Mapping a Path to Autonomous Flight in the National Airspace

Kenneth N. Lodding[1]
NASA Langley Research Center, Hampton, VA 23681

The introduction of autonomous flight, whether military, commercial, or civilian, into the National Airspace System (NAS) will present significant challenges. Minimizing the impact and preventing the changes from becoming disruptive, rather than an enhancing technology will not be without difficulty. From obstacle detection and avoidance to real-time verification and validation of system behavior, there are significant problems which must be solved prior to the general acceptance of autonomous systems. This paper examines some of the key challenges and the multi-disciplinary collaboration which must occur for autonomous systems to be accepted as equal partners in the NAS.

I. Introduction

Let me put it this way, Mr. Amor. The 9000 series is the most reliable computer ever made. No 9000 computer has ever made a mistake or distorted information. We are all, by any practical definition of the words, foolproof and incapable of error. ~HAL in 2001: A Space Odyssey by Arthur C. Clarke

IN fiction, autonomous systems do not have a good reputation. From the Jewish folk tale, *Golem of Prague*, through Mary Shelley's *Frankenstein*, to Fred Saberhaen's *Berserkers* series, as well as the infamous *Cylons* of *Battlestar Galactica*, automatons are generally not presented as society's benefactors. Even when not terrorizing the world, they are often seen as paranoid – *HAL* of 2001 and *Marvin* of *The Hitchhikers Guide to the Galaxy* – or worse, as bumbling tin men looking for a brain. These literature plots reflect their author's perspective of potential problems that can, or might result from the introduction of autonomous systems into the real world. Although the introduction of autonomous flight into the National Airspace will be technically challenging, the most difficult part may be getting good press for the automatons.

To accomplish the smooth introduction of autonomous aircraft into the National Airspace System (NAS) we need significantly more knowledge than simply knowing how to design, implement, test, and verify an autonomous system. We need to understand not only the individual autonomous system, but also how collections of tens or hundreds of them interact with each other, as well as with piloted traffic, all vying for maneuvering room within the NAS. Methods need to be developed to understand how unexpected interactions can occur within the control system of an autonomous aerial vehicle, and the potential for unexpected emergent behavior in the swarms of autonomous unmanned aerial systems (UAS) moving within the NAS. And when the understanding of how to build well-behaved autonomous systems is learned, it will be necessary to develop the capabilities to be able to formally verify and validate that these systems are safe: safe for themselves, safe to other aerial vehicles, and safe to anyone or thing on the ground. But the most daunting problem will be developing the skills and knowledge to make autonomous systems trustworthy. For at the bottom line, allowing autonomous aerial systems into the NAS is all about trusting our systems.

II. Autonomous Systems

At the most fundamental level an autonomous system is a system *capable of operating in the real-world environment without any form of external [i.e. human] control for extended periods of time.*[1] More completely, the system needs to be acting in a self-directed and self-governing fashion; able to perceive its environment and to act upon it so as to achieve pre-specified goals. The level of autonomy may vary during operation. This behavior is seen both in remotely piloted vehicles, as well as in crewed aircraft. In this scenario, human control of the system maybe absent for long periods of time, with the human only taking control at critical phases of the operation. This paper

---

[1] Computer Engineer, Flight Software Systems Branch, Mail Stop 472, NASA Langley Research Center.

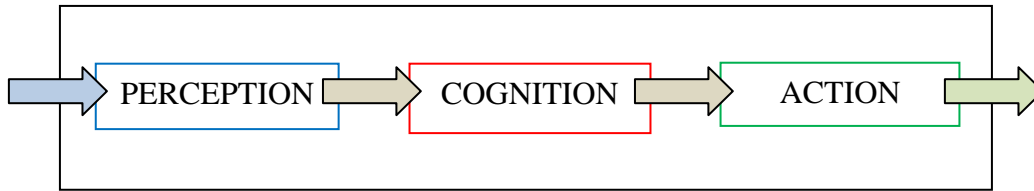American Institute of Aeronautics and Astronautics

Figure 1. Autonomous System Structure

will focus on the fully autonomous aircraft concept, viewing 'variable autonomy' architectures as a subset of a fully autonomous system.

The definition of autonomy just given maps nicely into a system architecture defined by three fundamental, high level capabilities.

Perception: *The system must be able to gather information about its environment.*
Cognition: *The system needs to be able to learn about its environment and make decisions based upon what is learned.*
Action: *The system must be able to act upon its environment.*

These requirements provide a basic set of high-level software and hardware capabilities necessary for the implementation of an autonomous aerial vehicle. All autonomous systems will contain to varying degrees these meta-capabilities, with the fundamental flow of information between them as shown schematically in Fig. 1. While the boundaries between primary functions are shown as being well defined, in actuality the functions could tend to blend together, often seemingly overlapping: is object recognition perception or cognition or perhaps a little of both? Each of the primary components further decompose into their own internal hierarchies of processing levels which are tuned for specific functionality within the general overarching functional structure.

III. Perception

The most fundamental UAS function for safe operation is the ability to perceive the active, real-time world of its environment. From taxi and runway operations, take-off, cruising flight, weather monitoring, and landing, a successful UAS must be able to perceive its environment in a manner which is minimally equivalent to a human pilot, and ideally superior. Sensors able to operate in parts of the spectrum and at distances not visible to humans will be able to provide raw data across multiple sensory channels to provide maximum information about the adjacent airspace. Figure 2 presents a highly schematic view of the perception process.

Machine vision consists of taking rapid snap shots of the local environment in the measurement space of the sensor. For visual sensors this is roughly 380 to 750 nm of the visible spectrum. As the fidelity of images in this part of the spectrum can be hampered by environmental conditions, sensors operating in other parts of the spectrum can be added to the sensor suite. IR sensors operating in the approximately 1 mm to 750 nm ranges will see the environment differently, as will UV and radar. The sensors are not restricted to the electromagnetic spectrum, as acoustic sensors could potentially provide additional information. This raw data is fused together to build a coherent picture of the local environment across the monitored spectrum.

However, data are not information and that conversion, *data to information*, is the goal of the remainder of the perception process. Once the environment is 'sensed', perception involves capturing the gist of the scene: what in the scene is important (salient) and what can be safely ignored at this time. In biological systems the perception process prioritizes scene components, moving primary focus to rapidly approaching, nearby objects, a desirable property for a UAS perception component.

The key capabilities required of the *perception* system are:

- **Detect hazards**: Sense for other aircraft, weather, and terrain.
- **Track hazards**: Gain data and confidence about the potential hazards, including validity of hazard and accuracy of the location and trajectory.
- **Evaluate threat of hazard**: Determine if immediate reaction is required, or simply add to global world model.
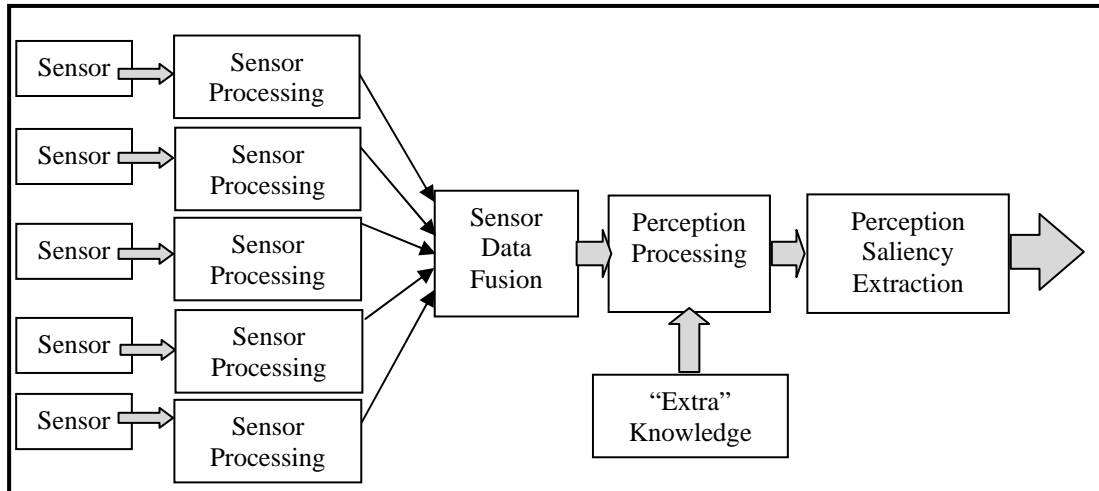
American Institute of Aeronautics and Astronautics

Figure 2. Perception Processing Structure

IV. Cognition

If perception is seen as the process of *converting data to information*, then cognition can be viewed as the continuation of that process: *converting information into plans*. The cognitive stage is hierarchically structured in typical *subsumption architecture*[2] style. See Fig. 3. For this paper, three primary levels are defined: *reactive, supervisory,* and *deliberative*.

Perception information enters the cognitive processing at the reactive layer, the lowest in the cognitive structure. The reactive layer provides for an immediate response to critical information by forwarding a copy of the information, without 'thinking about it,' to the action component for immediate reaction. One example of such behavior could possibly occur in a 'sense and avoid' situation. Conceptually, the reactive layer provides a similar function to the biological reflex action.

Situated above the reactive layer is the layer dealing with short term, tactical goals: the supervisory layer. At this level we are processing flight management issues as well as issuing action commands to perform local maneuvering, for example entry into a landing pattern. The supervisory level insures that action commands issued are within the current system constraints, such as vehicle maneuvering speeds.

The top level of the cognitive process performs the higher level functions mostly associated with strategic decisions. Man-machine and machine-machine communications also occur at this level in the generic architecture. Example strategic function performed at this level includes prioritizing goals relative to current and expected future vehicle fuel and power availability, predicted weather, and sensor capabilities.

The key capabilities required of the cognition system are:

- **React:** Respond to high priority events from the perception system. Note that these can be internally generated events signaling system problems (e.g. equipment failure), as well as external events, such as intruding aircraft on a crossing trajectory.
- **Supervise:** This is the flight management component of the system. This includes tracking targets, prioritizing them, issuing action commands to direct the vehicle trajectory, and in general perform the *piloting* of the vehicle. Previously this function would have been performed by a human pilot and termed *flight control.*
- **Deliberate:** UAS's are introduced into the airspace to achieve specific goals and tasks. These can range from autonomous package delivery to monitoring boarders for homeland security. The deliberation or strategic level manages the global goal/task functions relative to the current and predicted future states of the vehicle. As appropriate, the goals are expressed as tactical requirements and passed downward to the supervisory level.
- **Command:** The end result of the cognitive process is the generation of flight management commands which are passed to the action component for execution. The commands at this point in time are high level statements of intent: *climb to FL35 at rate of n FPM on heading… .*
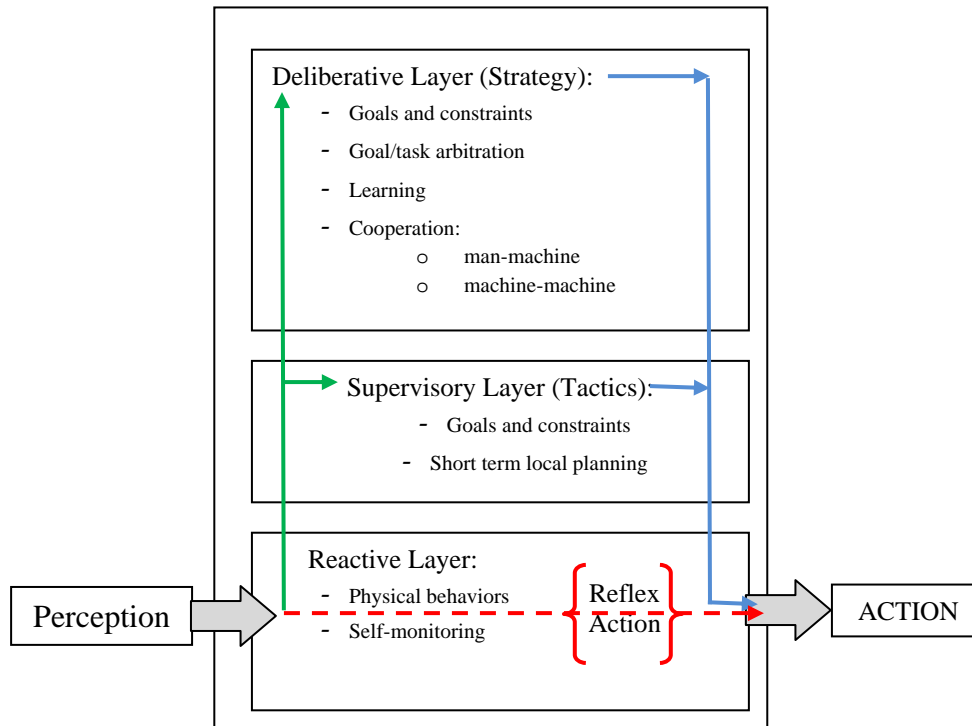
Figure 3. Cognitive Subsumptive Architecture

V. Action

Action refers to the ability of the autonomous system to act upon its environment. In the context of this paper it is simply the application of flight control inputs to alter the speed, and/or trajectory of the aircraft. The action component of the architecture accepts inputs from the cognition component, decomposes the request into appropriate actuator commands for engine and flight controls and issues the commands. The system self-monitors the action components (e.g. flight control surfaces) and performs localized corrections as necessary.

The key capability required of the action system is:

- **Execute:** Execute the commanded maneuver. This may require the decomposition of the command into a series of sub-commands which are actually executed by the vehicle flight control system.

VI. Scalability

A critical factor to the introduction of large numbers of UAS's into the NAS is the high complexity of multi-vehicle coordination. The design and implementation of a single UAS is challenging, but generally seen as a doable task. Even small *swarms* of autonomous vehicles, acting as a cohesive flight group can be envisioned working in the NAS. However, the *George Jetson* world view of untold numbers of autonomously flown air vehicles, scooting about the airspace like bumper-cars gone wild is not an easily digestible image. It is difficult to imagine ground based air traffic control for large collections of independent vehicles, so the answer must lay with onboard sensors and the development of suitable collision avoidance algorithms. In the one-on-one scenario, the implementation of sense and avoid may be directly understood: resolve the conflicted trajectory. As the number of agents increases in local, congested portions of the airspace, such as terminal areas, the accuracy, reliability, and predictability of conflict avoidance algorithms becomes critical as the number of potential conflicts increases and the available airspace for maneuver decreases. Verifying and validating algorithms of this complexity will be a critical component to the successful introduction of large numbers of autonomous aircraft into the NAS, but not the key one.

VII. Trusting Autonomous Systems

Autonomous systems by themselves are complex. Any non-trivial implementation intended to operate in the NAS will consist of numerous interacting components, coupled together in a labyrinth of interconnected software and hardware modules of unprecedented complexity. Expand the complexity by increasing the number of players, autonomous or human piloted, in the same game space and the complexity grows exponentially. Complexity permeates the game space. Any part of the system can potentially affect any other part of the systems, making the difficulty of predicting and controlling the interactions exceed our current control and safety abilities.

Key to successfully introducing UAS's into the NAS is that they are perceived by all of the other players in the NAS as being trustworthy. Trust will not be found in the formal methods of verification and validation; at least not the kind of trust that humans experience when trusting other humans. Formal methods to prove the correctness of algorithmic system will certainly be a part of any attempt to introduce UAS's into the NAS. However, that places the focus on the mathematical verification of a mathematical algorithm. But the final acceptance of autonomous systems into a space containing humans will be that the humans trust the system in a manner isomorphic to how they trust other humans. And that means we need to learn how humans understand the concept of *trust*. What signals flow between humans to allow a building of mutual trust? Certainly the specifics change depending upon the circumstances, but there are some overarching behaviors which facilitate the learning of trust between individuals. We will need to understand those and design ways to ease the process of introducing UAS's into the NAS.

The first components of building up trust are relatively simple. Is the system reliable? Can it perform the function or functions required? These are the questions asked by humans of any new team member. A more subtle question is whether the action of the system is predictable. Is the behavior exhibited expected and appropriate for the current circumstances, or has a new, unexpected behavior emerged? Can such unexpected behaviors by life threatening? This leads to the question of moral issues in learning to trust. Does a crashing UAS attempt to avoid hitting a school building, or does it simply crash and burn without concern for the building occupants? Society trusts its pilots to be make such ethical decisions under the stress of unexpected failure. Can we ask less of our autonomous systems?

The ability to specify, design, implement, test and validate the numerous software algorithms needed to control a UAS are critical parts of the puzzle, it is the human-machine trust that is key to acceptance of UAS's into the NAS. It is trust that will make UAS's a full-fledged member of the airspace. What is fundamentally required for humans to accept and trust autonomous systems is "…to be able to understand their reasoning process and the factors that precipitate certain actions. The machine needs to be able to communicate the reasoning behind its actions in an unambiguous manner that is also accessible to non-technical personnel – in other words, to *explain itself*."[3]

VIII. Conclusions

The conceptual architecture sketched in this paper illustrates a possible functional decomposition for autonomous systems. Consisting at the highest conceptual level of only three functional modules, any actual implementation will quickly expand these functions into a web of interconnected software and hardware modules of unprecedented complexity.

As hardware becomes more failure resistant, the complexity of coupled software is becoming more fragile because of the number of pathways with possible interactions which thread their ways through the system. Autonomous systems, by their nature, consist of large numbers of interacting modules, with inputs from numerous sensors, and issuing commands to a myriad of action devices. This results in a system that must not only handle hardware component failure and software anomalies ('bugs'), but also error conditions which potentially arise from the interactions of correctly functioning modules. The difficulty of the problem increases as the numbers and types of piloted and autonomous systems simultaneously interacting grows.

The complexity of the technical problems does not negate the desirable goal of integrating autonomous aircraft into the NAS. It is going to happen. Developing the autonomous flight management software will be the easy part; designing software to be trusted by all of the players will be the hard part. Developing methods to formally validate the accuracy and reliability of a mathematical algorithm is necessary, but not sufficient for developing trust. The final 'yes or no' acceptance of UAS's into a system containing humans will be made by the humans. And that means we need to understand how humans view the concept of *trust* and how machines can exhibit the necessary behaviors to win that trust.

References

[1] Beckey, G. A. *Autonomous Robots*, Cambridge, MA, MIT Press 2005.

[2] Brooks, Rodney A., *A Robust Layered Control System For A Mobile Robot*, A. I. Memo 864, Sept. 1985, Massachusetts Institute of Technology, Artificial Intelligence Laboratory.

[3] Kröske, J., O'Holleran, K. & Rajaniemi, H. *Trusted Reasoning Engine for Autonomous Systems with an Interactive Demonstrator*. 4th SEAS DTC Technical Conference 2009. URL: http://www.science.mod.uk/codex/issue5/journals/documents/codex5_journals_2.pdf