Design and Development of a Rapid Research, Design, and Development Platform for In-Situ Testing of Tools and Concepts for Trajectory-Based Operations

Matthew C. Underwood[1]
NASA Langley Research Center, Hampton, VA, 23681

To provide justification for equipping a fleet of aircraft with avionics capable of supporting trajectory-based operations, significant flight testing must be accomplished. However, equipping aircraft with these avionics and enabling technologies to communicate the clearances required for trajectory-based operations is cost-challenging using conventional avionics approaches. This paper describes an approach to minimize the costs and risks of flight testing these technologies in-situ, discusses the test-bed platform developed, and highlights results from a proof-of-concept flight test campaign that demonstrates the feasibility and efficiency of this approach.

Nomenclature

4DT	=	Four Dimensional Trajectory
ADS-B	=	Automatic Dependent Surveillance—Broadcast
ADS-C	=	Automatic Dependent Surveillance—Contract
ANSP	=	Air Navigation Service Provider
ARINC	=	Aeronautical Radio, Incorporated
ASTOR	=	Aircraft Simulation for Traffic Operations Research
ATM	=	Air Traffic Management
AvBus	=	Avionics Bus
DataComm	=	Digital Data Communications
DRNAV	=	Dynamic Area Navigation
DRNP	=	Dynamic Required Navigation Performance
EICAS	=	Engine-Indicating and Crew-Alerting System
EPP	=	Extended Projected Profile
FAA	=	Federal Aviation Administration
FMS	=	Flight Management System
HITL	=	Human-In-The-Loop
IFR	=	Instrument Flight Rules
ILS	=	Instrument Landing System
KLFI	=	Langley Air Force Base
KPHF	=	Newport News-Williamsburg International Airport
MCDU	=	Multi-Function Control Display Unit
MCP	=	Mode Control Panel
NAS	=	National Airspace System
NASA	=	National Aeronautics and Space Administration
ND	=	Navigation Display
NextGen	=	Next Generation Air Transportation System
PFD	=	Primary Flight Display
R2D2	=	Rapid Research Design and Development
RNAV	=	Area Navigation
RNP	=	Required Navigation Performance

---

[1] Research Aerospace Engineer, Crew Systems and Aviation Operations Branch, NASA Langley Research Center, Mail Stop 152, Hampton, VA 23681-2199, AIAA Member

RPFMS	=	Research Prototype Flight Management System
RTA	=	Required Time of Arrival
RWY	=	Runway
SMART-NAS	=	Shadow Mode Assessment using Realistic Technologies for the National Airspace System
TBO	=	Trajectory-based Operations

I. Introduction

Trajectory-based Operations (TBO) are a key component of the Next Generation Air Transportation System (NextGen) [1] [2]. TBO is based on the premise that, in the near future, aircraft operating in the National Airspace System (NAS) will be represented by a four-dimensional trajectory (4DT), which defines the lateral and vertical path of the aircraft with time components. TBO extends this premise by providing separation, sequencing, spacing, and merging services to flights based on a combination of their current and projected positions, thus providing efficiency, capacity, and safety benefits. [2]

Several Air Traffic Management (ATM) concepts, tools, and technologies that enable TBO rely on advanced technologies such as Automatic Dependent Surveillance-Broadcast (ADS-B), Digital Data Communications (DataComm), airborne and ground-based automation, and advanced controller and pilot interfaces. These concepts, tools, and technologies have been developed and investigated by the National Aeronautics and Space Administration (NASA), the Federal Aviation Administration (FAA), and others. Research and development to date for concepts that rely on the exchange of data between aircraft or between ground systems and aircraft have almost exclusively been conducted using modeling, analysis, and Human-in-the-Loop (HITL) simulation. To continue progress toward implementation of these advanced concepts, tools, and technologies, extensive validation and refinement must occur in operationally relevant environments.

Due to their inherently high cost, flight test activities associated with advanced ATM concepts are often limited in scope. However; many advanced concepts will require extensive validation involving multiple field tests. Results of an in-flight concept validation often result in changes to the equipment and standards, which results in a new cycle of expensive and time-consuming flight testing. Reducing these cost and time factors are critical factors in meeting modernization expectations for TBO concepts, tools, and their enabling technologies.

To enable rapid, efficient, low-risk flight testing of TBO concepts, tools, and technologies, the Rapid Research Design and Development (R2D2) Platform was developed at the NASA Langley Research Center. The R2D2 platform allows researchers to cost-effectively simulate, test, and demonstrate various elements of TBO without the need to purchase commercial off the shelf certificated avionics. The R2D2 platform was tested in a Beechcraft UC-12 (Beechcraft 200) in the Fall of 2016.

II. R2D2 Platform

A. R2D2 Platform Components

The R2D2 platform contains several components, each of which is discussed further in the following sections. The R2D2 platform is currently designed to be read-only from the aircraft, i.e., no data generated by the components of the R2D2 platform are supplied directly to the aircraft's autoflight systems or displays. For this set of operational trials, all guidance generated by the R2D2 platform was relayed to the pilot, who made manual inputs into the aircraft's flight systems, via the Mode Control Panel (MCP) or Flight Management System (FMS).

1. Researcher Displays, Interfaces, and Tools

The majority of components in the R2D2 system reside with a researcher seated in the cabin of the flight test aircraft. The researcher displays, interfaces, and tools shown in Figure 1 contain all of the required functionality needed to test TBO operations. These modules reside on a laptop computer that was connected to a hard-wired local area network on-board the aircraft.

The emulated displays, interfaces, and internal communication mechanisms are similar to those of a current commercial aircraft cockpit. These displays and interfaces were previously created for the NASA Langley-developed Aircraft Simulator for Traffic Operations Research (ASTOR), which is a medium-fidelity HITL computer workstation-based aircraft simulation used to test new ATM tools and procedures. [3]

The Multi-Function Control Display Unit (MCDU), located in the bottom right of Figure 1, is the user interface to the NASA Langley Research Prototype Flight Management System (RPFMS). The RPFMS is a simulated FMS, and includes the capabilities of a production FMS in addition to the research flexibility afforded by a software-based simulation. The RPFMS is capable of generating 4D Trajectories (4DTs) that are subject to multiple Required Time of Arrival (RTA) constraints or time windows constraints, receiving DataComm messages, and executing TBO concepts such as dynamic routing. [4] Furthermore, the RPFMS is capable of generating simulated Automatic Dependent Surveillance-Contract (ADS-C) Extended Projected Profile (EPP) messages, which provides a representation of the FMS-calculated 4DT for the aircraft to ground-based automation platforms.

An emulated data bus, known as the Avionics Data Bus (AvBus), serves as the inter-process communication backbone of R2D2. [5] The AvBus is a buffered shared memory that allows several simulated avionics processes to communicate in a flexible, efficient, and standardized (conforming to Aeronautical Radio, Incorporated (ARINC) 429 standard) manner. In the R2D2 platform, the AvBus is populated with state data obtained directly from the research aircraft's avionics systems, and the RPFMS was able to use this data in its 4DT computations. Furthermore, the state data received from the aircraft was used to drive the Primary Flight Display (PFD), shown at the left of Figure 1, and the Navigation Display (ND), shown in the center in Figure 1.

The Engine-Indicating and Crew-Alerting System (EICAS), shown in the upper right of Figure 1, is the user interface and display to review and load DataComm messages received by the aircraft into the RPFMS. Once the FMS receives a DataComm message, the flight crew had two options. The first option was to load the message in the RPFMS, execute the message within the RPFMS, and accept the message on the EICAS interface. The second option was for the flight crew to reject the incoming clearance by sending an unable message to the Air Navigation Service Provider (ANSP).
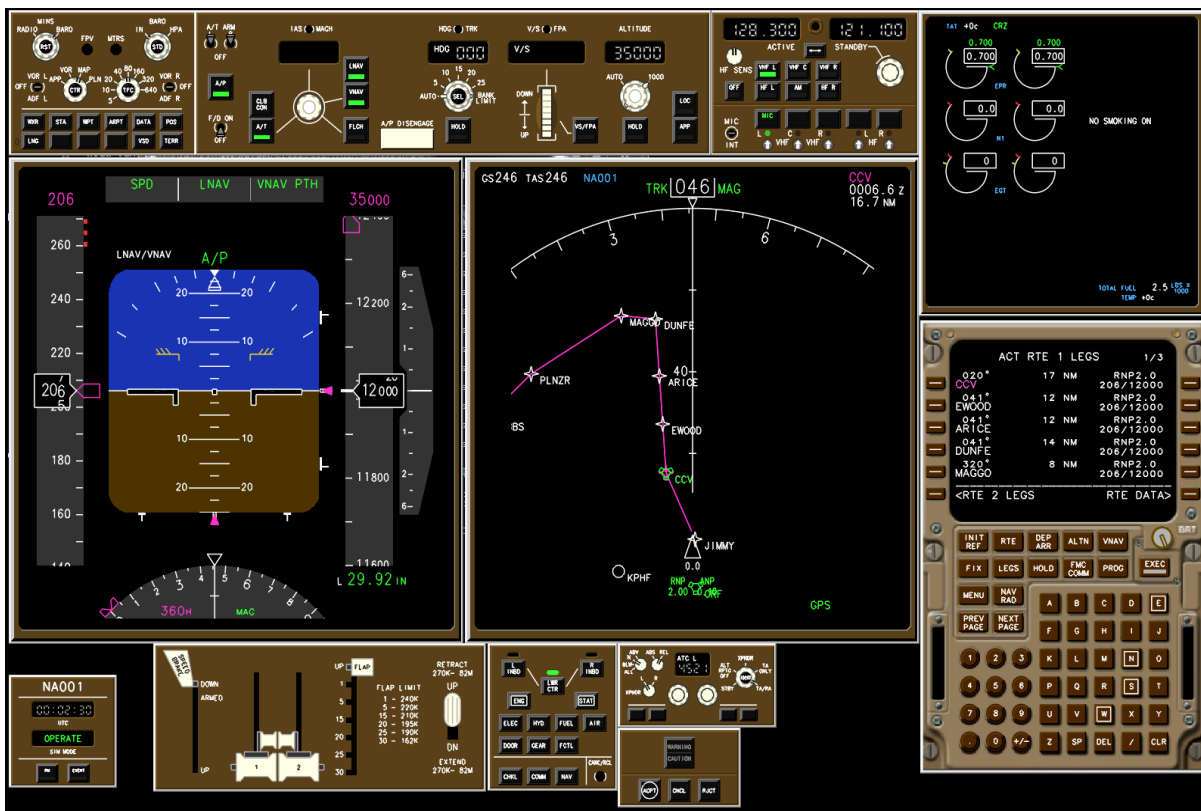

Figure 1: Researcher Displays, Interfaces, and Tools

2. Pilot Display and Interfaces

An EFB-based display and interface (Figure 2) for the pilot of the research aircraft was used to convey information from the researcher displays, interfaces, and tool located in the cabin. This tablet shows the pilot information relevant to the operation being conducted. The tablet computer is connected to the researcher laptop via a wireless local area network onboard the aircraft.

For this initial operational trial, two data fields and a display were included on the pilot display and interface. The first field was the RTA speed (top of Figure 2), which provided speed guidance to the flight crew to achieve the required timing at a certain waypoint. This field displayed a speed that the flight crew will either maintain manually or enter into the MCP of the aircraft. The speed will be shown in green with a black background if the aircraft is conforming to the speed command. If not, the speed will be shown in reverse video, as seen in Figure 2.

The second field is the Dynamic Route Guidance field, located below the RTA speed field in Figure 2. This field provides textual guidance to the flight crew of the route to be input into the aircraft's FMS in order to comply with the dynamic route clearance sent to the aircraft via DataComm. Additionally, the Dynamic Route Guidance field has "accept" and "reject" buttons that allow the flight crew to inform the researcher in the cabin of the aircraft that they have loaded the route into the aircraft's FMS ("accept") or will not load it ("reject").

The final component of the pilot display is a replication of the ND shown on the researcher's display, located at the bottom of Figure 2. Its intended function is threefold—to provide the flight crew with situation awareness of the route of flight, to ensure that the same route is entered into the aircraft's FMS as in the RPFMS, and to visualize the dynamic route guidance. The flight crew was also given the ability to change the range of the map (i.e., zoom in or out) as seen at the very bottom of Figure 2.

It should be noted that this display and interface is a component of a research and development platform—not an operational tool. Therefore, it was not subjected to human factors requirements for displays in the cockpit, nor was it evaluated for its usability and aesthetics by human factors specialists. Furthermore, the researcher and pilot were in constant voice communication. This was done to ensure that the pilot was aware that an operation was about to occur and as an avenue to discuss and resolve any confusion during the operation.

3. DataComm Surrogate

The R2D2 platform contains an application that allows for DataComm messages to be sent at a user's request to the RPFMS (seen in Figure 3). These messages conform to the message standard for DataComm [6], and are used to send clearances to the RPFMS such that it can act upon the information in the message. The DataComm surrogate was located on the researcher laptop.

The bottom field allowed the researchers to load a predetermined DataComm message from a list. The contents of the message are displayed in the message pane for review by the researcher. Once the message is reviewed, it can either be sent to the RPFMS or cleared so that a new message can be loaded. Once the message was sent to the RPFMS and acted upon, the response from the flight crew was shown in the top field.

4. ARINC 429 Dongle and Spider

A research data port on the aircraft was used to obtain the required state data (refer to Table 1) from the research vehicle's systems. However, the data coming from this port was raw ARINC 429 data from various systems on the aircraft. To convert it to a useable form, a Ballard ARINC 429 USB dongle
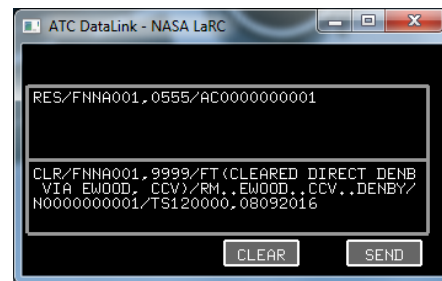

Figure 2: Pilot Display and Interface

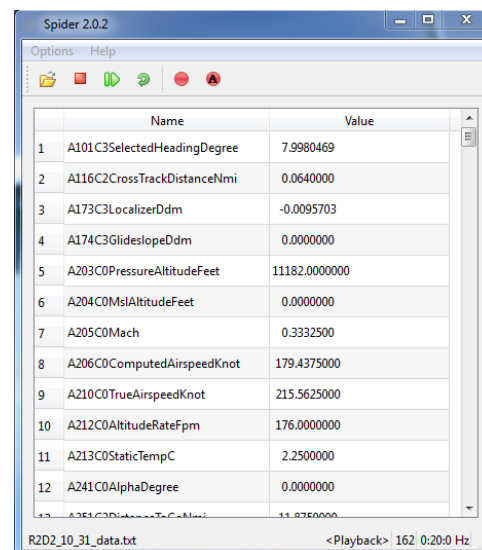**Figure 3: DataComm Surrogate**

Figure 4: Spider Application

translated the raw ARINC 429 data into practical engineering units. The data coming from this dongle was then pushed into a NASA Langley-developed buffered shared memory application known as Spider. The user interface for Spider is shown in Figure 4. From Spider, the data was loaded on to the AvBus, where it was read by the RPFMS. The ARINC 429 USB dongle was connected to the main research computer on the aircraft, and the Spider software resided on that computer. The data from Spider was transmitted from the research computer to the laptop via the hard-wired local area network on-board the aircraft.

5. Aircraft Performance Model

The R2D2 platform contained a medium-fidelity performance model of the UC-12 aircraft created from data within the Pilot Operating Handbook. This model included information about the aerodynamic coefficients, performance limitations, and engine performance of the research vehicle. These performance limits were used in the calculation of the 4DT performed by the RPFMS. The performance model was located on the researcher laptop.

Table 1: Ownship Data

Data Source	Data Element
Air Data Computer	Computed Airspeed (kts)
	True Airspeed (kts)
	Mach
	Pressure Altitude (ft)
	Vertical speed (ft/min)
	Static Temp (deg C)
Flight Management System	Cross Track Distance
	Latitude (deg)
	Longitude (deg)
	Ground Speed (kts)
	True Track (deg)
	True Heading (deg)
	Wind Speed (kts)
	Wind Direction (deg)
Primary Flight Display System	Selected Heading
Attitude Heading Reference System	Pitch (deg)
	Roll (deg)
	Pitch rate (deg/sec)
	Roll rate (deg/sec)

B. System Architecture

The system architecture for R2D2 is shown in Figure 5. The R2D2 platform requires state data from the research vehicle. For these tests, as previously mentioned, the research vehicle is outfitted with a data port that allows research systems to access these data. After the data is decoded by the Ballard ARINC 429 card and inserted into Spider, the R2D2 platform ingests these data via the AvBus. The researcher displays, interfaces, and tools read these data, and compute a 4DT while meeting any constraints imposed by the ANSP. Once the 4DT is computed, it is shown on the ND and MCDU on the researcher's laptop.
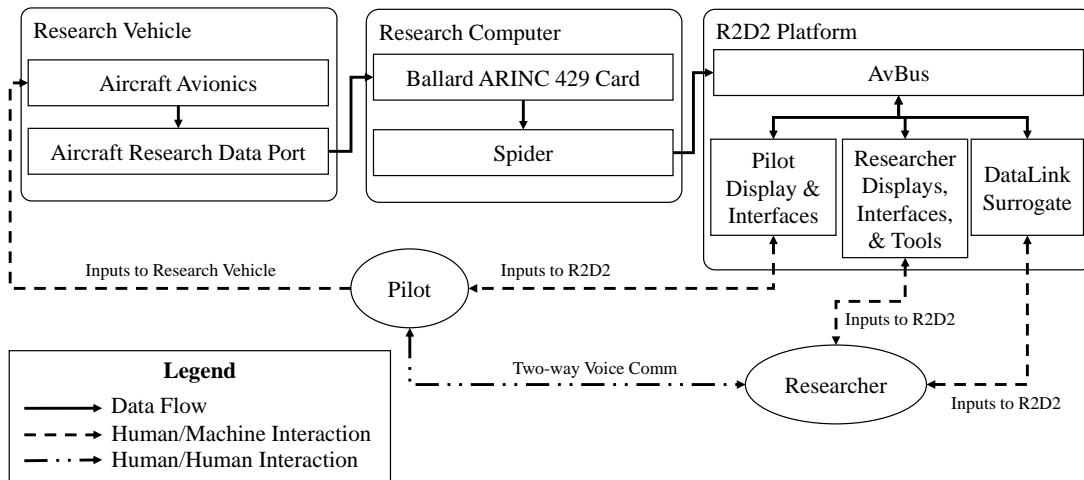

Figure 5: R2D2 Platform System Architecture

C. Anticipated Benefits of R2D2 Platform

The R2D2 platform bridges the gaps between research, development and implementation of advanced TBO decision support tools and their enabling communication and surveillance technologies. Anticipated benefits are:

- Costs of flight trials are mitigated for specific airborne functions that rely on advanced technologies or new standards for communication and surveillance. Cost reduction is achieved by reducing or eliminating the

requirements to procure and permanently install expensive equipment on aircraft to enable participation in the test.

- Flight trial preparation time is substantially reduced since aircraft are equipped with existing test hardware on a temporary basis, new functional capabilities are provided from outside the aircraft, and the new capabilities already exist in simulation laboratories in some instances.
- A virtual representation of communication and surveillance technology rather than reliance on actual hardware enables the exploration and validation of communication, navigation, and surveillance (CNS) requirements *in-situ*. Simulations of future hardware are easily modifiable to explore and test potential new standards for message content and transfer rate. Test hardware will not become obsolete as a result of CNS standards changes, and the virtual environment will enable rapid update of the test system as standards evolve.
- Simulated aircraft may augment the actual aircraft in the test, thereby increasing complexity of test scenarios with a minimal increase in costs.
- Research algorithms and crew interface software stay on the research platforms on which they were developed and tested, thereby reducing costs, preparation time, and flight trial execution risk.

In order to realize these anticipated benefits, the R2D2 system performing an advanced TBO operation must be tested in situ for both its feasibility and practicality. The results of this test, coupled with the results of numerical analyses and simulations, will determine if the results provided by the R2D2 system are both correct and valid. However, additional costs may be incurred for a tool tested using the R2D2 platform. These may include the cost for verifying, validating, and certifying the enabling technologies (e.g., ADS-B, DataComm, etc.) used by the tool, the hardware on which the tool resides, and the software and algorithms used in the tool. These costs are not trivial—however, testing the tool using the R2D2 platform provides initial data regarding whether the benefits of the tool outweigh the costs of validation, verification, and certification.

III. Proof-of-Concept Flight Test

To meet the aforementioned objectives, a flight test campaign was conducted to evaluate the operational capabilities and feasibility of the R2D2 system. This flight test campaign involved six flights (four check flights to test systems and procedures and two data collection flights) during which approximately six gigabytes of data was collected. In this section, the use-cases tested in the flight test are explained, the flight test routes are illustrated, results are provided and discussed, and the next steps for this research capability are presented.

A. Initial Use Cases

Two use case operations that specifically target the functionality of the R2D2 platform were chosen for this initial flight test campaign. These operations were specifically chosen because they trigger trajectory re-computations in the RPFMS.

1. Dynamic Reroute Operation

The first use case was a dynamic re-route of the aircraft. In an operational TBO context, a dynamic reroute may be needed for various reasons, including a path stretch for spacing, avoiding convective weather, areas of icing or turbulence, or at the user's (flight crew and/or dispatcher) request. Two types of dynamic reroutes are anticipated to be options to mitigate the aforementioned issues—a dynamic area navigation (DRNAV) route and a dynamic required navigation performance (DRNP) route. DRNAVs are dynamically generated area navigation (RNAV) re-routes - navigating by means of named fixes and navaids - but do not contain any navigational performance requirement. DRNPs are based on the similarly-named concept of operations by the FAA [7] and are targeted at improving the flexibility of the NAS. A DRNP [6, 8] is a re-route defined by a set of waypoints (which can include latitude/longitude points), RNP data for the re-route on a leg-by-leg basis, and fixed-radius-transitions or radius-to-fix legs to fully define the turn geometries along the re-route.

In this operational trial, the DRNAV option was chosen for three reasons. The first reason was that, for clarity in understanding the clearance request by the flight crew, a reroute with named fixes was chosen. The second reason was since the R2D2 platform does not write data to the aircraft's avionics, the pilots are responsible for manually entering this reroute into the FMS. Using named fixes was a mitigation for the risk of flight crew transposition error when entering the clearance into the FMS. Finally, since this operational trial occurred in controlled airspace with the potential requirement of an Instrument Flight Rules (IFR) flight plan, requesting the route deviation from the ANSP was made easier through the use of named fixes.

2. Required Time of Arrival Operation

The second use case was an RTA operation. An RTA operation may be required in a TBO environment to efficiently meter aircraft across a point or boundary. RTA time control relies on trajectory prediction to generate a flight trajectory that achieves the desired arrival time. The trajectory generator will iterate on possible trajectories until the estimated time of arrival at the RTA waypoint is within a pre-specified tolerance of the RTA. There are a number of methods for accomplishing this iteration. The RPFMS on the R2D2 platform uses the cost index[1] as the independent variable for this iteration. [4] Periodically, the trajectory generator will update the estimated arrival time, as well as the maximum and minimum arrival times, from the current aircraft location along the reference trajectory to the RTA waypoint. This update is done approximately every 60 seconds in the RPFMS. If the estimated time of arrival is earlier or later than the RTA by more than a set time error tolerance (30 seconds in these operations), the RPFMS will trigger a new trajectory iteration to meet the RTA.

B. Flight Test Vehicle and Data

The flight test vehicle chosen for the initial checkout flights of the R2D2 platform was the NASA Langley UC-12B King Air, call sign NASA528, seen in Figure 6. The UC-12B is a military version of a Beechcraft B200 King Air. It is capable of flying missions up to 28,000 feet and can fly for up to six hours, depending on the payload. It has a maximum airspeed of 260 knots and a range of 1,250 nautical miles. [9]

Figure 6: UC-12 King Air Aircraft, NASA528

The UC-12 test aircraft was equipped with a modern certified avionics, including:

- Dual Garmin G600 suite with Synthetic Vision
  - GDU 620 PFD, ND
  - GTN 750 Multifunction Display, with traffic display
  - GDL 88 Dual band ADS-B
  - GRS77 Attitude Heading Reference System
  - GDC 74 Air Data Computer
- TCAS I collision avoidance
- Applanix Position Orienting System (high resolution inertial data system)

Prior to the system checks, it was discovered that the autopilot system installed in the UC-12B was not compatible with the new avionics suite. Due to timing of the flight campaign and constraints imposed by the maintenance schedule of the aircraft, it was deemed impractical to procure and install a compatible autopilot system in the UC-12B prior to the initial check flights. As a mitigation, the flight crew would follow the guidance of the Garmin FMS and/or the R2D2 Pilot Display and Interface, but hand-fly the aircraft. The resulting impacts of this mitigation are discussed in the results section of this document.

C. Flight Test Routes and Procedures

This section of the document describes the routes and procedures that were used in the flight test campaign. The flights departed from Langley Air Force Base (KLFI), were conducted over the Hampton Roads area and the eastern shore of Virginia, and terminated at Newport News-Williamsburg International Airport (KPHF), where additional, unrelated research conducted during the campaign was performed. After the research at KPHF was conducted, the aircraft returned to KLFI. The flight trial was split into two portions—an outbound leg and an inbound leg—during which clearances for both use case operations were issued.

1. Flight Test Routes

Four routes were created for the test—two that took into account the departure at KLFI (either runway (RWY) 08 or RWY 26) and two that considered the arrival at KPHF (either the Instrument Landing System (ILS) approach to

---

[1] The cost index for a flight in commercial transport operations is typically set by an airline dispatcher, and is a number used to specify that airline's preference between saving flight time and reducing fuel burn.

RWY 07 or the ILS approach to RWY 25). The combinations are shown in Table 2, and the waypoints for each route are shown in Figure 7.

Table 2: Route Options

	ILS Approach KPHF RWY 25 (Option A)	ILS Approach KPHF RWY 07 (Option B)
Depart KLFI RWY 08 (Option 1)	Route 1.A	Route 1.B
Depart KLFI RWY 26 (Option 2)	Route 2.A	Route 2.B

2. Flight Test Procedures

Prior to each flight, the researcher briefed the flight crew on the flight plan and the two types of operations planned to be conducted during the flight. Additionally, to mitigate negative effects due to the lack of an autopilot system in the aircraft, the researcher requested that the flight crew stay within certain bounds for altitude deviation, lateral deviation, and speed deviation, which are discussed in Section D.1. A briefing was issued by the flight crew that included the weather information for the flight, the expected time en route, and any constraints on the aircraft regarding maintenance issues.

After departure, when the research aircraft completed the turn after JIMMY (Route 1.A and 1.B) or RIPPS (Route 2.A and 2.B) and was established on course to the CCV, the outbound leg began. The first operation performed on the outbound leg was a DRNAV clearance. The researcher in the cabin controlling the R2D2 platform issued a DRNAV clearance to the flight crew via the DataComm surrogate, which instructed the flight crew to proceed direct from their current position to a navaid named JRAYE, then proceed to a navaid named MELFA, then rejoin their route at ARICE. After the flight crew obtained approval to deviate from their planned route from the ANSP (if the flight was on an IFR flight plan), loaded the re-route into the aircraft's FMS, and accepted the re-route clearance on the Pilot Display and Interface, the re-route was executed in the R2D2 RPFMS by the researcher.

Figure 7: Waypoints for Each Route

Once the aircraft was established on the new route, the second clearance on the outbound leg—an RTA operation—was issued by the researcher in the cabin. The RTA waypoint for the outbound leg was ARICE. The RTA Speed data field on the Pilot Display and Interface became active, providing speeds that the flight crew followed to achieve the RTA.

Once the aircraft sequenced ARICE, the outbound leg was complete and the flight crew prepared for the inbound leg. The flight plan included a tear-drop turn at DUNFE to assist in setting up the inbound leg. Once the turn at DUNFE was made and the aircraft was established on course direct to ARICE, the inbound leg began.

The first operation performed on the inbound leg was a DRNAV clearance. The clearance instructed the flight crew to proceed direct from their current position to a navaid named EWOOD, then proceed to a fix named CCV, then rejoin their route at DENBY. After the flight crew followed the procedures for obtaining approval from the ANSP if applicable, loaded, and accepted the clearance, the re-route was executed in the R2D2 RPFMS by the researcher.

Similar to the outbound leg, once the aircraft was established on the new route, an RTA clearance on the inbound leg was issued by the researcher in the cabin. The RTA waypoint for the inbound leg was DENBY. After the aircraft sequenced DENBY, the inbound leg was completed and the flight crew prepared for the approach to KPHF.

D. Flight Test Metrics

Two primary metrics and a secondary metric were used to evaluate the success of the operational trials. It should be noted that performance metrics for the use-case operations—i.e., conformance to lateral path in the case of the DRNAV clearance or time error at the RTA waypoint—were not included in this evaluation. This was done for two reasons. First, this was the first trial of a prototype system in which the main objective was to test the injection of actual aircraft state data into the RPFMS, not the performance of the use-cases. Second, the performance model of the

aircraft was not at the correct level of fidelity to test the performance of the use cases, especially with respect to the RTA use-case.

1. Description of Primary Metrics

The first primary metric for this operational trial was the number of trajectory computation errors during each of the operational flight trials. These errors were categorized based on the process in the trajectory generator in which they were encountered (i.e., vertical trajectory generation errors, lateral trajectory generation errors), as well as which phase of the trajectory the error occurred in (climb, cruise, descent). The goal for this metric was an average of less than five trajectory errors per flight leg.

The second primary metric was the number of successful use-case operations conducted per leg per flight. This was a binary metric—either the operation was initiated and executed or it failed. The goal for this metric was that at least one (50%) of the use-cases on a given leg were successful.

2. Description of Secondary Metric

As previously mentioned, to mitigate negative effects due to the lack of an autopilot system in the aircraft, the researcher requested that the flight crew stay within the following bounds for altitude deviation, lateral deviation, and speed deviation:

- ±100 feet of vertical path,
- ±1 nautical mile of lateral path, and
- ±10 knots of indicated airspeed.

These bounds were used to evaluate how well the flight crew remained on path and speed throughout the operational flight trial, and the values of the bounds were derived from known issues on prior check flights. Their intended function was not to judge the skill of the flight crew, but to help provide root causes for failure to meet the primary metrics— i.e., it helped answer the question: "If the goal was not met for number of trajectory errors or number of successful operations, was it a result of the mitigation for lack of autopilot or was it due to another factor?"

To provide the researcher and development team with an answer to the secondary metric, a cost function was designed. This cost function sought to answer two main questions:

1. How well did the aircraft follow the vertical path, lateral path, and speed profile generated by RPFMS?
2. When the aircraft deviated outside of the bounds set by the researcher, what was the impact of those excursions on the flight?

The cost function was calculated as follows in Eq. (1):

$$J = \sum_{i=1}^{m} \left(A_i \left(B_i X_{1_i} + (1 - B_i) \left(\frac{1}{n} \right) \sum_{j=1}^{n} X_{2_{i,j}} \right) \right) \tag{1}$$

where:

- J is the value of the cost function,
- m is the number of error dimensions,
- A_i is the weighting factor based on impact of the i[th] error dimension,
- B_i is the weighting factor based on impact of full flight error versus peak errors for the i[th] error dimension,
- X_{1_i} is the normalized full flight error term for the i[th] dimension,
- n is the number of times that the aircraft deviated out of the containment bounds on a given flight, and
- $X_{2_{i,j}}$ is the normalized peak error term of the i[th] dimension for the j[th] occurrence that the flight deviated out of the containment bounds.

X_{1_i} is given by Eq. (2):

$$X_{1_i} = C_i \left(\frac{RMS(err_i)}{bound_i} \right) + (1 - C_i) \left(\frac{t_{err\_total_i}}{t_{flight\_total}} \right) \tag{2}$$

where:

- C_i is the weighting factor based on impact of magnitude of error versus duration of error for the i[th] error dimension,
- $RMS(err_i)$ is the root mean square of the i[th] dimension's error signal,
- $bound_i$ is the absolute value of the containment bounds of the i[th] dimension,
- $t_{err_{total_i}}$ is the total amount of time that the aircraft spent outside the containment bounds of the i[th] dimension during the flight, and
- t_{flight_total} is the total flight time.

$X_{2_{i,j}}$ is given by Eq. (3):

$$X_{2_{i,j}} = \frac{\int_{t_{j_{start}}}^{t_{j_{end}}} err_i \, dt}{\left(t_{j_{end}} - t_{j_{start}}\right)\left(max\left(err_i\big|_{t_{j_{start}}}^{t_{j_{end}}}\right)\right)} \qquad (3)$$

where:

- $t_{j_{start}}$ is the time of the start of the j[th] occurrence of a deviation out of the containment bounds,
- $t_{j_{end}}$ is the time of the end of the j[th] occurrence of a deviation out of the containment bounds, and
- err_i is the i[th] dimension's error signal.

For these flights, three error signals (err_i, $i = 3$) were used as inputs to the cost function in order quantify the behavior of the flight—cross-track error, vertical error, and speed error. The aircraft was determined to be out of its conformance bounds when the value of these error terms exceeded or fell below the upper or lower bounds respectively.

Cross-track error was calculated following the procedure set forth by Ryan, et al. in [10] for the "closest segment" alternative. Figure 8 demonstrated how this calculation is performed. The first step was to find the trajectory segment closest to the aircraft's current position. In general, the closest segment was the segment with the shortest perpendicular from the track point to the segment. In Figure 8, Q1 is the aircraft's current position (track point) and the line segment Q2-Q3 is the closest trajectory segment. The perpendicular is shown to intersect the segment at point Q4. The cross track error is then defined as the length of the line Q1-Q4.
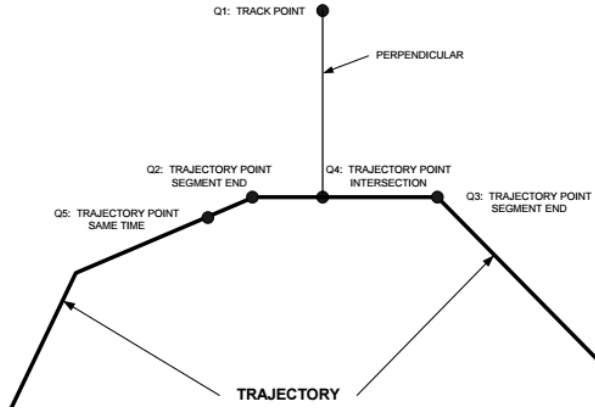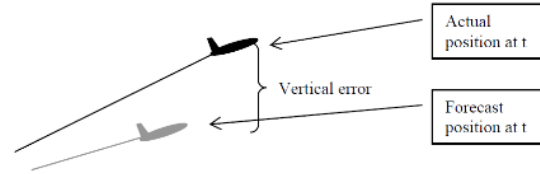
Vertical error (shown in Figure 9) was defined as the difference between the aircraft's current altitude at a given time and the altitude defined by the trajectory generated by the RPFMS at the same time. Finally, the speed error was defined as the difference between the aircraft's current speed at a given time and the speed delineated by the trajectory generated by the RPFMS at the same time.

Table 4 shows the weightings for the cost function. The weighting terms used in the cost function were set by the researcher based on subject matter expertise regarding which of the various terms would cause the RPFMS to fail. From discussions with subject matter experts, it was determined that the RPFMS is most sensitive to the vertical profile. The RPFMS attempts to null out the vertical error in cruise flight by generating a trajectory with either a cruise climb or cruise descent whenever the aircraft is more than 150 feet above or below the vertical path specified in the trajectory. For this reason, the researcher set the vertical error term and the peak vertical errors with the highest sensitivities, as can be seen in Table 4. The RPFMS is more robust to lateral and speed excursions, thus those weightings are set relatively low, and the impacts peak and full-flight errors are treated as equal for both error dimensions.

Figure 8: Cross-Track Error Calculation [10]

Figure 9: Vertical Error Calculation

Table 3: Cost Function Weightings

	Vertical	Lateral	Speed
Impact of Error Dimension (A_i)	0.75	0.15	0.1
Impact of Full Flight Error (B_i)	0.25	0.5	0.5
Impact of Peak Error ($1-B_i$)	0.75	0.5	0.5
Impact of Magnitude of Error (C_i)	0.5	0.5	0.5
Impact of Duration of Error ($1-C_i$)	0.5	0.5	0.5

Table 4: Cost Function Value Quantifiers

Value Range	Qualifier
[0 0.2)	Perfect
[0.2 0.4)	Great
[0.4 0.6)	Good
[0.6 0.8)	Marginal
[0.8 1.0]	Poor

Each of the error terms in the cost function are normalized to provide a sensible value to the researchers reviewing the data. If the values of the cost function equals 0, then the flight crew flew the exact guidance that the FMS dictated (i.e., no cross-track, vertical, or speed error) and maintained the aircraft within the bounds for the entirety of the flight. As the values of the cost approach a value of 1, it indicated that the aircraft was either not adhering to the FMS guidance for a majority of the flight, the aircraft had significant deviations outside of the containment bounds set by the researcher, or both conditions existed simultaneously. The quantifiers in Table 3 were associated with the values of the cost index.

E. Flight Test Results and Discussion

Overall, the flight test campaign was a success. The R2D2 platform was utilized for approximately 1 hour (approximately 30 minutes per flight) and generated 220 (111 and 109 in the first and second flights, respectively) trajectories during the two operational trials.

1. Number of Trajectory Computation Errors Results and Discussion

During the operational trials, no trajectory computation errors occurred. Thus, the stated goal of an average of less than 5 trajectory errors per flight leg was achieved. Figure 10 below shows a successful trajectory computation from the second operational trial.
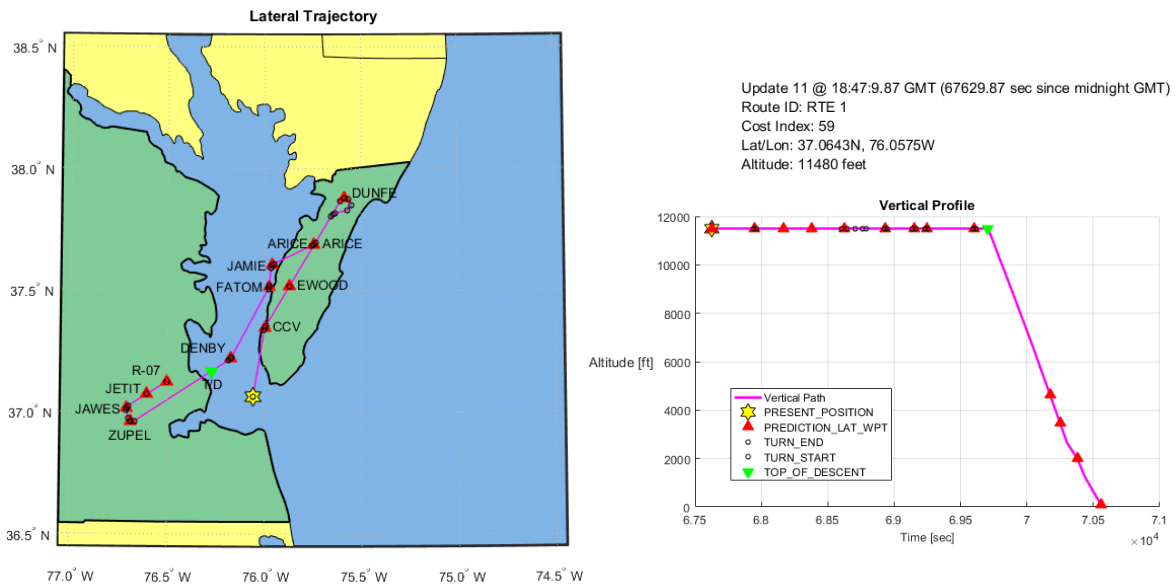

Figure 10: Successful Trajectory Computation. *The lateral trajectory is shown on the left, and the vertical trajectory is shown on the right. The yellow 6-pointed star indicates the aircraft's position when the trajectory was computed, the red triangles depict waypoints in the route, the black circles indicate the beginning and end of turns, and the green upside-down triangle represents the top-of-descent point.*

However, in the check flights prior to the operational trials, significant numbers of trajectory computation errors occurred. In one check flight there were 21 instances where the trajectory generator in RPFMS failed, and 134 instances where errors occurred in the cruise portion of the trajectory generator. Upon examining the data after this particular flight, vertical deviations from the planned cruise altitude (due to the lack of an autopilot system), while not extreme, caused the majority of these errors. As mentioned previously, the RPFMS attempts to null out the vertical error by building a cruise climb or cruise descent in the vertical profile. Additionally, since the RPFMS was originally designed for a large transport aircraft simulation, a limit for the minimum distance that an aircraft must be at cruise prior to starting its descent is set. However, the UC-12 is not a large transport aircraft, and the flight routes for the operational trials were not very long. Therefore, when the RPFMS in R2D2 attempted to build a cruise climb or cruise descent in the vertical profile in an attempt to null out the altitude deviations, the minimum cruise distance was not met, thus causing errors in the vertical component of the trajectory generator. To mitigate this particular error, the minimum cruise distance parameter in the R2D2 RPFMS was modified to be 10 nautical miles instead of 50. The results of this modification were observed immediately in the next flight—no trajectory errors occurred.

2. Number of Successful Use-Case Operations Results and Discussion

During the operational trials, all instances of the use-case operations (the DRNAV use-case and the RTA use-case) were completed successfully. Thus, the stated goal that at least 1 (50%) of the use-cases on a given leg were successful was met. Figure 11 below illustrates a successful DRNAV use-case operation. During check flights prior to the operational trials, no issues occurred with this use-case operation.
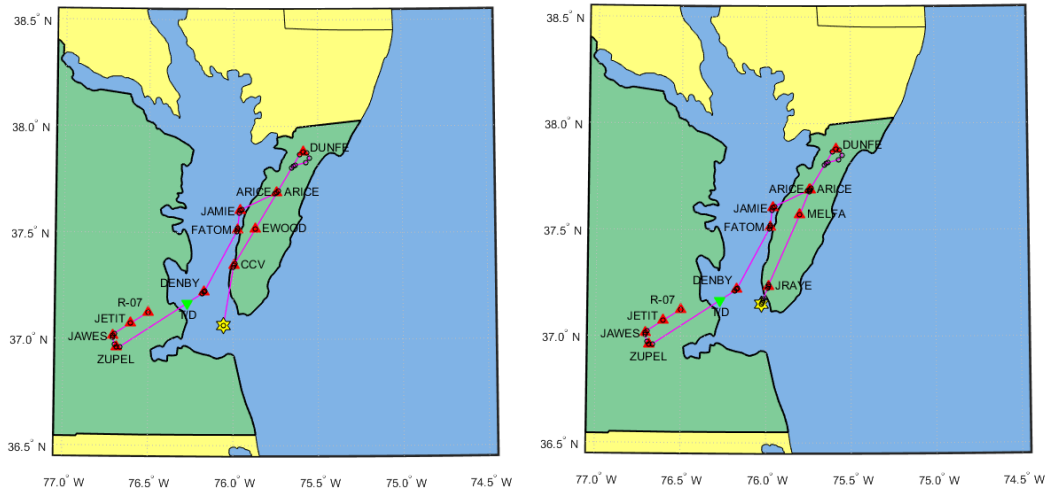

Figure 11: Original vs. DRNAV Trajectories. *The original trajectory that follows Route Option 1.B is shown on the left, and the trajectory computed after performing the DRNAV operation (Direct JRAYE, MELFA, rejoin at ARICE) is shown on the right.*

Additionally, Figure 12 demonstrates a successful RTA operation that was conducted during the second leg of the second operational flight trial. The figure shows the time error that the RTA algorithm was trying to null versus the time-to-go to the RTA point. As is evident in the figure, the time error was very small due to the flight crew's ability to closely follow the speed guidance provided by the R2D2 RPFMS and shown on the Pilot Display and Interface.
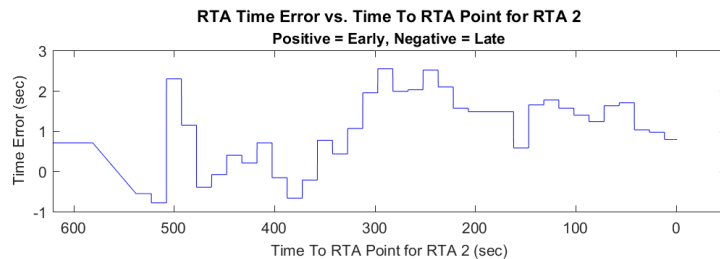

Figure 12: RTA Operation

However, in the check flights prior to the operational flight trials, several incidents occurred where the R2D2 platform was unable to perform an RTA operation. These problems were coupled to the issues associated with the trajectory generation errors mentioned in the previous section. After the modification to the R2D2 RPFMS minimum cruise distance was made, no issues with the RTA operation were experienced.

An operational issue with the RTA use-case was discovered when the RTA use-case operations were succesfully performed in both the check flights and the operational trials. During a few operations, the speeds presented to the flight crew to fly to achieve the required arrival time were outside of the flight crew's comfortable operating speed limits. This issue is attributed to the fidelity of the aircraft performance model in the R2D2 platform, and will be addressed in future work.

3. Results and Discussion of Cost Function Data

The cost function data for the two operational trials is shown in Table 5. As is evident by the results of the cost function and the associated qualifier for each flight, the aircraft was within the vertical, lateral, and speed bounds for a significant portion of the flight with a few deviations in each error dimension. The output of the cost function

supported the notes that the researcher took during the flight with respect to how well the aircraft conformed to the bounds that the researcher set.

F. Next Steps

Based on the results and lessons learned from the flight test campaign, two next steps have been identified to make this platform more useable in the future.

1. Integration with Ground-based TBO Tools

Several TBO tools and concepts involve the use of ground-based tools used by the ANSP in conjunction with airborne tools used by the flight crew. These concepts require communication links between the flight deck and ground systems such as ADS-B and DataComm. To test these concepts fully in-situ, a means by which to test both the concepts and communication requirements needs to exist. One proposed method, known as the Networked Air Traffic Infrastructure Validation Environment, proposes to emulate these data links by using software simulations of the communication links and in-flight Internet as the

Table 5: Cost Function Results

	Flight 1	Flight 2
Vertical Full Flight Component:	0.28073	0.22855
Vertical Peak Component:	0.64249	0.62979
Vertical Component:	0.55205	0.52948
Lateral Full Flight Component:	0.15604	0.17362
Lateral Peak Component:	0.9114	0.54777
Lateral Component:	0.53372	0.3607
Speed Full Flight Component:	0.46225	0.36421
Speed Peak Component:	0.70409	0.79717
Speed Component:	0.58317	0.58069
Cost Function Value For Flight:	0.55241	0.50929
Qualifier:	Good	Good

communication backbone. [11, 12] Furthermore, the NASA Shadow Mode Assessment using Realistic Technologies for the National Airspace System (SMART-NAS) Test Bed promises to "fill important gaps in the air traffic community's simulation and testing needs for allowing more efficient acceleration and acceptance of NextGen and far-term concepts and technologies." [13] The SMART-NAS Test Bed uses distributed communication to connect various data sources (e.g., SWIM, Weather Providers) with various ATM laboratories (containing ANSP simulators and Flight Deck simulators) and flight assets to validate concepts using multiple operational domains and investigate concepts related to revolutionary operations. [14] Finally, researchers at NASA Langley have developed a prototype capability that allows for demonstrations of some of the functionality of advanced TBO concepts, such as time-based metering, merging, and spacing and DRNP and DRNAV re-routing. This capability uses the SMART-NAS Test Bed to communicate with an ASTOR for concepts that require a flight deck component.

To integrate the R2D2 platform with ground-based tools, an in-flight Internet system must be installed on the aircraft. Once the in-flight Internet system is installed and tested, development can begin regarding the transmission of data from the aircraft to the ground system and vice versa using a combination of the Networked Air Traffic Infrastructure Validation Environment concept (emulations and simulations of the datalinks) and the SMART-NAS Test Bed (communication protocol and networking to ATM labs). Finally, for advanced TBO concepts that require both ANSP and flight crew interactions, data can be shared between the R2D2 platform and the TBO prototype in a similar manner to how data is currently shared between an ASTOR and the TBO prototype.

This step is viewed by the author as the most critical step to implement to realize the benefits of the R2D2 platform.

2. Aircraft Performance Model Refinement

The performance model incorporated in the R2D2 platform is of medium-low fidelity, as a result of the research and development team's decision to modify an existing business jet model to reduce the development effort. Two major differences between the jet performance model and the flight test aircraft are: the UC-12 is a twin-turboprop rather than a jet, causing issues with engine modeling, and the UC-12 does not have an autothrottle system like the business jet model, potentially causing errors due to inaccurate trajectory generation in the climb and descent phases of flight. Furthermore, the performance model lacked high-quality performance data; the only data available for developing the performance model for the UC-12 was the pilot's operating handbook.

Future testing of development activities and operational procedures, such as testing a new RTA algorithm and the procedures associated with it, will require a more accurate and higher-fidelity performance model. This is the second-highest priority modification that needs to be made to the R2D2 platform.

IV. Conclusion

This paper describes the design and development of an in-situ flight testing platform—R2D2—as well as the operational trials and results regarding the feasibility of the platform. The R2D2 platform provides the features of an advanced 4-dimensional flight management system, an avionics suite comparable to a modern large transport category

aircraft, and multiple communication systems while reducing the costs associated with these systems by using simulations and emulations. It provides the flight crew with guidance to perform TBO concepts, and provides a research capability to test these concepts in-situ. The R2D2 platform does not interact, or interfere, with flight control or safety systems, which ensures greater safety and reliability while reducing risk during a flight test.

The check flights and operational trials confirmed issues that were presumed to exist when integrating a software simulation with actual flight data; however, these issues were mitigated through modifications in the software and with refined flight test procedures. The two operational trial flights resulted in zero trajectory computation errors and the successful completion of the TBO use-case operations. An initial cost function was developed to quantify the conformance of the aircraft to the procedures set by the researcher; however, it must be refined to obtain more meaningful results.

The operational trials described successfully demonstrated that the R2D2 platform provides a timely and efficient means by which to test TBO tools and concepts in-situ by using emulations and simulations of avionics and communication networks.

Acknowledgments

References

[1] NextGen Joint Planning and Development Office, "Concept of Operations for the Next Generation Air Transportation System, Version 3.2," JPDO, Washington, D.C., 2011.

[2] NextGen Joint Planning and Development Office, "JPDO Trajectory-Based Operations (TBO) Study Team Report," JPDO, Washington, D.C., 2011.

[3] D. Finkelsztein, T. Lung, R. A. Vivona, J. Bunnell, D. Mielke and W. Chung, "Airspace and Traffic Operations Simulaiton for Distributed ATM Research and Development," in *AIAA*, 2005.

[4] M. G. Ballin, D. H. Williams, B. D. Allen and M. T. Palmer, "Prototype Flight Management Capabilities to Explore Temporal RNP Concepts," in *Digital Avionics Systems Conference*, St. Paul, 2008.

[5] M. T. Palmer and M. G. Ballin, "A High-Performance Simulated On-Board Avionics Architecture to Support Traffic Operations Research," in *AIAA Modeling and Simulation Technologies Conference*, Austin, 2003.

[6] RTCA, "DO-350A; Safety and Performance Requirements Standard for Baseline 2 ATS Data Communications," RTCA, Inc., Washington, D.C., 2016.

[7] Federal Aviation Administration, "Dynamic Required Navigation Performance Preliminary Concept of Operations, Version 2.0," Washington, D.C., 2014.

[8] RTCA, "DO-351A; Interoperability Requirements Standard for Baseline 2 ATS Data Communications," RTCA, Inc., Washington, D.C., 2016.

[9] NASA, "B-200 (UC-12B) - LARC," 3 February 2015. [Online]. Available: https://airbornescience.nasa.gov/aircraft/B-200_UC-12B_-_LARC. [Accessed 3 October 2016].

[10] H. F. Ryan, M. M. Paglione and S. M. Green, "Review of Trajectory Accuracy Methodology and Comparison of Error Measurement Metrics," in *AIAA Guidance, Navigation, and Control Conference*, Providence, 2004.

[11] M. McCormack, M. Underwood, A. Gibson, L. Miller, N. Dennis and M. Ballin, "Feasibility of a Networked Air Traffic Infrastructre Validation Environment for NextGen Concepts," in *Tenth USA/Europe Air Traffic Management Research and Development Seminar*, Chicago, 2013.

[12] J. Grisham, N. Larson, J. Nelson, J. Reed, M. Suggs, M. Underwood, Y. Papelis and M. Ballin, "Proof-of-Concept of a Networked Validation Environment for Distributed Air/Ground NextGen Concepts," in *AIAA Aviation Technology, Integration, and Operations*, Los Angeles, 2013.

[13] K. Palopo, G. B. Chatterji, M. D. Guminsky and P. C. Glabb, "Shadow Mode Assessment using Realistic Technologies for the National Airspace System (SMART NAS) Test Bed Development," in *AIAA Modeling and Simulation Technologies Conference* , Dallas, 2015.

[14] K. Palopo, "NASA Technical Reports Server," 24 October 2016. [Online]. Available: https://ntrs.nasa.gov/search.jsp?R=20170000219. [Accessed 21 March 2017].