

Determinism and Autonomy in the National Airspace System (NAS)

James C. Neidhoefer*, Christopher S. Gibson†, and Maithilee Kunda‡
Aerotomy, Inc., Lithia Springs, GA 30122

and

Eric N. Johnson§
Georgia Institute of Technology, Atlanta, GA, 30332-0150

DOI: 10.2514/1.32696

Collaborative efforts between U.S. government agencies and industry have been underway to address certification and operational issues associated with integrating Unmanned Aircraft Systems (UAS) into the National Airspace System. This fact, combined with the obvious potential synergies between UAS and high-level autonomous systems, indicates a likelihood that in the near future, pressure from industry, the DoD, and research organizations will emerge to certify increasing levels of autonomy for UAS operations in the NAS. The complex and sensitive nature of the NAS will require that the characteristics of and requirements for these autonomous systems, some of which may involve sophisticated decision-making or interacting with humans, be carefully evaluated. One characteristic of autonomous systems that is highly relevant for NAS applications is that of determinism. This note defines and discusses two types of determinism. “Ideal determinism” is exhibited when, given exactly repeatable inputs and initial conditions, a system has exactly repeatable outputs. “Functional determinism” is exhibited when, given inputs and initial conditions that are repeatable within the measurement limitations of an observer, the system has repeatable outputs. Functional determinism in autonomous systems is desirable for their safe and effective integration into the NAS, both to maximize the potential benefits of such systems and to ensure that the operational environment of the NAS is not degraded for any stakeholders with respect to safety, organization, or ease of operation.

I. Introduction

RECENT work has been underway to integrate Unmanned Aerial Systems (UAS) into the National Airspace System (NAS). Currently, UAS are typically limited to operating in Special Use airspace, primarily in military operational or test areas. Only a few UAS operations in civil airspace have been allowed in recent years, with each mission evaluated and accepted on a case-by-case basis and tightly controlled by the Federal Aviation Administration (FAA). However, given: 1) very recent advancements in the FAA such as the issuance of “Experimental Airworthiness Certifications” intended specifically for UAS, 2) progress in collaborative efforts between industry and the U.S.

Received 08 June 2007; revision received 00 Month 2007; accepted for publication 28 September 2007. Copyright © 2007 by James C. Neidhoefer, Christopher S. Gibson, Maithilee Kunda and Eric N. Johnson. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 1542-9423/04 \$10.00 in correspondence with the CCC.

* CEO, Aerotomy, Inc., Lithia Springs, GA, 30122, AIAA Associate Fellow.

† CTO, Aerotomy, Inc., Lithia Springs, GA, 30122, AIAA Senior Member.

‡ Research Engineer, Aerotomy, Inc., Lithia Springs, GA, 30122, AIAA Student Member.

§ Lockheed Martin Associate Professor of Avionics Integration, Aerospace Eng. Dept., Georgia Institute of Technology, Atlanta, GA, 30332-0150.

government (FAA, DoD, and NASA), such as the now defunct Access 5 program,² the ASTM F-38 committee,^{3,4} and the RTCA Special Committee (SC) 203¹ towards developing federally sanctioned certification and operational procedures, and 3) advances in UAS technology itself such as advanced sensing devices, high-level autonomous algorithms, and new and sophisticated modes of communications, the near future will likely see more widespread military and civilian use of UAS in the NAS.

To date, the SC-203's work has been focused on three main areas: 1) developing and publishing a set of "Guidance Material" for integrating UAS into the NAS,¹ 2) addressing issues specific to UAS Command, Control, and Communications (C3) systems, and 3) addressing issues associated with UAS Detect, Sense, and Avoid (DSA) systems. The SC-203 Guidance Material addresses a broad spectrum of UAS-specific NAS-related practices, from certification to operational issues. The development of the Guidance Material has been led by SC-203 Working Group 1, while the C3 and DSA issues have been addressed by Working Groups 2 and 3, respectively.

One major area that, to date, has not been addressed in detail by the SC-203 is the use of autonomous systems in the NAS. While the majority of the SC-203 work has focused on human-controlled unmanned systems, continued progress in the development of advanced, autonomous unmanned systems has made the issue of autonomy in the NAS a salient one. Of course, there are many highly automated systems already used regularly with manned systems in the NAS, such as fly-by-wire autopilots; as technology advances, it is likely that the role of autonomy will become increasingly central to day-to-day operations in the NAS. Moreover, the upsurge of interest in high-level autonomous systems, for instance those that involve complex decision-making or interacting with humans, requires that the characteristics of and requirements for these systems be carefully evaluated before they can be introduced into the safety-critical NAS environment.

One such characteristic of autonomous systems that has not been addressed previously in literature or in the work of major committees is that of *determinism*. This note presents a case for the importance of determinism in autonomy for applications in the NAS. In particular, the primary goals of this paper are as follows:

- 1) To provide useful definitions of the abstract notions of ideal and functional determinism.
- 2) To show why determinism is an important quality for autonomous systems (both manned and unmanned) intended for applications in the NAS.
- 3) To illustrate the impacts of deterministic versus nondeterministic behavior on performance, safety, and human acceptance, using two exemplars of advanced guidance systems.
- 4) To confirm the feasibility of implementing deterministic, high-level autonomous systems on real-world aircraft.

II. Ideal vs. Functional Determinism

The definition of "ideal determinism" in a system is that, given full knowledge of some initial conditions, which may include the system's current state as well as a set of control inputs, the outputs of the system can be determined exactly. This definition of ideal determinism corresponds with the more general notion of predictability; that is, if a system is completely predictable, then we say it is deterministic, while if its outputs exhibit some measure of randomness, we say the system is nondeterministic.

In any complex environment composed of multiple, interacting, autonomous systems, a system that is deterministic provides additional information to other systems (namely the knowledge of what it is likely to do next), which can then be used to improve decision-making and other behaviors and interactions among all of the systems that are involved. For example, consider the simple scenario of a human pilot in the NAS whose aircraft is approaching another aircraft head-on, faced with the choice of turning either right or left to avoid a collision. If the pilot knows that the other aircraft will turn towards the right, then he can also turn right and successfully avoid a collision. On the other hand, if the pilot knows only that the other aircraft will turn right or left with equal probability, then the best he can do is guess a uniform random direction for himself. In a hypothetical situation restricted to constant altitude, both aircraft will turn either right or left half of the time and thereby remain safe, but the other half of the time, they will have chosen corresponding opposite directions and be headed for a collision.

This scenario, of course, is one of the numerous situations for which desired actions are explicitly encoded in aviation regulations: "When aircraft are approaching each other head-on, or nearly so, each pilot of each aircraft shall alter course to the right" (FAR Section 91.113e).⁵ In essence, many rules and regulations such as this one define

specific, deterministic modes of behavior for operators in the NAS (e.g. for trained pilots, air traffic controllers, etc.), which, because they are known to everyone, enable individuals both to predict what others will do in order to react appropriately and to follow prescribed patterns in order to be predictable to others. Without this prior knowledge, an individual's actions would be arbitrary and situation-specific, which would greatly reduce the levels of safety, organization, and coordination in the complex, interaction-rich environment of the NAS.

This line of reasoning provides compelling motivation for preferring that autonomous systems operating in the NAS behave in a deterministic manner. Consider the same situation given above of two aircraft approaching head-on except with the agent in question being an autonomous UAS instead of a human pilot. If the UAS were operating using a non-deterministic collision avoidance algorithm, then it might turn left and continue to be on a collision course with the other aircraft. Of course, the UAS would hopefully have more advanced collision avoidance capabilities than just executing a single turn, and so it might turn left and then turn right, successfully avoiding a collision after all. But if the other aircraft were piloted by a human, this maneuver would no doubt appear confusing. Furthermore, even if the other aircraft were another UAS with a similar collision avoidance algorithm, having to do a double swerve would likely be both less safe and less efficient than if both UAS were following a consistent, deterministic pattern of actions.

The point illustrated by this example is that, when evaluating the behavior of a system, there is more to consider than just the outcome of the behavior, e.g. while two UAS may perform collision avoidance with equal success, if one UAS uses a non-deterministic algorithm while the other uses a deterministic algorithm, then there may well be substantive reasons to prefer one over the other. Of course, the deterministic system may not always be the “best” choice; one can easily imagine situations in which determinism would be undesirable in an autonomous system, for instance in a UAS operating in a war zone filled with adversarial agents.

However, for the particular environment of the NAS, where autonomous systems (manned or unmanned) must cooperate extensively with other autonomous systems as well as with people, utilizing determinism will likely facilitate certification of such systems and increase levels of safety, organization, user comfort, and performance.

An important point is the subtle difference between ideal and functional determinism. For applications within the NAS, it would be desirable for a system to not only satisfy the definition of ideal determinism (as stated at the beginning of Section II) but also to be “functionally deterministic,” i.e. deterministic within the measurement limitations of relevant observers. This distinction is made because, while an autonomous algorithm with certain inputs and initial conditions may give exactly repeatable outputs in ideal, laboratory-like conditions, a complete autonomous system operating in the real world that employs this same algorithm may not exhibit deterministic behavior to external observers, due to the influence of process and measurement noise. The key to understanding this distinction lies in the fact that in the real world, it is never possible to have full knowledge of the *exact* state of a system (and thus *exactly* repeatable inputs and initial conditions are not possible). Observations made by an observer will likely have some associated error, for both human and autonomous observers. For example, if a person were visually observing a UAS hovering at some point, it would be impossible to define the exact 3D coordinates of that point. Even if the UAS' own state measurement could be communicated directly, there would be some noise in that measurement itself.

How, then, does the reality of imperfect observation affect the notion of whether a system is functionally deterministic? If a system exhibits significantly different behaviors for different inputs or initial conditions that are not distinguishable in the presence of measurement noise, then even though the system is ideally deterministic (exactly the same inputs result in exactly the same outputs), it may appear nondeterministic to an observer (indistinguishably small changes in the inputs result in large changes in the outputs). Consider, for example, a UAS that turns right to avoid a collision if its relative heading to an obstacle is greater than or equal to zero degrees and turns left otherwise. Now suppose that the UAS encounters two obstacles in succession, with the first obstacle at $+0.01^\circ$ and the second at -0.01° . To an observer who is able, perhaps, to distinguish heading angles up to a $\pm 1.0^\circ$ resolution, these two encounters appear to have identical initial conditions, with each obstacle *appearing* to be at $\sim 0.0^\circ$ degrees in front of the UAS. However, the UAS' choice of different actions in each situation could lead the observer to conclude that the UAS exhibits nondeterministic behavior. Thus, one can say that the UAS is functionally nondeterministic with respect to this particular observer, since it is impossible for the observer to successfully predict its behavior, given that whether the UAS has a heading of $+0.01^\circ$ or -0.01° is itself dependent on many random factors in the real world.

Therefore, in evaluating whether an autonomous system is deterministic, it is important to take into account any relevant observers along with their specific observational capabilities. Given the context of the NAS, this means that it is desirable for autonomous systems to be functionally deterministic within the measurement limitations of other NAS users, both human and autonomous, such as UAS operators, other pilots, Air Traffic Control (ATC), or other autonomous systems. In particular, it is undesirable for an autonomous system to exhibit apparent *chaotic* tendencies, i.e. where small changes in the initial conditions can lead to large changes in the system's output behavior.

Of course, to say that a deterministic system exhibits or lacks chaotic tendencies is a subjective judgment that depends entirely on how one defines "small" changes in initial conditions (how small?) and "large" changes in outputs (how large?). As the use of autonomous systems becomes more prominent and widespread, quantitative assessments of the level of determinism exhibited by these systems will likely have to be made, taking into account both the behavior of the various systems as well as their intended functions, in order to properly establish criteria and preferences for safely and successfully integrating them into the NAS.

III. Determinism in an Autonomous Guidance System

The need for functional determinism in autonomous systems used in the NAS will likely span a wide range of applications, from algorithms for guidance and control functions to interface tools designed to reduce pilot or air traffic controller workload. In this section, we use a concrete example from the realm of autonomous guidance to illustrate the implications of functional determinism and how it can be achieved in autonomous systems. However, guidance is by no means the only area in which considerations of determinism can have an impact on the effective integration of autonomous systems into the NAS.

Guidance is, though, a good area with which to illustrate the importance of functional determinism, because the delineations between deterministic and nondeterministic guidance systems are easy to visualize. Also, autonomous guidance systems are highly relevant, as the recent NASA funded National Research Council Decadal Survey of Civil Aeronautics⁷ recommended "advanced guidance systems" (such as formation flying, collision avoidance, swarming, etc.) as the highest national priority research and technology topic in the area of dynamics, navigation & control, and avionics.

We focus in this discussion on the specific application of *autonomous guidance for collision avoidance*. This problem is an interesting one to consider for two reasons. First, it requires achieving the basic guidance functions of trajectory planning and collision avoidance and thus is a good benchmark for advanced guidance systems. Second, it is useful for illustrating that autonomous systems developed for operations in the NAS are equally relevant to both manned and unmanned aircraft. Contemplation of the "human" factor introduced when autonomous systems must interact with manned vehicles enables a more intuitive and pronounced understanding of why determinism is needed in these systems.

In the following subsections, we first present two specific approaches to autonomous collision avoidance—one uses a popular algorithm that is ideally deterministic but functionally non-deterministic, and another method that is equally effective at avoiding collisions but is functionally deterministic. Simulation-based results illustrate the significant differences in behavior generated by these two approaches. Finally, we present experimental flight results that indicate the practical feasibility of designing and implementing deterministic autonomous systems on real-world aircraft.

It is noted that the presentation of the following examples is not to promote a particular collision avoidance approach but rather to facilitate comparison of deterministic and nondeterministic approaches to solving a specific problem and to demonstrate that autonomous systems can be designed to operate in a functionally deterministic manner and be implemented successfully in the real world.

A. Nondeterministic Versus Deterministic Approaches to Autonomous Collision Avoidance

The following potential-field-based, multi-ship collision avoidance example was implemented to demonstrate how a guidance technique that may be ideally deterministic may not also be functionally deterministic (deterministic within the measurement limitations of an observer). This example involves five identical vehicles with symmetric starting conditions, four of which are trying to reach a goal on the far side of a collision point in the middle. A simulation was run four separate times using all of the same initial parameters but with a small amount of noise added to the initial conditions of the starting locations of the vehicles. Each of Figure 1 through Figure 4 shows

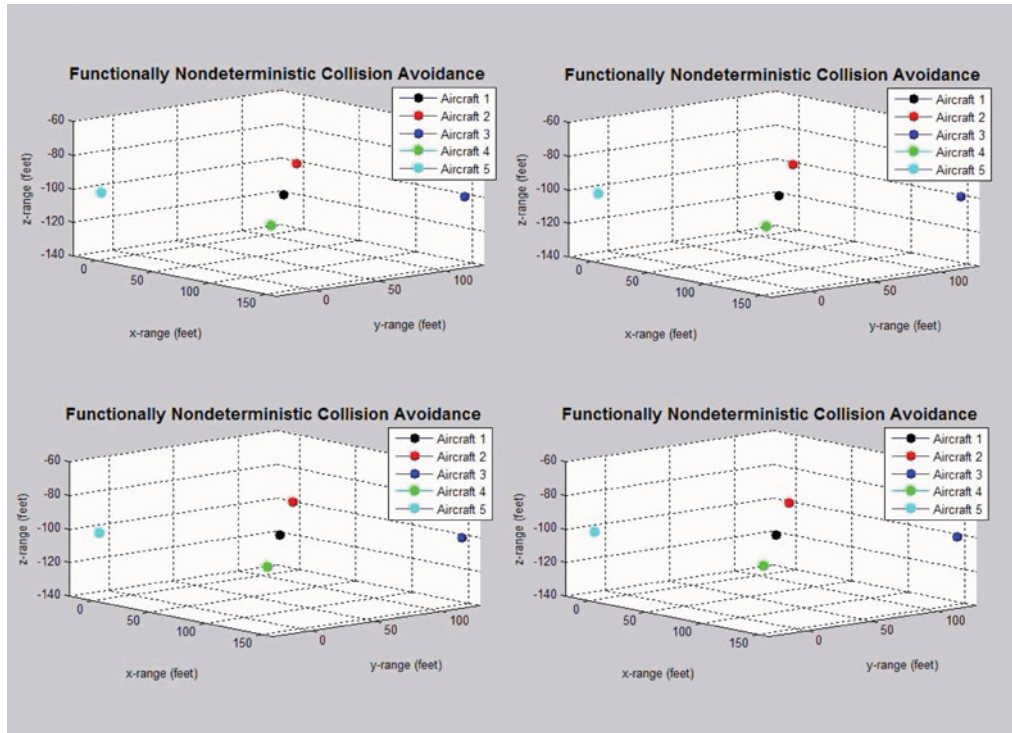


Fig. 1 Potential field-based collision avoidance no. 1.

snapshots of these four simulated runs taken at various times. Clearly, even though the simulation was begun with nearly identical initial conditions in each case (see Figure 1), the behaviors resulting from the potential-field-based collision avoidance method are drastically different (see Figure 4). To a pilot sitting in one of the vehicles, or to an air traffic controller observing the situation, the starting situation among the various cases would be indistinguishable (Figure 1). The early behavior of the system would also be fairly predictable (Figure 2). However, as the vehicles close and autonomously avoid collision using the potential-field-based method, any possibility of prediction is overwhelmed by the extreme variations in the vehicles' behavior (Figure 3 and Figure 4).

To fully understand the significance of functionally nondeterministic autonomous systems interacting with manned systems, imagine being a pilot onboard one of these vehicles. Your vehicle is on a collision course with several other vehicles, each under the control of a functionally nondeterministic autonomous system. Now, imagine that although you may have been in this same situation before, even dozens of times, you still have no idea how the other vehicles in the situation will react. The other vehicles might fly over your vehicle, under it, or fly around it with no forewarning of intent. Would this be an acceptable situation?

In contrast to this type of functionally nondeterministic example, it is possible to implement autonomous guidance systems using a functionally deterministic approach. Figure 5 shows four snapshots in time of a multi-ship collision avoidance maneuver similar to the one described above, in which each vehicle is equipped with a functionally deterministic autonomous guidance system. Not only are the results extremely predictable and repeatable, even with the addition of significant noise to the initial conditions, but the red and green vehicles are also following the "rules of the road" observed by manned aircraft (i.e., when approaching another aircraft head-on, one should pass on the right).⁶ This behavior is highly deterministic and completely controllable.

B. Feasibility of Functionally Deterministic Approaches

The primary objective of this section is to present results from an actual UAS flight test experiment of an autonomous collision avoidance system that exhibits functionally deterministic behavior. The intent of this

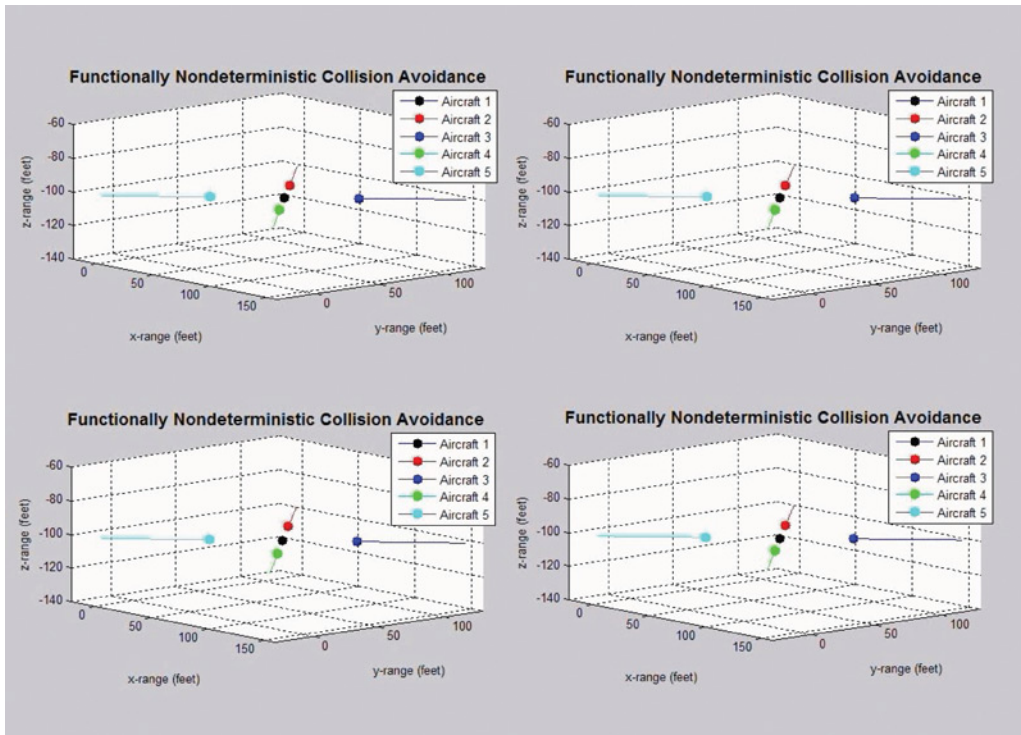


Fig. 2 Potential field-based collision avoidance no. 2.

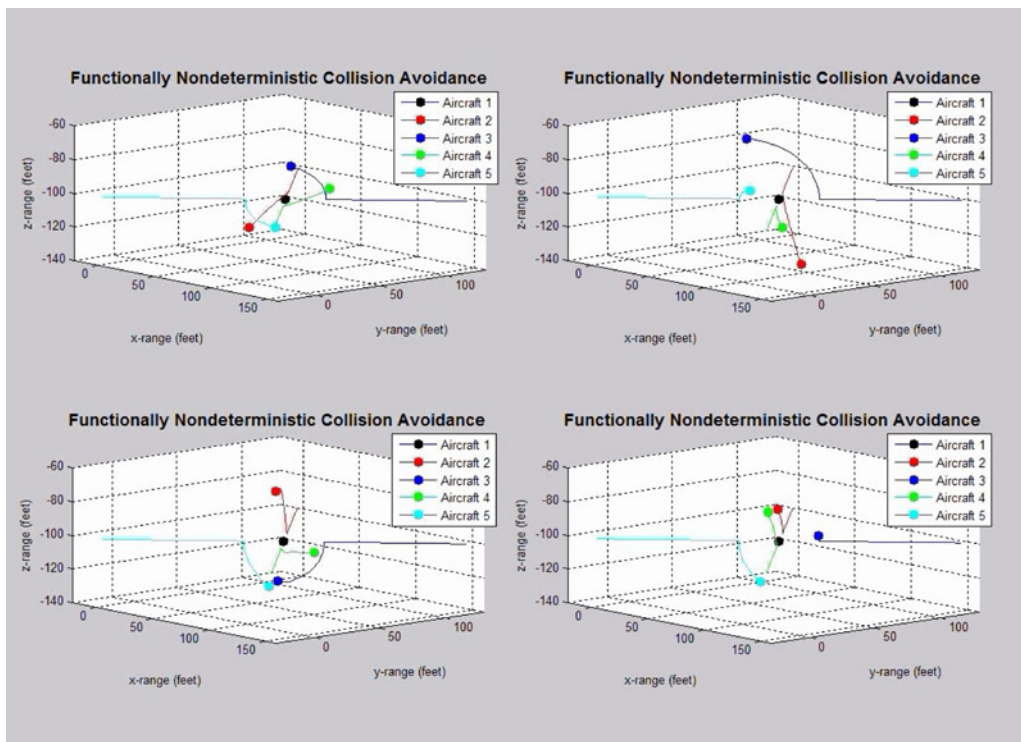


Fig. 3 Potential field-based collision avoidance no. 3.

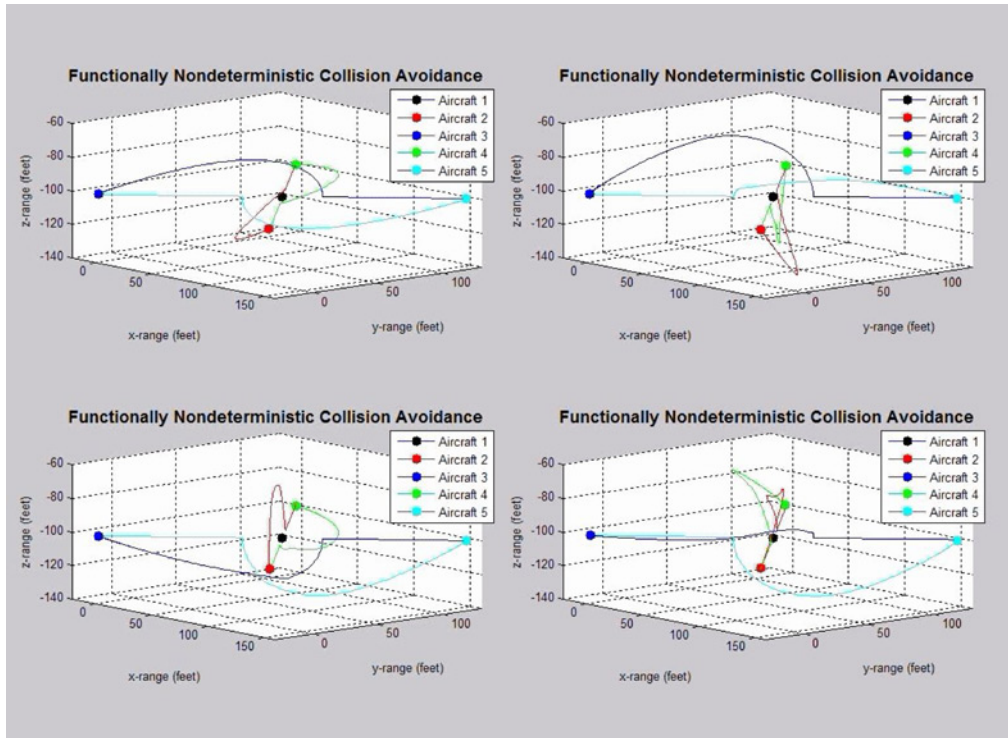


Fig. 4 Potential field-based collision avoidance no. 4.

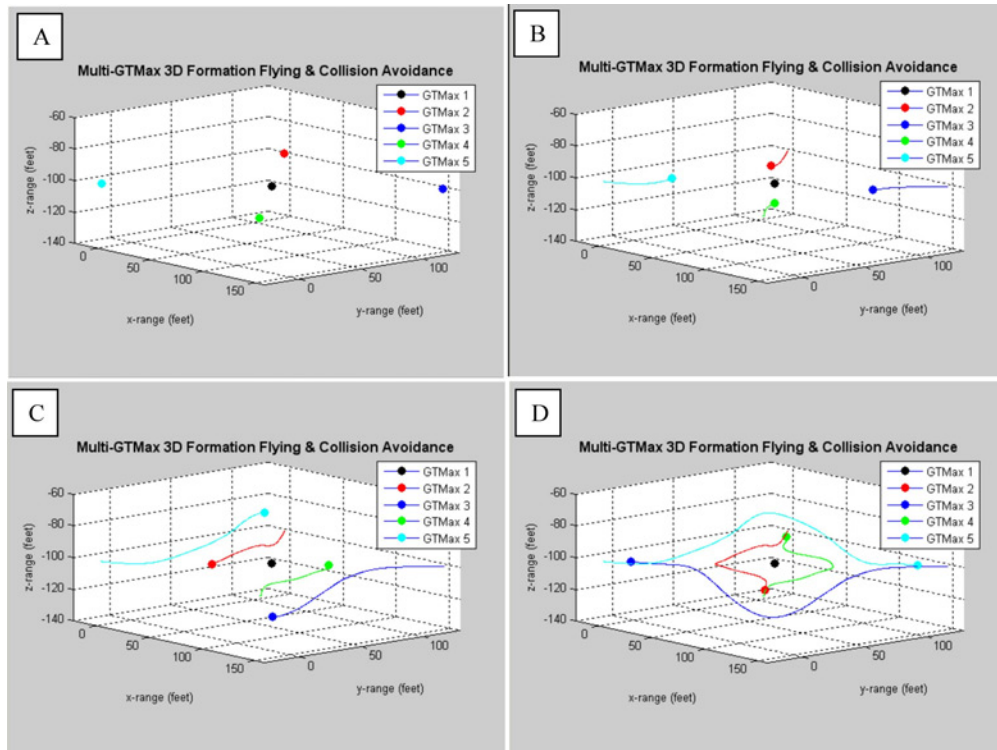


Fig. 5 A multi-ship collision avoidance maneuver with the AFFS.




Fig. 6 Flight test data: functionally deterministic autonomous collision avoidance.

presentation is not to promote a particular collision avoidance algorithm or approach but rather to demonstrate that practical implementations of highly autonomous functionally deterministic systems are possible.

A scenario was created in which a desired baseline trajectory for the test UAS would cause a collision with another simulated aircraft. The autonomous guidance system then had to predict the collision and generate a safe, efficient modified trajectory to avoid the collision and still achieve the waypoint goal. For this experiment in particular, the simulated aircraft was commanded to remain in a stationary hover, while the UAS⁸ was commanded to autonomously move from a position 75 ft. to the right of the hovering aircraft to a position 75 ft. to the left of the aircraft. The straight, baseline trajectory generated for this waypoint of course intersected the hovering aircraft's position, requiring the UAS to fly directly through the hovering vehicle. Therefore, as the UAS began to move down the baseline commanded trajectory, the autonomous collision avoidance system predicted a collision and generated a safe, modified trajectory to avoid the hovering aircraft while still achieving the final commanded position.

An additional objective of this experiment was to demonstrate that the modified trajectory commanded by the autonomous collision avoidance system could be completely controlled using different parameter settings in the guidance algorithm. That is, the UAS could be commanded to take alternate routes in order to avoid a collision in a deterministic (but still autonomous) manner.

To demonstrate this functionality, a series of five different maneuvers were commanded using the same collision scenario as the initial test described above. Each maneuver was specified so that the UAV would avoid the hovering aircraft by autonomously taking a route 1) directly behind the aircraft, 2) above and behind the aircraft, 3) directly above the aircraft, 4) above and in front of the aircraft, and 5) directly in front of the aircraft, respectively. Figure 6 shows flight test data for all five maneuvers performed as part of this flight test plan. It can be seen that, in addition to avoiding the hovering aircraft while achieving waypoint goals, the guidance system successfully generated a trajectory according to the specified route in each test run.

IV. Conclusions

It is likely that in the near future, pressure from industry, the DoD, and research organizations will emerge to certify increasing levels of autonomy for UAS operations in the NAS. One characteristic of autonomous systems

that is highly relevant for NAS applications is that of determinism. This note defined and discussed two types of determinism: “ideal determinism,” which is exhibited when, given exactly repeatable inputs and initial conditions, a system has exactly repeatable outputs, and “functional determinism,” which is exhibited when, given inputs and initial conditions that are repeatable within the measurement limitations of an observer, the system has repeatable outputs.

It was concluded that functional determinism in autonomous systems is desirable for their safe and effective integration into the NAS, both to maximize the performance and potential benefits of such systems and to ensure that the operational environment of the NAS is not degraded for any stakeholders with respect to safety, organization, or ease of operation. In other words, it is undesirable for autonomous systems in the NAS to exhibit *apparent* chaotic tendencies, i.e. where small changes in the initial conditions can lead to large changes in the system’s output behavior.

In addition, it was concluded that it is possible to implement autonomous systems that are functionally deterministic. Flight test results were presented to support this conclusion.

Acknowledgments

This work was supported in part by NASA under contract #NND06AA32C and the U.S. Army under contract #W911NF-05-C-0092. The authors would like to acknowledge and thank Paul Mays for his contributions to the work presented in this paper.

References

- ¹RTCA A. Available at: <http://www.rtca.org/comm/Committee.cfm?id=45>, last accessed on 2/25/2007.
- ²The Access 5 Program, available at: http://www.nasa.gov/centers/dryden/research/HALE_ROA/index.html, last accessed on 2/25/2007.
- ³The ASTM F38 Technical Committee. Available at: www.astm.org/COMMIT/COMMITTEE/F38.htm, last accessed on 2/25/2007.
- ⁴ASTM Committee F38, “Standard Specification for Design and Performance of an Airborne Sense-and-Avoid System”, Standard F 2411-04, ASTM International, July 2005.
- ⁵Federal Aviation Regulations, Regulatory and Guidance Library. Available at: http://www.airweb.faa.gov/Regulatory_and_Guidance_library/rgFAR.nsf/, last accessed on 2/25/2007.
- ⁶Aeronautical Information Manual. Available at: <http://www.faa.gov/ATpubs/AIM>, last accessed on 2/25/2007.
- ⁷The National Research Council, “Decadal Survey of Civil Aeronautics”. Available at: <http://www.nap.edu/catalog/11664.html>, last accessed on 2/25/2007.
- ⁸Johnson, E., Schrage, D., Prasad, J., Vachtsevanos, G., “UAV Flight Test Programs at Georgia Tech”, Proceedings of the AIAA Unmanned Unlimited Technical Conference, Workshop, and Exhibit, 2004.

Michael Hinchey
Associate Editor