Presented at 17th Aviation Technology, Integration, and Operations Conference, Denver, CO, June 2017.

A Wrapper Paradigm for Trusted Implementation of Autonomy Applications

Sally C. Johnson[1] and Jesse C. Couch[2]
Adaptive Aerospace Group, Inc., Hampton, VA 23666

Autonomous systems inherently have characteristics that make them difficult to verify and validate, such as nondeterminism, adaptation, and learning. A paradigm for implementation of autonomous systems monitored by a high-integrity run-time assurance wrapper is proposed. This paper examines the verification, validation, and certification challenges of autonomous systems, then presents the wrapper paradigm and discusses how the paradigm can be applied to address those challenges. An example wrapper design is described that is being developed for the Safe Autonomy Flexible Innovation Testbed (SAFIT^{TM}), which will enable safe flight operation of a wide range of autonomy applications by providing integrated flight protection, including traffic and obstacle avoidance, flight envelope protection, and geospatial containment.

I. Introduction

Autonomous flight operations have the potential to enable an explosion in growth of national and global aviation, with concurrent advances in the standard of living for millions of people worldwide. One of the key barriers to the introduction of autonomy into the National Airspace System (NAS) is lack of trust in autonomy by pilots, controllers, NAS regulators, and the public. The FAA is understandably cautious in certifying autonomous flight systems, since these systems have characteristics that make them especially difficult to verify and validate (V&V), such as being nondeterministic and possessing learning capabilities that modify their behavior to address non-routine, unanticipated situations.

Adaptive Aerospace Group, Inc., (AAG) is developing the Safe Autonomy Flexible Innovation Testbed (SAFIT^{TM}), which will enable safe flight operation of autonomy applications, such as complex and/or nondeterministic systems. SAFIT^{TM} is an Unmanned Aircraft System (UAS) testbed designed under a grant from NASA to support NASA's autonomy research and employs a new paradigm for trusted design, implementation, and real-world testing of autonomy applications, based on state-of-the-art wrapper technology. The SAFIT^{TM} wrapper, which is called SAFIT-Wrap^{TM}, will ensure safe flight operation of unproven prototype applications by providing integrated flight protection, including traffic and obstacle avoidance, flight envelope protection, and geospatial containment. While numerous geofencing applications have been developed[1], SAFIT-Wrap^{TM} provides a unique set of rigorous protections. Although originally designed as a NASA research testbed to support safe flight testing of unproven test applications, a high-integrity version of SAFIT-Wrap^{TM} is now being developed to enable implementation, certification, and operational use of a wide range of civil UAS applications. AAG has completed the design of the functional logic for the wrapper and created a simulation prototype[2] and has plans to develop and V&V a high-integrity version as a core flight management system to support operational use of autonomous flight applications.

II. V&V Challenges for Autonomous Systems

There are several characteristics that are typical of autonomous systems for aviation that present challenges for trusted implementation. The National Academy of Sciences conducted a comprehensive study of autonomy research for civil aviation[3] and identified a number of barriers to the development and certification of autonomous aviation applications. Certification of avionics requires demonstration that the system meets a rigorous set of safety, reliability, and operational performance requirements; however, it can be difficult to clearly define the system-level performance requirements or to detail the functionality required for an autonomous system that interacts with an evolving mission

[1] Senior Research and Development Engineer, member.

[2] Aerospace and Computer Engineer, member.

and operational environment. Even if the requirements can be defined, autonomous systems typically entail highly complex software with non-deterministic and adaptive behavior, so verification that the requirements are met is difficult. Additionally, most civil aviation applications will likely be implemented as a combination of autonomous software with real-time mission management and oversight by a human operator, necessitating V&V of the total human/machine system. Thus, certification must address the ability of the software, interacting with the human operator, to perform its intended function across a defined range of missions and operational environments, and to identify changes in function required within a bounded level of uncertainty about the evolving operational environment.

The US Department of Defense (DoD) conducted a detailed study of the role of autonomy in military systems[4], which provides additional insight into the effective design and evaluation of autonomous applications for aviation. Two key desirable design characteristics from the DoD report that are also applicable to civil autonomy applications are resilience and effective replanning. Resilience is a measure of the ability to adapt behavior to handle new missions or unexpected conditions, including operating effectively when surprises occur and gracefully degrading to a safe, lower level of performance in the presence of failures. For autonomous applications with human oversight, resilience includes providing real-time situation awareness to the operator, alerting when unexpected conditions are arising or when problematic performance trends are occurring, and graceful handover to the operator if necessary. It is important that the system not react with brittleness, such as sudden, unexpected failure to perform or abruptly turning control over to the operator after reaching an unsafe or unstable state. Effective replanning requires the system to recognize when the current plan is becoming invalid and to perform replanning that meets important constraints and is effective for the current situation. This may include recognizing the need for operator intervention and alerting the operator to the need for replanning. The DoD report also discusses the unique, challenging aspects of V&V of autonomous systems. Since autonomous systems are designed to operate in a dynamic environment and react non-deterministically to evolving conditions, the enumeration of test cases covering all combinations of conditions and then generating the full range of acceptable non-deterministic responses is unrealistic. Thus, new testing and evaluation techniques are needed.

III. Wrapper Paradigm

A wrapper is a state-of-the-art concept in the design of highly reliable systems, for example to allow use of Commercial-Off-The-Shelf (COTS) components that have not been subjected to rigorous V&V. The wrapper executes independently from the application being monitored and checks the outputs of the unproven software for one of the following:

- Correctness: The wrapper is able to ensure through simple analysis that the autonomy application has produced a correct solution, i.e., a solution that meets the specific full correctness criteria. An example would be an application that employs autonomy characteristics to find a solution to a complex mathematical problem, where the validity of the solution produced can be checked using a simple algorithm.
- Reasonableness: The wrapper is not able to fully check the validity of the solution, but is able to determine whether the solution meets certain criteria that should be true if the solution is reasonable. An example would be an algorithm that produces a series of outputs that should have a certain relationship with each other, where the wrapper could be monitoring for outliers or inconsistency in trends.
- Adherence to a set of safety properties: The wrapper is able to ensure through simple analysis that the autonomy application has produced a solution that is consistent with a set of predefined safety criteria. The optimality of the path would not be assessed by the wrapper. An example would be an application that employs autonomy characteristics to find an optimal path through a room , avoiding a set of obstacles. The wrapper would use simple algorithms to check that the path produced is a valid path that does not conflict with the obstacles. If the obstacles are stationary, the wrapper could check the full route before proceeding to move. However, if the obstacles can move or the route is planned incrementally, then the route should be checked in real-time, continuously verifying the safety or validity as the mission progresses. Continuous monitoring is typically more appropriate for aviation applications of autonomy because the validity and optimality if the solution will vary as conditions change and the mission evolves.

Not every application lends itself to the wrapper runtime assurance paradigm. The application must be one with outputs that can be checked using simple code that can be verified for correctness or bounded, and there must be a means for partitioning the unproven code from proven code. Additionally, timing issues must be addressed; i.e., if the application fails to complete execution, the wrapper must take control in a timely manner and produce a satisfactory or fail-safe solution using an alternate safe and deterministic algorithm. Many aviation applications of autonomy do in fact lend themselves to the wrapper paradigm. For these systems, the certification requirement of demonstration that the system meets a rigorous set of safety, reliability, and operational performance requirements may be able to be met by certification of a wrapper that monitors the complex autonomous software for compliance with a set of performance and safety properties and takes over as needed using simple, verifiable algorithms that are safe bu non-optimal. It may not be necessary to clearly define the functionality required for the autonomous system to interact with an evolving mission and operational environment as long as bounds can be placed on the mission and environment and performance and safety properties can be enumerated that cover the full range of missions and environmental parameters.

For an autonomous system that interacts with a human operator, a wrapper can support the desirable human/machine interface features discussed above, such as resilience and effective replanning. In addition to providing situation awareness and alerting information to the operator and receiving control commands, a wrapper can monitor performance trends and make the operator aware of anticipated problems. In some applications, a wrapper may also be able to aid in problem recovery by ensuring that the system is in a safe and stable state before handing off control to the operator. In a previous research project that AAG conducted for the FAA[5,6], a concept for a runtime assurance wrapper was developed to allow certified implementation of an uncertified general aviation autopilot. In the concept, the wrapper monitored the trends of the vehicle state parameters to determine whether the vehicle was trending toward an unstable state, such as an unrecoverable attitude. If so, the wrapper would alert the pilot and would use unsophisticated but highly reliable control algorithms to return the aircraft to straight and level flight before hand-off control to the pilot. In some cases, a wrapper can also be employed to determine when the current plan is becoming invalid, to recognize the need for operator intervention, and to alert the operator to the need for replanning.

IV. Application of the Paradigm in the Design of SAFIT™

The wrapper paradigm is being employed in the design and implementation of SAFIT™, a multipurpose UAS testbed designed to support a variety of research projects. The SAFIT^{TM} system requirements, vehicle, ground control station, and wrapper are described in the following subsections.

A. System Requirements

The system requirements for SAFIT^{TM} were defined with inputs from a wide range of NASA research projects, including UAS Traffic Management, various autonomy projects, and adaptive controls and controls upset research, and included a number of unusual requirements, including:

- Tolerance of +3 g to -2 g normal accelerations
- Inclusion of communications and operator interface to support Beyond Visual Line of Sight operations
- Ability to follow predefined waypoints or waypoints generated in real-time by test software
- Support for manual navigation from test software control commands
- Flight envelope, ground collision, and geospatial containment protections tunable to allow high-g maneuvers and maneuvering near the flight envelope boundary
- Redundant control surfaces that can be used to simulate stuck or damaged control surfaces and resilient controls configurations
- Ability to flight test user-supplied test wings, empennage, and tail

To meet these system requirements, the design for SAFIT^{TM} features four innovations that make it uniquely suited to supporting NASA's autonomy research:

- A reconfigurable vehicle design that enables a wide range of mission scenarios
- A robust aero-propulsive control system to ensure stability and controllability, with the ability to mimic a range of UAS test vehicle performance by setting parameters to limit the turn rate, climb rate, and power to be used for a given test period
- Safe flight evaluation of unproven prototype applications by providing integrated flight protection, including traffic and obstacle avoidance, flight envelope protection, and geospatial containment

- Support for variable levels of autonomy, ranging from full waypoint-based route plans to waypoints generated in real-time by test software to direct control inputs, with all the above being subject to the integrated protections

B. Vehicle Design

The vehicle design, shown in Figure 1, is a quad tilt-rotor with a high wing featuring an H-tail and redundant control surfaces. It features a Vertical Take-Off and Landing (VTOL) configuration with good cruise endurance supported by pivoting all four motor/propellers or a Conventional Take-Off and Landing (CTOL) configuration with more efficient cruise by optionally removing two of the motor/propellers, enabling a wide range of mission scenarios. While the design is scalable, the prototype design has a wingspan of nine feet to meet anticipated payload requirements and is powered by four 15-inch fixed mid-pitch propellers. Onboard electric motors and rechargeable batteries will support carrying a 6-lb payload for up to 30 minutes of CTOL flight at a cruise speed of 40 miles per hour. In the VTOL configuration, a 3-lb payload is supported, with up to 10 minutes of hovering flight. If a larger-scale vehicle is produced in the future to support higher payload weights, an onboard generator may be required. Construction materials include carbon fiber joiner and trunnion tubes, high-density foam and fiberglass surfaces and panels, thin-wall aluminum fuselage tubes, and poplar and birch plywood bulkheads.

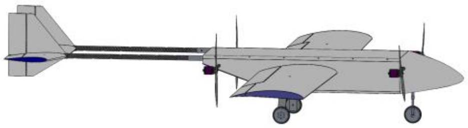

Figure 1. SAFIT^{TM} vehicle design

The highly modular design, as shown in Figure 2, supports reconfiguration of the vehicle to tailor performance for a specific mission, enables testing of user-supplied wing and tail panels, and facilitates transport of the vehicle. Expeditious access to test systems and ship systems in the fuselage is provided via a large removable panel on the top of the fuselage as well as removable nose and tail cones. This also provides for rapid replacement of batteries, which are mounted in the fuselage, between missions. Payload test systems are mounted on fuselage rails, which will hold the test modules securely in the event of high-load flight conditions during test missions.

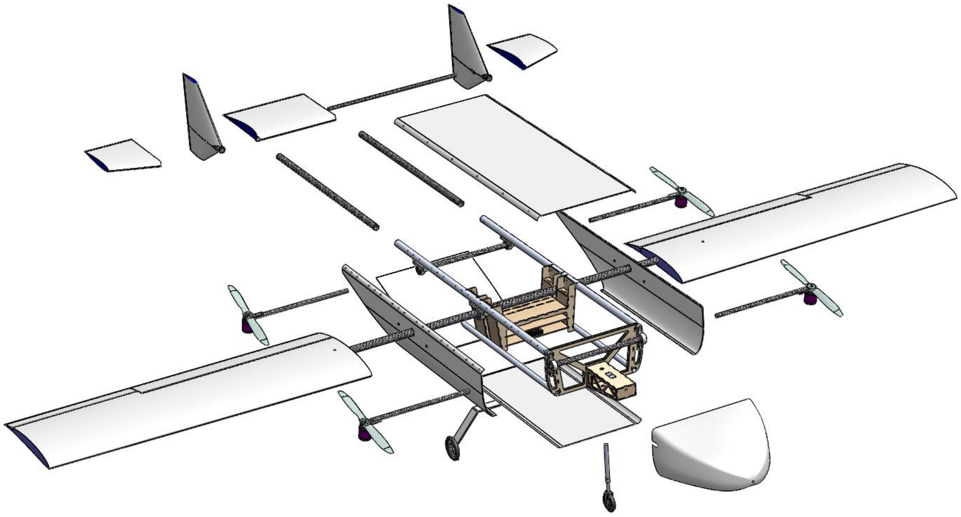

Figure 2. Modular vehicle design

C. Ground Control Station

The SAFIT^{TM} system includes a Ground Control Station (GCS), which supports the researcher in inflight communication of operational commands to the vehicle and displays real-time flight information. The GCS is hosted on a laptop computer with reconfigurable displays of the following: flight information in the format of Primary Flight Display/Multi-Function Display, forward-facing camera image, alerting and status information from UAS core and test modules, limited research data display capability, and a message board to support entering and transmitting information and commands to the vehicle, including inflight changes in autonomy control modes. The GCS will support implementation of simple automated checklists before and during a flight and will alert the operator of problems or status changes in the UAS system and test modules. The researcher can trigger an emergency landing via the GCS and can also send manual control commands to the vehicle. Advanced GCS functions are currently being designed, including alerting the operator that the vehicle is entering an area where its maneuvering will be severely restricted due to obstacles and supporting the operator in coordinating the management of multiple UAS.

D. SAFIT-Wrap^{TM} Design

SAFIT-Wrap^{TM} will ensure safe and economical flight evaluation of unproven prototype applications by providing integrated flight protection including traffic and obstacle avoidance, flight envelope protection, and geospatial containment. Variable autonomy levels will be supported, ranging from full autonomy with auto-takeoff and auto-land test systems to manual commands from the GCS, enabling testing of autonomous applications that produce either waypoint-based route plans or direct control inputs, with all the above being subject to the protections of SAFIT-Wrap^{TM}. During a flight test, the wrapper will monitor the current vehicle position and state information and check flight commands to provide envelope protection as well as traffic and obstacle avoidance and geospatial containment.

The Safety Critical Avionics Systems Branch at NASA Langley Research Center has been developing traffic and obstacle avoidance algorithms and applying formal methods to validate their safety properties for nearly ten years[7,8], and AAG has learned much from their work. Most of NASA's early algorithms were tailored to the performance of civil air transport and high-end general aviation aircraft, although NASA has recently developed and publicly released an integrated functionality suitable for UAS applications. Key SAFIT-Wrap^{TM} functionality for traffic conflict detection and avoidance is based on a set of open-source algorithms developed by NASA, which were combined with AAG-developed algorithms for waypoint management, obstacle detection and resolution, and geospatial containment tailored to a small UAS maneuvering in an urban environment, and integrated with envelope protection based on parameterized vehicle performance. AAG plans to build on NASA's formal methods work to develop a high-integrity implementation of SAFIT^{TM}.

The algorithms developed by NASA operate by calculating bands defining headings, altitudes and altitude change rates that result in conflicts, and AAG has continued to use this techique, creating combined bands that integrate the traffic and obstacle avoidance and geostationary containment components. Parameterized envelope protection is then applied to downselect to achievable conflict resolution maneuvers, based on command limitations for heading, altitude, vertical rate, and speed.

The effective integration of envelope protection, waypoint management, traffic and obstacle avoidance, and geospatial containment involves fairly complex logic. The individual functions require prediction that an unacceptable state is imminent, followed by suggesting a reaction to avoid the unacceptable state. States involving combinations of factors may require earlier detection to ensure all factors can be resolved. The logic must also address cases where a satisfactory solution cannot be found, prioritizing flight envelope protection vs. buffer violation (i.e., crossing into the buffer and perhaps even incurring a separation violation but avoiding a Near Mid-Air Collision violation) vs. airspace violation (i.e., traversing slightly outside the boundary of the geofenced area). Each avoidance protection is implemented with a buffer of additional protection space, which can be traversed when necessary to avoid a more important protection.

The waypoint management algorithm must determine a new path to reach a waypoint that is on the other side of an obstacle, ensuring delivery of the vehicle to all defined waypoints in the assigned order, even when the closest path may be blocked by the geospatial containment boundary or traffic aircraft. The logic must also consider when entering a constrained area between obstacles or traffic aircraft may be inadvisable.

V. How V&V Challenges Are Addressed in SAFIT^{TM}

A high-integrity version of SAFIT-Wrap^{TM} is now being developed as a core flight management system to enable implementation, certification, and operational use of a wide range of civil UAS applications. For most autonomous UAS applications, the primary safety concerns are that people and property are protected. SAFIT^{TM} will provide real-time assurance that each application meets a rigorous set of safety, reliability, and operational performance requirements, including flight envelope protection, traffic and obstacle avoidance, and geospatial containment. Additionally, SAFIT^{TM} will provide the interface with a human operator, ensuring resilience by providing real-time situation awareness to the operator, alerting when unexpected conditions are arising or when problematic performance trends are occurring, and graceful handover to the operator if necessary for replanning. For many applications, SAFIT^{TM} could obviate the insurmountable task of certifying complex, non-deterministic autonomous software.

The wrapper itself must be certified, which means that it must be demonstrated that the software, interacting with the human operator, will perform its intended function across a defined range of missions and operational environments, including functioning correctly within a bounded level of uncertainty about the evolving operational environment. Since SAFIT^{TM} will operate in a dynamic environment and react to evolving conditions, the enumeration of test cases covering all combinations of conditions and acceptable responses is still unrealistic. Thus conventional testing and evaluation techniques will be insufficient. AAG's vision is to build on NASA's formal methods work as a component of developing and conducting V&V on the high-integrity implementation of SAFIT^{TM}, including generation of a formal mathematical specification for the key high-integrity components of the software and formally verifying that the specification satisfies a limited set of safety properties necessary for safe multi-UAS operations.

V. Concluding Remarks

A key barrier to widespread use of autonomous systems in the operation of vehicles in the NAS is the inherent difficulty in verifying and validating these systems. A paradigm for implementation of autonomous systems monitored by a high-integrity run-time assurance wrapper has been proposed, and the use of the paradigm has been demonstrated with an example wrapper design that was created for SAFIT^{TM}, which is designed to enable safe flight operation of unproven autonomy applications by providing integrated flight protection including traffic and obstacle avoidance, flight envelope protection, and geospatial containment. The verification strategy for SAFIT^{TM} relies on judicious use of formal methods combined with partitioning and extensive testing, bringing a high level of rigor to verification of the core algorithms. The FAA has not yet adopted a certification standard for UAS maneuvering autonomously in the NAS; however, AAG plans to continue working with the FAA Small Airplane Directorate to ensure that the verification methods employed for SAFIT^{TM} are sufficient to meet any standard likely to be adopted for UAS, including for vehicles maneuvering autonomously to avoid traffic and obstacles, for a single operator handling multiple vehicles, and for Beyond Visual Line of Sight operations.

Acknowledgments

This work was conducted under NASA grant NNX15AV70G. Our Grant Technical Officer, Dr. Christine Belcastro of NASA Langley Research Center, provided us with valuable advice and enthusiastic support on this project. We are also grateful to the members of the Safety Critical Avionics Systems Branch at NASA Langley Research Center who have generously provided advice and feedback.

References

[1] Atkins, E., "Autonomy as an Enabler of Economically-Viable, Beyond-Line-Of-Sight, Low-Altitude UAS Applications with Acceptable Risk," *Association for Unmanned Vehicle Systems International 41st Annual Symposium*, May 2014.

[2] Johnson, Sally, Petzen, Alexander, and Tokotch, Dylan, "Exploration of Detect-and-Avoid and Well-Clear Requirements for Small UAS Maneuvering in an Urban Environment," *17th Aviation Technology, Integration, and Operations Conference*, June 2017.

[3] Committee on Autonomy Research for Civil Aviation, "Autonomy Research for Civil Aviation: Toward a New Era of Flight," National Academies Press, Washington, D.C., 2014.

[4] Defense Science Board, "Task Force Report: The Role of Autonomy in DoD Systems," Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, Washington, D.C., July 2012.

[5] Hook, Loyd R., Clark, Matthew, Sizoo, David, Skoog, Mark A., and Brady, James, "Certification Strategies Using Run-Time Safety Assurance for Part 23 Autopilot Systems," *IEEE Aerospace Conference*, Big Sky, MT, March 2016.

[6] Fuller, Justin, Hook, Loyd, Hutchins, Nathan, Maleki, K. Niki, and Skoog, Mark A., "Toward Run-Time Assurance in General Aviation and Unmanned Aircraft Vehicle Autopilots," *IEEE/AIAA 35th Digital Avionics Systems Conference*, Sacramento, CA, September 2016.

[7] Munoz, Cesar, Narkawicz, Anthony, Hagen, George, Upchurch, Jason, Dutle, Aaron, Consiglio, Maria, Chamberlain, James, "DAIDALUS: Detect and Avoid Alerting Logic for Unmanned Systems," *IEEE/AIAA 34th Digital Avionics Systems Conference*, Prague, Czech Republic, September 2015.

[8] Hagen, G., Butler, R., and Maddalon, J., "Stratway: A Modular Approach to Strategic Conflict Resolution," *11th AIAA Aviation Technology, Integration, and Operations Conference*, September 2011.