figure 33 reveals that the expected savings erosion is minimal. This may vary for different weather patterns. For instance, a convective stormy weather pattern may be less predictable than a calm day. However, it is assumed that airline operators utilize weather data that is constructed from data receied within the last hour. As such, a forecast 6 hours in advance represents a flight whose weather model is populated with a realistic operational worst case scenario.

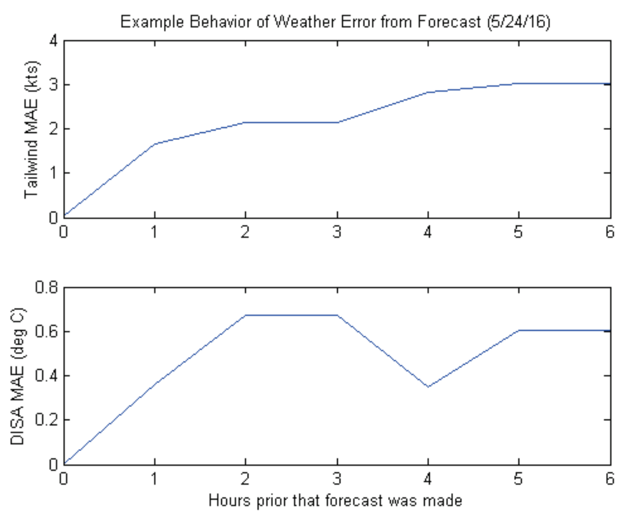

Figure 34. MAE and forecast lead times.

Impact of forecast lead time on observed weather mean absolute error.

Figure 35 depicts an example of the climb profile for one flight predicted to the optimum cruise altitude as calculated by three methods: Legacy, UCCD and Gen B. The grey line represents the flight profile as predicted by Legacy. The red line is the profile generated by UCCD technologies without consideration of weather data. The large difference in climb profile results largely from use of variable thrust command. Also noticeable is the slightly higher cruise altitude chosen when the cost of climbing and descending are considered in the cruise altitude selection. The green line shown in Figure 35 is the profile generated by Gen B technology that accounts for weather. The climb profile harnesses the available tailwinds that extends the distance in climb, but it also chooses a higher cruise altitude.

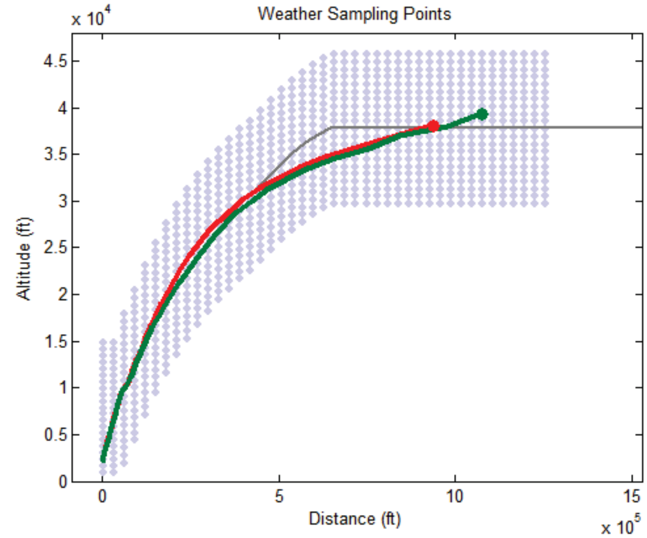

Figure 35. Climb profile from three methods.

Impact of accounting for weather on climb profile and cruise altitude.

6.4 TRL 5 Benefit Assessment

A TRL 5 Monte Carlo test set is run using the CLEEN II ATT to assess the benefit of the Gen B technology. This test set consists of many cases with varying aircraft type, weight, weather, and route. See paragraph B.2.1 for a detailed description of the case setup.

6.4.1 Fuel Savings (Cost Index 0)

For cases where Cost Index is set to zero, the total direct operating cost (DOC) is weighted to be entirely fuel cost (weighting of time cost is zero).

Outliers and unflyable path cases are omitted from this aggregate analysis. When removed, cumulative fuel savings of 0.53% and 0.43% are observed over the Legacy and Legacy+ methods, respectively. This metric represents the total fuel savings from all flights as a percentage of the total fuel burn for all flights. The savings over Legacy 99% confidence interval is 0.49% - 0.57%. In other words, there is a 99% chance that the true mean fuel savings lies in the interval 0.49-0.57%. The 99% confidence interval for the Legacy+ FMS is 0.38% - 0.48% fuel savings. Histograms for savings are shown in Figure 36.

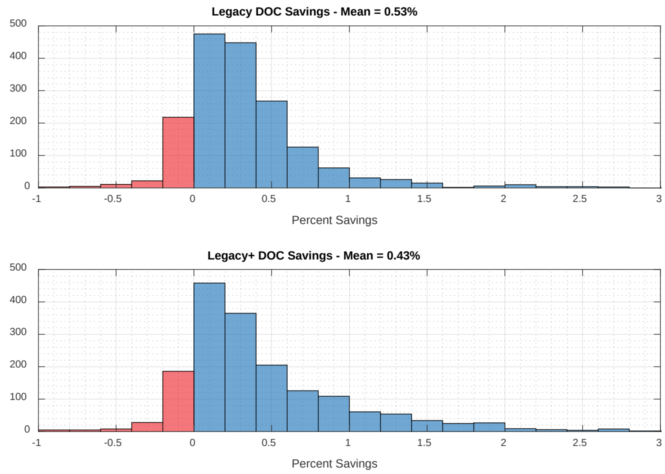

Figure 36. Histogram of DOC savings for Cost Index 0.

In practice, an operator would compare the output of the CLEEN algorithm with their Legacy approach and choose the control that produced the lowest cost. In other words, the cases where the CLEEN algorithm fails to save fuel over the Legacy methods would not be flown by the aircraft operator, and instead, the Legacy profile would be used, obviously resulting in a 0% savings over the Legacy method. Thus, the true fleetwide benefit an operator would observe is calculated by replacing the negative savings trial cases from the TRL 5 assessment with zeros. The resulting histograms and mean savings for the operator fleetwide benefit is presented in Figure 37. From this modified dataset, the cumulative fuel savings of CLEEN Generation D technology for Cost Index 0 are **0.54%** and **0.44%** over Legacy and Legacy+, respectively.

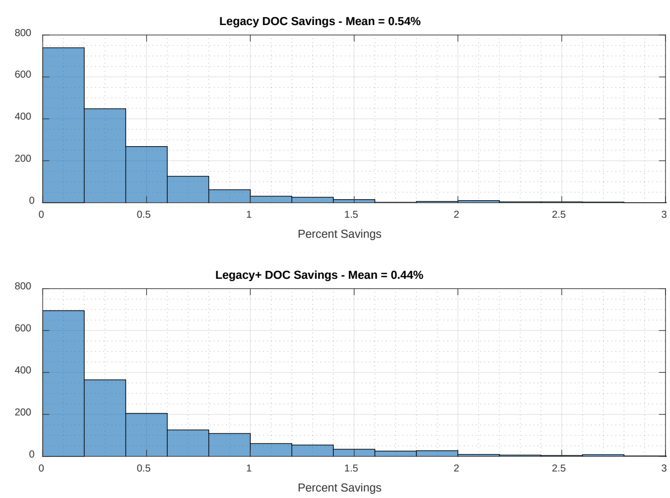

Figure 37. Gen B operational savings for Cost Index 0.

Cases with negative predicted savings would use Legacy control, considered as zero savings

6.4.2 Cost Savings (Cost Index 25)

With outliers and unflyable paths removed, cumulative Direct Operating Cost savings of **0.44%** and **0.36%** are observed over Legacy and Legacy+, respectively. This metric represents the total direct operating cost savings from all flights as a percentage of the total direct operating cost. The savings over Legacy 99% confidence interval is 0.41% - 0.48%. The 99% confidence interval for the Legacy+ FMS is 0.32% - 0.39% direct operating cost savings. Histograms for the data set with outliers removed are shown in Figure 38.

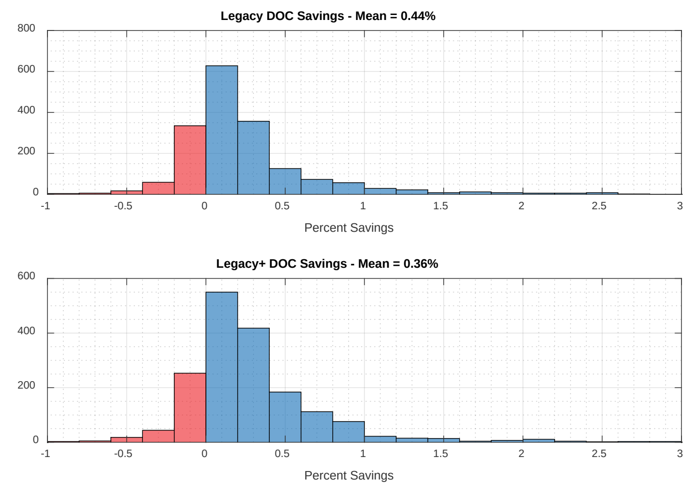

Figure 38. Histogram of DOC savings for Cost Index 25.

The true fleetwide benefit an operator would observe is calculated by replacing the negative savings trial cases from the TRL 5 assessment with zeros. The resulting histograms and mean savings for the operator fleetwide benefit is presented in Figure 39. From this modified dataset, the cumulative cost savings of CLEEN Generation D technology for Cost Index 25 are **0.46%** and **0.37%** over Legacy and Legacy+ respectively.

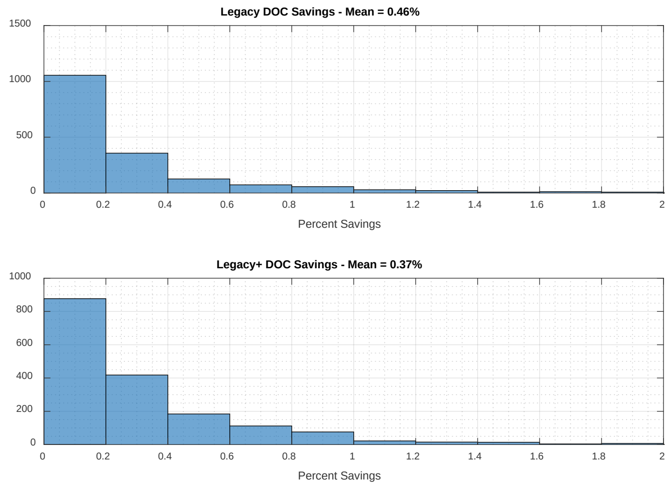

Figure 39. Gen B operational savings for Cost Index 0.

Cases with negative predicted savings would use Legacy control, considered as zero savings.

6.4.3 Example: Detailed Analysis

In this section, individual cases that both saved and lost fuel are examined to determine the root source of benefit (or loss).

Figure 40 and Figure 41 depict altitude profiles of flights and the resulting cumulative fuel savings as a function of distance. The savings buildups represent a case where CLEEN saves fuel and a case where CLEEN loses fuel respectively. These specific cases were selected to illustrate how savings are developed or lost throughout a flight.

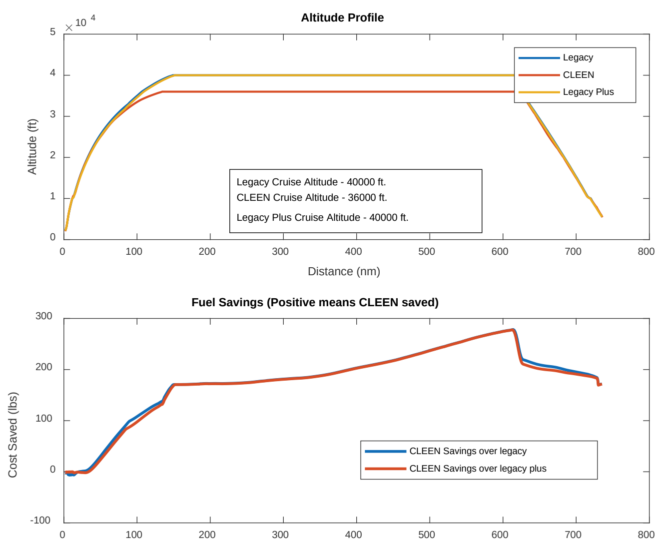

Figure 40. Example of flight that saves fuel.

To understand how savings are generated, it is important to look at each phase of flight individually as well as the combined effect of all flight phases. In the case that saved fuel (Figure 40), it is shown that CLEEN saves fuel throughout the climb phase. The reason for this is twofold. First, a variable speed and thrust climb allow for more efficient climbing flight compared to a max thrust constant speed climb. Second, in this case, CLEEN selects a lower cruising altitude that has the benefit of a shortened high-thrust climb phase and a more optimal cruising altitude. For each case, the CLEEN technology assesses many different altitudes and determines the optimal altitude based on many factors including weather, performance, and the tradeoff of the altitudes with other flight phases (i.e., climb and descent). When combining climb, cruise, and descent, an overall most optimal flight profile is generated that results in the lowest cost flight possible given the implemented constraints.

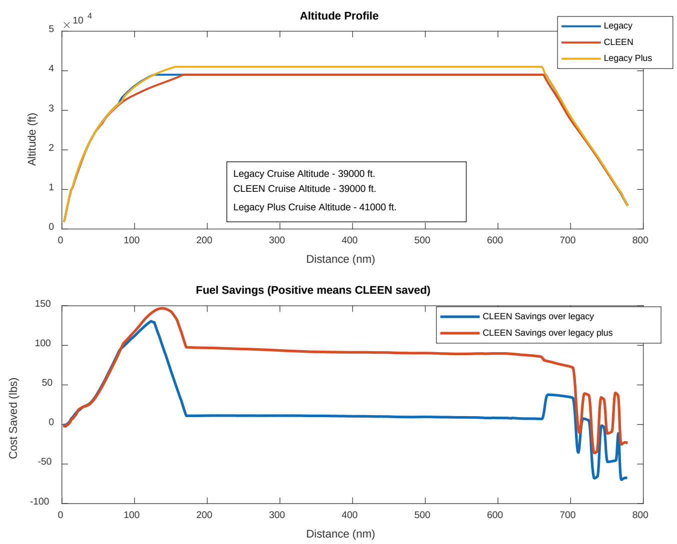

Figure 41. Example of flight that loses fuel.

Even though CLEEN should always find the optimal path, there are still cases that lose fuel (shown in Figure 41) due to unaccounted aircraft effects. The reasons for this include difference between observed and predicted weather and neglecting higher order terms in the simulation or optimizer algorithms. Investigation reveals that in most cases that lost fuel, descent is the primary culprit. As seen in Figure 41, in both Legacy and Legacy+ comparisons, CLEEN saved fuel through the point that all flights were descending. Once in descent, CLEEN had reversion logic activated (note the spikes in fuel saved) that lost fuel until it no longer provided any cumulative savings. This is generally caused by neglection of the higher order wind effects in the simulator physics. Given a proper descent path without these thrust reversion spikes, it is expected that CLEEN would have saved fuel in this example flight.

6.4.4 Sensitivity Study

While parsing through the benefit assessment data, various trends can be observed in the data based off aircraft type, route, weight, weather, and altitude quantization; this section summarizes those trends.

By combining the savings for route length and aircraft type groups, some trends can be seen in savings based on flight distance. This data is show in Figure 42. The greatest percent cost savings for CLEEN flights compared to Legacy come from the shortest and longer flights; the most benefit over Legacy+ is gained during shorter flights. Benefit in short flights is generated mostly from UCCD and variable speed and thrust climb. This effect is discussed in more detail in paragraph 6.4.3.

By separating the flights by relative route length and then sorting those routes by aircraft, the effect CLEEN has on each flight length with respect to aircraft type can be observed. For short flights, narrow and wide body aircraft benefit the most from CLEEN technology. This can be attributed to the lower cruise altitudes selected for shorter cruise distances – a direct result of the unification of climb, cruise and descent. Cost savings increases with the altitude delta between CLEEN and Legacy flights. Short flights have a large altitude delta and therefore greater savings. This trend is also due to climb efficiencies. For shorter routes, climb is larger portion of the flight, causing the fuel saved in climb to have a larger impact on the overall percent fuel savings.

Another apparent trend is clear for long flights, where the savings increase with the increase in aircraft size. Wide body aircraft save the most while regional aircraft save the least. The length distribution of wide body flights causes this trend, which includes routes up to 7000 nmi in length. The long distance allows for significant savings to accumulate throughout cruise when different altitudes are picked as compared to those of the Legacy method.

By creating a correlation matrix between percent savings, absolute value of altitude difference between CLEEN and Legacy profiles, and route length, relationships between the 3 can be observed; see Figure 43. The diagonal of the plots are histograms of each individual comparison element and the rest of the plots are correlations of the x and y specified elements. For example, the top right plot contains a comparison of route length versus savings.

By analyzing the plots, a few trends can be observed. First, short and long routes have the largest altitude discrepancies between Legacy and CLEEN flights. Additionally, short and long flights tend to save more fuel than the mid-range flights. For short flights, this is achieved mainly due to the UCCD technology; CLEEN assesses the tradeoff between cruising and climbing and decides to cruise at a lower altitude. This altitude is suboptimal for just the cruise phase but is less so than the cost of climbing to that cruise altitude. For long flights, the larger fuel savings are mainly achieved by selecting a significantly more optimal cruising altitude. Finally, matching the other trends discussed, larger altitude differences typically result in higher cost savings. Note that there are cases where these relationships do not hold but the trend exists for most flights.

Analysis of the savings as a function of passenger loading (PAX factor), weather day, and minimum vertical separation limit reveals no correlation.

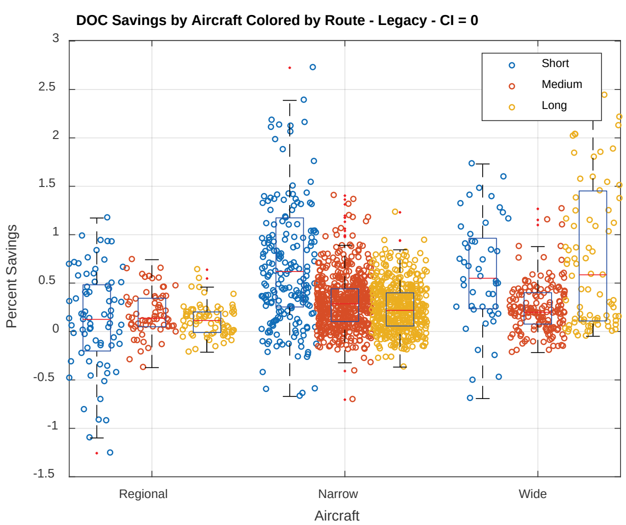

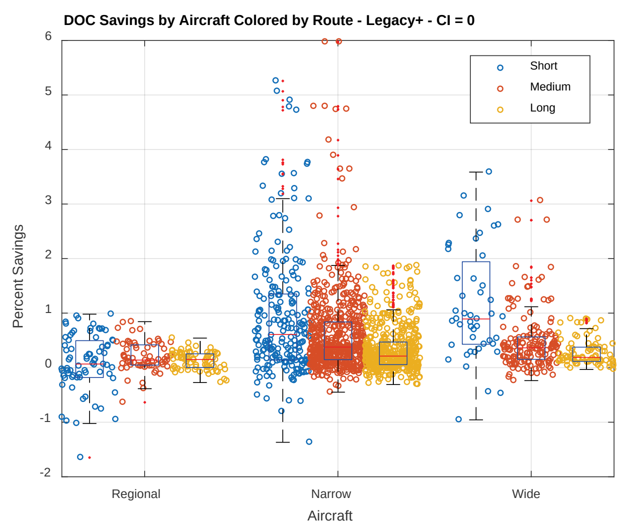

Figure 42. DOC percent savings separated by aircraft type, colored by route.

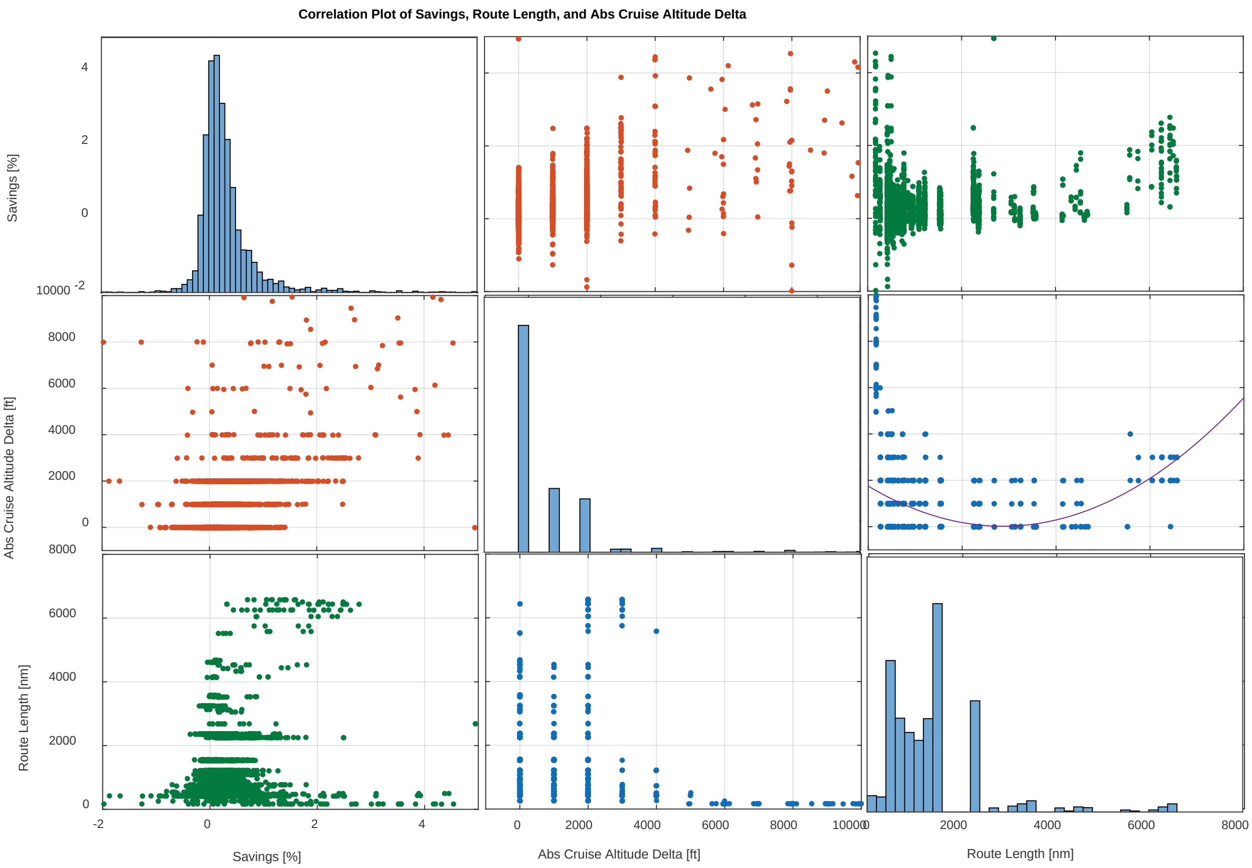

Figure 43. Correlation matrix.
Correlation plot of savings, route length, and absolute value of cruise altitude difference between Legacy and CLEEN.

7. Generation C: Step Climbs and Descents

7.1 Summary

Current methods used by Flight Management Systems to compute optimal cruise step locations are deficient in at least one of several ways: they do not exist in the native FMS functions and require the pilot to manually enter and execute them; or if they do exist they are overly simplistic in their calculation and produce only one optimal step that must be a climb or they do not account for known weather data or use a low-fidelity representation of the weather. The CLEEN Generation C technology provides a method to compute more than one optimal step, including step climbs or descents of varying magnitude, all optimized with respect to high-fidelity weather data; see Figure 44. Multiple steps enable tracking of both aircraft performance changes with weight and favorable weather patterns, such as tailwinds at different cruise altitudes. This Generation C technology is built on the Generation B technology that incorporates high-fidelity data of both wind and temperature into the CLEEN Unified Climb, Cruise, and Descent (UCCD) framework.

In its final form, the Generation C technology performs the following steps:

1. Collect weather data along a 2D grid (altitude and distance) that follows an approximation of the entered flight plan's lateral path
2. Conduct a brute force cruise altitude search to determine the optimum control commands that produce a profile with the lowest flight cost for any origin and destination combination using the following process:
    (a) Calculate a multi-variable climb path optimization to the initial cruise altitude, and generate a Legacy descent path estimation so that cruise distance is known
    (b) Perform a graph search for optimal cruise step climbs and descents throughout the cruise phase, including an integration of the final cruise trajectory (this is the heart of Generation C technology)
    (c) Combine the three phases to determine the profile cost
3. Predict the flight plan using the optimal control

TRL 5 benefit assessment testing (described in paragraph 7.4) demonstrates that with the Generation C technology, an operator can achieve **1.02%** fuel savings as compared to Legacy and **0.40%** as compared Legacy++ (see paragraph 3.3 for a description of these baseline methods). This result is shown in Figure 45. This technology is implemented in a real-time TRL 6 prototype described in later sections of this document.

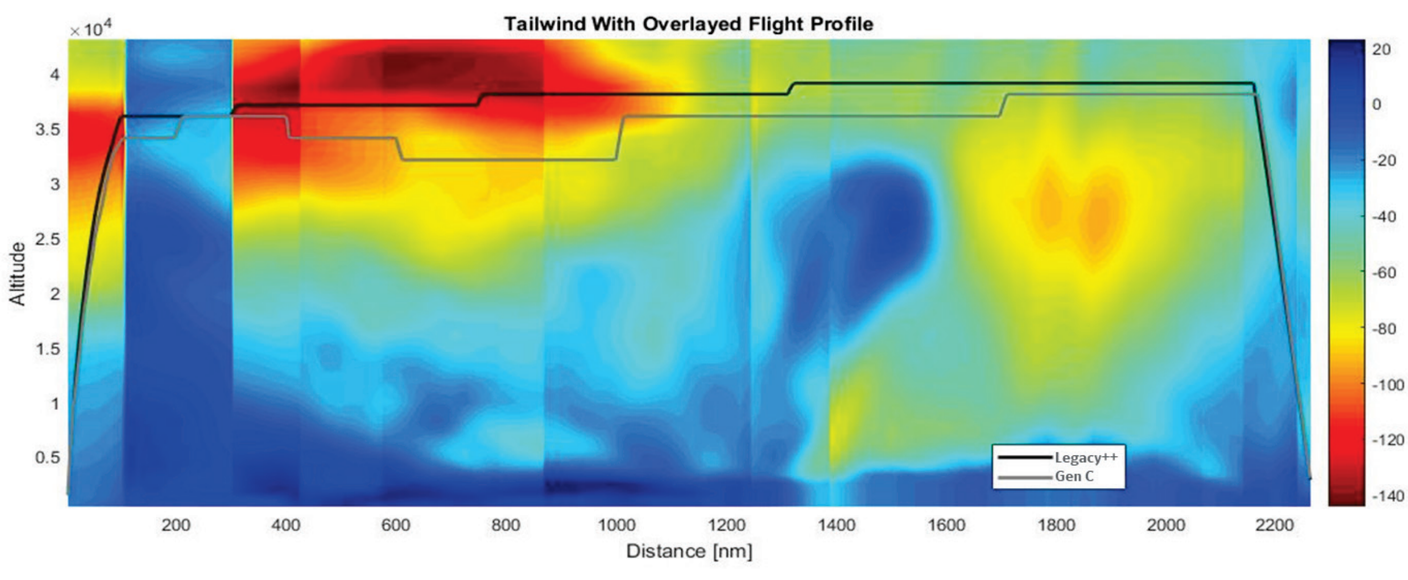

Figure 44. Generation C profile.

This features multiple optimal climbs and descents among weather patterns to produce cost savings. In this example, the large headwind during the early cruise phase is avoided by stepping down to a lower cruise altitude.

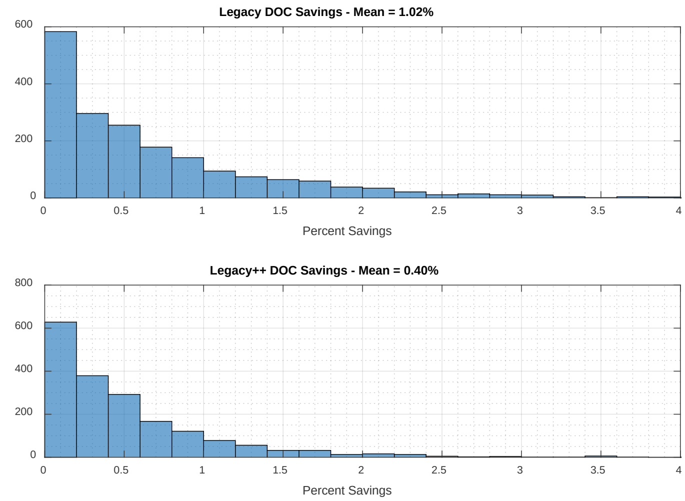

Figure 45. TRL 5 benefit.

Monte Carlo study in TRL 5 simulation testbed shows significant savings from Generation C technology

7.2 Development

7.2.1 Variable Cruise Altitude and Speed

The Gen C Cruise Step problem is solved with Dijkstra's algorithm, a shortest path-finding algorithm similar to what Google Maps[††] uses to determine an optimal route between two locations. The traditional application of the algorithm is to consider distance as the cost function, thus indicating that the algorithm should find the shortest path between vertices. However, any measurement of cost can be used, including the Direct Operating Cost (DOC) of flying each edge.

Dijkstra's algorithm is exhaustive, which means all search paths are expanded and it is guaranteed to find the lowest-cost path, given enough execution time. Given V vertices and E edges in a graph, the total run time of the algorithm is at worst $O(V \log V + E \log V)$. This is true for the specific variant of Dijkstra's Algorithm used in UCCD Gen C; namely, that where a min-priority queue is used to store the vertices.

In order to set up Dijkstra's Algorithm to solve the Gen C Cruise Step problem, it is necessary to define a graph of vertices and edges that represents the problem. To begin, consider the simple distance-altitude grid that represents the vertical and longitudinal flight path of the aircraft, like that of Figure 46.

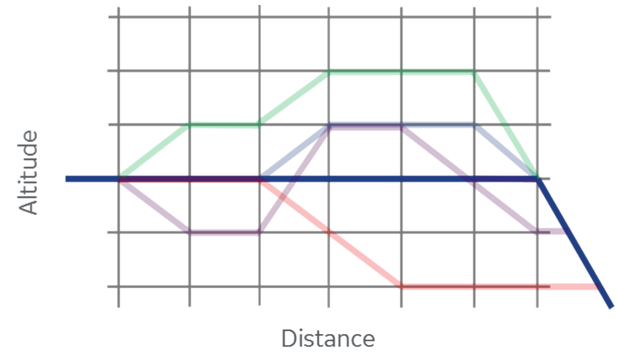

Figure 46. Simple representation of vertical flight path grid.

This grid depicts distance on the x-axis and altitude on the y-axis. Shown overtop the grid are potential cruise paths through the vertical space, most of which show step climbs and/or descents at certain distances. For our application, the grid intersection points (distance-altitude pairs) are considered the graph vertices. The edges are the set of segments that connect a given vertex (the parent node) to the set of connected vertices (the child nodes) that all are one "distance step" more than the parent and have a variety of different altitudes. The initial node is the Top-Of-Climb (TOC) point, and the final node is the Top-Of-Descent (TOD) point. A path between the initial and final node is formed by connecting several nodes together with cruise segments that monotonically increase in distance.

For instance, for a parent node at distance X and altitude A, and assuming the distance step is parameter dx and the altitude separation is da, the set of child nodes is:

$$\{ (X + dx, A), (X + dx, A + da), (X + dx, A + 2 * da), \dots, (X + dx, A + n * da),$$

$$(X + dx, A - dA), (X + dx, A - 2 * dA) \dots (X + dx, A - n * dA) \}$$

Where n controls the number of steps above and below the current altitude that are considered in the search. Thus, in this simple example, if n is 3, then each node has 7 child nodes: 3 above, 3 below, and the original altitude indicating no step climb or descent is required, all at distance $X + dx$.

This example has only one control variable: altitude. However, this idea can be extrapolated to an extra dimension so airspeed can also be a control variable. This can be done by imagining that a third axis comes out of the page in Figure 46, on which Mach is selected. Then, the graph vertices are distance-altitude-speed tuples, where altitude and Mach are the control

[††] Google Maps is a trademark of Google

variables and thus each child node is a variation in both altitude and speed at the next distance step. This expands the number of child nodes per parent node multiplicatively: if 13 speed steps are analyzed (6 above and below the current speed, and one at the current speed), then each node has $13 * 7 = 91$ child nodes. These parameters can be set to control the size and resolution of the search problem. The Gen C optimizer technology uses this two-control variable approach that is referred to as the 2D Dijkstra when distinguishing from the 1D (altitude only) counterpart (described above).

Throughout the rest of the document, this 2D application of Dijkstra's algorithm to the cruise step problem is generally called the Dijkstra cruise step optimizer or Dijkstra optimizer.

7.2.2 Cost Function Calculation

The graph definition above provides the lists of edges and vertices for the cruise optimization problem, leaving the "cost" of each vertex-connecting edge as the only parameter requiring calculation. This cost is that of the cruise segment cost, defined by the cost to fly the aircraft from one distance-altitude-speed tuple (parent node) to a sequential distance-altitude-speed tuple (child node). If the child node has the same speed and altitude as the parent node, then the cost is simple – it is a cruise segment of constant speed, and the cost is easily calculated using aircraft model functions to determine the fuel flow required to produce thrust that counteracts the aircraft's drag.

If the child node demands an altitude change, then a step climb or step descent must be calculated. Step climbs are generally computed at maximum climb thrust and step descents at idle thrust. For a cruise segment that specifies an altitude change, the climb or descent is performed starting at the parent node's location and at the appropriate thrust level. The distance of the climb or descent is calculated using the equations of motion, and then the remaining distance of the segment is performed in cruise mode with required thrust. This is diagrammed notionally in Figure 47, where node A is the parent node and nodes B, C, and D are child nodes. Speed changes at each node are assumed to be instantaneous and occur at the end of the segment; thus, the speed for the entire segment is the Mach specified at node A.

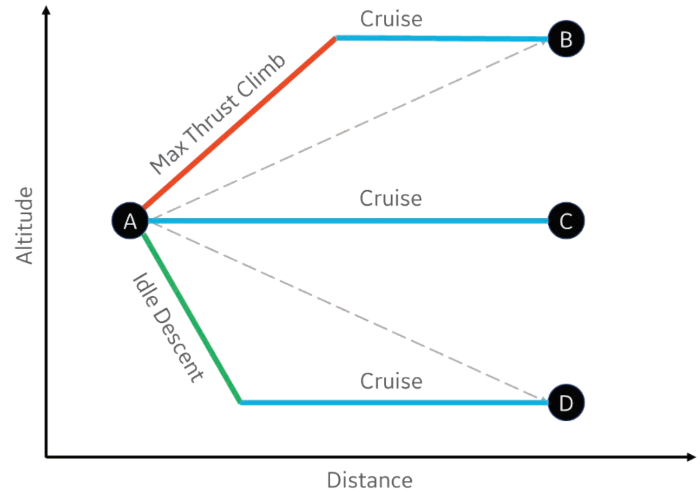

Figure 47. Climb and descent calculation.
Cruise segment for a step climb (A to B), no change (A to C), and step descent (A to D).

As with the rest of the UCCD features, the cost of any given flight segment is the combination of the cost of fuel and the cost of time, which are related via the Cost Index.

Several design considerations were performed during the development phase of the Gen C optimizer. This included determining how many iterations of the graph search algorithm were necessary to generate an optimal path and determining the grid resolution used in each graph search.

7.2.3 Optimization Algorithm

The implementation of Gen C: Step Climbs and Descents builds on the previous implementation of Gen B software, described in paragraph 6.2.3.

In the few examples below, the prototype Dijkstra cruise step optimizer is used to optimize the flight path through different weather patterns. The first example in Figure 48 shows a "contrived weather" scenario, where the weather data is faked to produce a checkerboard pattern of tailwind pockets. The goal of this case is to observe that the Dijkstra optimizer follows the obvious best route and avoids heavy headwinds in favor of heavy tailwinds. As expected, the optimal path from TOC to TOD in this case is to perform step climbs and descents when appropriate to keep the aircraft in the tailwind regions.

The second example case is shown in Figure 49, and contains real weather data, including tailwind and DISA. Again, the optimal path includes step descents to track a favorable tailwind between 1000 nmi and about 1400 nmi, and then returns to a higher altitude when the second half of the flight encounters heavy headwinds throughout the altitude profile (due to a turn in the flight path).

7.3 Algorithm Validation Testing

7.3.1 Test Methods

Both the legacy systems (Legacy and Legacy++) and Gen C algorithms are implemented in a TRL 4 MATLAB test suite using simplified point-mass integration of the trajectory. Once a planned control trajectory is calculated, separate point-mass simulation is performed to compute the cost of each profile. See paragraph B.1 for a detailed description of the TRL 4 testing method.

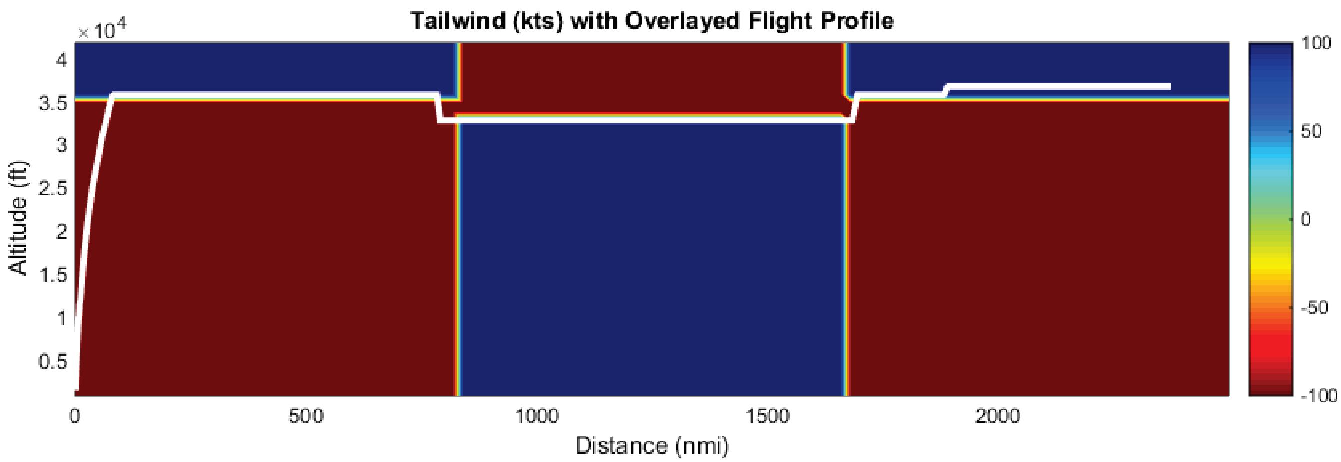

Figure 48. Contrived weather cruise profile example.
The cruise profile (white line) is shown with contrived weather, where blue is beneficial weather pattern.

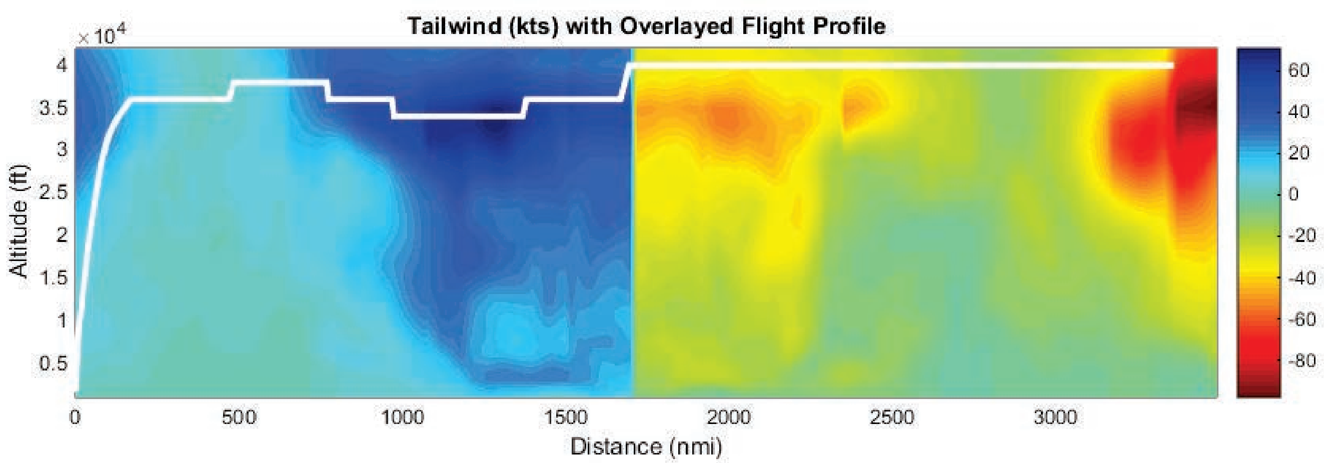

Figure 49. Real weather cruise profile example.
The cruise profile (white line) is shown with real weather, where blue is beneficial weather pattern.

The four test sets shown in Figure 50, totaling 6540 flights, were developed to evaluate the Gen C technology. In all, these flights drove 2180 comparisons between the three optimization methods.

#	Description	Cases
1	Standard Atmosphere	420
2	Vary Algorithm Resolution	1152
3	Contrived Weather Pockets	1512
4	Real Weather	3456

Figure 50. TRL4 test set flight breakdown.

Within each test set is a nominal variation of more basic flight parameters that is shown in Figure 51. Note that each test set has a subset of this variation according to the purpose of the test set, and not the complete variation. For instance, test set 2 only contains flights of 4000 nmi in order to determine the effect of distance step on cost and computation time.

Parameter	Nominal variation in each test set
Takeoff Weight	110 – 170 klbs
Flight Distance	500 – 4000 nmi
Altitude Quantization	1000, 2000 ft
Cost Index	0, 25, 50, 100

Figure 51. Nominal variation of basic flight plan parameters in each test set.

Each test set is designed to either exercise a certain aspect of the optimizer, or determine the sensitivity of the cost savings to various external effects:

1. Test set 1 is the basic test set that is designed to show the general concept of step climbs. Large savings over Legacy are expected; but performance may be close to Legacy++ that makes step climbs similar to UCCD Gen C when weather is not a factor.
2. Test set 2 is designed to analyze the algorithm sensitivity to distance step length and minimum cruise distance parameters.
3. Test set 3 is a contrived weather scenario designed to stress the optimizer. Each case has fake weather data with “pockets” of tailwinds and headwinds in a “checkerboard” pattern. The Gen C optimizer is expected to track the tailwind patterns and produce a profile that utilizes them as best as possible. An example weather pattern from this set is shown in Figure 48.
4. Test set 4 is the benchmark test set used to report final savings numbers. It includes the largest number of cases with a variety of six different real weather days and is expected to be representative of the benefit achieved in service.

7.3.2 Test Results

7.3.2.1 Overview

The benefit achieved by the CLEEN technologies for each test set is presented relative to the Legacy and Legacy++ implementations. For each test set, cases where the aircraft total weight dropped below 95 klbs were filtered out – the reasoning being that the operating empty weight of a standard single-aisle transport is about 91.3 klbs, and some fuel must be left for reserves. Figure 52 summarizes the results.

As expected, the savings achieved compared to the Legacy method are greater than that compared to the Legacy++. Test set 4 in the chart above is important as the primary numbers reported for the benefit of this technology in TRL 4. Test set 3 savings are omitted from the chart as that test set is used only for verification that the Gen C profiles follow the expected path.

Test set 1 shows that, with a lack of weather variation, the Legacy++ method performs remarkably close to the Gen C method. This indicates that the Legacy++ method of calculating step climb locations based solely on the decreasing weight of the aircraft is sufficient in the presence of standard atmosphere. Test set 2 shows large savings over Legacy because this set contained only long flights of 4000 nmi, but the data from this test set will be used later to perform more algorithm resolutions analyses. Finally, test set 4 shows expected savings in real-weather scenarios, assuming perfect forecast.

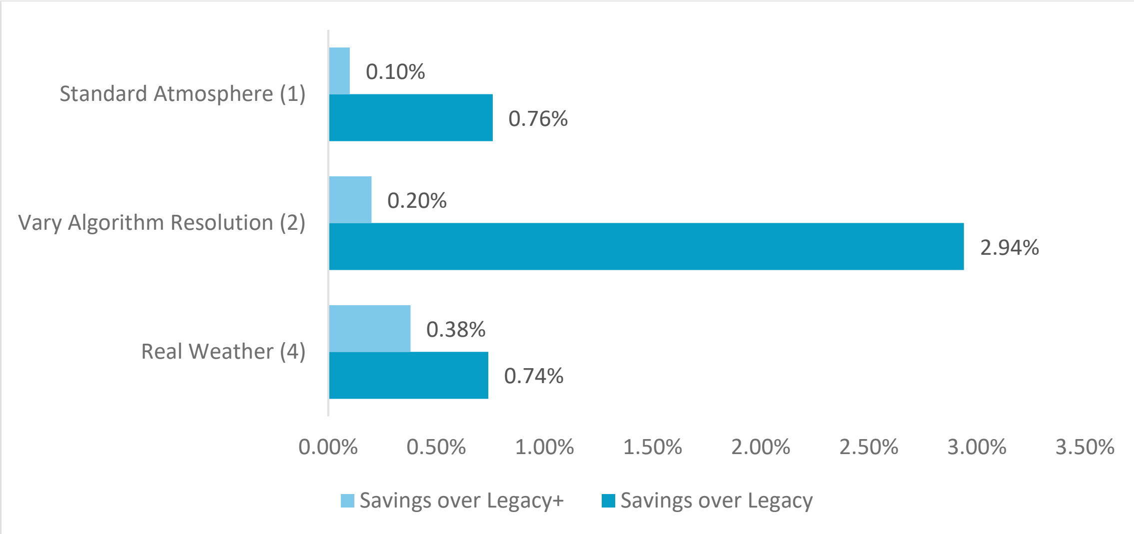

Figure 52. Gen C Fuel Savings for Test Sets 1 - 4 (not including 3).

7.3.2.2 Test Set 4 (Real Weather) Results

In this section, the results from test set 4 are broken out in more detail. Figure 53 and Figure 54 show histograms of the percent savings over each stated method, with the histograms separated by flight distance. The translucent bars in the graph show the histogram data itself overlapping each other, while the bold lines represent a notional fit of a gamma distribution to the data. This helps to reveal more clearly the trends present in the data. The dashed lines indicate the mean of each histogram, which is useful for observing how the savings correspond to flight distance in each subset.

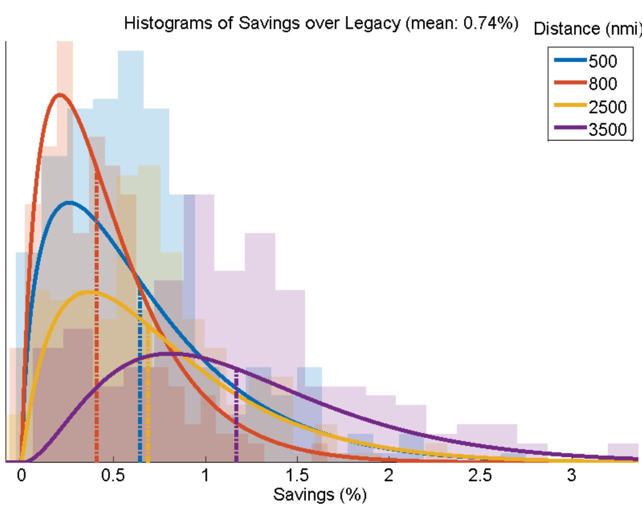

Figure 53. Savings over Legacy.

Test set 4 savings data over Legacy and correlation to flight distance.

It is first worth noting that, when reporting cost savings numbers in a percent, the flight distance plays an important role. The same absolute savings corresponds to different percent values for different length flights. For instance, for a short flight and a long flight, savings of 100 lbs. corresponds to a larger and smaller percent, respectively. This effect is present in all the following results.

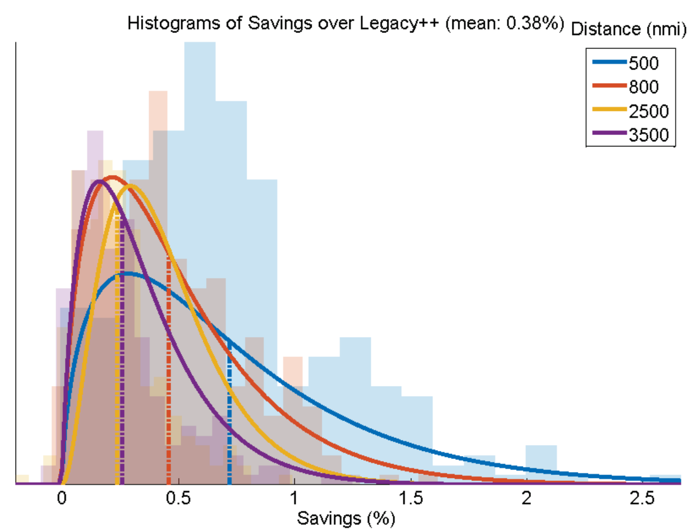

Figure 54. Savings over Legacy++.

Test set 4 savings data over Legacy++ and correlation to flight distance.

As seen in Figure 53, the average savings over Legacy generally increase significantly with increasing flight distance. This makes sense, as with longer flights the cruise phase becomes the dominant portion of the longer flight and the Legacy method does not utilize step climbs to adjust the cruise altitude. Accordingly, the Legacy method has a very suboptimal altitude at the end of long cruise phases, when a good portion of the aircraft fuel weight has been burned off. The one exception to this generality is that the 500 nmi routes show more savings than the 800 nmi. This is due to the UCCD optimizer choosing an altitude *lower* that the Legacy altitude (an altitude that is technically more suboptimal) because the UCCD framework determined that it was not cost effective to burn the climb fuel to get to a higher altitude where little time was to be spent cruising. Figure 55 shows that this not solely the effect of flight distance inflating percent values as mentioned in the above paragraph – indeed, the shorter 500 nmi flights save more absolute cost than the 800 nmi flights. This effect has been known since Generation A and is document in Section 4.

For the savings over Legacy++ (shown in Figure 54), the trend is opposite: longer flights generally result in the Gen C optimizer saving less percent over the Legacy++ method. There is an opposite effect at work here from the savings over Legacy: for shorter flights, the climb phase dominates, and the cruise phase of Gen C is similar in many cases to the cruise phase of Legacy++ as both perform step climbs. This effect prevents the benefit from accumulating throughout cruise. Instead, the variable-thrust variable-speed climb phase provided by the UCCD framework generates the majority of the savings for the flight. Thus, longer flights result in a smaller percent saving number.

Figure 55 provides the same mean cost savings data as shown in Figure 53 and Figure 54 but in absolute terms (converted to pounds) instead of percentages. The absolute cost savings over Legacy is most extreme at the longer flight ranges, where the Gen C technology takes advantage of step climbs and descents. On the other hand, for the savings over Legacy++, the 90-100 lbs. of savings on the shorter flights is a larger percent than the 140-200 lbs. of savings on the longer flights, resulting in the percentages seen in Figure 54.

Distance (nmi)	500	800	2500	3500
Mean savings (lbs.) over Legacy	79.1	74.1	308.1	701.8
Mean savings (lbs.) over Legacy++	94.2	92.7	143.6	214.2

Figure 55. Test set 4 mean savings in pounds instead of percent.

Figure 55 shows one unexpected trend. For shorter flight distances, the absolute savings over Legacy++ are greater than the savings over Legacy, indicating that Legacy outperforms Legacy++ on average for those distances. Investigation shows that this trend is due to two inaccuracies in the Legacy++ algorithm. The first is that, in some cases, the Legacy++ algorithm chooses a different initial cruise altitude than Legacy to harness potential benefits from wind. However, in a subset of those cases, the different altitude is suboptimal, due to either the short duration of cruise or poor modeling of the weather – the true weather may be different in between the data points that were entered into the Legacy++ algorithm. The second inaccuracy is that some cases show the Legacy++ commanding a step climb at a distance partway through the short flight, with the intention of accumulating fuel savings at the more optimal altitude. However, the Legacy++ step climb algorithm does not consider the remaining distance available to cruise at the new altitude. The resulting altitude switch is suboptimal because there is not enough distance over which to accumulate the savings, and the savings are offset by the cost of making the step climb. Both shortcomings are deficiencies in the current market method of computing step climbs.

7.3.2.3 Test Set 2 (Algorithm Resolution) Results

Test set 2 contains flights exclusively of 4000 nmi, with the intention of further understanding the effect of step length in Dijkstra's algorithm in terms of optimality and computation time. This set also provided for studying the optional minimum cruise distance parameter that specifies how long the aircraft is required to stay at a certain cruise altitude before being allowed to step. This parameter had minimal effect on the outcome, and thus is ommitted from the plot below.

Each boxplot in Figure 56 contains a variation in takeoff weight and weather day. The red bar on each boxplot indicates the median of the data, while the blue box covers the 25th to 75th percentile and the dashed whiskers extend to the furthest data point not considered an outlier. As can be seen from the top subplot, an increase in step length decreases the computation time; but the bottom subplot shows that there is a corresponding loss of optimality (measured as savings over Legacy++ here). It is notable that decreasing the step length to a value smaller than about 125 nmi does not significantly improve the optimality (savings improve from about 0.19% to 0.21%, but more than doubles the computation time. This suggests that an acceptable value for the distance step length parameter is between 100 and 125 nmi.

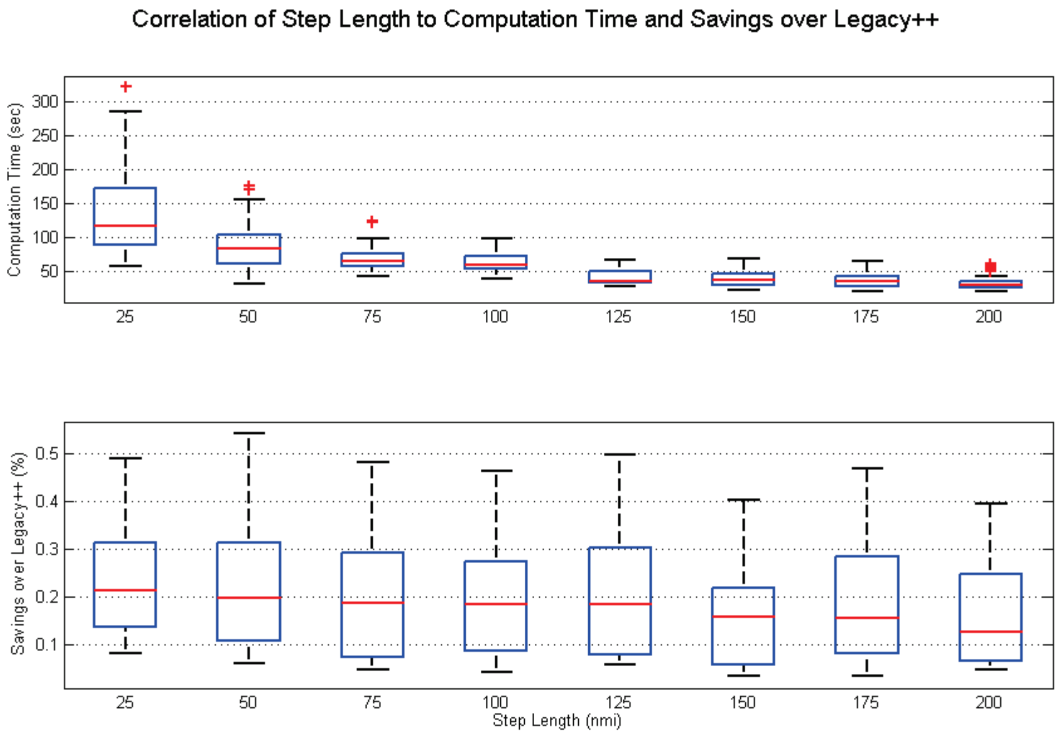

Figure 56. Test set 2 results.
Test set 2 boxplots show how distance step length affects computation time and optimality.

7.4 TRL 5 Benefit Assessment

7.4.1 Fuel Savings (Cost Index 0)

With outliers removed from the dataset, the cumulative fuel savings of CLEEN Generation C technology for Cost Index 0 are 1.00% and 0.39% over Legacy and Legacy++ respectively. This metric represents the total fuel savings from all flights as a percentage of the total fuel burn from all baseline flights. The 99% confidence interval for Legacy savings is 0.96% - 1.05%. This means that there is a 99% confidence that the true mean will fall within the range specified. The 99% confidence interval for Legacy++ is 0.34% - 0.44%. The histogram in Figure 57 is the distribution of percent savings for all Cost Index 0 cases.

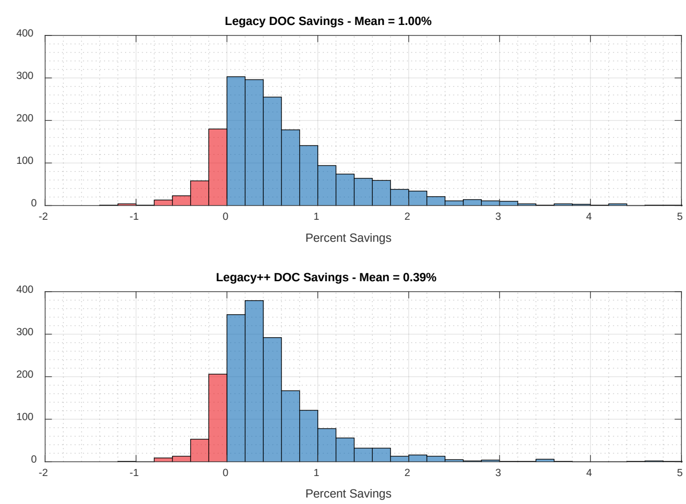

Figure 57. Gen C cost savings histogram for Cost Index 0 with outliers removed.

In practice, an operator would compare the output of the CLEEN algorithm with their Legacy approach and choose the control that produced the lowest cost. In other words, the cases where the CLEEN algorithm fails to save fuel over the Legacy methods would not be flown by the aircraft operator, and instead, the Legacy profile would be used, obviously resulting in a 0% savings over the Legacy method. Thus, the true fleetwide benefit an operator would observe is calculated by replacing the negative savings trial cases from the TRL 5 assessment with zeros. The resulting histograms and mean savings for the operator fleetwide benefit is presented in Figure 58. From this modified dataset, the cumulative fuel savings of CLEEN Generation D technology for Cost Index 0 are **1.02%** and **0.40%** over Legacy and Legacy++ respectively.

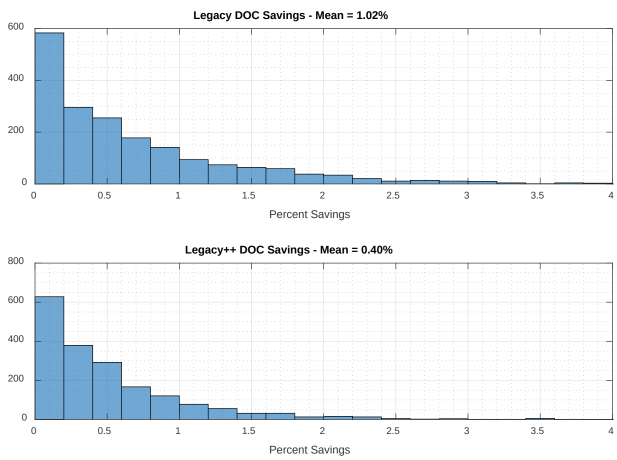

Figure 58. Gen C operational savings for Cost Index 0.

Cases with negative predicted savings would use Legacy control, considered as zero savings

7.4.2 Cost Savings (Cost Index 25)

With outliers removed, the cumulative cost savings for cases with a Cost Index of 25 are 0.89% and 0.37% over Legacy and Legacy++ respectively. This metric represents the cumulative cost savings as a percentage of the sum of all baseline costs. The 99% confidence interval for savings over Legacy is 0.85% - 0.94%. For Legacy++, the 99% confidence interval is 0.33% - 0.41%. The histogram in Figure 59 is the distribution of percent savings for all Cost Index 25 cases.

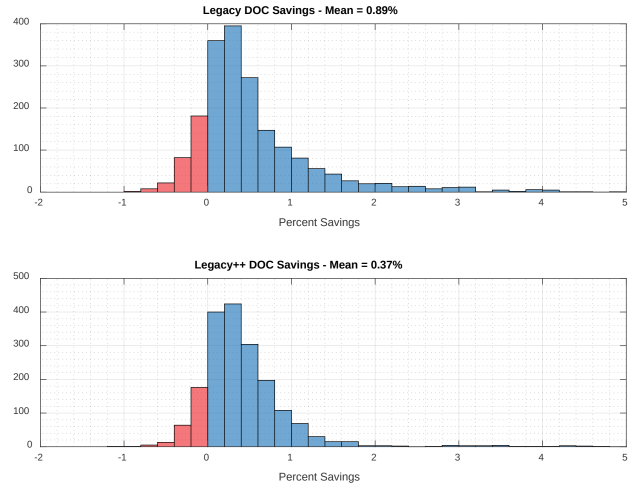

Figure 59. Gen C DOC savings for Cost Index 25 cases with outliers removed.

The true fleetwide benefit an operator would observe is calculated by replacing the negative savings trial cases from the TRL 5 assessment with zeros. The resulting histograms and mean savings for the operator fleetwide benefit is presented in Figure 60. From this modified dataset, the cumulative cost savings of CLEEN Generation D technology for Cost Index 25 are **0.91%** and **0.38%** over Legacy and Legacy++ respectively.

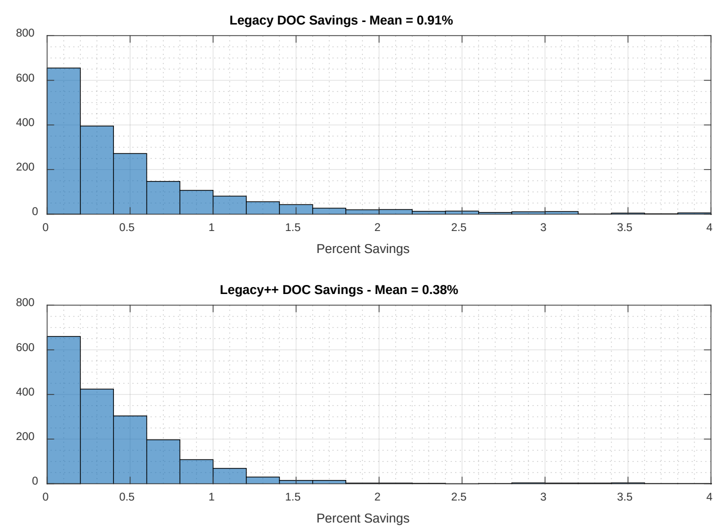

Figure 60. Gen C operational savings for Cost Index 25.

Cases with negative predicted savings would use Legacy control, considered as zero savings

7.4.3 Execution Time

For Gen C, long flights can cause the computational complexity of the graph search to exponentially increase due to the increased number of nodes in the search. However, the execution time of the algorithm is quick enough to satisfy operator desires. Figure 61 shows a histogram of execution time of the CLEEN algorithm for the full TRL 5 test set, indicating a maximum execution time of just under 15 seconds. Furthermore, the mean execution time is under five seconds, and the 95th percentile is just under 9 seconds. Thus, the optimization algorithm execution time will satisfy operator desires for real-time applications or for batch processing across a fleet of aircraft.

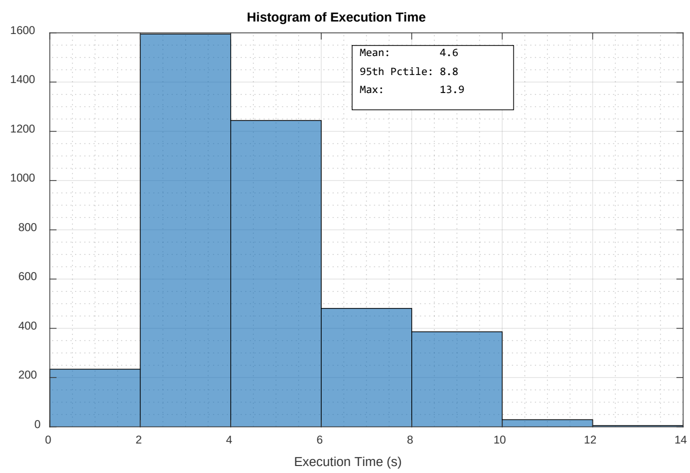

Figure 61. Execution time histogram of the Gen C algorithm in TRL 5 software.

7.4.4 Example: Detailed Analysis

In contrast to previous sections, where the aggregate results of all cases combined has been examined, this section details the savings of individual flights and examines the root causes of CLEEN fuel savings. Figure 62 displays a comparison of flight path with the cumulative fuel savings throughout the flight.

Gen C software is developed as an improvement to the previous generations that continue to exist within Gen C. The increasing slope of the CLEEN Savings over Legacy line confirms that each CLEEN step reaches an increasingly optimal altitude, with the sawtooth pattern emerging from the cost to climb to the new altitude.

It is less obvious from this case study that the altitude selection of CLEEN is better than the Legacy++ profile. Since both technologies climb at separate times, savings goes up and down based on which one has climbed most recently; a step climb initially burns more fuel but switches to a more optimal altitude. By correctly assessing the tradeoff between the cost of climbing and the benefit of a more optimal cruise altitude, an optimal profile is created. CLEEN's correct assessment is proven by the net fuel savings over Legacy++ at the end of the flight.

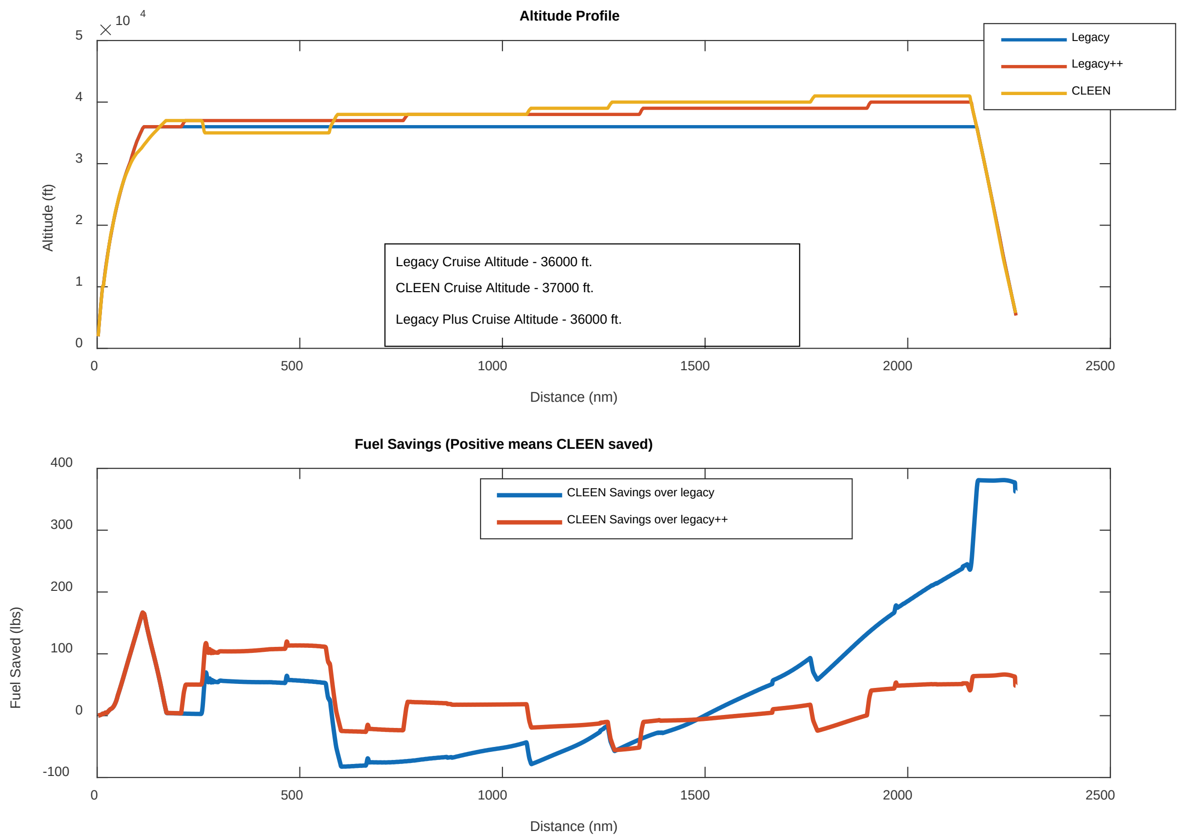

Figure 62. Detailed flight comparison.

7.4.5 Sensitivity Study

Akin to the previous discussion of correlation between CLEEN savings and aircraft type, there is a clear trend that large aircraft benefit significantly more than smaller aircraft from the addition of steps in the Gen C technology. This effect is more drastic when comparing against Legacy flights where the mean savings for B777 is more than double that of smaller aircraft. Figure 63 shows this effect. The savings of each test case is plotted as a point above the aircraft type, and boxplots summarize the data for each aircraft. The red marks for each aircraft represent the mean percent savings with the outliers removed. Finally, the mean values are annotated beneath each boxplot. The greater benefit for larger aircraft is due to the route lengths flown; larger aircraft fly longer routes allowing for more opportunity to benefit from steps. A mean of over 2% fuel savings – the largest mean observed – can be achieved for the B777 aircraft model over Legacy.

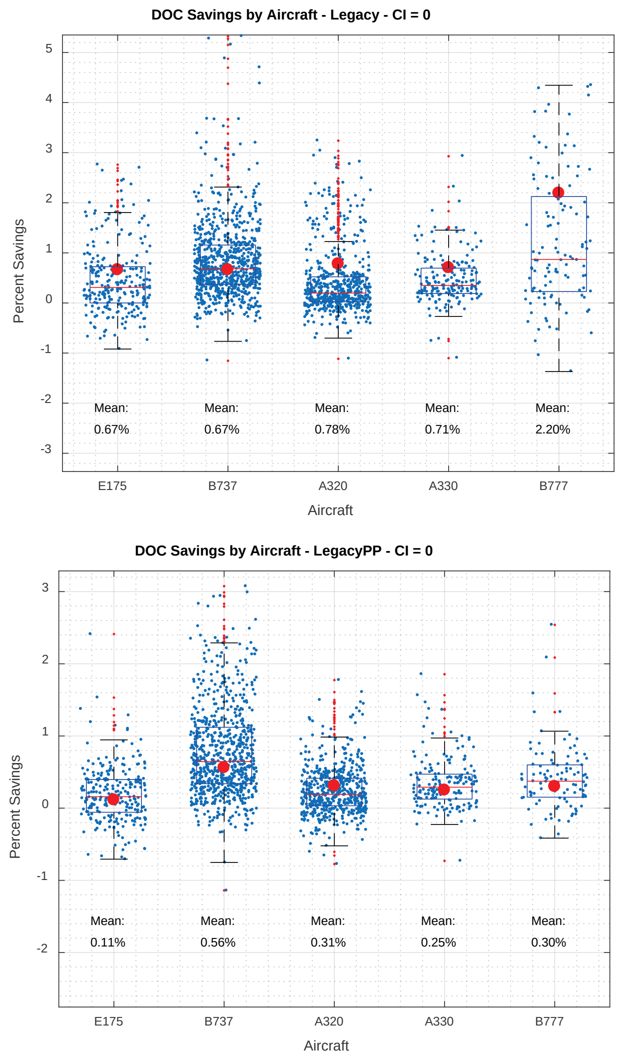

Figure 63. DOC Savings by aircraft.

Percent savings by aircraft type, with mean for each aircraft marked by a red dot.

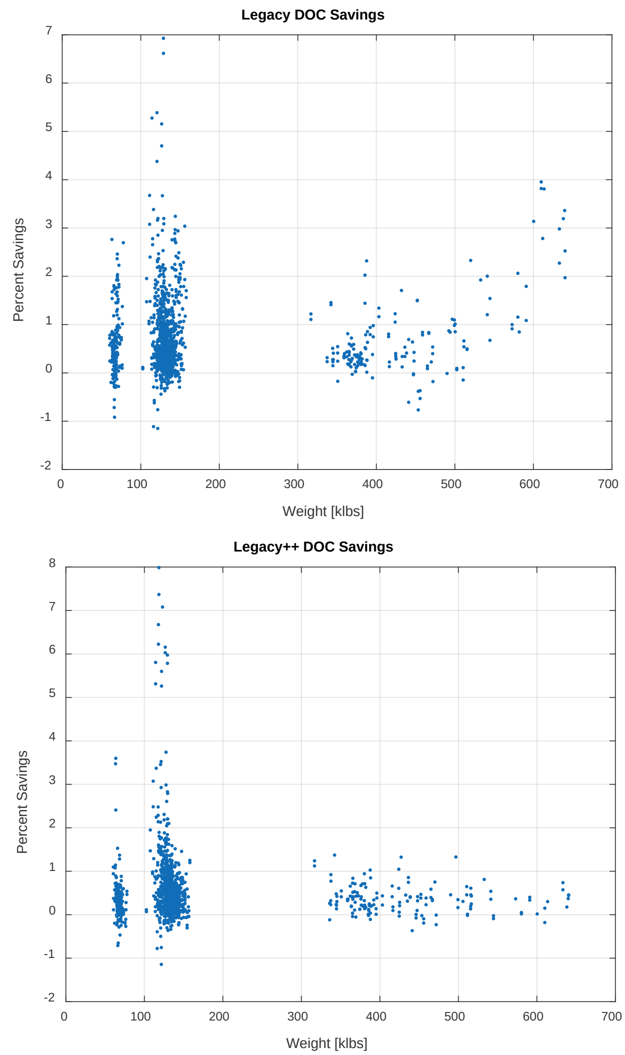

Figure 64. Legacy DOC savings.

Leverage plot for Cost Index 0 shows three distinct aircraft groups (regional, narrow, and wide) and their savings.

For Legacy comparison flights, the percent savings is the highest for the heaviest aircraft. These heavier aircraft have long flights with higher takeoff weights and thus more step climbs and descents during cruise, resulting in more fuel savings. The suboptimality of a constant altitude (Legacy) flight worsens with longer flight distance and becomes increasingly detrimental.

This high percent savings for heavy aircraft has a large effect on the overall mean savings shown in paragraph 7.4.1. Generally, large aircraft have more leverage on the mean savings due to having an overall higher aircraft weight which typically results in more fuel burn. E175 has low leverage, due to its low weight, when calculating the overall mean savings so it has a small effect on the mean. Wide body aircraft have higher leverage, due to a larger weight, and therefore, have a high influence on the mean. Since wide body aircraft have a higher savings relative to the other aircraft types, the mean is increased due to that influence. The leverage plot of aircraft weight versus fuel savings is shown in Figure 64 for Cost Index 0.

Twenty routes were chosen to represent average operable routes flown by aircraft operators. The routes were then chosen to represent trends in historical data for each aircraft.

The greatest percent cost savings for CLEEN flights compared to Legacy come from short and long flights; the most benefit over Legacy++ is gained during short flights. Benefit in short flights is generated mostly from UCCD and variable speed and thrust climb. This effect is present in the previous Gen B and is discussed in more detail in paragraph 6.4.3. Long flights benefit greatly from Gen C steps. Since Legacy++ technology has step climbs, the benefit of CLEEN is reduced when compared to the benefit for the same flights over Legacy technology. By splitting the flights up by aircraft type, the trend is observed for each subset of aircraft (see Figure 65).

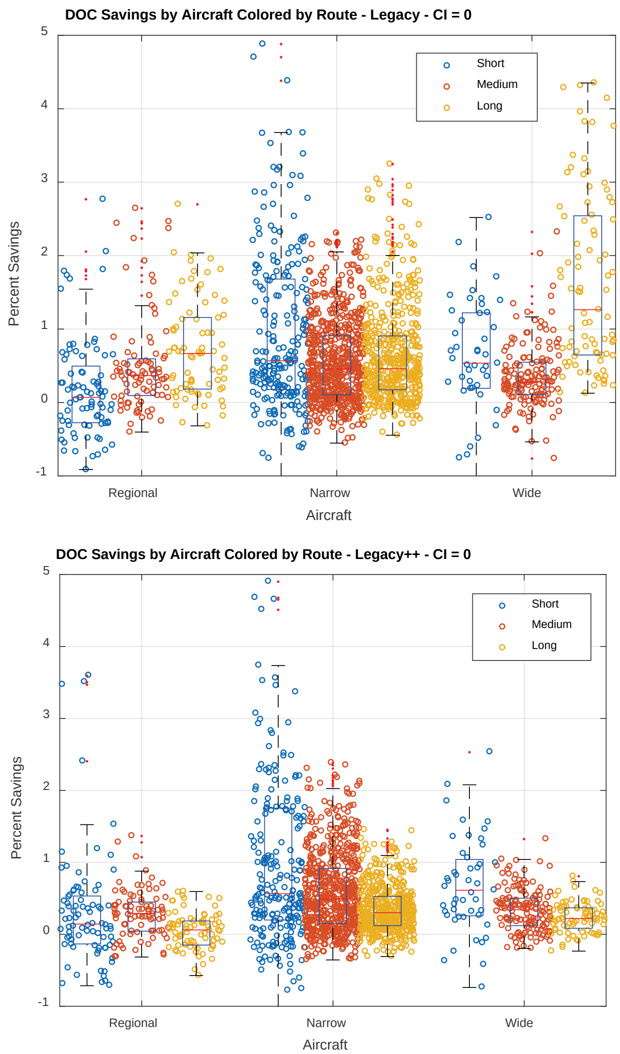

Figure 65. DOC percent savings separated by aircraft type, colored by route.

Observe that in Legacy flights, short and long routes save the most fuel (with exception of the regional Legacy short flights). CLEEN benefits from short routes the most over Legacy++ technology. The same trend, while not displayed here, is observed for Cost Index 25 cases.

Legacy methods have an a priori assumed distance baked into the calculations used to generate the economy tables. This assumed distance falls into the medium length flight range, meaning that the tables providing economy data should perform the best for those flight distances. This is another reason why CLEEN technologies save the most over Legacy flights for short and long routes. CLEEN technologies work for all flight lengths allowing the most optimal profile to be found in all flight scenarios. This further emphasizes the need and opportunity of a non-tabular based approach to flight optimization.

Weight plays a vital role in aircraft dynamics for determining operating limits. Therefore, the initial weight of the aircraft has one of the largest impacts on flight profile. Legacy profiles are forced to one altitude (that can be optimal for the given initial weight) while Legacy++ and CLEEN profiles change altitudes, adapting to changing weight as flight distance increases. Due to this effect, savings over Legacy profiles are also correlated with the number of steps the CLEEN flight takes (see Figure 66).

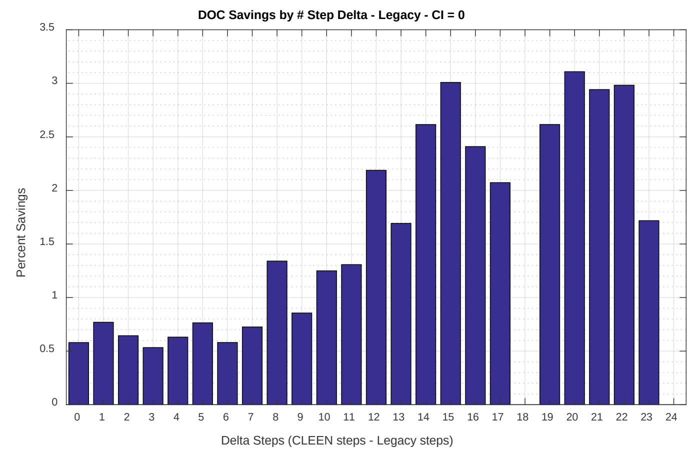

Figure 66. Legacy DOC savings by step number.

Average savings over Legacy generally increases with the number of steps performed.

Observe that as the number of steps increase, cost savings generally increases as well. The number of steps is correlated to the weight of an aircraft. For heavier aircraft, aircraft start at more efficient lower altitudes. Then as fuel weight is burned off throughout the flight, CLEEN profiles step to account for the changing weight. More steps tend to occur for heavier aircraft as shown in Figure 67.

Each grouping of points corresponds to a different aircraft class; red is regional, green is narrow body, and orange is wide body. Within each grouping, as weight increases, so do the number of steps that occur. This corresponds with the trend in Figure 64 that shows savings increasing with aircraft weight.

Some flights have as many as 24 steps. This large number of steps only happen in very long flights. This effect can be seen by overlaying the flight on the weather data as is done in Figure 68.

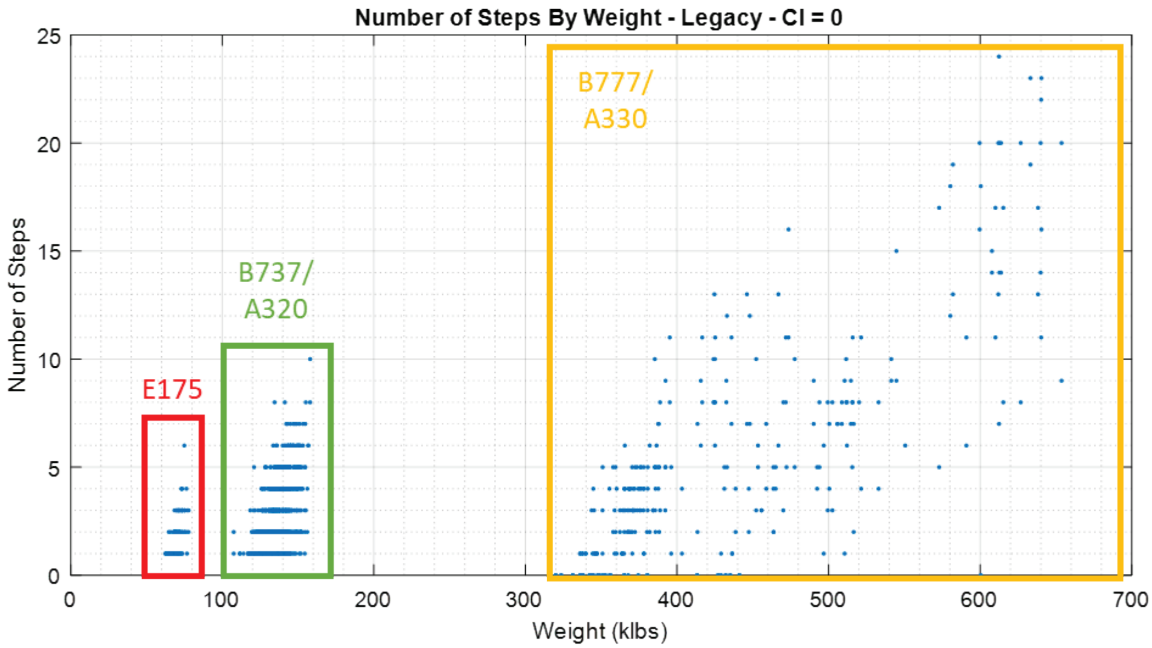

Figure 67. Effect of weight on number of steps.

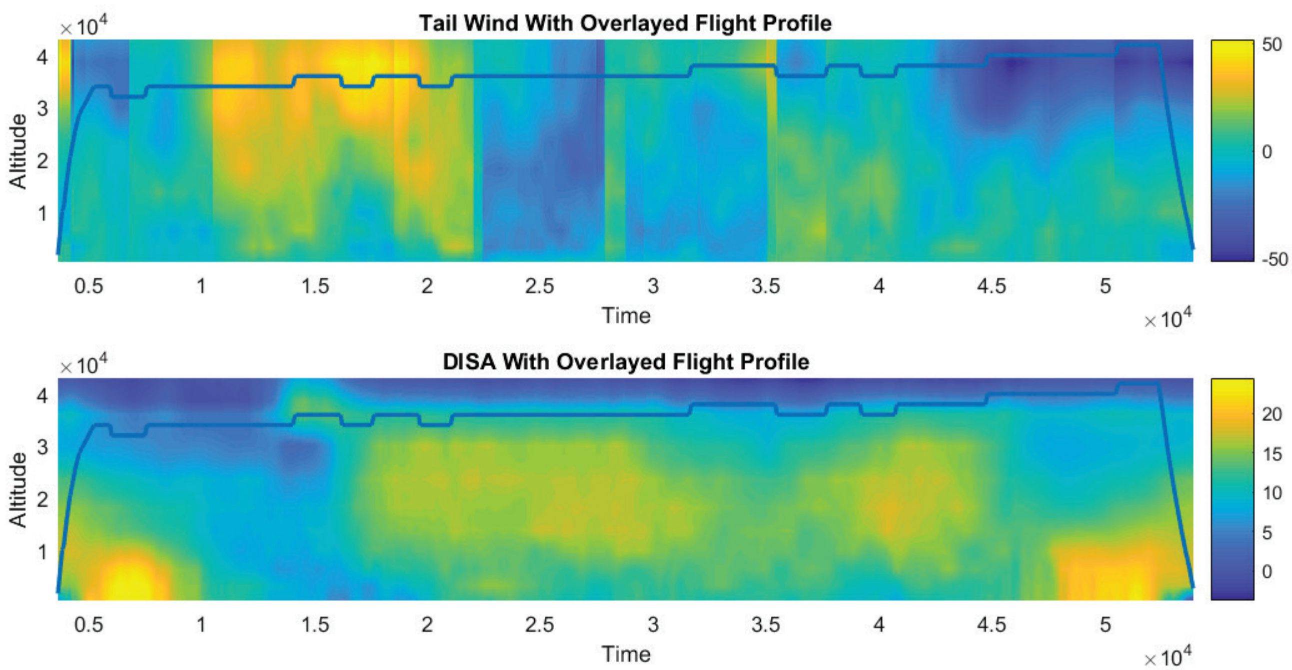

Figure 68. Flight path overlaid on weather profile.

While this case study doesn't always show the nuances that cause each step descent and climb, by analyzing the temperature and wind patterns associated with each descent, it can be seen that descents typically occur to mitigate headwinds. Lower headwinds often allow for a more optimal flight profile; however, it should be noted that the optimizer picks the most optimal profile based off many factors. Therefore, it can't be assumed that the optimal flight path is guaranteed to track wind or temperature patterns along the flight path.

For the Legacy++ comparison, the difference in number of steps can be used to infer how number of steps affects savings (see Figure 69).

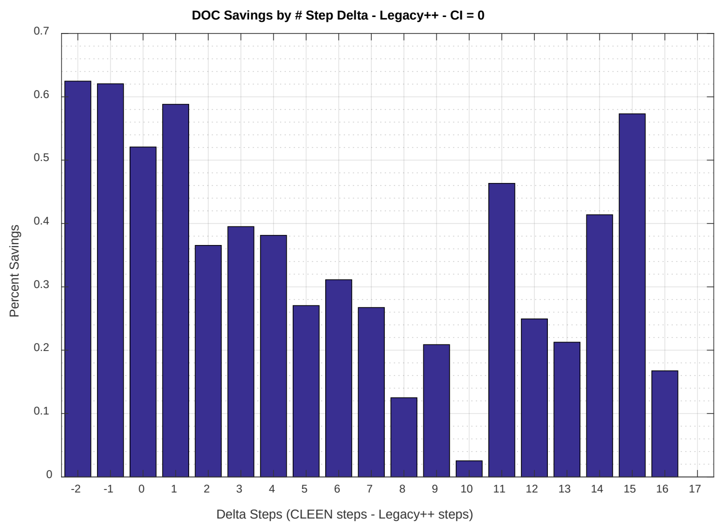

Figure 69. Legacy++ DOC savings by step number.

Average savings over Legacy++ does not show a consistent trend with the number of steps performed but is always positive.

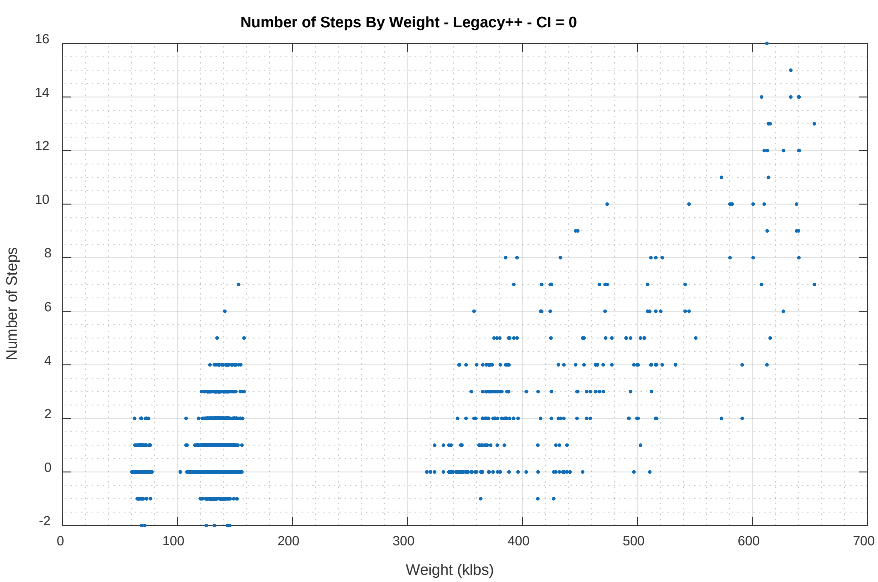

Figure 70. Number of steps by weight for Legacy++.

Note that negative numbers mean that CLEEN had fewer steps than Legacy++. Generally, the difference in number of steps is not a primary factor on the amount of savings generated when compared to Legacy++ cases. As can be concluded from the range of step deltas, CLEEN technologies not only could step more to follow weight changes and weather patterns but can also eliminate unnecessary steps that are not beneficial to perform. The difference in number of steps between CLEEN and Legacy++ is not dependent on weight except for heavier weights where large positive differences in number of steps happen frequently; Figure 70.

8. Generation D: Required Time of Arrival Constraint

8.1 Summary

As the national airspace becomes increasing filled with commercial flights, improved air traffic management is becoming an increasingly important issue. To make flying safer, more efficient, and more predictable, the FAA has encouraged the development of several technologies and processes under the banner of The Next Generation Air Transportation System, or NextGen. This initiative, which began rolling out improvements in 2007, and plans to have all major components in place by 2025, encompasses dozens of innovative technologies and processes, including those focused on air traffic management and decongestion. As the FAA continues to implement the NextGen procedures, Required Time of Arrival (RTA) capabilities will see increased use as a tool to assist with this effort; see Figure 71. According to the FAA's NextGen[‡‡]:

> "Expanding the use of Time-Based Metering and advancing Required Time of Arrival (RTA) capabilities will enable a new level of predictability that will greatly enhance collaborative planning."

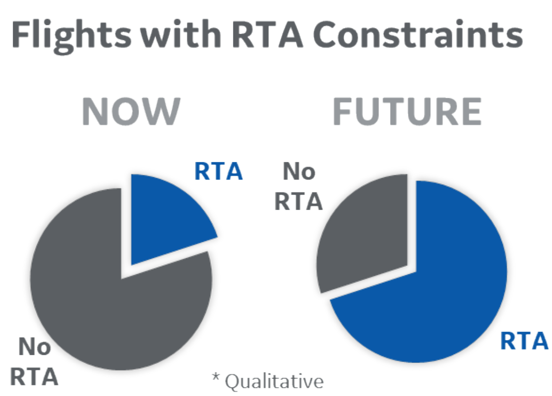

Figure 71. RTA constraints.

RTA capabilities will see increased use with the continued implementation of NextGen.

RTA capability allows ATC to schedule aircraft arrivals in specified intervals to decongest the airspace around airports and prevent overlapping arrival times. In addition to the use of RTA constraints at destinations, RTA constraints may be applied at any waypoint in the flight. Effective RTA capabilities produce many desirable effects: an increase in airspace capacity allowing higher traffic volumes, improved flight efficiency with less controller involvement, more flexibility in traffic deconfliction options, and precise flight predictions for trajectory synchronization and negotiation.

Current methods used by Flight Management Systems to compute RTA profiles are deficient in at least one of several ways: they lack the ability to vary the altitude to meet the RTA, they do not account for known weather data or use a low-fidelity representation of the weather, or they rely on ATC-specified manual lateral offsets of the flight path to slow down arrival time. The Generation D technology developed provides a method to compute a flight profile that can meet a wider range RTAs with modifications to both cruise speed and altitude (such as step climbs and descents), while taking advantage of favorable weather patterns, such as tailwinds at different cruise altitudes; see Figure 72. This ability is accomplished primarily through augmenting the graph search described in Generation C to keep track of time-of-arrival predictions during the optimization to select a path that meets the RTA.

The existing means of profile optimization to which Generation D is compared are Legacy and Legacy++. These methods are similar to those tested against in Generation C and are described in paragraph 3.3, but both are augmented with the Legacy method of meeting RTAs, described in paragraph 3.3.1.

‡‡ "FAA NextGen", January 2010. [Online]. Available: https://www.faa.gov/nextgen/media/FAA_TASKFORCE_RESPONSE_1-31-2010.pdf.

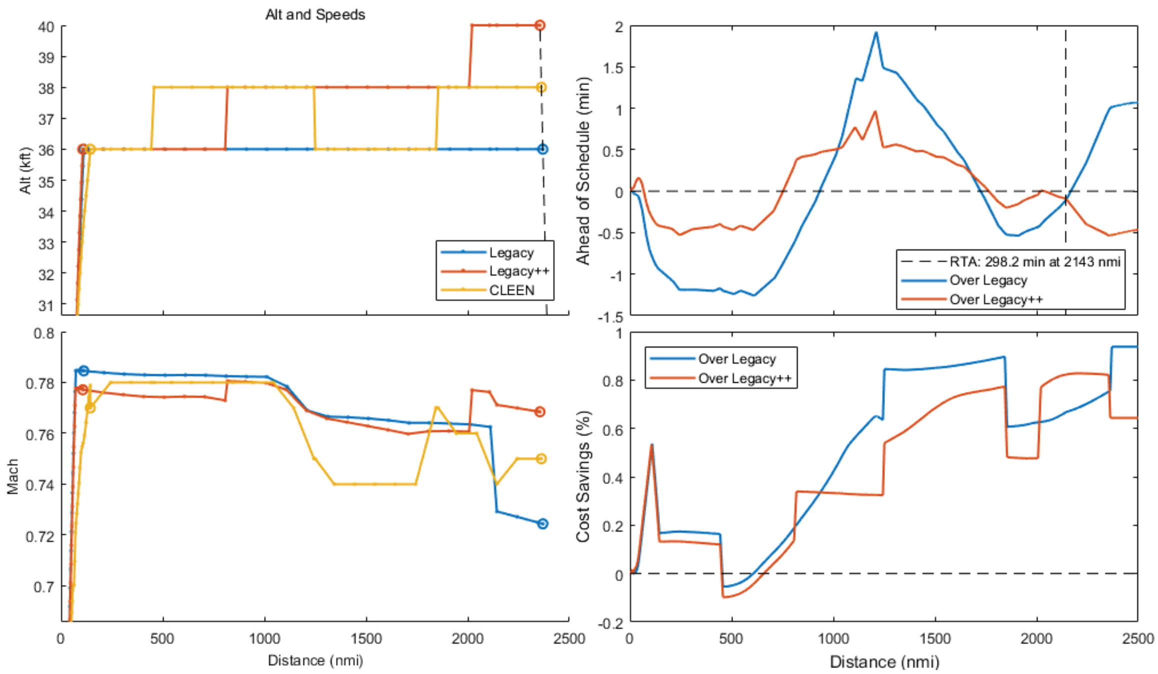

Figure 72. CLEEN Generation D technology.

This technology uses altitude and speed changes to optimize profile and achieve an RTA target – top right subplot shows the time-of-arrival difference at a given distance between CLEEN and Legacy (blue) and CLEEN and Legacy++ (red).

In the TRL 5 benefit assessment testing, Generation D demonstrates that an operator can achieve **0.81%** fuel savings as compared to Legacy and **0.47%** as compared to Legacy++. This result is shown in Figure 73. In addition, the benefit of Generation D technology extends beyond cost savings. Several other metrics of improvement can be considered when dealing with flights containing RTA constraints, as indicated in Figure 74. TRL 5 testing demonstrated that the Generation D technology performs about **2.7 times fewer re-predicts** than the Legacy methods, improving profile precision and increasing efficiency of trajectory synchronization and negotiation with Air Traffic Control (ATC). Generation D also demonstrated a **52% wider range of achievable RTA constraints**, allowing more flexibility for flight arrival times and thus improving ATC ability to schedule flights. Figure 74 indicates that Generation D has a greater than 93% success rate of meeting the RTA constraint in simulation. This rate can be improved via modifications to the algorithm that will increase computation time.

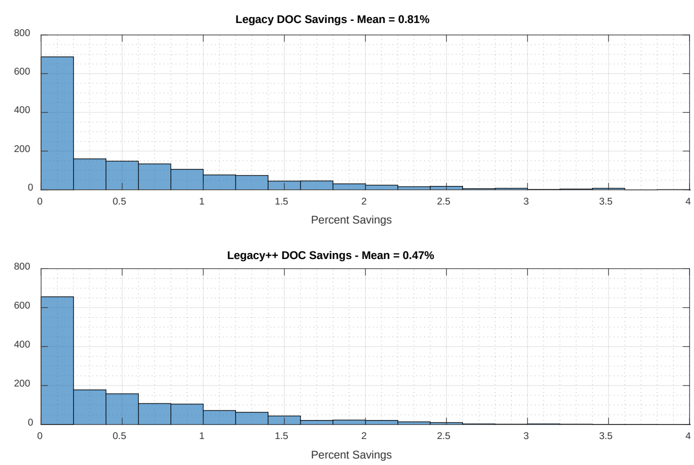

Figure 73. TRL 5 benefit assessment testing.

Monte Carlo study in TRL 5 simulation testbed shows significant savings from Generation D

Metrics	Value
Mean DOC savings over Legacy method	0.81%
Mean DOC over Legacy++ method	0.47%
Maximum cost savings over Legacy method	Up to 3.5%
Percentage of cases with improved RTA accuracy AND fuel (over Legacy)	67%
Percentage of cases with improved RTA accuracy OR fuel (over Legacy)	98%
Improved solution stability measured in reduction in re-predicts	2.7x
RTA failure rate	Less than 7%
Increased RTA window	52%

Figure 74. Key takeaways from the Generation D TRL 4 testing and TRL 5 benefit assessment.

8.2 Development

The TRL 4 implementation of Gen D builds on the previous TRL 4 implementation of Gen C software.

In the example below, the Gen D optimizer is used to optimize the flight path subject to an RTA constraint and with real weather data. In this example, the RTA location is set around 2200 nmi, as indicated by the vertical dashed line in the upper right plot of Figure 75. The left two plots show the selected speed and altitude profiles for each method. Observe that the Gen D profile performs a step descent about halfway through the cruise, then climbs up again before reaching the top of descent. Accordingly, the Gen D profile uses a lower speed during this segment.

The top right plot shows the time accumulation of the Gen D profile compared to both Legacy methods. This is done by treating each Legacy method's distance-time profile as a "schedule", and determining if Gen D is ahead or behind (positive indicating ahead). The dashed line crosshairs indicate the RTA constraint in both distance and time. Finally, the bottom right plot shows the accumulated cost savings for Gen D over each method. This example shows that all three methods meet the RTA within the specified tolerance of 10 seconds, while the Gen D method does so with 0.94% savings over Legacy and 0.64% savings over Legacy++. In this example, a part of these savings come from the step descent that saves 91 pounds of fuel by harnessing the larger tailwind at 36,000 ft (Figure 76), during which the Gen D profile cruises at a slower speed of Mach 0.74 (that is more efficient) while maintaining the ability to reach the RTA. The majority of the remaining savings come from the optimal climb phase.

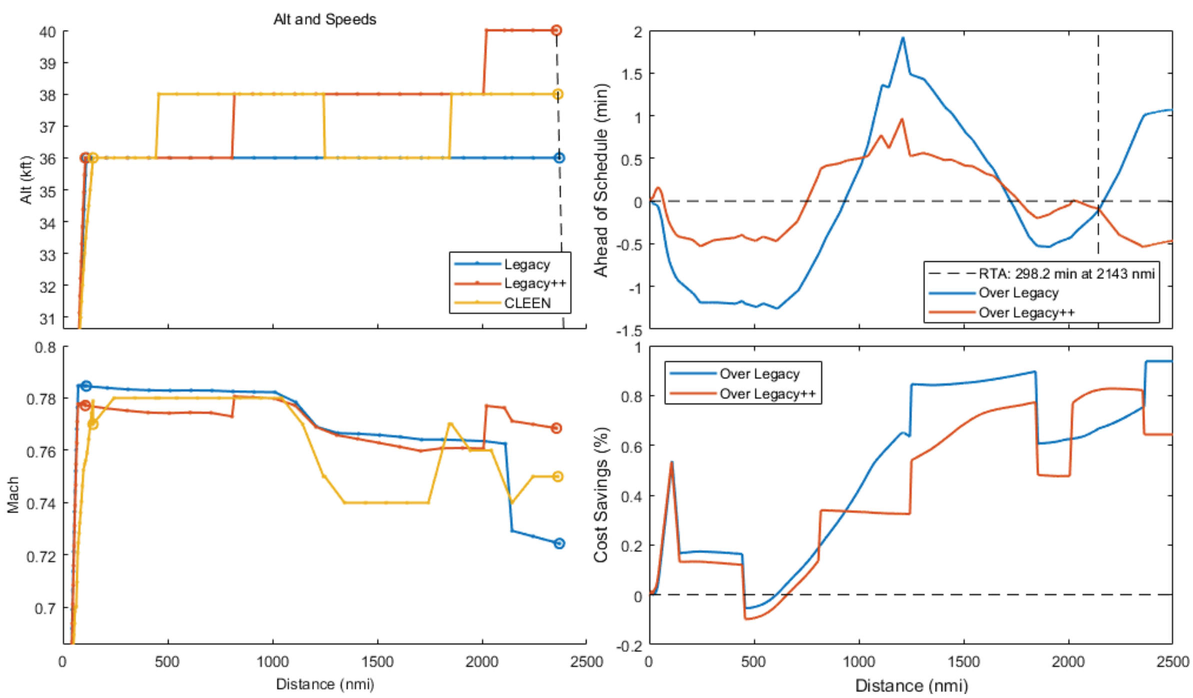

Figure 75. CLEEN achieves more optimal flight while maintaining ability to hit an RTA.

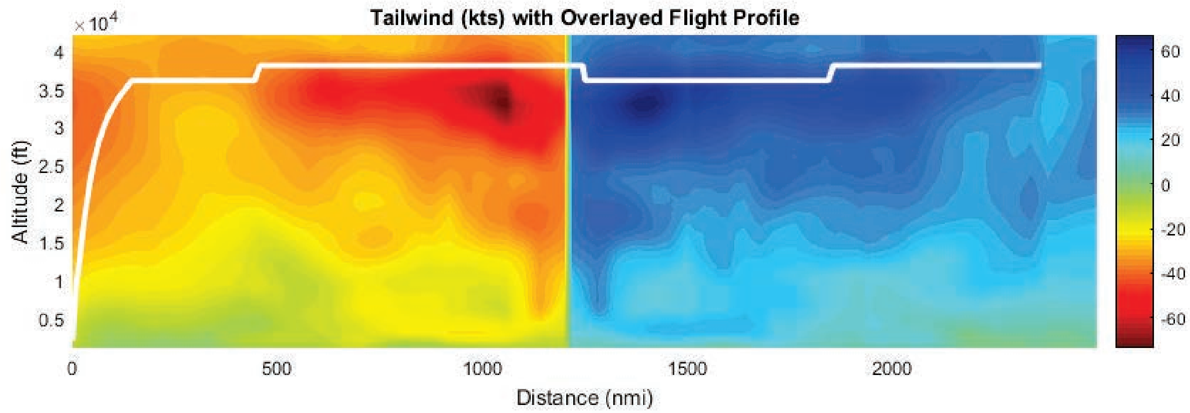

Figure 76. Weather overlaid on CLEEN profile shows tracking to favorable winds.

8.3 Algorithm Validation Testing

8.3.1 Test Methods

Both the legacy systems (Legacy and Legacy++) and Gen D algorithms are implemented in a TRL 4 MATLAB test suite using simplified point-mass integration of the trajectory. Once a planned control trajectory is calculated, separate point-mass simulation is performed to compute the cost of each profile. See paragraph B.1 for a detailed description of the TRL 4 testing method.

The single test set shown in Figure 77, totaling 2880 flights, is developed to evaluate the Gen D technology. The test set performs a sweep of all combinations of several parameters, as shown in the table.

Parameter	Variation of Parameter
Takeoff Weight	110, 130, 150, 170 klbs
Flight Distance	500, 800, 2500, 3500 nmi (both directions)
Altitude Quantization	2000 ft
Cost Index	0
Weather Days	4 normal and 2 inclement
RTA Distance	60% of CRZ, 90% of CRZ, 50% of DES
RTA Time	10%, 25%, 50%, 75%, 90% of window

Figure 77. Variation of basic flight plan parameters in the test set.

The weather data for TRL 4 testing comes from the NOAA National Operational Model Archive and Distribution System (NOMADS); specifically, the Rapid Refresh System (RAP) format. The weather system utilized throughout CLEEN work was established during Gen B development (see paragraph 6.2.1 for more details). The tested weather days are identical to those selected in Gen C TRL 4 testing, in Test Set 4 (paragraph 7.3.1). They are 06/23/16, 03/08/16, 05/24/16, 10/28/16 (normal days), and 3/17/16, 4/15/16 (inclement days).

As indicated in Figure 77, the RTA distance and time parameters are determined by a percent value of the distance travelled in a Legacy++ computed profile. See Figure 78 for a visual depiction of this process.

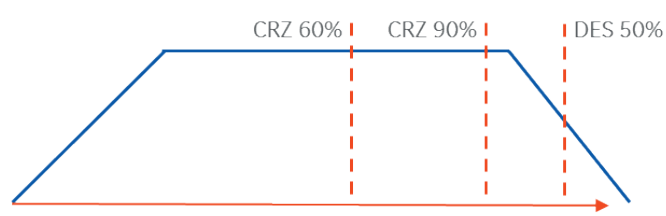

Figure 78. Visualization of RTA Distance parameter in test set.

Further Legacy++ profiles are generated with a minimum Cost Index (-30) and again with the maximum Cost Index (200) to give a full window of arrival time at the chosen RTA location, where the maximum Cost Index gives a minimum arrival time (the fastest flight), and vice versa. By treating the maximum arrival time (from the minimum Cost Index and thus a slow flight) as 100% and the minimum arrival time (from the maximum Cost Index and thus a fast flight) as 0%, the RTA constraint time for the case is set according to the RTA Time parameter. This is depicted in Figure 79.

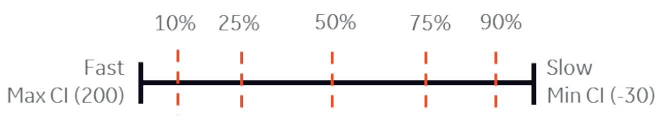

Figure 79. Visualization of RTA Time parameter in test set.

8.3.2 Test Results

The Gen D optimizer is judged on four main criteria:

1. **Optimality:** Average percent savings over the Legacy and Legacy++ methods. Note that metrics in this TRL 4 assessment are only relative to one another and do not accurately represent true savings in service due to descent angle assumptions used to calculate benefits at this early development stage. Later TRL 5 benefit assessment (paragraph 8.4) provides a more accurate measure of true savings.
2. **Computation Time:** Both average and maximum execution time of the optimization algorithm
3. **RTA Success:** The percent of cases for which the optimizer returns a solution that meets the RTA within a certain tolerance
4. **Cost Improvement Success:** The percent of cases for which the optimizer returns a solution that saves cost compared to the Legacy method

Figure 80 and Figure 81 collectively summarize the benefits using these four metrics, as well as show the development history and how certain features and improvements contributed to these benefits.

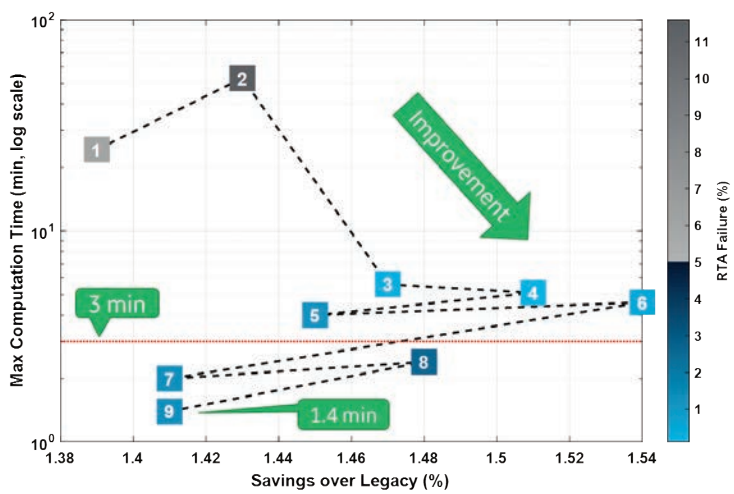

Figure 80. Gen D development history.
This shows an iterative approach to reduce computation time and improve optimality and RTA success rate.

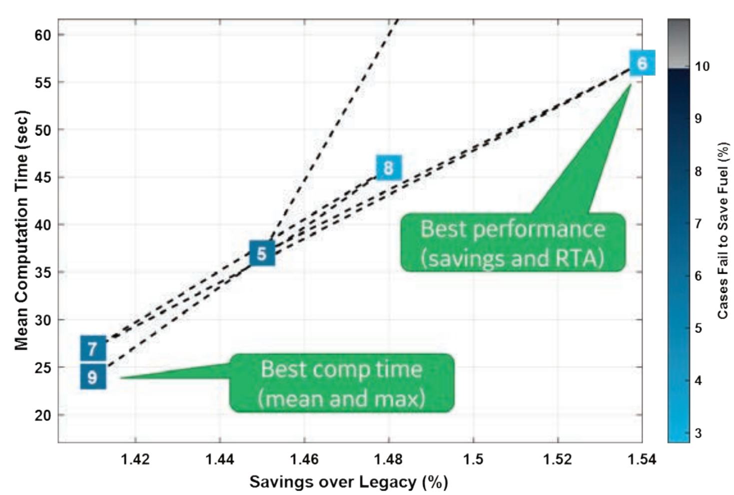

Figure 81. Gen D development history (zoomed in).
This shows that Algorithm 9 provides the best computation time.

Figure 80 shows each algorithm version, numbered in order of development. Each algorithm is plotted on the chart with the savings over Legacy, maximum computation time (on a logarithmic scale), and each is colored by RTA failure rate. As is shown by the green arrow, overall algorithm improvement is indicated by proximity to the lower right edge of the chart, and lighter blue color – higher savings, lower computation time, lower RTA failures.

Observe that the original attempt for solving the Gen D problem, is listed first as Algorithm 1 and had an undesirably large RTA failure rate of above 5%.

The goal for maximum computation time, under 3 minutes, is called out with a dashed red line on the plot, and as can be seen, the latest improvements to the Gen D optimizer achieve that goal. In the final form, a solution is produced in under 1.5 minutes for all cases, while only sacrificing about 0.13% optimality compared to the highest performing option studied.

Figure 81 shows results for the same algorithms as Figure 80, but with a few changes: focus is given to the final five iterations, the y-axis is switched from maximum to mean computation time, and the algorithm boxes are now colored by cost improvement failure rate. This figure shows the trade-off between optimality and computation time for various sets of features in the Gen D optimizer.

For the remainder of this analysis, the low computation time is prioritized and thus Algorithm 9 is considered the final Gen D method. The distribution of savings over each method is indicated in Figure 82 and Figure 83. In these plots, a positive number indicates that the Gen D optimizer outperformed the relevant Legacy method. The cases where Gen D underperformed are indicated by red highlighting. Theoretically, a perfect implementation of the optimization techniques used in Gen D should not produce a profile worse than the Legacy profile. However, in practical applications, losses can still occur. There are a few causes of these anomalies, such as algorithm discretization or prior simplifications made to speed up execution time. However, in a TRL 6 implementation with the EFB and Connected FMS, these cases will be mitigated: If Gen D underperforms the current Legacy profile, the Gen D solution is not sent to the FMS, and the Legacy profile continues to be used by the FMS, resulting in no savings or losses.

Histograms separated by flight distance, weight, and RTA time reveal more clearly the trends present in the data. The dashed lines indicate the mean of each sub histogram that is useful for observing how the savings correspond to the indicated parameter in each subset.

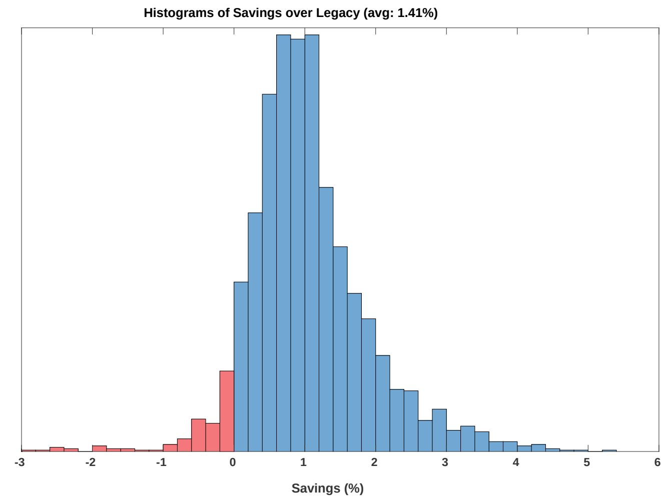

Figure 82. Gen D savings over Legacy.

Savings averages 1.41%, with most cases saving fuel.

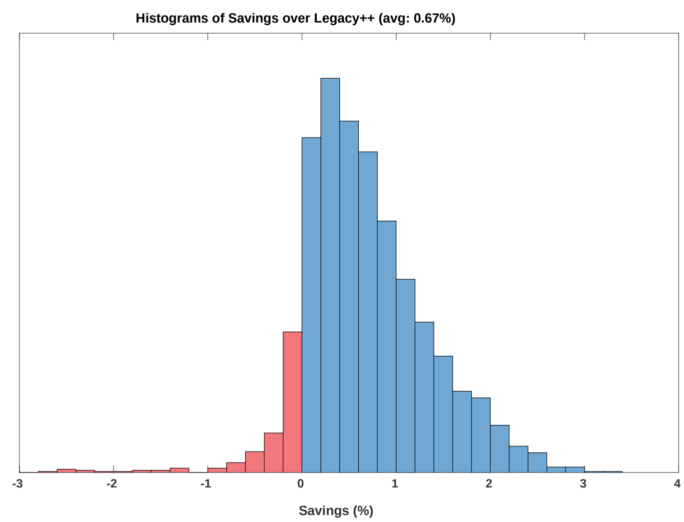

Figure 83. Gen D savings over Legacy++.

Savings averages 0.67%, with most cases saving fuel.

It is first worth noting that, when reporting cost savings numbers in a percent, the flight distance plays an important role. The same absolute savings corresponds to different percent values for different length flights. For instance, for a short flight and a long flight, savings of 100 pounds of fuel corresponds to a larger and smaller percent, respectively. This effect is present in all the following results.

The average savings over Legacy generally increase significantly with increasing flight distance (Figure 84). This makes sense, as with longer flights the cruise phase becomes the dominant portion of the longer flight and the Legacy method does not utilize step climbs to adjust the cruise altitude. Accordingly, the Legacy method has a very suboptimal altitude at the end of long cruise phases, when a good portion of the aircraft fuel weight has been burned off.

The one exception to this generality is that the 500 nmi routes show more savings than the 800 nmi. This is due to the Gen D optimizer choosing an altitude *lower* that the Legacy altitude (an altitude that is technically more suboptimal when viewed in isolation) because the UCCD framework determined that it was not cost effective to burn the climb fuel to get to a higher altitude where little time was to be spent cruising.

For heavier aircraft, the savings generally increase. This correlation is essentially the same effect as the correlation with distance, due to the filtering applied to the cases.

For the savings over Legacy++ (shown in Figure 54), the trend is opposite and not near as strong: longer flights generally result in the Gen D optimizer saving slightly less percent over the Legacy++ method. There is an opposite effect at work here from the savings over Legacy: for shorter flights, the climb phase dominates, and the cruise phase of Gen D is similar in many cases to the cruise phase of Legacy++ as both perform step climbs. This effect prevents the benefit from accumulating throughout cruise. Instead, the variable-thrust variable-speed climb phase provided by the UCCD framework generates the majority of the savings for the flight. Thus, longer flights result in a smaller percent saving number.

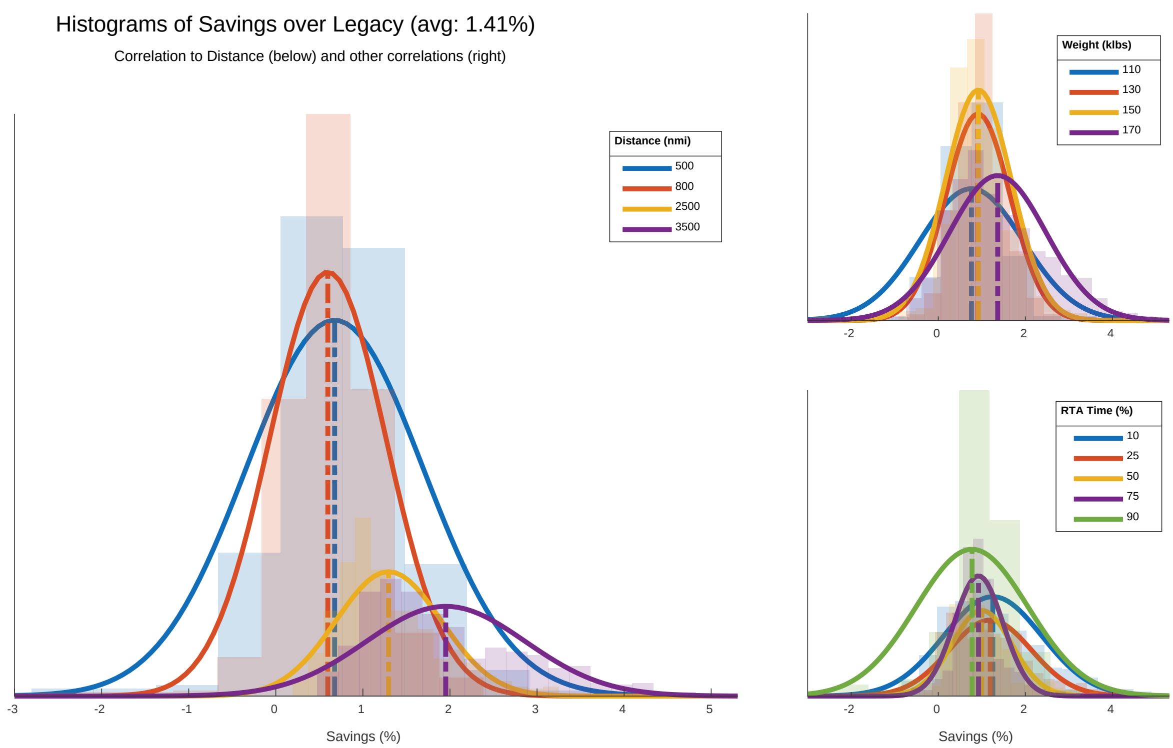

Figure 84. Savings over Legacy.
Heavier aircraft travelling longer range exhibit more CLEEN Gen D savings over Legacy.

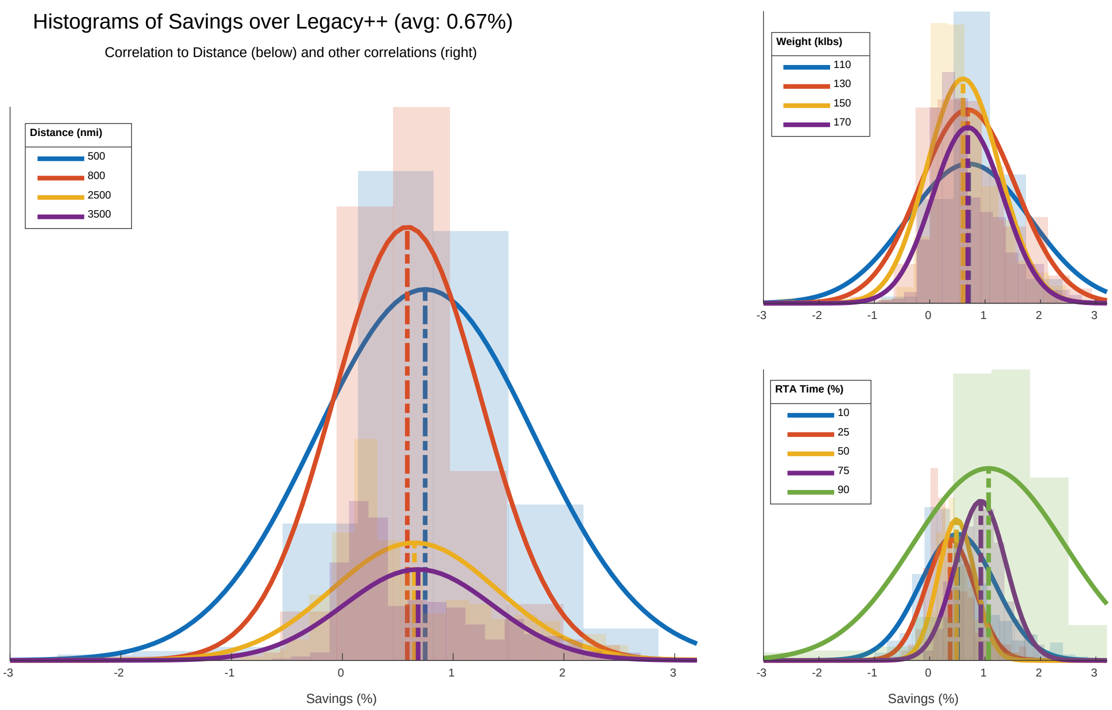

Figure 85. Savings over Legacy++.

CLEEN Gen D savings over Legacy++ are not strongly correlated with weight or distance.

8.3.3 Ancillary Benefits

Unlike in prior generations of the CLEEN software, cost savings is not the only benefit that Gen D exhibits. When dealing with flight profiles containing RTAs, RTA achievement accuracy, profile predictability, and achievable RTA range are key operational concerns. These benefits are assessed using an early version of the Gen D optimizer. The analysis has not been repeated for the final version of the optimizer, but the same general trends and magnitude of benefits are qualitatively observed in both.

In general, a prediction (from either the Gen D optimizer or from a Legacy method) that successfully meets the RTA within the required tolerance of 10 seconds does not necessarily meet the RTA in simulation (due to unknown environment or aircraft parameters at the time of control calculation). If the given method had a poor representation of the true tailwinds, then the true tailwinds combined with the profile controls from predictions of inaccurate tailwinds might cause the actual arrival time to be different than the ETA that was predicted.

To quantify this, Figure 86 shows fuel burn and RTA accuracy benefits in a consolidated view. On the x-axis is the fuel savings that Gen D produces in pounds of fuel, where a positive number indicates Gen D saved fuel. On the y-axis is the RTA accuracy improvement that Gen D provides, characterized as the difference between Gen D and Legacy in absolute RTA error, where RTA error is the difference between the RTA time and the actual arrival time computed from cost calculator simulation. The location of each data point in the scatter plot indicates if Gen D improved fuel consumption compared to Legacy, improved RTA accuracy compared to Legacy, improved both, or improved neither. All data points in the top right quadrant showed a Gen D improvement in both fuel and RTA accuracy, corresponding to 67% of the data when comparing to Legacy and 56% of the data when comparing to Legacy++. Likewise, only a small handful of cases (2% and 3% respectively) showed an improvement in neither. The green and gray arrows on the axes demonstrate what each quadrant implies about that data point.

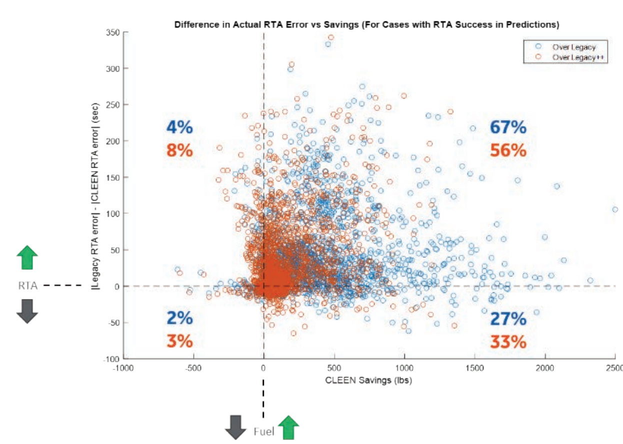

Figure 86. Fuel burn and RTA accuracy.

Majority of cases against both Legacy and Legacy++ show CLEEN improvement in both fuel savings and RTA error.

In service, the Legacy methods rely upon frequent re-prediction (that update control targets) to adjust the profile to meet the RTA while in the presence of unexpected winds. If the initial flight profile that is predicted is untrustworthy and will require a re-predict at some point enroute, that places a burden on pilot to inform ATC of the adjusted profile. In turn, ATC will need to revise local estimates of aircraft locations and schedules, causing potential conflicts that will need to be addressed by negotiating new schedules with nearby flights. This process, known as trajectory synchronization and negotiation, will be made inefficient by the variability in the flight profile caused by frequent re-predicts. In a crowded airspace, the impact of this inefficiency may be significant. However, if the flight profile that the crew initially filed with ATC is more predictable, stable, and trustworthy, then work performed by ATC to ensure adjacent trajectories are synchronized will be more efficient. This will effectively minimize controller workload. The Gen D optimizer achieves this level of predictability mainly by the high-fidelity weather model that is much more accurate than models in the Legacy methods.

In practice, the FMS utilizes a rate logic to determine when a re-predict of aircraft profile is necessary to meet the RTA. The logic is shown in Figure 87: for the flight time until the RTA is encountered (in minutes on the x-axis), the function shows the RTA error tolerance (in seconds on the y-axis) that is allowed before a re-predict is triggered. Typical limits are shown at 15 seconds on one end and 120 seconds on the other, and the essential rate is one second of RTA error tolerance per minute of flight time until the RTA.

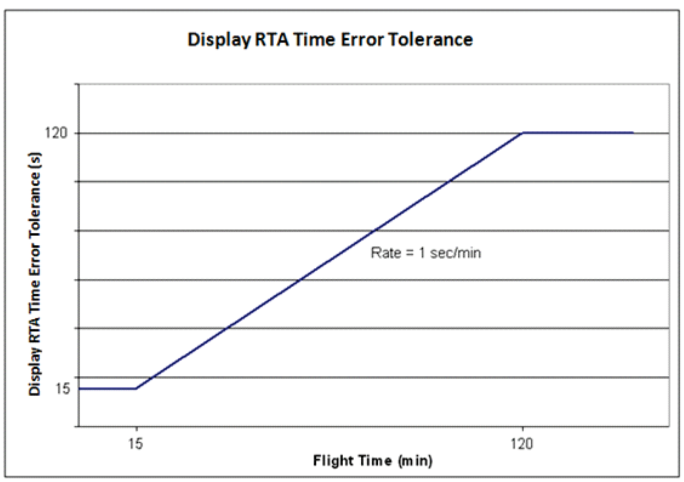

Figure 87. Legacy logic determining when to trigger re-predicts.

Under this logic, the FMS would attempt to re-predict the profile to meet the RTA if the error between the ETA and the RTA becomes larger than 15 – 120 seconds, depending on the distance away from the RTA. Taking a more stringent 30-second rule as a constant, it is possible to determine how many cases from each method would require *at least one* re-predict during the flight. Figure 88 provides counts for the number of cases where the RTA was met successfully in predictions (within a 10 second tolerance), and then how many of those cases failed the RTA in simulation (more than 30 seconds of error). Using these data, the final column shows the percent of cases that would require at least one re-predict for each method. As seen in the final column of Figure 88, the Gen D optimizer requires around five to six times fewer re-predicts than both Legacy methods. This is under the assumption that only a single re-predict would be needed to realign the inaccurate profile to meet the RTA, which might not be the case when the weather model in the Legacy method is especially inaccurate.

Method	Predictions RTA Success (10 sec)	Simulation RTA failure (30 sec)	Cases needing re-predicts (%)
Legacy	2865	1033	36.1%
Legacy++	2863	942	33.9%
Gen D	2717	181	6.6%

Figure 88. RTA successes.

Gen D flights are 5-6x less likely to require a speed optimization recalculation during flight than Legacy approaches.

Another benefit Gen D brings about is an increased range of available RTAs. In practice, the ultimate range of achievable RTAs is limited by the aircraft flight envelope, especially on the fast side: typical commercial transport aircraft cannot travel much faster than Mach 0.8. However, within those limitations, the flight path generation process is responsible for creating profiles that meet the selected RTA. Because of the Legacy method's limitation to only vary speed to meet RTAs, the range of achievable constraints is rather limited.

In the case where the RTA time constraint requires the aircraft to slow down (hereafter called a slow constraint), the Legacy method may reduce speed; however, that solution can only go so far before reaching the lower end of the flight envelope. If the RTA is still not achievable, operators will create modifications to the lateral path in order to arrive at the correct time. These path-lengthening solutions create many inefficiencies, as would be expected in any solution where the distance traversed is increased from optimal. According to a baseline test by The MITRE Corporation, each additional minute of flying time created by a lateral offset results in an additional

71 pounds of fuel consumption[§§]. Furthermore, these solutions require intervention from ATC and thus increase pilot and ATC workload.

The ability to meet RTAs outside of the traditional Legacy range without modifications to the lateral path is highly desirable. While earlier arrivals may seem to be the most important given the desires of passengers, the ability to create later arrival times and meter traffic flow in just the vertical path and aircraft speed is certainly attractive to operators and controllers, especially in congested airspace.

To determine this benefit, a separate test set was created with RTAs outside of the 0-100% RTA time window described in paragraph 8.3.1. This test set utilizes the same weather days, true Cost Index (0), and altitude separation (2000 ft). However, only the route distance of 1000 nmi with weight of 135 klbs is tested. The RTA distance parameter is fixed at the 90% of cruise location. Finally, the RTA time ranges are set outside of the usual window, to the ranges depicted in both the faster and slower regions in Figure 89. The RTA time parameter is varied in 10% increments; thus, the test range of RTA times is [-40, -30, -20, -10, 110, 120, 130, ..., 240, 250] %. The left limit of -40% was chosen experimentally: with RTA times faster than that, none of the three methods were able to meet it. The range extends out far into the slow end, with 250% chosen by engineering intuition.

Figure 89. Test range for RTA time window expansion.

For this analysis, the allowable search range for the RTA Cost Index in the two Legacy methods is expanded to match the Gen D range of [-50, 500], ensuring that all methods have an equal opportunity to reach the RTA. The results of the analysis are depicted in Figure 90 that shows the percent of cases that successfully met the RTA for each of the selected RTA times, with the three methods side-by-side in separate colored bars.

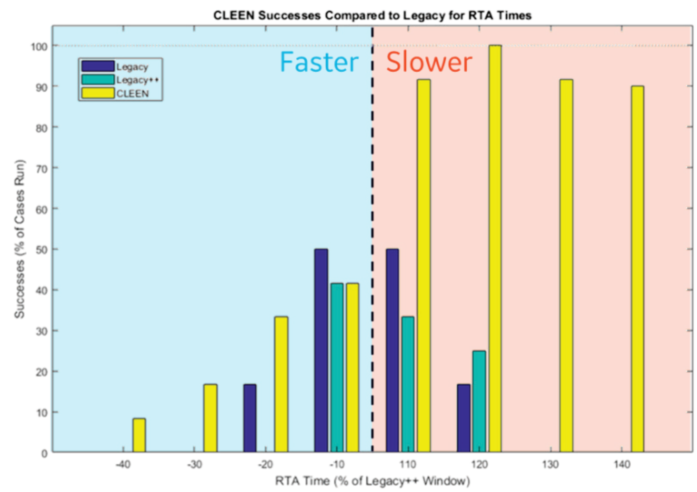

Figure 90. Successes in RTA.

Improvement in available RTA range for Gen D (truncated at 140%) shows 52% wider RTA achievable range.

This shows that on the faster end, the Gen D method can meet a handful of RTA times that neither Legacy method can meet, in the -30% and -40% area, while all three methods can meet an equal number of -10% RTA times. This is expected, as the RTA time of 0% is already near to the aircraft flight envelope limits, and there is not much room available for the aircraft to speed up any further. The slower end of the spectrum is where the Gen D optimizer brings a major improvement. Without the need for costly lateral offsets, the Gen D optimizer can meet most slow RTAs. While the plot in Figure 90 is truncated at 140% for the sake of clarity, the trend continues onwards towards the 250% RTA time value. Gen D is able to create profiles that meet a whole host of slow RTAs, while both legacy methods can go no slower than 120% without the need for a lateral offset. To determine the increased range, the mean RTA time value for all the cases with RTA successes for each method is shown in Figure 91.

From the data, the average value for achievable RTA time for the Legacy methods is not much more than 10% outside of the usual window. However, the Gen D optimizer improves that range to -19.2% on the fast end, and 171% on the slow end. Comparing the size of these intervals, it can be seen that the Gen D optimizer has a 52% wider range of available RTAs, the majority of which are on the slow end.

[§§] P. Ostwald, "Impacts of ATC Related Maneuvers on Meeting a Required Time of Arrival", The MITRE Corporation, Egg Harbor Township, NJ, 2007.

Method	Fast	Slow
Legacy	-12.5%	112.5%
Legacy++	-10%	114.3%
Gen D	-19.2%	171.1%

Figure 91. Mean of RTA time of success cases for each method.

8.4 TRL 5 Benefit Assessment

8.4.1 Fuel Savings (Cost Index 0)

With outliers removed from the dataset, the cumulative fuel savings of CLEEN Generation D technology for Cost Index 0 are 0.74% and 0.39% over Legacy and Legacy++ respectively. This metric represents the total fuel savings from all flights as a percentage of the total fuel burn from all baseline flights. The 99% confidence interval for Legacy savings is 0.68% - 0.81%, meaning that there is a 99% chance that the true mean will fall within the range specified. The 99% confidence interval for Legacy++ is 0.32% - 0.45%. The histogram in Figure 92 shows the distribution of percent savings for all Cost Index 0 cases. This analysis includes all test cases where both the CLEEN and baseline sets of optimal control result in a flight plan that meets the RTA. While a rare occurrence, cases where either set of controls miss the RTA are excluded from this analysis.

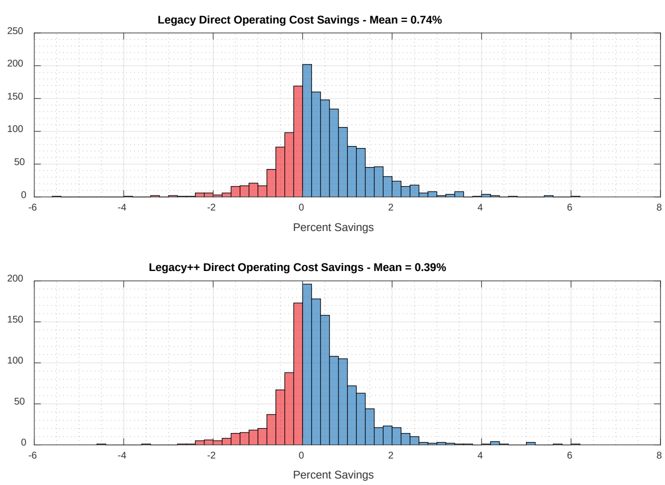

Figure 92. Cost Index 0 savings.

Gen D technology saves significant fuel over Legacy and Legacy++ baselines.

In practice, an operator would compare the output of the CLEEN algorithm with their Legacy approach and choose the control that produced the lowest cost. In other words, the cases where the CLEEN algorithm fails to save fuel over the Legacy methods would not be flown by the aircraft operator, and instead, the Legacy profile would be used, obviously resulting in a 0% savings over the Legacy method. Thus, the true fleetwide benefit an operator would observe is calculated by replacing the negative savings trial cases from the TRL 5 assessment with zeros. The resulting histograms and mean savings for the operator fleetwide benefit is presented in Figure 93. From this modified dataset, the cumulative fuel savings of CLEEN Generation D technology for Cost Index 0 are **0.81%** and **0.47%** over Legacy and Legacy++ respectively.

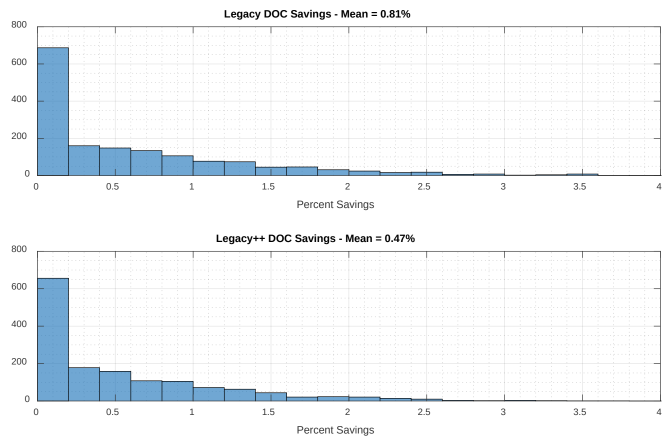

Figure 93. Gen D operational savings for Cost Index 0.

Cases with negative predicted savings would use Legacy control, considered as zero savings.

8.4.2 Cost Savings (Cost Index 25)

With outliers removed from the dataset, the cumulative fuel savings of CLEEN generation D technology for Cost Index 25 are 0.54% and 0.20% over Legacy and Legacy++ respectively. This metric represents the total fuel savings from all flights as a percentage of the total fuel burn from all baseline flights. The 99% confidence interval for Legacy savings is 0.49% - 0.60%. This means that there is a 99% confidence that the true mean will fall within the range specified. The 99% confidence interval for Legacy++ is 0.14% - 0.25%. The histogram in Figure 94 is the distribution of percent savings for all Cost Index 25 cases. As expected, the benefit is lower for Cost Index 25 than Cost Index 0, because the time component of the flight is fixed for both CLEEN and Legacy methods. With Cost Index 25, this represents a large portion of the overall DOC equation.

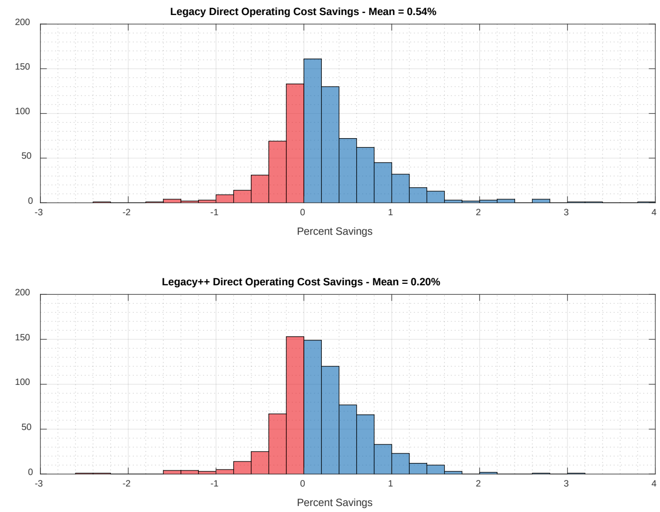

Figure 94. Cost Index 25 savings.

Gen D technology reduces DOC over Legacy and Legacy++ baselines with Cost Index 25.

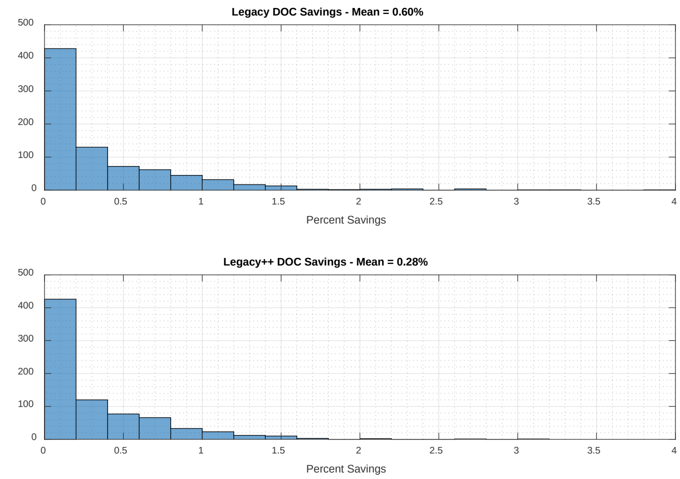

Figure 95. Gen D operational savings for Cost Index 25.

Cases with negative predicted savings would use Legacy control, considered as zero savings.

It should also be noted that Cost Index has no factor on the flight profile prior to the RTA. Normally, without an RTA constraint, Cost Index is the mechanism for relating the cost of time to the cost of fuel. For a case under the presence of an RTA constraint, time is fixed; therefore, prior to an RTA, only fuel matters. For that reason, cases where the RTA waypoint location is located at the end of descent are only run for a Cost Index of 0. This means that Cost Index 25 analysis only includes cases where the RTA is located at the last waypoint in cruise.

The true fleetwide benefit an operator would observe is calculated by replacing the negative savings trial cases from the TRL 5 assessment with zeros. The resulting histograms and mean savings for the operator fleetwide benefit is presented in Figure 95. From this modified dataset, the cumulative cost savings of CLEEN Generation D technology for Cost Index 25 are **0.60%** and **0.28%** over Legacy and Legacy++ respectively.

8.4.3 Execution Time

For Gen D, long flights or difficult-to-meet RTAs can cause the computational complexity of the graph search to exponentially increase due to the increased number of nodes in the search. Work is performed to improve the execution time of the algorithm during development to keep it below a desired maximum time of three minutes.

In the TRL 5 software, compiled code improves execution time over the MATLAB prototype, meaning that even longer flight distances (beyond the 3500 nmi maximum distance tested in TRL 4) will execute relatively quickly. Figure 96 shows a histogram of execution time of the CLEEN algorithm for the full TRL 5 test set, indicating a maximum execution time of just under three minutes. Furthermore, the mean execution time is under seven seconds, and the 95th percentile is just over 13 seconds, meaning that execution times longer than several seconds are outliers, occurring less than 5% of the time. Thus, the optimization algorithm execution time will satisfy operator desires for real-time applications or for batch processing across a fleet of aircraft.

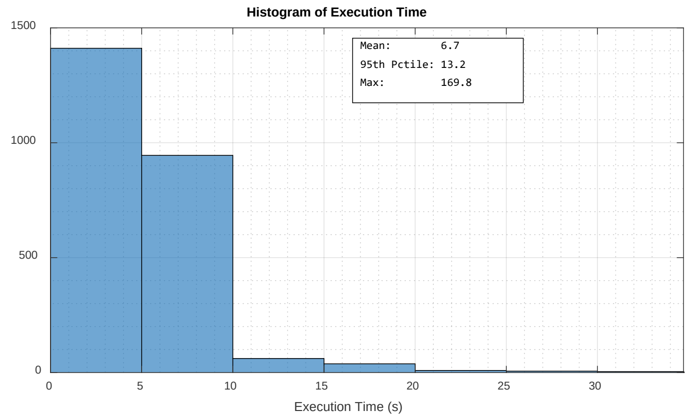

Figure 96. Execution time histogram of the Gen D algorithm in TRL 5 software.

This shows the x-axis limited to remove outliers.

8.4.4 Example: Detailed Analysis

An example of a Gen D flight is shown in Figure 97. Note that the bottom plot in the figure shows the difference between the simulated time of the Legacy and the CLEEN flights at each distance. The CLEEN flight saves fuel by continually varying altitude and airspeed to meet the RTA, rather than a constant speed with a pre-defined altitude profile. Through this method, the CLEEN technology takes advantage of favorable environmental factors to meet the RTA.

The same benefit-extracting path modifications from previous generations can be seen above; that is, the elongated climb path from Gen A and the speed profile and steps determined by Gen B and Gen C lead to cost savings. In addition, time now has a significant effect on the final trajectory. The bottom plot shows how the CLEEN flight time typically compares to the Legacy flight time throughout the flight. The CLEEN flight has a slower climb time and therefore must make up the time through the cruise phase of flight such that the RTA is achieved. Once passing the RTA distance (the vertical dashed line), the CLEEN flight is no longer constrained to a time and can fly whatever is most optimal for the remainder of the flight. Overall, an optimal profile that meets the RTA is achieved.

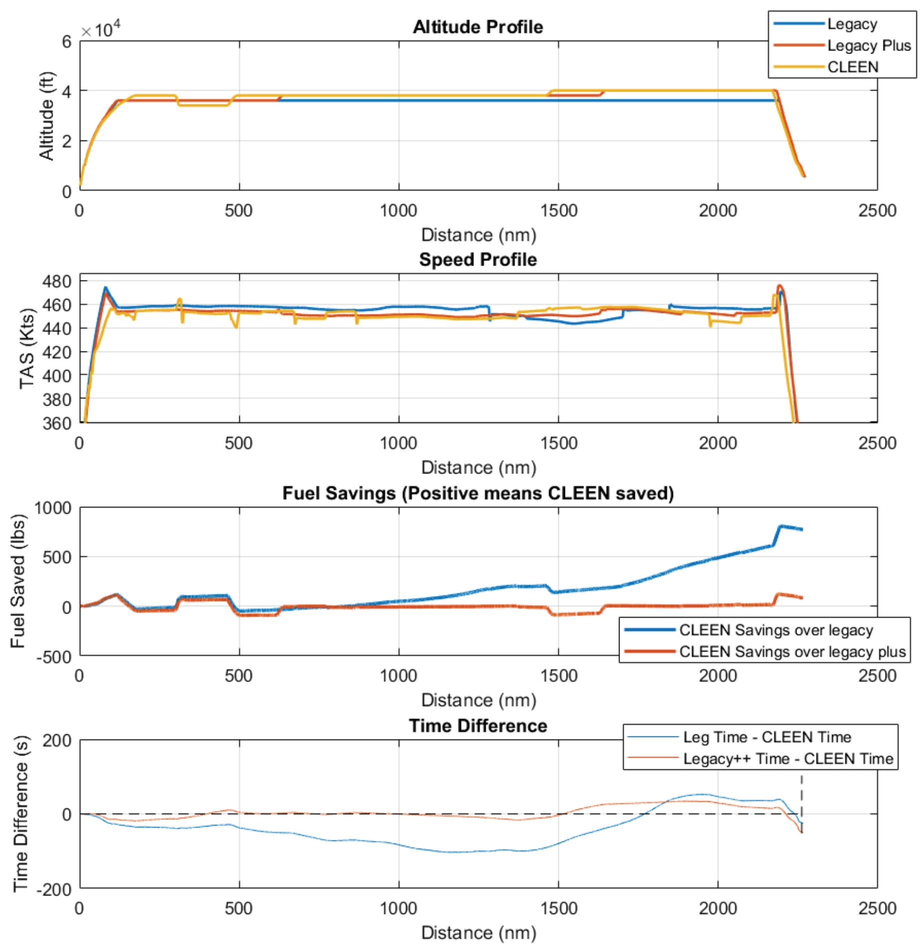

Figure 97. Example RTA flight with Cost Index 0 and RTA location at the end of cruise.

8.4.5 Sensitivity Study

The Gen C benefit assessment reveals that shorter and longer routes saved considerably more cost over Legacy than the middle-distance routes. For the Gen D assessment, these route-based increased savings are only present for longer routes over Legacy; shorter routes savings are varied over both Legacy and Legacy++. The significant cost savings observed over Legacy in longer routes are not attained over Legacy++. Since the Legacy++ technology includes step climbs, the margin of benefit for longer flights is reduced. This trend is further supported by the route-based savings for each aircraft shown in Figure 98.

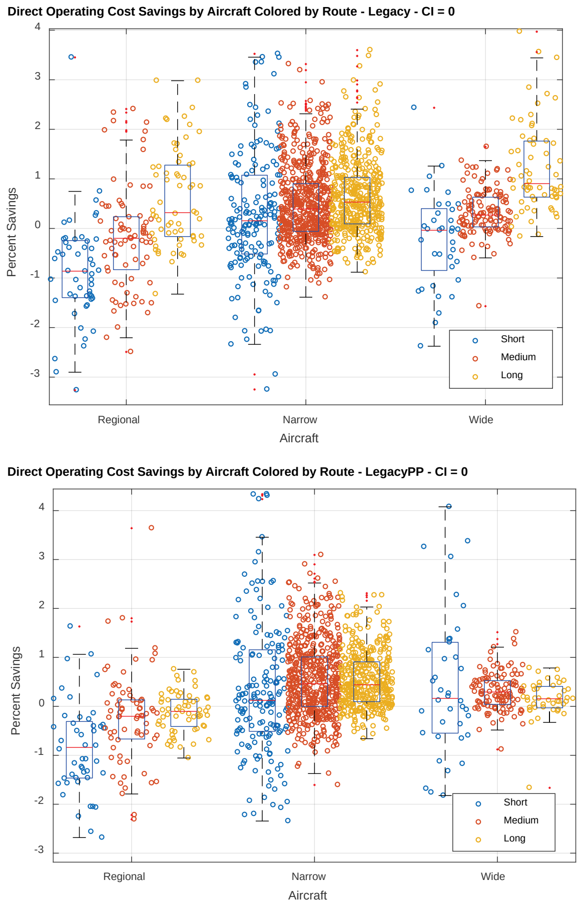

Figure 98. Route-based savings by aircraft.

Gen D performs better for longer routes over Legacy, but this trend does not hold for Legacy++.

As discussed in paragraph B.2.1.6, the Gen D assessment adds the complication of a relatively equal number of fast and slow RTAs with varying distributions at two separate distinct locations. The cost savings results for each flight varies based off the RTA location and what time the RTA is relative to the nominal flight time. The top and bottom histograms shown in Figure 99 depict the savings for the RTA waypoint located at the end of cruise and end of descent, respectively.

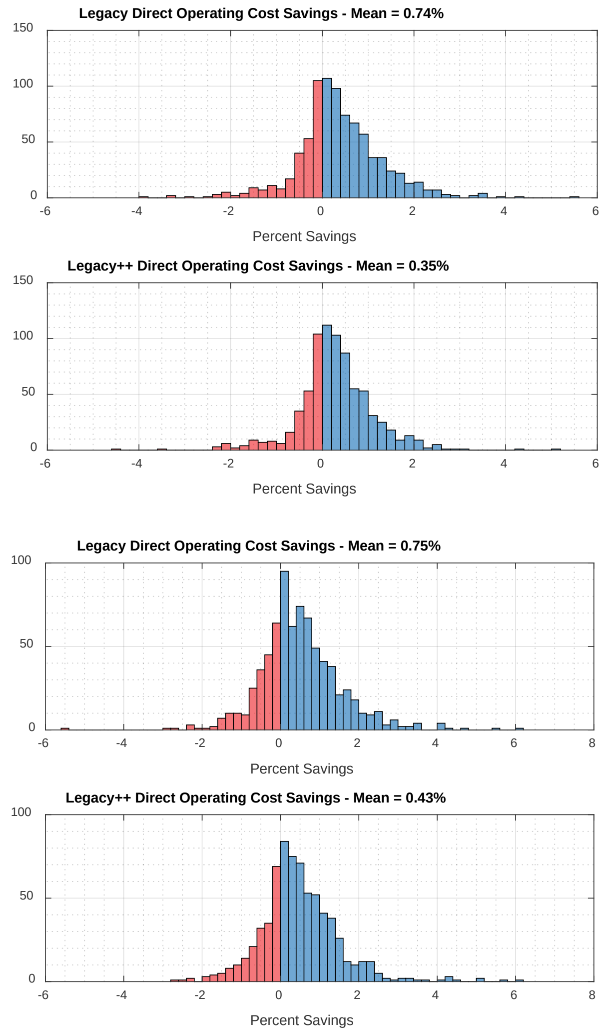

Figure 99. RTA waypoint savings.

End of cruise (top) and end of descent (bottom) RTA location-based savings show no correlation between performance and RTA location.

When compared to Legacy flights (shown on top), the savings for end of cruise and end of descent RTA location are 0.74% and 0.75% respectively; Legacy++ flights saved 0.35% and 0.43% respectively. The data shows that CLEEN technologies generally perform better when compared to Legacy functionality for RTAs located at the end of descent.

Figure 100 separates the results according to fast and slow RTAs. Recall that a fast RTA requires faster flight (less time enroute to the RTA waypoint) than the nominal flight time and vice versa for a slow RTA.

Figure 100. Slow (top) and fast (bottom) RTA based savings for CI 0.

These histograms show increased benefit for fast RTA against Legacy.

When compared against Legacy flights, the cost savings for slow and fast RTAs for Cost Index 0 is 0.62% and 0.86% respectively. The CLEEN technologies for fast RTAs perform significantly better than Legacy flights. While not as substantial, the trend is strong and opposite when compared to Legacy++ flights, where slow and fast RTAs save 0.41% and 0.36% respectively.

8.4.6 Ancillary Benefits

In addition to the significant cost savings extracted by the CLEEN technologies, fewer re-predicts are also required to meet the RTA. This means that the original profile prediction/estimation of Gen D is more accurate than that of both Legacy technologies. Figure 101 displays various re-predict and RTA metrics from the benefit assessment results.

	Legacy	Legacy++	CLEEN
Number of Flights	2881	2870	2893
RTA Met	2682 (93%)	2556 (89%)	2677 (93%)
Number of Re-predicts (For cases that met RTA)	2964 (~1.1 per flight)	2783 (~1.1 per flight)	1111 (~0.4 per flight)

Figure 101. Ancillary benefits.

CLEEN re-calculates the control trajectory to meet the RTA 2.6x more often than Legacy and Legacy++ algorithms.

CLEEN flights performed significantly better in terms of number of re-predicts needed to meet the RTA, as a re-predict occurred around less than half of the number of times compared to both Legacy technologies. The rate of 1111 re-predicts over a total of 2677 flights results in about a 2.7 times reduction over the Legacy and Legacy++ re-predict rates.

CLEEN and Legacy performed similarly in terms of meeting the RTA but Legacy++ flights performed noticeably worse. The altitude profile for Legacy++ flights is determined outside of the RTA speed calculation. Higher altitudes give access to a smaller speed envelope meaning that a smaller range of RTAs can be achieved compared to flying at a constant lower altitude; recall that the RTA range is determined by Legacy ETA extremes. Due to this effect, Legacy++ flights have trouble meeting some of the more extreme RTA cases. This effect is shown in Figure 102.

Note that the plot represents only flights that ran without error. The plot shows how Legacy++ (middle bar chart) meets significantly fewer RTAs than CLEEN and Legacy flights as the RTA becomes more extreme.

Additionally, while not tested as part of this benefit assessment, CLEEN technologies can meet a wider range of RTAs than Legacy technologies due to having the freedom to step up and down while utilizing the full speed envelope. This effect is analyzed in TRL 4 in paragraph 8.3.3.

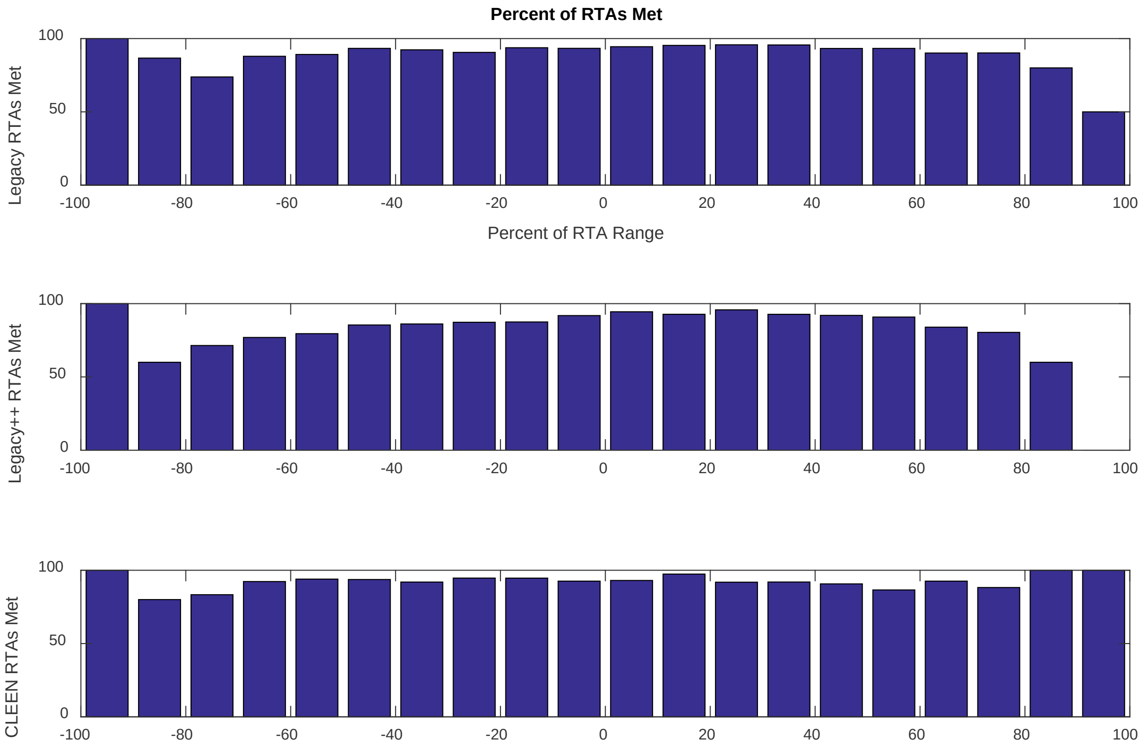

Figure 102. Percent RTA met for each technology.

9. TRL 6 Real Time System Implementation

9.1 TRL 6 System Design

The TRL 6 system implementing the CLEEN technologies consists of a Flight Management System (FMS) and an Electronic Flight Bag (EFB). The FMS is the flight hardware that executes the modified Operational Flight Plan (OFP) software, responsible for creating a flight plan and producing guidance outputs. The EFB is a tablet executing the CLEEN optimizer. Running the optimizer on the tablet provides the greater hardware performance required to generate an optimal control within an acceptable amount of time. Having the optimizer outside of the OFP software also provides the ability to have the software at a lower Design Assurance Level (DAL) that leads to lower cost and more rapid development and updates.

When deployed in the field, the system will be as shown in Figure 103. The FMS (top left) executes the OFP software that follows the optimal controls generated by the optimizer in the EFB (bottom left), by transmitting them as commands to the other flight hardware via ARINC 429. The EFB communicates to the FMS wirelessly through the Aircraft Interface Device (AID, middle), where the protocol is converted to ARINC 429 and routed through the Communications Management Unit (CMU, top right) to the FMS. The lab environment that simulates the field equipment shown is known as FM Workstation and is used for verification and validation of GE's production flight management systems.

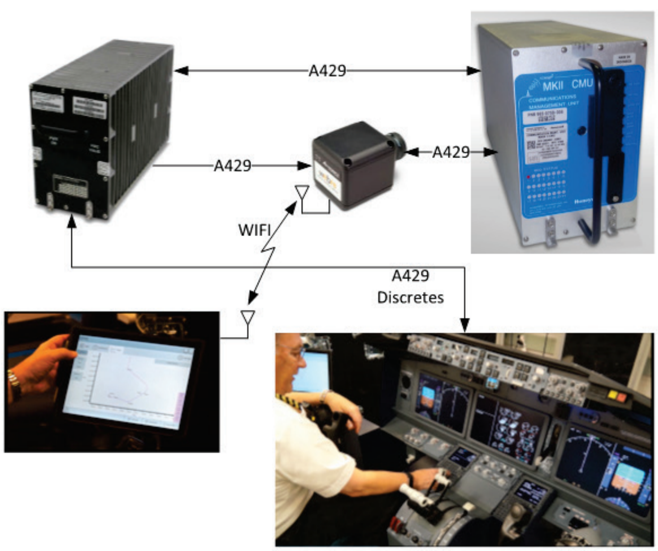

Figure 103. TRL 6 system design.

Figure 104 shows normal operation of the TRL 6 system. Operation begins with the pilot enabling pairing in both the EFB and the FMS. The EFB displays a key that is used to encrypt the communication between the units. After the pilot enters the key into the FMS, optimization can be enabled, and the FMS sends the flight plan and state data required for optimization to the EFB. Depending on the optimization type selected in the EFB, (Gen B, C, or D) the TRL 5 optimizer will generate a profile and compare that to the Legacy profile to produce a displayed estimate of the cost savings of the optimized profile over the Legacy.

To provide the best weather to be used in the lower fidelity model within the FMS, a weather augmenter processes the high-fidelity weather data in the EFB to produce a best-fit representation of the data using the FMS's lower fidelity model. The weather augmentation process is based on "Flight Management System Weather Input Optimizer Final Report" [***].

If the estimated savings are desirable, the pilot can have the optimal profile, control, and augmented weather data sent to the FMS for review before accepting it in the higher DAL system. The software used to establish the link and communicate between the FMS and EFB is the Connected FMS Software Development Kit (CFMS SDK). It provides an API that handles the pairing, encryption, and communication protocol used.

Figure 105 shows the closed-loop guidance of the system, including the guidance when under an RTA constraint. As the aircraft travels along the predicted path using the optimal control, it will deviate from the predicted time of arrival due to errors propagated as it flies along the path. When the FMS performs re-predicts during the flight, the optimal control is adjusted ensuring that the RTA is met. The adjusted control is then used by the guidance to issue targets that the auto pilot will track to control the plane. To inform the pilot of the current savings following the adjusted controls, the controls are fed back to the EFB and used to calculate the savings of the controls that are currently being used by the FMS. This allows the pilot the ability to check the savings of the current controls against controls generated by a new optimization run.

[***] GE Aviation Systems LLC, "Flight Management System Weather Input Optimizer Final Report", FAA CLEEN, 2013.

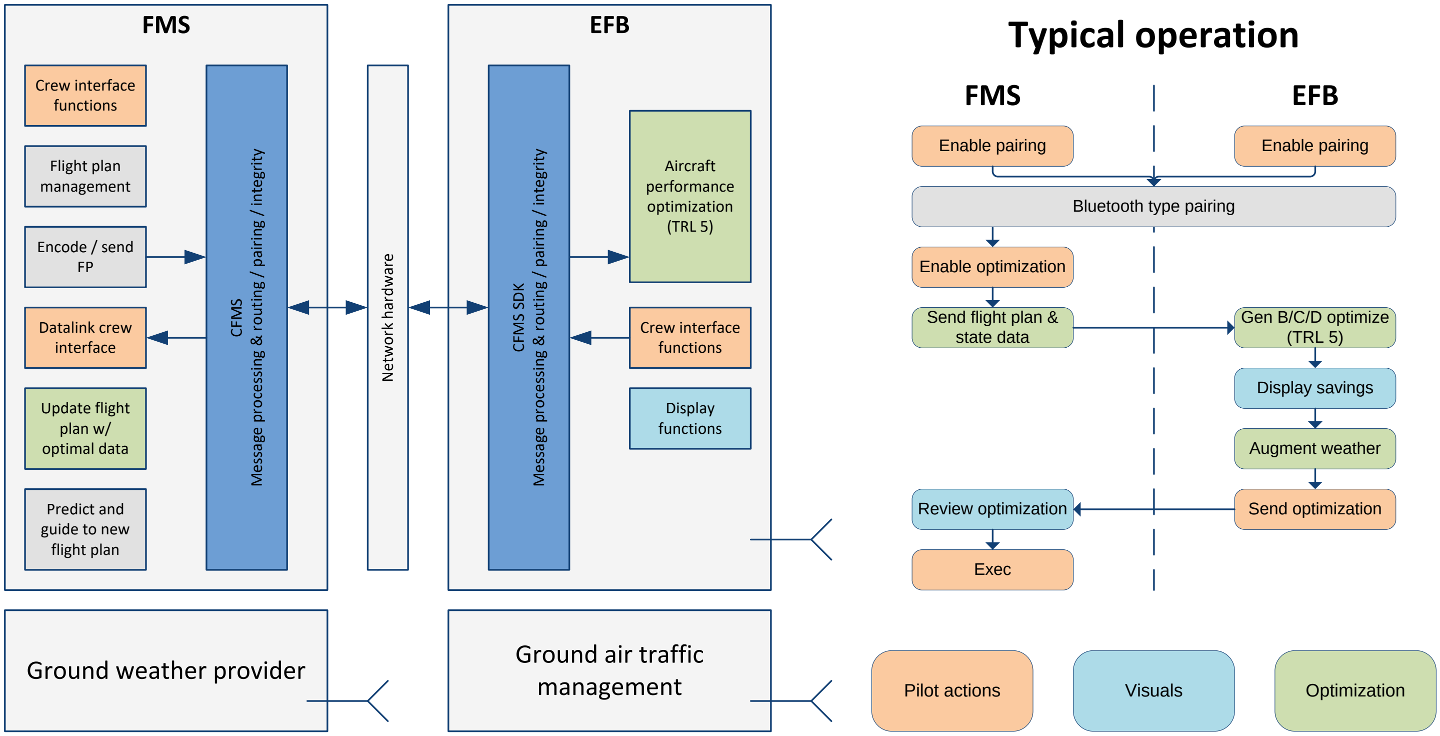

Figure 104. TRL 6 system operation diagram design.

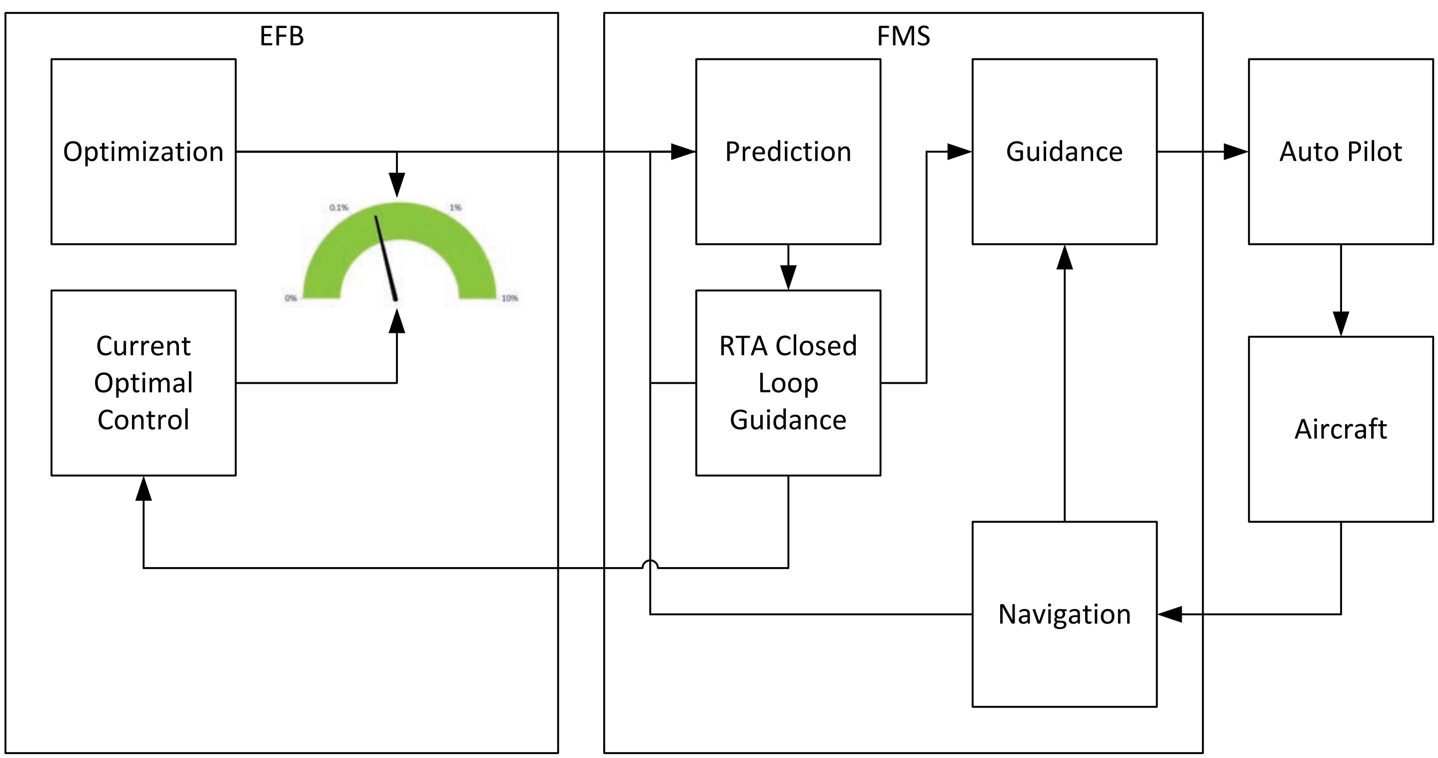

Figure 105. RTA closed loop guidance.

9.2 TRL 6 Benefit Assessment

9.2.1 Methodology

A fundamental chicken or the egg problem arises when considering benefit assessment for performance optimization technologies. The benefit of these technologies must be proven before expending significant effort to produce a flight-worthy certified system, and the most accurate way to assess the benefit is to collect data from a fielded fleet of aircraft.

A solution to the problem of demonstrating benefit would be to perform a small number of controlled flight tests comparing outputs of pre-production versions of the new technology against the Legacy production system. In these tests, the aircraft is equipped with an advanced sensor package and all variables are held at prescribed values during the test interval (on the order of minutes). In engine testing for fuel efficiency, even the physical position of crew members aboard the aircraft is fixed to ensure the aircraft center of gravity doesn't shift during the test. Despite efforts to control the aircraft environment, small variations arise due to factors outside of crew control (wind gusts, turbulence, etc.). The uncertainty in test results caused by these effects is further compounded by sensor measurement errors, making it very difficult to measure a benefit around 1% with statistical significance.

The aircraft trajectory optimization benefit assessment cannot be performed over a short, highly controlled experiment. To measure the benefits of this technology, one must look at an entire flight, end-to-end. The experiment must be statistically significant over conditions relevant to a fleet of aircraft for their service life.

The total fuel burn and time of arrival (the direct operating costs) are dependent on environmental conditions and aircraft performance. In this type of experiment, issues arise when attempting to duplicate flight conditions between two flights.

In an experiment, one could fly the same aircraft with the same weight and balance, flying the same filed flight plan between the same two cities within a window of several hours. Even in this scenario, weather patterns vary, in addition to discrepancies due to the timing of air traffic control clearances. Another option would be to fly two aircraft (one with new technology and one with old) simultaneously along the same route. In this hypothetical experiment, slight differences between engine performance and aerodynamics would introduce fuel and time differences between the flights.

Statistical hypothesis testing may be used to estimate the number of flight trials (observations) required to demonstrate fuel savings. The concept is to use an estimator for fuel burn and pose the detection of fuel burn improvement as a statistical hypothesis testing problem. If there is fuel burn reduction with the improved FMS technology (relative to the Legacy), the estimator applied to a set of flight trials with the improved FMS will have a different mean error than the estimator applied to a set of flight trials with the Legacy FMS.

The problem may be formulated as follows: let μ_0 be the mean of the estimator error when applied to trials with the Legacy FMS and μ_1 be the mean of the estimator error when applied to trials with the GE FMS. Then, to show at least δ benefit, the following null and alternate hypotheses may be formulated:

- $H_0\colon\ \mu_1 < \mu_0 + \delta$
- $H_1\colon\ \mu_1 \geq \mu_0 + \delta$

In words, the null hypothesis is that at least δ amount of cost isn't saved and the alternate hypothesis is that at least δ amount of cost is saved. Several assumptions are made:

- The estimator error is normally distributed
- The standard deviation of the estimator error, σ_e, is known and is the same when the estimator is applied to either set of trials
- The observed difference between the means is: $D \coloneqq \mu_1 - \mu_0 > \delta$

Note that the D must always be strictly greater than δ; that is, it can never be shown that more cost is saved than was observed. Moreover, the closer δ is to D the more trials that will be required.

Currently, the estimation methods are accurate within 0.4% - 1.0% standard deviation, including modelling and sensor measurement errors. We optimistically assume the $\sigma_e \coloneqq 0.4\ \%$ standard deviation case for the remainder of this analysis.

TRL 5 Benefit Assessment of the Generation C (non-RTA) optimization technology shows a mean of 1.00% fuel savings over the Legacy system (see paragraph 7.4.1). Further processing of the fuel savings data reveals that the savings distribution has a standard deviation of 0.82%. Estimation methods suggest errors in estimation of tail-specific performance variation of 0.3%. Statistically combining error sources and reframing into the hypothesis testing framework results in a population mean of $\mu_D \coloneqq 1.00\%$, and standard deviation of the $\sigma_D \coloneqq \sqrt{0.82^2 + 0.3^2} = \ 0.87\%$ of the experimental results population D.

In a flight test experiment, consider each flight trial as random draws from the distribution described above. As such, there will be flights that exceed the mean savings, and some that do not. There may even be observations where the new technology is costlier than the Legacy system. Figure 106 shows that the probability that the experiment is successful in proving a benefit at least as large as δ given a specified number of flights, defined as showing a 0.75% benefit with statistical significance (simply chosen as three quarters of the estimated benefit from simulations). The figure shows that on the order of tens to a hundred flight trials are required to achieve above a 50% chance of success of such an experiment and up to a thousand flight trials will not improve the success rate above 60%. This large number of trials is an economically infeasible proposition. To perform a flight trial of this scale, the trials would need to be performed as part of normal commercial revenue flight, which leads back to the original problem of incurring cost to produce a certified system without proof of benefit.

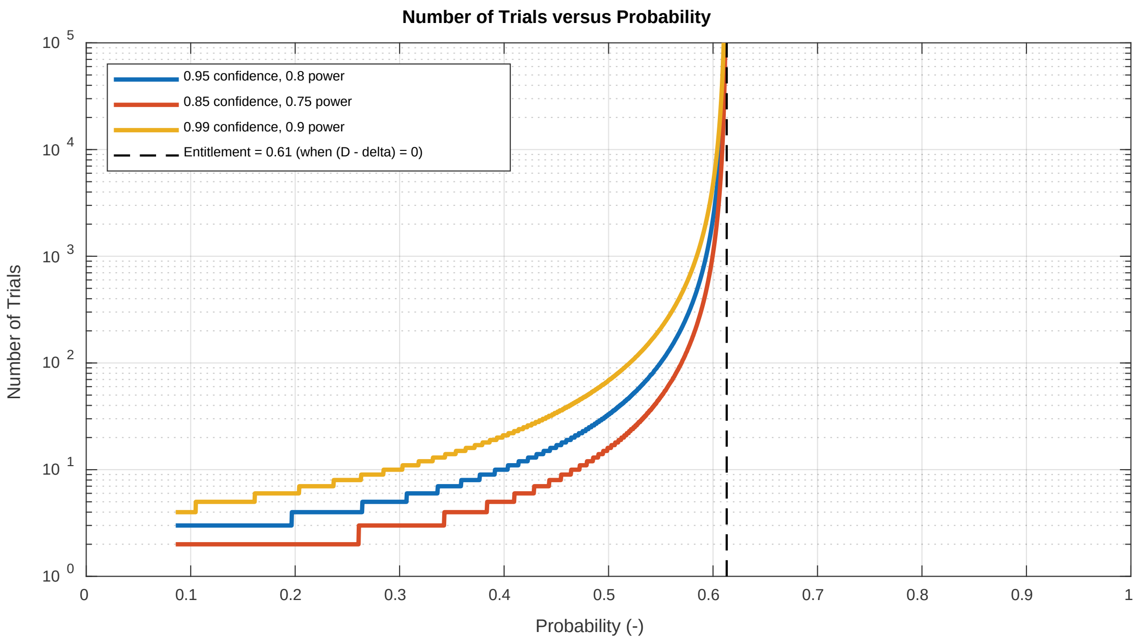

Figure 106. Likelihood of statistical significance.

Based on this analysis, the TRL 5 fleet-wide benefits presented in paragraphs 7.4 and 8.4 for Gen C (no RTA) and Gen D (with RTA), respectively, provide the most realistic estimate of DOC savings. To validate these savings, a benefit assessment in the TRL 6 system is performed. The TRL 6 Benefit Assessment is run on a small subset of cases from the TRL 5 Benefit Assessment to ensure that the results from the large-scale Monte Carlo assessment hold without any unexpected deviations due to real-time processing aspects. See paragraph B.3 for a full description of the testing methodology.

9.2.2 Results

A comparison of Gen C benefit cases between TRL 6 and TRL 5 is presented in Figure 107. In the figure, the bars show the Direct Operating Cost savings achieved by the CLEEN method over the Legacy method for the given case, in pounds (both on the y-axis and annotated in each bar). Note that, in cases where the Cost Index is non-zero, the savings achieved from the cost of time is converted to the equivalent pounds using the Cost Index and added to the fuel savings, resulting in a Direct Operating Cost in pounds that includes the cost of time. Each case number is noted on the x-axis; this number is an identifier for each case and has no intrinsic meaning. Important input variables and observed cost savings for each case are detailed in Figure 108.

Six of the nine TRL 6 runs exhibited cost savings greater than those of the TRL 5 runs. The three cases that saved less than their TRL 5 counterparts all were within 15 pounds of fuel (or equivalent, as defined by the Cost Index) and are considered very close matches. In total, six of the nine flights matched very closely between TRL 5 and TRL 6 (cases 137, 222, 931, 1318, 2723, and 3885). The three flights that were not close matches showed TRL 6 savings that exceeded 15 pounds more TRL 5 (see cases 1477, 1969, and 2287). This is indicative of either minor improvements made in the development of the TRL 6 system or of the CLEEN control profile performing better in the dynamic TRL 6 simulator than in the TRL 5 one. The mean absolute error over all nine cases is 17.2 pounds.

One case (2287) exhibits a much larger difference between TRL 5 and 6 than any of the other cases. Investigation revealed that this is caused by the simulated guidance in the descent phase of flight; namely, thrust reversion logic in the lower TRL system. For this case, the TRL 5 CLEEN method performed more thrust reversions during descent than the TRL 5 Legacy method, in order to track the descent speed. This caused a larger fuel burn, and thus lower savings numbers. Comparing this case at top of descent instead of bottom (excluding the difference caused by this reversion) resulted in a close match of TRL 6 and TRL 5 – about 10 pounds. The reason for the initial discrepancy is described more in paragraph B.3.2.

This analysis concludes that the CLEEN software is sufficiently consistent between TRL 5 and TRL 6 implementations to validate the TRL 5 fleet-wide benefit assessment results.

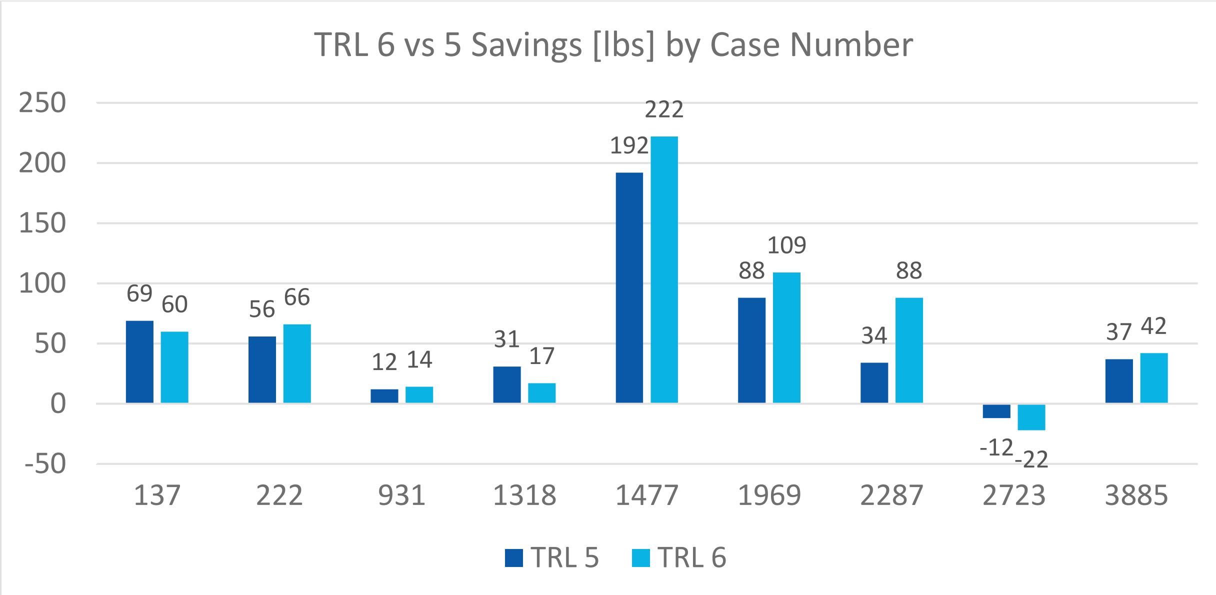

Figure 107. Chart comparing cost savings (in lbs.) between TRL 5 and 6 for a subset of cases.

Case #	Route (nmi)	Weight (klbs)	Cost Index	Altitude Sep. (ft)	Weather Day	TRL 5 Cost Savings (%, lbs.)		TRL 6 Cost Savings (%, lbs.)	
137	300	113.8	0	1000	10/28	1.46%	69	1.32%	60
222	800	135.6	0	2000	4/10 (incl)	0.52%	56	0.61%	66
931	1000	137.4	25	1000	8/19	0.05%	12	0.06%	14
1318	800	130.9	0	2000	3/8	0.40%	31	0.22%	17
1477	2500	147.0	0	1000	2/19	0.80%	192	0.93%	222
1969	1000	138.5	0	1000	12/9	0.58%	88	0.71%	109
2287	1500	134.1	25	1000	6/23	0.12%	34	0.31%	88
2723	600	118.7	25	1000	8/13	-0.11%	-12	-0.19%	-22
3885	500	126.0	0	1000	4/15 (incl)	0.63%	37	0.72%	42

Figure 108. Case parameters and cost savings used in producing TRL 5 to 6 comparison.

9.2.3 Example TRL 6 Result

In this section, an example case (number 1969) from the Gen C TRL 6 testing is analyzed and compared to the corresponding results from TRL 5 testing. Figure 108 shows the relevant parameters for the example flight shown, as well as the cost savings the CLEEN flight provides over the Legacy flight in TRL 5 and TRL 6, for comparison. Each plot depicts four flights, CLEEN and Legacy for both TRL 5 and 6. For clarity, the CLEEN flights are colored red and orange and the Legacy flights with blue and teal. Dashed lines indicate predicted values, and solid lines indicate actual values.

The altitude profiles are compared in Figure 109. As expected, the Legacy flight in TRL 5 and 6 have the same cruise altitude, and the CLEEN flights have the same cruise altitude step locations. During the last portion of the climb phase for the CLEEN flight, the TRL 5 flight more closely follows the initial predicted altitude-distance profile than the TRL 6 one does. This is due to the more accurate weather in the prediction of TRL 5, which is explained more fully in paragraph B.3.2.

The descent paths in TRL 5 (for CLEEN and Legacy) are nearly identical, and the TRL 6 descent paths are likewise nearly identical, but they do not match each other exactly – the TRL 6 methods deviate from their initial prediction and begin their final descent earlier, whereas the TRL 5 simulator is designed to follow the prediction exactly and does not deviate from the initial prediction. This is not concerning, as the TRL 6 simulator has a more robust descent system, and furthermore any fuel saving effects caused by this would cancel out, as the Legacy and CLEEN perform the same descent path.

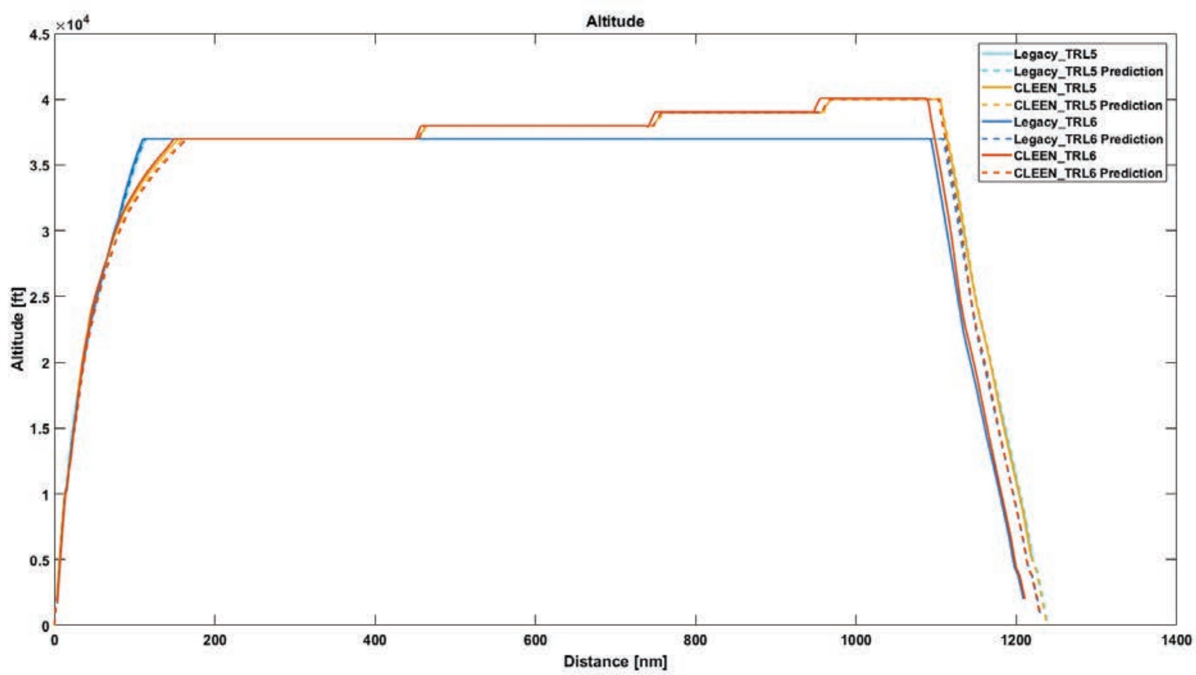

Figure 109. Altitude comparison of TRL 5 and 6 example case.

Figure 110 through Figure 112 compare the lateral track of the case that reveal that the lateral profile is nearly identical in all methods. A slight difference between the profiles' turn radii can be observed in Figure 111 and Figure 112 that is caused by small foreseeable speed variations during climb and descent. In addition to small differences in weather models between the systems, the TRL 5 simulation uses a simulated version of the high-fidelity lateral function that exists in the TRL 6 simulation.

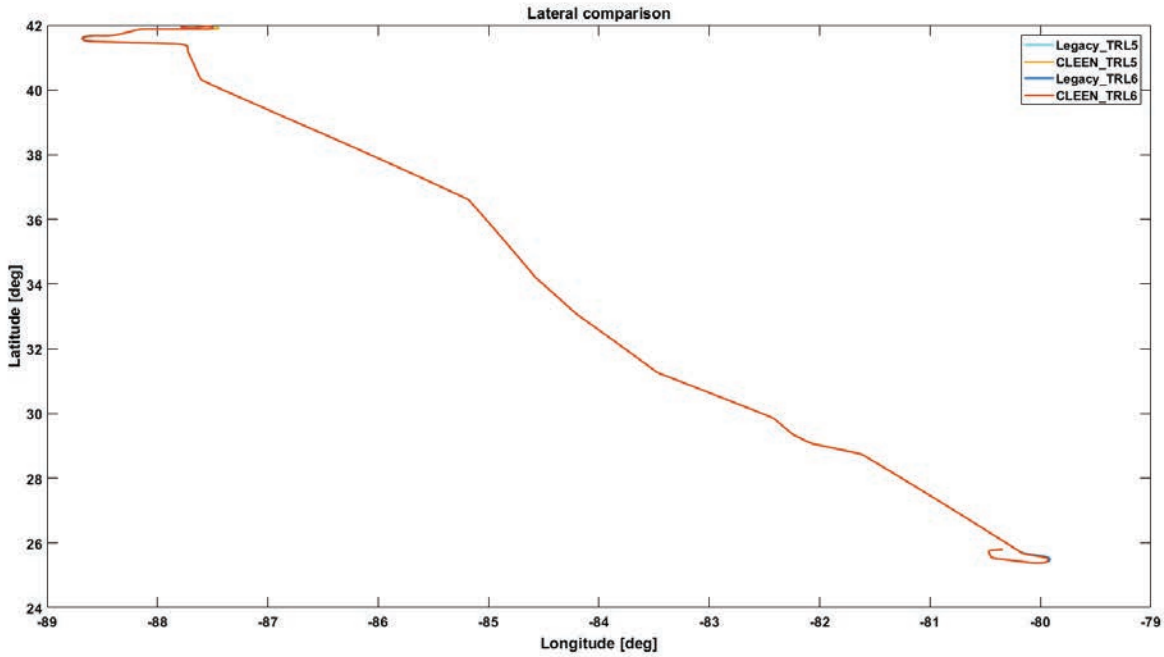

Figure 110. Lateral comparison of TRL 5 and 6.
The example case shows nearly identical match.

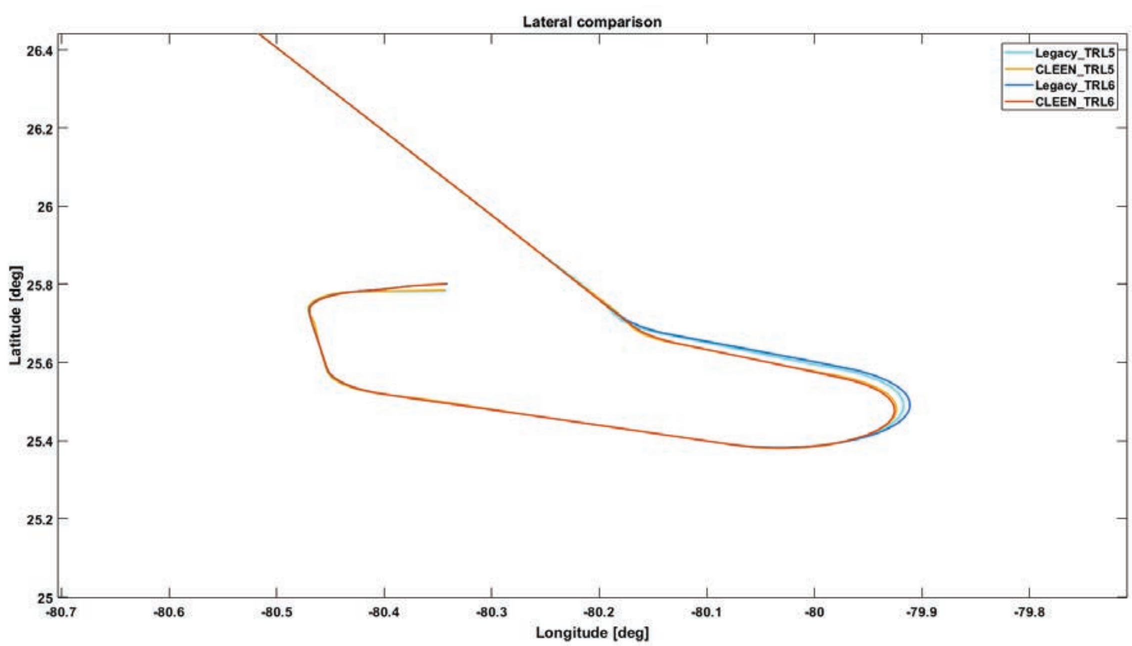

Figure 111. Lateral comparison (zoomed in on departure) of TRL 5 and 6 example case.
Differences in lateral path curve radius due to difference in commanded airspeed.

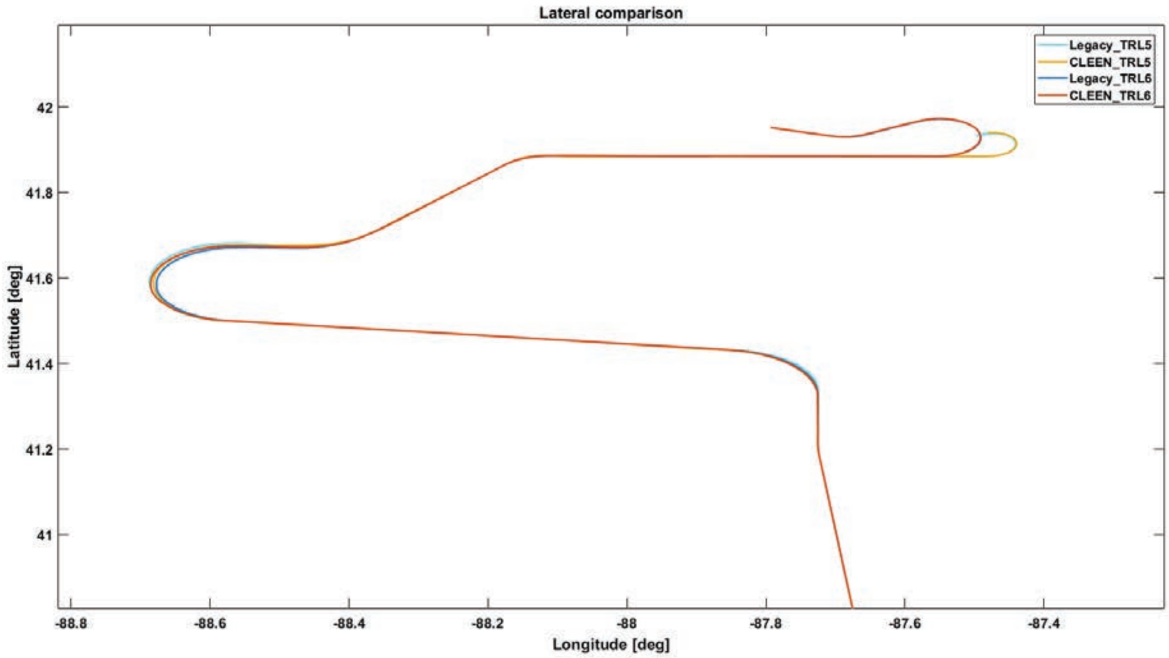

Figure 112. Lateral comparison (zoomed in on arrival) of TRL 5 and 6 example case.

Figure 113 compares the speed profiles of each method for this example case, in CAS, TAS, and Mach versus distance. The CAS profiles show that the CLEEN Climb Optimizer is producing a variable-speed climb profile that is being tracked by the simulation. The Mach plot is zoomed in to more clearly show the cruise phase. Observe that the Legacy profiles (in blue shades) in TRL 5 and 6 have a similar Mach, as do the CLEEN profiles.

Figure 114 compares the thrust and drag. The CLEEN profiles show a lower thrust at the end of the climb phase (between 50 and 150 nmi) than the max-thrust Legacy profiles, indicating the proper execution of the variable-thrust profile that the Climb Optimizer generates. Small dips and bumps in the thrust are shown during the cruise phase, occurring whenever the CLEEN profiles perform a step climb or step descent, or when either method performs speed changes. The drag plots show erratic drag during climb and descent, as is to be expected with the changing angle of attack during those phases.

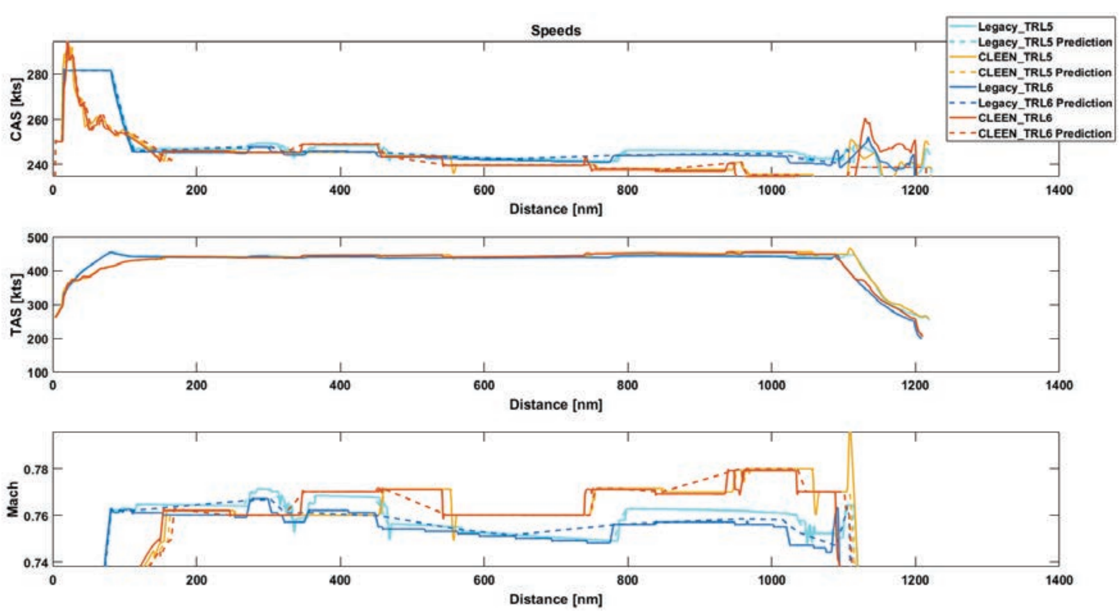

Figure 113. Speed comparison of TRL 5 and 6 example case.

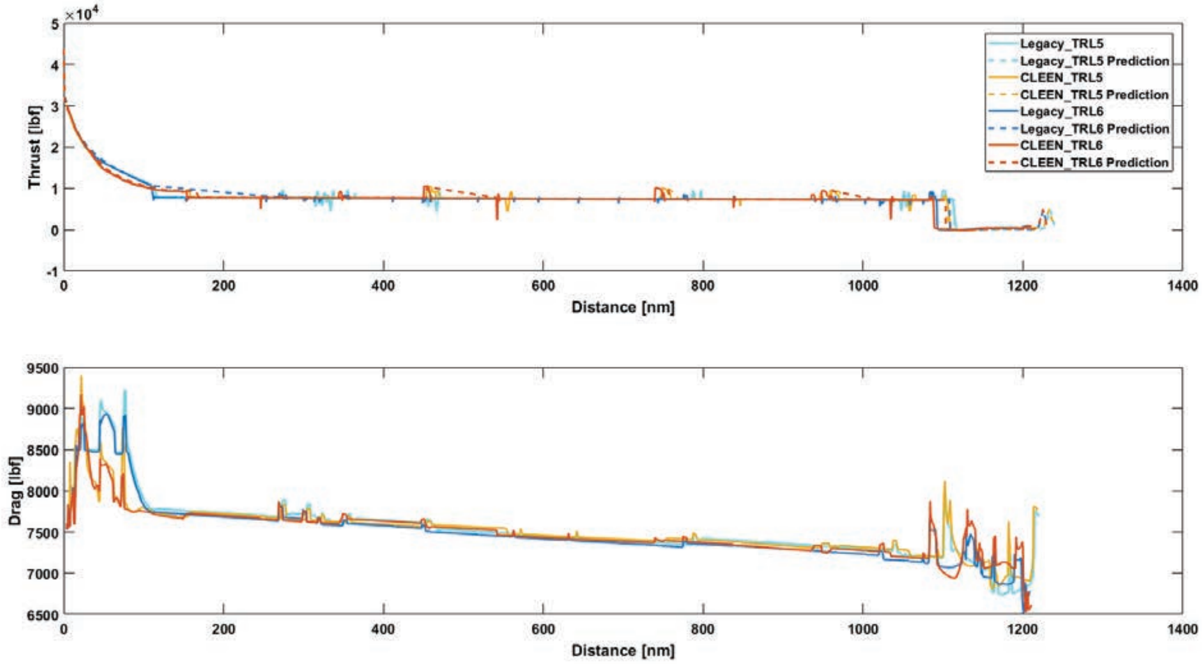

Figure 114. Thrust and drag comparison of TRL 5 and 6 example case.

Figure 115 shows the weather data for the case, both simulated (solid lines) and predicted (dashed lines). The CLEEN TRL 5 test has the most accurate predicted weather when compared with its actual, simulated weather. The reason for this is described in detail in paragraph B.3.2. The weather used in both simulations – the "actual" or experienced weather in the test flight – are consistent between TRL 5 and 6. This is appropriate given that all flights are following the same lateral track and have similar altitude profiles.

Finally, Figure 116 compares the aircraft weights, fuel flow, and fuel savings. For this case, the Cost Index is 0, meaning that the cost savings are entirely comprised of fuel savings, and time has no cost. The top two plots show that all four aircraft began at the same weight and the last plot shows fuel flow during the cruise phase that tracks well with the thrust (shown previously).

The Fuel Saved plot (second from bottom) shows an accumulation of fuel saved, measured as the difference in total fuel burned between each flight and the CLEEN TRL 6 flight. The dark blue line represents the Legacy TRL 6 flight and serves as the best measure of the benefit achievable with the CLEEN technologies. The annotated point at the right end indicates that 109 pounds of fuel were saved.

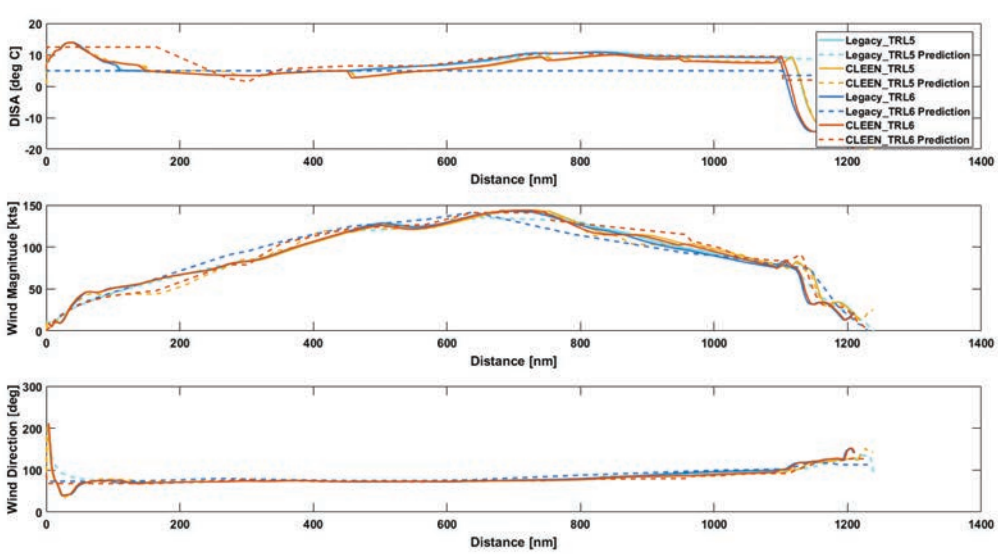

Figure 115. Weather comparison of TRL 5 and 6 example case.

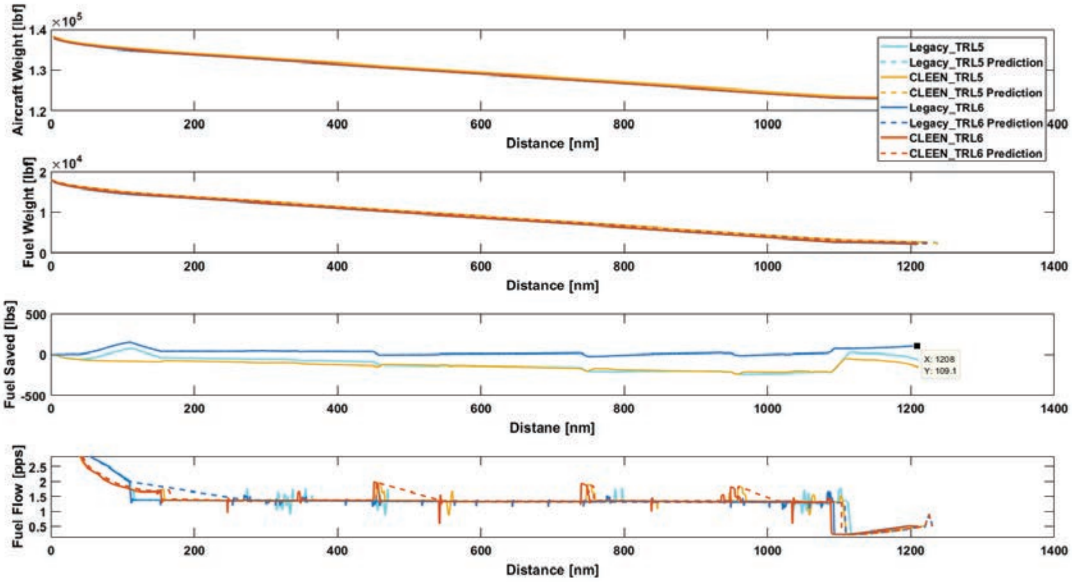

Figure 116. Weight and fuel comparison of TRL 5 and 6 example case.

10. Conclusion and Recommendations

In conclusion, large, fleet-wide simulation of the GE CLEEN II technology demonstrates significant fuel burn reduction over both the Legacy system in service today and a representation of best-in-class ground tools. Beyond GE's work on CLEEN and CLEEN II, the field of flight path optimization is crowded with many similar technologies and tools proposed by various industry and academic teams. Each team has their own baseline for assessment leading to operator confusion and often over-inflated benefits metrics. GE proposes the method described within this document as a standardized means for assessment to provide an "apples-to-apples" comparison removing uncontrollable effects such as air traffic control vectoring, the effect of weather forecast inaccuracy, and the quality of vehicle model, and compares against the best a legacy system can do in the absence of external errors. Each of these external effects is important and can be improved, but ultimately are outside of the control of a flight path optimization algorithm.

As flight optimization tools continue to mature and enter service, GE recommends further engagement with airline operators and OEMs in partnership with the FAA to promote the benefits of the technology and ensure minimal operational and training impact. To facilitate a smooth transition to new optimal vertical control policies, GE recommends deployment of high-accuracy 4D physics-based prediction capabilities (such as those deployed on the Electronic Flight Bag in this project) within Airline Operation Center (AOC) ground tools and Air Traffic Control (ATC). With these capabilities mirrored, the airborne and ground infrastructure will "see" the same picture of the overall airspace; providing a means of trajectory synchronization with reduced trajectory negotiation.

GE recommends continued development to exploit the benefits of the flight path optimization tools described in this report to include additional control parameters (such as lateral path planning) and cost functions (such as community noise, aircraft scheduling, traffic, throughput, etc.). This work should be performed in conjunction with advancing core optimization technology for run-time and memory improvements that allow for practical implementation of optimization.

GE recommends the work performed for the CLEEN II program is complimented with future work in airspace optimization, improved weather forecasting, and adaptive modelling technology to provide better inputs to the optimization engine, and thus produce better results with lower variation in benefit.

Appendix A Summary of Benefits Metrics

A Benefits Metrics

% Cost Savings	Gen B – Constant Altitude / HiFi Weather		Gen C – Cruise Steps / HiFi Weather		Gen D – Cruise Steps / HiFi Weather with RTA	
Cost Index	0	25	0	25	0	25
Legacy	0.53%	0.44%	1.00%	0.89%	0.74%	0.54%
Legacy+	0.43%	0.36%	N/A	N/A	N/A	N/A
Legacy++	N/A	N/A	0.39%	0.37%	0.47%	0.20%

Figure 117. Summary of benefits including negative savings cases.

% Cost Savings (Operational)	Gen B – Constant Altitude / HiFi Weather		Gen C – Cruise Steps / HiFi Weather		Gen D – Cruise Steps / HiFi Weather with RTA	
Cost Index	0	25	0	25	0	25
Legacy	0.54%	0.46%	1.02%	0.91%	0.81%	0.60%
Legacy+	0.44%	0.37%	N/A	N/A	N/A	N/A
Legacy++	N/A	N/A	0.40%	0.38%	0.47%	0.28%

Figure 118. Summary of benefits when operationally negative savings are set to zero.

Appendix B Benefit Assessment Methodology

This section describes the benefit assessment methodology introduced in paragraph 3.4.

B Benefit Methodology

B.1 TRL 4

In a TRL 4 environment, all Legacy systems and the proposed generations of CLEEN algorithms are coded in MATLAB along with all inherent limitations imposed by commercial flight (i.e., horizontal and vertical speed restrictions and acceleration limits). Simplified point-mass integrations and entitled instantaneous performance are used to produce a flight profile, following each method's optimal control. Once produced, separate higher-fidelity point-mass integrations are used to accurately simulate the trajectories resulting from the previously calculated control. These simulated trajectories produce the data required to compute the cost of each profile. See paragraph 3.1 for a formulation of the cost of each integration step in climb and cruise. These costs are combined with an estimated cost of descent to compute an overall flight cost.

The MATLAB suite used to simulate the flight resulting from the optimum control is collectively referred to the Cost Calculator (CC), as it simultaneously computes the cost of each flight and stores important information. Each simulated optimal profile is compared to that of the Legacy method, with calculations beginning at the same distance and takeoff gross weight.

The cost calculator approach does not simulate any actual aircraft dynamics, but rather makes use of simplified integrations to determine the fuel consumed by each method's optimal solution. More specifically, the calculations simplify the model of fuel flow in that speed changes are assumed to be instantaneous. This permits the calculation of the fuel flow for one segment of the discretized cruise profile as the optimal fuel flow at the step start multiplied by the ratio of step distance to optimal speed at the step start. This simplification also assumes that slow vehicle dynamics dominate the behavior of the aircraft and the resultant forces and fuel flow.

The TRL 4 Assessment is conducted for each generation of technology using a multitude of test sets. Within each test set is a nominal variation of basic flight parameters, mainly including takeoff weight, flight distance, altitude separation, weather day, and Cost Index. An example variation of these parameters is shown in Figure 119. Note that each test set has a subset of this variation according to the purpose of the test set, and not the complete variation. Furthermore, for each generation, custom-tailored test sets are employed to either exercise a certain aspect of the optimizer or determine the sensitivity of the cost savings to various external factors. These test sets are described more fully in their corresponding sections.

Parameter	Nominal variation in each test set
Takeoff Weight	110 – 170 klbs
Flight Distance	500 – 4000 nmi
Altitude Quantization	1000, 2000 ft
Weather Day	Normal: 06/23/16, 03/08/16, 05/24/16, 10/28/16 Inclement: 3/17/16, 4/15/16
Cost Index	0, 25, 50, 100

Figure 119. Nominal variation of basic flight plan parameters in each test set.

B.2 TRL 5

B.2.1 Monte Carlo

Historically for Part 25 transports, the air vehicle manufacturer designed and specified the performance and management of the aircraft via database tables. These data and algorithms were then delivered to a supplier (such as GE) for inclusion in an embedded FMS. These legacy functions are widely used throughout the commercial airline industry and are included on all aircraft types with advanced flight management systems.

The work of the CLEEN II program sought to create new technologies that utilize data available within the FMS to improve the capability to produce an optimized profile using real-time optimization software rather than a tabulated approach. Proving the viability of this technology over such a wide application range requires a large selection of simulated test flights. As such, test flights are chosen with differing combinations of aircraft type, takeoff gross weight, lateral routes, and other typical operational demands (for instance, flight level must be at 2000ft intervals).

A Monte Carlo method is employed to generate scenarios for the individual test flights, beginning in Generation B (paragraph 6.4) and onward. This random sampling of probabilistic distributions allows for a relatively small number of test flights to be representative of the target fleet, as opposed to an exhaustive set of all combinations of variability.

B.2.1.1 Aircraft

The CLEEN II technology is an algorithm improvement that can be employed on any airframe model. Five different aircraft are chosen for this assessment to be characteristic of the commercial airline fleet in service today. Figure 120 contains an estimated fleet breakdown for the top eight largest passenger airlines of North America in terms of enplaned passengers, fleet size and number of destinations. These fleet estimates classify over 70% as narrow-body, 14% as wide-body, and 13% as regional carriers[†††].

SIZE Rank	Airline	Number Narrow	Number Wide	Number Regional	Total Number	% Narrow	% Wide	% Regional
1	American	720	146	77	943	76%	15%	8%
2	Delta	410	150	272	832	49%	18%	33%
3	Southwest	722	0	0	722	100%	0%	0%
4	United	561	179	0	740	76%	24%	0%
5	Air Canada	72	70	25	167	43%	42%	15%
6	Alaska	218	0	67	285	76%	0%	24%
7	JetBlue	155	0	100	255	61%	0%	39%
8	WestJet	115	4	0	119	97%	3%	0%
	Totals	2973	549	541	4063	73%	14%	13%

Figure 120. Representative fleet airframe classifications.

[†††] GE Aviation Systems LLC, "Flight Management System Weather Input Optimizer Final Report", FAA CLEEN, 2013

Collectively, these three aircraft categories represent the vast majority of all commercial aircraft in operation today. To that point, the Boeing 737 series (a narrow body transport) is the best-selling jet commercial airliner in history[‡‡‡]. There are over 2,000 Boeing 737s airborne at any given time, with one departing or landing somewhere every two seconds[§§§]. The Airbus counterpart to the 737 – the A320 – is ranked as the world's fastest-selling jet airliner family according to records from 2005 to 2007, and as the best-selling single-generation aircraft program (also a narrow body transport)[****][††††]. There are more than 6,000 Boeing 737 and Airbus 320 aircraft in today's service mainline fleets[§§§].

Similarly for the wide bodies, the Boeing 777 has received more orders than any other wide-body airliner with over 1900 ordered and 1500 delivered, and the A330 is not far behind with over 1600 orders placed and over 1300 delivered and in operation[‡‡‡‡].

For this benefit assessment, the Boeing 737-800 and Airbus A320-200 are chosen in the narrow body category, the Boeing 777-300 and Airbus A330-200 as wide bodies, and the Embraer 175 AR as a regional jet. The test set consists of routes flown by aircraft types that match the fleet breakdown percentages of Figure 120, with a slight preference given to the 737 for narrow-body flights based on distribution of these aircraft in service.

An analysis of test results confirms that the airframe distribution matches that described above. Figure 121 depicts the percentage of test runs ascribed to each airframe.

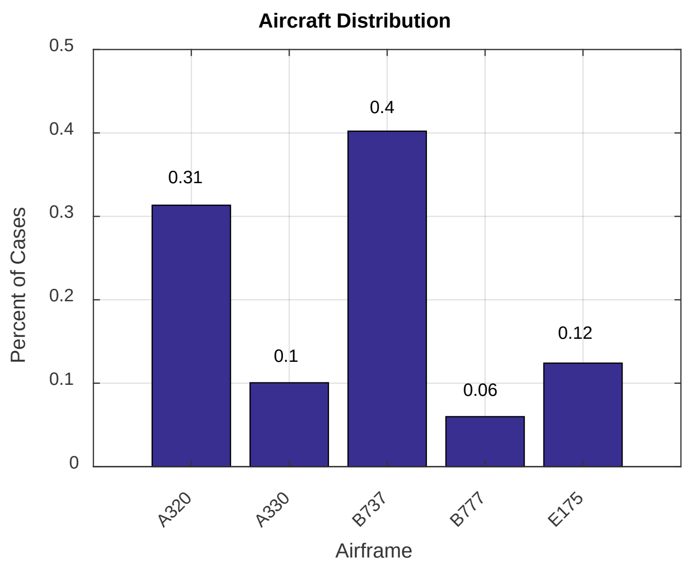

Figure 121. Airframe distribution as percentage of total cases.

Performance and physical characteristics for each aircraft are contained within a Model Engine Database.

‡‡‡ "Plane Spotters", 2019. [Online]. Available: https://www.planespotters.net. [Accessed 10 March 2017].

§§§ M. Kingsley-Jones, "6,000 and counting for Boeing's popular little twinjet", FlightGlobal, 22 April 2009. [Online]. Available: https://www.flightglobal.com/news/articles/pictures-6000-and-counting-for-boeings-popular-little-325472/. [Accessed 10 March 2017].

**** "Boeing 737 Facts", Boeing, April 2014. [Online]. Available: http://www.boeing.com/farnborough2014/pdf/BCA/fct%20-737%20Family%20Facts.pdf. [Accessed 18 October 2018].

†††† "Airbus steals the Paris air show", Hellocompany.org, 19 June 2007. [Online]. [Accessed 7 March 2017].

‡‡‡‡ "Airbus orders and deliveries", Airbus S.A.S., 31 January 2017. [Online]. [Accessed 7 March 2017].

B.2.1.2 Routes

There are over 20,000 flights occurring daily within the United States traversing an immense number of unique routes. World-wide, that estimate grows to over 100,000 flights covering distances as long as 8,600 nautical miles and as short as a few miles.

An analysis of the daily flights by each chosen airframe is conducted to determine what distances are representative of in-service operation. Figure 122 contains a summary of this data. Note that the minimum and maximum distances are not absolute, but rather represent the observed extrema found in the selected date range.

Airframe	Minimum (nmi)	Median (nmi)	Maximum (nmi)
Boeing 737-800	25	1125	3325
Airbus A320-200	25	975	2775
Boeing 777-300	100	3500	10900
Airbus A330-200	50	1750	6650
Embraer 170 AR	25	600	2075

Figure 122. Flight distance statistics.

The final route length distribution in the Monte Carlo analysis is generated to match those observed in the historical data. The test set is discretized into twenty routes, with distances varying from 200 to 7000 nautical miles. Each route is created by selecting origin and destination airports whose separating distances fall within the ranges of Figure 122. Jepsen-documented airways and navigable waypoints are then chosen between airports to elongate paths and create realistic lateral tracks for simulation; several routes also included standard instrument departures and standard terminal arrival routes. Figure 123 depicts the lateral path of several planned routes.

Limited weather data availability demands that all routes be constructed within the United States. In order to prevent the effective cancellation or stabilization of weather, routes are constructed that traverse large latitude and longitude spans. Traveling around large portions of the contiguous US increases the technologies exposure to varying weather types and severity.

Route distances for each test flight are chosen via random sampling of a probability distribution function that is constructed to fit the historical data. Once a length is chosen, the closest matching discretized route is selected for the flight. This Monte Carlo methodology ensures a test set that matches that to be expected in commercial operation. Route direction was randomly selected to provide more variance in weather effects.

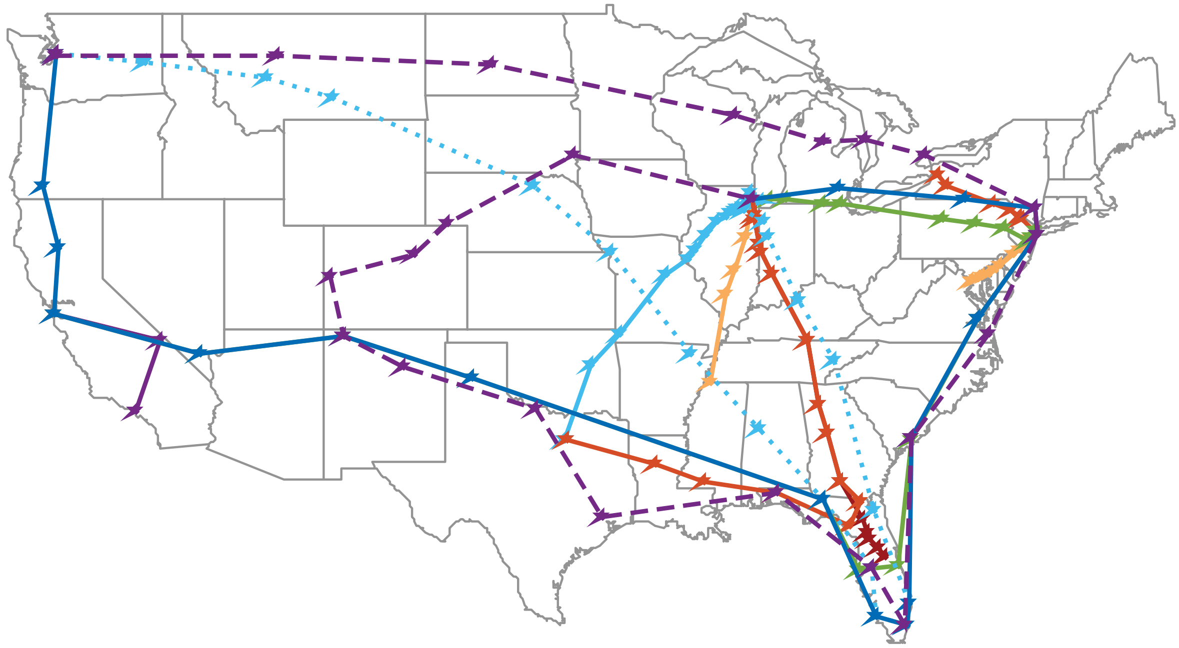

Figure 123. Benefit Assessment Route Mapping.
Stars along each route represent the waypoints used in their creation – no direction is shown, as routes were flown both 'forward' and 'backward'.

An additional consideration when choosing airport pairs is each's frequency of use. Preference is given to airports that are considered "hubs". Figure 124 summarizes the airport combinations for each route. Runways are selected for each airport to ensure takeoff and landing into the wind. Where appropriate, each runway also has an associated standard instrument departure, possibly with an enroute transition, and a standard terminal arrival with relevant approach and transition.

Route	Origin	Destination
200 NMI	KJFK	KDCA
300 NMI	KJFK	KBUF
400 NMI	KORD	KMEM
500 NMI	KLAX	KSFO
600 NMI	KORD	KLGA
800 NMI	KDFW	KORD
900 NMI	KORD	KMCO
1000 NMI	KORD	KMIA

Route	Origin	Destination
1500 NMI	KORD	KDFW
2500 NMI	KJFK	KSEA
2600 NMI	KSEA	KMIA
3000 NMI	KJFK	KSFO
3500 NMI	KSEA	KORD
4000 NMI	KSEA	KDFW
4500 NMI	KSEA	KORD
5000 NMI	KJFK	KJFK
5500 NMI	KSEA	KSFO
6000 NMI	KSEA	KORD
6500 NMI	KSFO	KSFO
7000 NMI	KSEA	KORD

Figure 124. Airport pairings for routes flown “forward”.

B.2.1.3 Weights

A weight range for each aircraft is chosen to represent the in-service operational range and an estimate of the rate of occurrence of each weight. This is accomplished using the following equation:

$$W_{Gross} = W_{operating\ empty} + W_{fuel} + PAX * \ Seats * \overline{W}_{passenger}$$

subject to the limitation of maximum takeoff weight.

A survey of publicly available data is conducted to determine each airframe’s number of available seats. Piano source data provided operational empty weight and maximum takeoff weight. FAA advisory circular AC120-27E provides for estimation of the average passenger weight to be approximately 200 lbs.

Fuel weight is determined for each airframe through simulation of several flights for each of the above-specified routes, with a range of aircraft weights. Analysis of this data provides for an accurate estimation of the average fuel required to fly each route. After estimating the fuel weight, predictions are completed for the flight and iterated upon until fuel weight converges on an unchanging value. The amount of fuel necessary to fly a 200 nmi flight is then added as an estimate for the required reserve fuel.

The passenger load factor (PAX) is the independent variable used to create a gross weight range representative of in-service operation. A random passenger weight between 0% PAX and 100% PAX is selected from a truncated normal distribution centered at 80% [N(80,20)]. This weight is then used in conjunction with the route, aircraft type, and weather to determine the overall weight of the aircraft.

An analysis of aircraft passenger loading confirms the normal distribution about each's mean of 80% loading. Figure 125 depicts the final PAX loading with the truncated normal distribution overlaid on each plot.

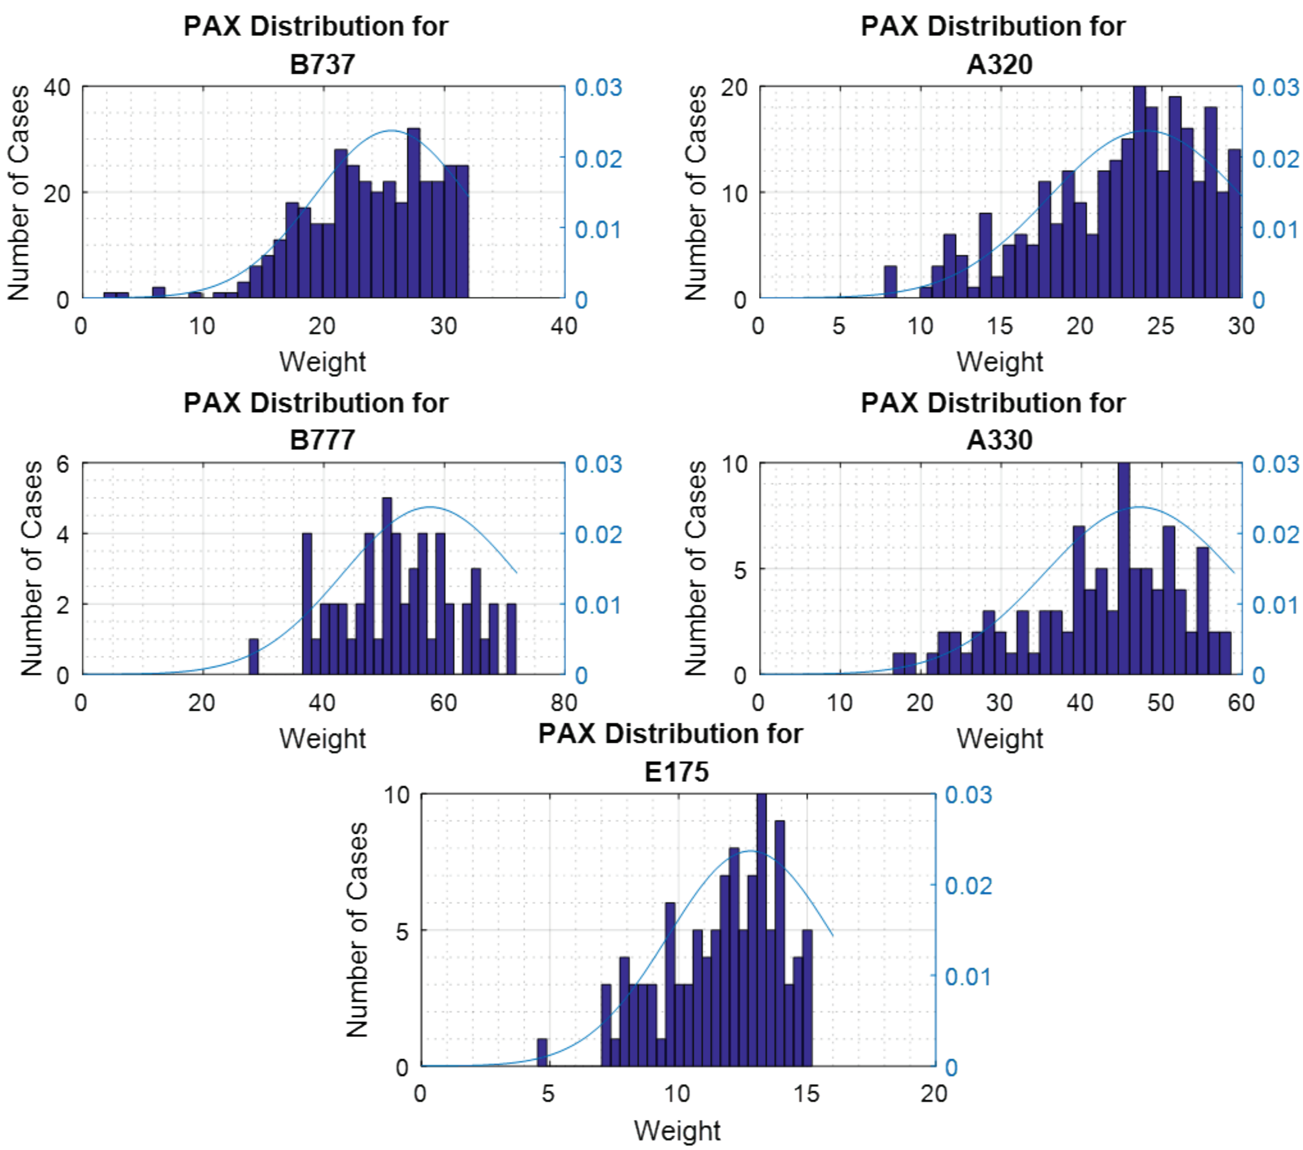

Figure 125. PAX weights in kilo-pounds for test runs with truncated normal distributions overlaid.

B.2.1.4 Weather Conditions

The NOAA National Operational Model Archive and Distribution System (NOMADS) is a network of data that accesses and integrates models and other data stored in geographically distributed repositories into heterogeneous formats. One such format, the Rapid Refresh system (RAP), is chosen to provide data for the benefit assessment due to its collection from multiple data sources, including commercial aircraft weather data, balloon data, radar data, surface observations, and satellite data.

Several representative days are chosen around climate phenomena observed in 2016 to be used in the benefit assessment. RAP data for these days exists as actual observed data (recorded every hour) and as forecasted data (as predicted at each hour, for the next 18 hours). The PPDriver generates predicted flight profiles that utilize the freshest forecast data relative to the takeoff time and date. Random takeoff times between midnight and noon are chosen, which ensures all flights begin and complete in a single day, thereby limiting the processing required for gathering weather data. The simulation of each of the flights utilizes only the ‘actual’, simulated truth weather data, which updates from the RAP forecast for every hour the flight progresses.

A subset of 7 ‘normal’ and 3 ‘inclement’ weather days were used for random selection, with approximately 87% of the days being ‘normal’ and 13% ‘inclement.’ Figure 126 contains the listing of weather days and their classification.

Day in 2016	Classification & Reasoning
February 19	Normal
March 8	Normal
March 17	Severe Storm in Central Plains
April 10	Severe Winter Weather in Central Plains
April 15	Severe Winter Weather in Central Plains
June 23	Normal
August 13	Normal
August 19	Normal
October 28	Normal
December 9	Normal

Figure 126. Weather selections.

The effect of inclement weather phenomena is regional. As such, many of the flown routes are only partially affected while still others – the 200, 300, 400, 500, 600, 900, 1000, 2500, 5000, and 5500 nmi flights – are too distant to be significantly affected.

B.2.1.5 Additional Parameters

Several other parameters are varied as part of the CLEEN Benefit Assessment. ‘Altitude quantization’ or ‘minimum vertical separation’ is used to assess the impact of complying with FAR 14 CFR 91.159 and 91.179, which requires cruise at flight levels that are multiplicative by 1000. To assess the full range of impact, quantization is selected from 1000 and 2000 feet to represent one-way and two-way airways in both the East and West primary directions. Only 2000 ft intervals are used for the RTA benefit assessment since past assessments have shown that altitude quantization does not have a large impact on cost savings.

Carriers use cost indices to account for the cost of the aircraft operating crew by specifying the relation between the cost of time and fuel. Every generated case is conducted with cost indices of both 0 (fuel cost only) and 25 (typical fuel and time blend) to assess direct operating cost.

B.2.1.6 RTA Time and Location

CLEEN technologies are assessed using a randomly generated RTA assigned at a specified waypoint for each test flight. This additional constraint is only applied to flights when judging the benefit of the Gen D technology; all other test efforts are executed without time constraints. For this assessment, two waypoint locations are chosen where the arrival time is controlled; the last waypoint in cruise and the first waypoint below 10000 ft in descent.

Each test flight's required time of arrival is determined by first calculating a nominal flight time as the time it takes to fly the profile to the constrained location at a Cost Index of 25. After determining the nominal flight time, an RTA type is randomly drawn (uniformly) as either 'fast' or 'slow'. A fast RTA is any time constraint that requires a flight time that is less than the nominal time, and a slow RTA requires more time.

With the RTA type selected, the magnitude of deviation from the nominal flight time is chosen using two, truncated normal distributions, each truncated at the nominal flight time. The final RTA is determined using the pre-selected takeoff time, the RTA type, and the time deviation. The distribution for determining the RTA deviation is shown in Figure 127.

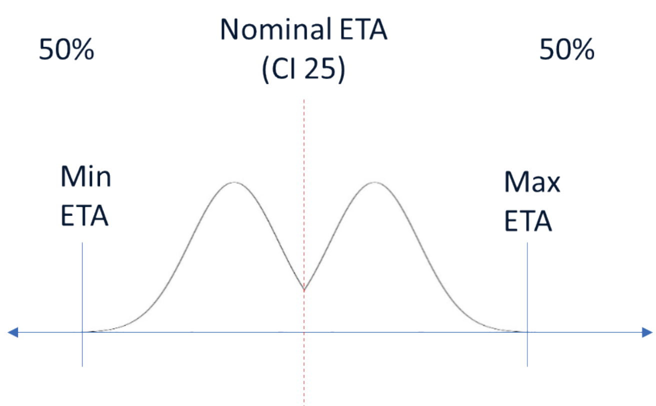

Figure 127. RTA time constraint distribution.

Note that the min, max, and nominal ETAs are determined by a Legacy profile generated using Cost Index limits defined by the Legacy system. This distribution changes with every different weight, aircraft type, and route combination. An example of what the distribution might look like for one combination of those parameters is shown in Figure 128; the red line represents the nominal ETA.

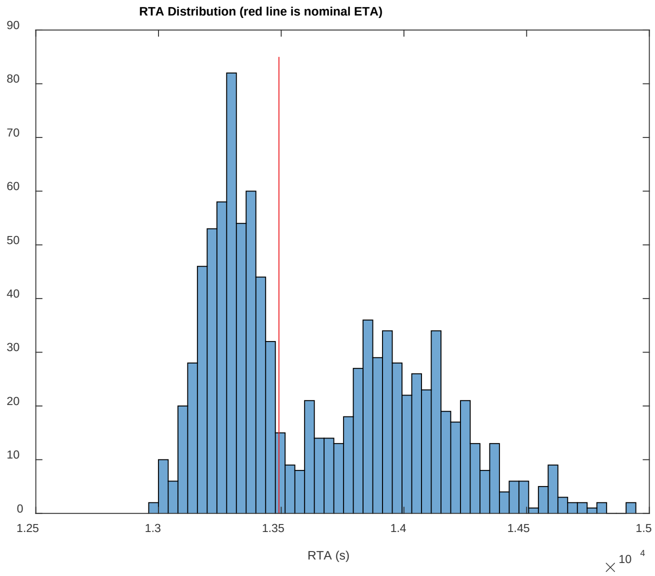

Figure 128. Possible RTA distribution.

The RTAs to the left of the red line are fast RTAs and to the right are slow RTAs. While the general shape of this distribution will remain the same for different cases, there are variations in min, nominal, and max ETA for every case based on the values derived from Legacy Cost Index. Generally, the range of fast RTAs will always be significantly smaller than the range of slow RTAs; that is, faster RTAs will have a smaller standard deviation than slow RTAs.

B.3 TRL 6

B.3.1 Description

The full TRL 5 Benefit Assessment described in paragraph B.2 of this document is performed in a TRL 5 development environment, using the ATT to perform fast-time simulations to assess the expected fleetwide benefit of the CLEEN technologies. This assessment is thorough in describing the expected fleet-wide benefit through analysis of many possible operating conditions and aircraft types, but it does make some simplifying assumptions about performance in the embedded avionics OFP (Operational Flight Plan) software with a real-time laboratory simulator. These assumptions include simplifications of lateral and vertical systems, descent path, re-predictions, etc., and are listed in paragraph B.3.2.

In the TRL 6 environment, the CLEEN Optimizer is implemented in a CFMS (Connected FMS) application on an Electronic Flight Bag (EFB) that communicates with the FMS hosted on an ADE processor via a network. See Section 9 for details on the EFB. A real-time simulator (FMWorkstation) is used on a lab PC to simulate the aircraft flight and provide an interface for pilot inputs to the FMS. A tool called AutoTestUtility automates the process of entering flight data into the FMS through the FMWorkstation interface and controls the CFMS app. A representation of the TRL 5 and TRL 6 environments are shown in Figure 129 and Figure 130.

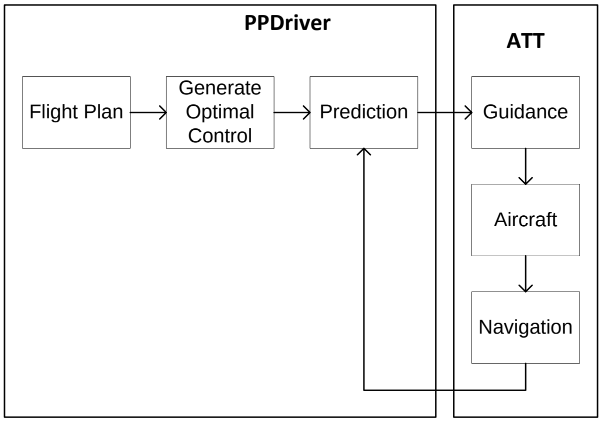

Figure 129. TRL 5 system for benefit assessment.

This assessment uses simulated FMS functions not central to optimization performance. PPDriver is a portion of the OFP, rehosted for desktop system.

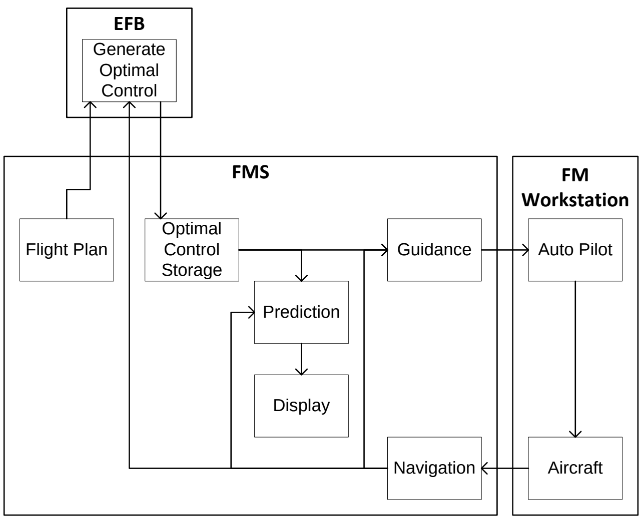

Figure 130. TRL 6 system for benefit assessment.

The EFB contains the PPDriver from Figure 129. All FMS functions are hosted in embedded avionics (real-time).

Due to the TRL 6 Benefit Assessment being performed on a real-time simulator, it is practically infeasible to simulate all the same cases that were performed in TRL 5 assessment. Thus, the TRL 6 Benefit Assessment contains only a small subset of the cases from TRL 5 and has a primary goal of validating the benefit numbers shown in the TRL 5 assessment. Therefore, the general strategy for the assessment is as follows:

1. Select a well-distributed subset of cases from the TRL 5 Benefit Assessment, using only the Boeing 737 aircraft model
2. For each case:
   - Use the input parameters (aircraft weight, route, weather day, Cost Index, etc.) to set up a flight in FMWorkstation via the AutoTestUtility
   - Fly the real-time simulation using Legacy control software and record relevant flight data such as speed, fuel burn, distance travelled, etc.
   - Fly the real-time simulation using the CFMS application to optimize the flight profile and record relevant flight data such as speed, fuel burn, distance travelled, etc.
   - Process the data into the format used for TRL 5 Benefit Assessment, and use the same methods performed for TRL 5 to calculate the cost benefit of the CLEEN profile
3. Compare the cost benefit calculated in TRL 6 for each case to the corresponding TRL 5 benefit for the same case
4. Analyze each case individually, compare the observed benefit, investigate any significant discrepancies, and assess whether the benefit can be extrapolated to make conclusions about expected total benefit in a TRL 6 environment.

This assessment is performed in two parts: testing of the cost savings from the Gen C software and testing of the RTA capabilities of the Gen D software. Based on this assessment, the fleetwide expected TRL 5 benefit is reliable and holds in a TRL 6 environment, and thus it is reasonable to assume the observed benefit will be maintained as the technology matures beyond TRL 6.

In this report, the TRL 6 Benefit Assessment is performed for only Generation C. See paragraph 9.2 for the results of this assessment.

B.3.2 Differences Between TRL 5 and TRL 6

The benefit achieved using the TRL 5 system is not identical as the TRL 6 system for the following reasons:

1. **Predicted Weather:** The predicted weather constitutes the largest difference between TRL 5 and 6. While each system uses the same weather gathering tool to gather data for predictions, the amount of data gathered the location and time from where the data is gathered, and its subsequent use varies. For the TRL 6 systems, the weather population method is designed to reflect how the process would occur in operation for each system. The predicted weather could have a large effect on the descent phase, where no optimal controls are generated, and instead the Legacy FMS computes the proper descent speed and angle (given the weather prediction) in both the CLEEN and Legacy methods. The details of each system regarding weather is described below.
   - TRL 5 CLEEN: The predictions query weather data from the weather gathering tool that are stored in each prediction block for its exact 4-D location (latitude, longitude, altitude, time). Additionally, the cruise segment is predicted using a purposefully higher density grid of prediction blocks (one approximately every 10 nmi), each of which contributes to a weather model resolution higher than would otherwise be available. This results in the most accurate predicted weather.
   - TRL 5 Legacy: The predictions query weather from the same source data as the CLEEN technology but are forced to store that weather in a much lower resolution model, with data only being stored at each waypoint in the prediction. This results in less-accurate predicted weather than TRL 5 CLEEN.
   - TRL 6 CLEEN: The flight profile is generated by the same system used in TRL 5, but the TRL 6 FMS is incapable of entering weather data at every prediction block. In an operational TRL 6 system, only the following weather data are enterable: top-of-climb wind and temperature, four winds of varying altitudes and one temperature at each waypoint in cruise, and three descent winds along with a descent temperature. The high-fidelity weather model used to generate the optimal control was reduced by an internally developed Weather Augmenter that determines the weather entries that result in a best-fit representation of the true full weather from the Optimizer. The FMS then uses the reduced model to predict the test case's flight profile.
   - TRL 6 Legacy: The predictions in this system utilize the most rudimentary weather model that is populated using an automated system (the AutoTestUtility) that mimics the methods a pilot would use to enter data. The utility queries the high-fidelity online model for data that is directly above the departure and destination airports, at the cruise altitude. For the departure airport, the data retrieved is that at the top-of-climb altitude, forecasted at the time of its sequence. Similarly, the destination airport data is retrieved at the top-of-descent altitude, as forecasted at the time of its sequence. The cruise data is limited to a single wind entry at each waypoint (not each prediction block). No temperature entries are available, so the top-of-climb temperature is propagated throughout the entire cruise phase. This results in the lowest-accuracy predicted weather that is confirmed by the data shown in Figure 115.
2. **Re-predictions:** In TRL 5 Gen C, the ATT simulation is designed to track the initial predicted path and does not perform any re-predictions of the path. In TRL 6 Gen C, the Legacy flight may perform a re-predict during flight, and produce new speeds if the true aircraft weight or experienced weather vary from what was initially predicted. Technically, the CLEEN flight may perform a re-predict as well, but it will not result in new speeds (or cruise altitude steps) as those are not re-optimized, and the flight will still follow the original optimal controls. The expected impact of these re-predicts in Gen C TRL 6 Legacy is small. In Gen D, re-predicts are performed in all TRL 5 and 6 systems, in order to adjust the profile to meet the RTA constraint.
3. **Optimizer Inputs/Outputs:** The Optimizer code base is the same PPDriver used in TRL 5 and 6. In TRL 5, the PPDriver directly produces the prediction that the ATT simulates. In TRL 6, the PPDriver is hosted in the cFMS App, and generates a prediction and optimal controls that are transmitted to the FMS. Ideally, the identical case produces the same inputs to the optimizer in TRL 5 and 6; however, there are slight numerical discrepancies that can arise due to the dynamic nature of the TRL 6 simulator.
4. **Controls:** The TRL 6 simulation (FMWorkstation) has a different control source than the ATT does. In TRL 5, the predicted path is used directly as the source of controls. In TRL 6, the FMS produces flight commands to the Autothrottle model that in turn generates commands for the engine model. The effect is that the TRL 6 system has more realistic lags in commands.
5. **Descent:** Descent paths can vary between TRL 5 and 6. This is partially due to different weather predictions, but also different thrust reversion logic between the real FMS guidance and the simulated version in TRL 5.
6. **Lateral/Vertical:** The ATT has less-advanced lateral and vertical tracking than the full FMWorkstation simulation does. Generally, this results in a minor difference in things like turn radius or angle of attack. However, this effect is essentially negated as it affects CLEEN and Legacy in the same magnitude, resulting in an offset.

Appendix C Previous Work

C GE's Previous Work

C.1 GE CLEEN I Climb Optimization

C.1.1 Summary

The Flight Management Systems in service today determine constant climb speed, constant cruise speed, and constant descent speed to minimize DOC based on takeoff weight and range and assuming maximum thrust for climb and idle thrust for descent. Typically, software look-up tables derived from flight trials or simulation define the control that is then limited to comply with performance and airspace requirements.

On some high-performance airplanes, the optimization method is derived from the calculus of variations and Pontryagin's minimum principle. The most common methods determine variable climb and descent speeds to achieve performance closer to optimum relative to the constant speed method. However, many simplifying assumptions have been applied to enable a practical design. These assumptions — motivated by the computer technology available at the time the methods were developed — introduce errors and thus yield suboptimal performance.

Accordingly, the problem was reformulated for climbing flight without these assumptions, and modern numerical methods were applied to achieve a control closer to optimum. Since the pilot has more discretion controlling the airplane during the climb phase, the study focused on climbing flight to prove the concept.

Relative to the legacy constant speed method, the following improvements were implemented:

1. Higher order equations of motion that include mass as a state variable and a model of the installed performance of each engine.
2. Since transient motion is a small share of the flight time, the fast dynamics have a negligible effect on fuel savings. Thus, the fast dynamics may be eliminated from the system equations. The result is a set of differential algebraic equations (DAE) that represents only the slow vehicle dynamics that enables a larger time step for integration.
3. Flight path angle (instead of speed) and throttle lever position (thrust) are the control variables to minimize the cost function.
4. Because the arrival time is free to vary, the length of the state trajectory is unknown. To avoid solving a mixed-integer programming problem, the independent variable is changed from time to altitude and is a known boundary condition.

The problem was formulated using a Non-Linear Programming (NLP) method, and an optimization solver developed by GE Global Research for model-predictive control was used to solve the problem numerically. However, because the cost surface is not smooth, the convergence time and robustness of the algorithm are unsatisfactory. Therefore, a simpler approach based on[§§§§] was developed as an alternative method. In this approach, called the Energy State Approximation (ESA):

1. The independent variable in the cost function is changed from time to energy.
2. The control variable is simplified to be just speed (thrust is added as a control variable during later development; see paragraph 5.2.1).
3. The vehicle state is approximated by specific energy.
4. The Golden Search method is used to find the minimum cost.

Because the vehicle state is approximated, the ESA method is suboptimal relative to the NLP method. However, the solver is more robust and converges in a shorter time, which makes the ESA method more suitable for embedded applications. Accordingly, the fuel savings achieved by both methods (relative to the Legacy method) was compared to validate the theory and to quantify the penalty introduced by the ESA method.

Relative to the Legacy constant-velocity optimization method and based on a small number of simulation trials, the NLP method reduces the fuel required (in the climb phase only) to climb to a specified cruise altitude by 1%. The ESA method reduces the amount of fuel required by 0.81%. Both the NLP and ESA methods were implemented in a prototype FMS. Due to superior robustness and convergence time, the ESA method is the best method for an airworthy design.

[§§§§] H. Erzberger and H. Lee, "Constrained Optimum Trajectories with Specific Range", *Journal of Guidance and Control,* vol. 3, pp. 78-85, 1979.

C.1.2 Technical Summary

In commercial Aviation, Direct Operating Cost (DOC) is defined as:

$$DOC\ [\$] = Fuel\ Cost\ [\$] + Time\ Cost\ [\$] \quad (1)$$

Past work of the CLEEN program has utilized a variation of this definition within a flight profile optimization algorithm, which leads to the following of climb cost:

$$\min_V J_{climb} = \min_V \int_{t_0}^{t_f} \rho(\pi, V, h) dt\,, \quad (2)$$

Where t_f is the time the climb terminates, called the top of climb, given the system equations:

$$\begin{aligned} \dot{E} &= \frac{V}{m} f_x^s(h, V, \pi) \\ \dot{h} &= V\gamma \\ \dot{x} &= V + V_W \\ \dot{m} &= -\sigma(h, V, \pi) \\ 0 &= f_z^s(h, V, \pi) - mg \\ 0 &= \tau_y^b(h, V, \pi) \end{aligned} \quad (3)$$

In eq. (2), $\rho(\pi, V, h) := c_f w_f(\pi, V, h) + c_t$, where w_f is the fuel flow rate, c_f is the cost per pound of fuel, and c_t is the cost rate (i.e., the cost of time). See "Optimal Variable-Speed Climb for a Fixed-Wing Aircraft"[*****] for a more thorough description of this climb formulation and its derivation.

C.2 GE CLEEN I Vertical Path Optimization

The basis for the Generation C technology is described in a previous GE CLEEN technology called Vertical Path Optimization. This work analyzes a few methods of computing optimally stepped cruise profiles, only considering cruise altitudes (not speed). The method focuses on using the GE Approximate Global Optimizer (GEAGO) to optimize steps. Dijkstra's algorithm is also implemented with a fine grid as an entitlement method with GEAGO and is referred to as a Dynamic Programming solution in that report.

While the report concludes that the GEAGO solution produces a profile within 0.1% of the cost of the entitlement Dynamic Programming/Dijkstra method in a much faster computation time, the GEAGO is not pursued for the Gen C technology due to the complexity of its code base, and substantial software memory and run-time improvements since the original implementation. Dijkstra's algorithm is very simple and can be implemented in only a few lines of code; this is desirable from both a developer and a certification standpoint.

[*****] R. Ghaemi, D. Lax, E. Westervelt, M. Darnell and N. Visser, "Optimal Variable-Speed Climb for a Fixed-Wing Aircraft". in *AIAA Aviation Forum*, Dallas, TX, 2019.

C.3 MIT-LL CASO

MIT Lincoln Laboratories produced a similar technology to the Gen C Optimizer and performed a similar study in a project they called CASO (Cruise Altitude and Speed Optimizer). The report can be found in "Optimal Variable-Speed Climb for a Fixed-Wing Aircraft"[†††††]. The study develops several implementations of cruise step optimization, distinguished by how the cruise altitude steps are performed. Namely, a drift-up cruise climb pattern, 1,000 and 2,000 ft increment step climbs (that is akin to the Legacy++ method), and what they call Flexible VNAV (that is akin to the Gen C optimization technology, providing for step climbs and descents).

Their report indicates a savings aggregate of 1.96% for Flexible VNAV, and 1.75% for 2,000 ft Step Climb. These numbers are concerning because they are much higher than the savings deemed achievable with Gen C. However, the MIT report's basis of comparison is to compare the methods with an as-flown trajectory, using radar track data from real flights and a fuel weight estimator. As mentioned in paragraph 3.4, this artificially inflates fuel savings. A more astute observation is that the difference between the Flexible VNAV and 2,000 ft Step Climb methods is 0.21% that is much closer in magnitude to the 0.38% the Gen C technology saves over the Legacy++ method.

The Gen C optimization technology implements several enhancements beyond the CASO design: variable speed/variable thrust climb, unified climb cruise design, and accounting for step cost in cruise. These advances provide additional cost savings over alternate modern step optimization methods proposed in competing products.

[†††††] R. Ghaemi, D. Lax, E. Westervelt, M. Darnell and N. Visser, "Optimal Variable-Speed Climb for a Fixed-Wing Aircraft", in *AIAA Aviation Forum*, Dallas, TX, 2019.