

Share the Sky: Concepts and Technologies That Will Shape Future Airspace Use

Mark G. Ballin¹

NASA Langley Research Center, Hampton, VA, 23681

William Cotton²

National Institute of Aerospace, Hampton, VA, 23666

and

Parimal Kopardekar³

NASA Ames Research Center, Moffett Field, CA, 94035

The airspace challenge for the United States is to protect national sovereignty and ensure the safety and security of those on the ground and in the air, while at the same time ensuring the efficiency of flight, reducing the costs involved, protecting the environment, and protecting the freedom of access to the airspace. Many visions of the future NAS hold a relatively near-term perspective, focusing on existing uses of the airspace and assuming that new uses will make up a small fraction of total use. In the longer term, the skies will be filled with diverse and amazing new air vehicles filling our societal needs. Anticipated new vehicles include autonomous air vehicles acting both independently and in coordinated groups, unpiloted cargo carriers, and large numbers of personal air vehicles and small-scale point-to-point transports. These vehicles will enable new capabilities that have the potential to increase societal mobility, transport freight at lower cost and with lower environmental impact, improve the study of the Earth's atmosphere and ecosystem, and increase societal safety and security by improving or drastically lowering the cost of critical services such as firefighting, emergency medical evacuation, search and rescue, border and neighborhood surveillance, and the inspection of our infrastructure. To ensure that uses of the airspace can continue to grow for the benefit of all, a new paradigm for operations is needed: equitably and safely sharing the airspace. This paper is an examination of such a vision, concentrating on the operations of all types of air vehicles and future uses of the National Airspace. Attributes of a long-term future airspace system are provided, emerging operations technologies are described, and initial steps in research and development are recommended.

I. Introduction

From the ground to the edge of space, the airspace over our nation is one of our most valuable treasures. This enormous natural resource is available to all citizens to share for scores of purposes, both terrestrial and airborne. Since the dawn of powered flight, access to the National Airspace has been a right. This freedom of access to the skies has been responsible for incredible innovation and it has fueled the economic engine of commerce. The National Airspace System (NAS), which provides regulation and safety services to those who use it, is massive and complex, providing services to over fifty thousand air transportation flights carrying over 1.6 million passengers each day. To provide the required services, today's air navigation service provider employs over 14000 air traffic controllers and managers at 600 air traffic control facilities, utilizing seventy-one thousand pieces of ground

¹ Project Scientist, NASA NextGen Concepts and Technology Development Project, M.S. 156A, NASA Langley Research Center, Hampton, VA, Associate Fellow, AIAA.

² Consultant, National Institute of Aerospace, Hampton, VA.

³ Project Manager, NASA NextGen Concepts and Technology Development Project, M.S. 210-8, NASA Ames Research Center, Moffett Field, CA.

equipment. While the services provided by the NAS are critical, the current approach to airspace management may impede our airspace's utility to society over the long term. Constraints include limited flexibility of operations, a lack of scalability in both accommodation of demand and in costs for services provided, and an inability to predict safety impacts of changes. The Joint Planning and Development Office (JPDO) Next Generation Air Transportation System (NextGen) modernization plans¹ promise some relief from current costly constraints to aviation's growth, but even that effort is limited in its ability to truly unleash the innovative forces for beneficial use of the airspace.

There are also many enthusiastic innovators of new air vehicles as well as several new business models for operating them within the National Airspace. An early adopter market can exist for these vehicles and their operations. By combining the use of new technologies applied to aeronautical challenges with new approaches to modernization, aviation's future can truly fulfill the expectations of a new generation of Americans. However, there are several highly challenging technological bottlenecks that may slow or even prevent such growth. Because these bottlenecks can only be removed through long-term research and development (R&D) commitments, the R&D community must play a critical role.

In this examination of aviation's future, we present the vision of shared use of the airspace by all. Attributes of a future system based on this vision are provided, and emerging operations technologies that can be applied to aviation are described. The paper describes some fundamental R&D issues that should be initiated or given additional emphasis, and it discusses some emerging approaches to investigating the issues. Many of the future operations described rely on advanced air vehicles. This paper does not concentrate on the future vehicles or their enabling technologies. Instead, the paper concentrates on the operations of vehicles in the National Airspace, focusing on technologies that directly support those operations.

II. The Vision: Share the Sky

The airspace challenge for the United States is to protect national sovereignty and ensure the safety and security of those on the ground and in the air, while at the same time ensuring the efficiency of flight, reducing the costs involved, protecting the environment, and protecting the freedom of access to the airspace. Many visions of the future NAS hold a relatively near-term perspective, focusing on existing uses of the airspace and assuming that new uses will make up a small fraction of total use. Scheduled airline operations can be envisioned in this environment, and they will almost certainly continue to exist over the long term. However, scheduled airline operations may make up a much smaller fraction of total airspace operations in the future. The long-term goal must be to accommodate both those who wish for a scheduled passenger service as well those who have other uses.

The skies of tomorrow will be filled with diverse and amazing new air vehicles filling our societal needs. In addition to the vehicle types we have today, anticipated new vehicles include autonomous unpiloted air vehicles acting both independently and in coordinated groups, unpiloted cargo carriers, and very large lighter-than-air freight-carrying vehicles. The airspace must also accommodate an anticipated explosion in growth of personal air vehicles and small-scale point-to-point transports. These vehicles will enable new capabilities that have the potential to increase societal mobility, transport freight at lower cost and with lower environmental impact, improve the study of the Earth's atmosphere and ecosystem, and increase societal safety and security by improving or drastically lowering the cost of critical services such as firefighting, emergency medical evacuation, search and rescue, border and neighborhood surveillance, and the inspection of our infrastructure. Further, the future airspace must accommodate uses that do not involve vehicles at all. Tethered aerial windmills may take advantage of high-altitude winds to provide a significant fraction of our future electrical energy needs. Therefore, the number of vehicles and other systems using the National Airspace in the long term will probably be orders of magnitude greater than the number we have today.

Therefore, we must develop new concepts of operation and technologies that reduce or eliminate a vehicle's burden on other vehicles and the supporting infrastructure. Without reducing this burden, future airspace access may be limited and the airspace may not be able to accommodate the anticipated new uses. We call this vision *Share the Sky* because, taken together, the enabling capabilities permit to the extent possible the safe coexistence of all conceivable flight missions.



Fig. 1. Safe coexistence of many airspace users will be critical in the future.

III. New Airspace Uses

Consider a typical day in the future. A flight to work begins by backing the vehicle out of the garage and moving to the center of the cul de sac for a vertical transition to flight toward the office. The driver enters the destination and commands the start of flight, but the automation takes the vehicle where it needs to go, avoiding obstacles and other airborne vehicles as necessary on the way. After landing on one of the several pads in the parking lot, the vehicle taxis to a parking spot. Precipitation and visibility are irrelevant to the operation, and the driver needs no special skill to “pilot” this vehicle. Some of the neighbors flew from home to airports at the edge of town to transfer to airliners flying individually optimized trajectories to distant destinations. Rental flying vehicles at these destination airports are used to complete the journey to places of business, vacation spots or the homes of friends and relatives.

Back at work, supplies and products are received and delivered by small package freight companies sending unpiloted craft to origin locations for package pick up and, after transfer to high speed flight among regional airports, delivering the packages to their individual destinations in more unmanned craft. As illustrated by Fig. 1, these personal aircraft will routinely share the airspace with each other as well as with vehicles carrying out other types of operations. The figure illustrates future airliners in climb and cruise phases of flight, small autonomous morphing “organic” aircraft of the future that will serve in a variety of roles, and future recreational air vehicles.

High altitude aerial windmills turn high-speed winds into massive power generation that is fed through their tethers into the power grid for use at home, work and play. The tether locations are dynamically tracked and distributed in real time to all airborne aircraft in the vicinity whose tactical conflict detection and resolution systems avoid them with the highest integrity, just as they do mountain peaks and other aircraft. Some of the generated power lights the tethers and the windmills, creating new stars in the sky. Figure 2 illustrates one concept of an aerial windmill.² The day’s commerce is supported by massive data communications economically transmitted over extremely high bandwidth pathways between aerial data relay stations. Some of these are tethered in place above hilly terrain, kiting and generating power when the wind is blowing and using the stored energy for sustaining lift to keep the platform on station when the wind is calm. It continuously remains motionless in inertial space despite

changing and gusty wind conditions. The wind energy is transmitted to ground storage through the tether and drawn from the storage to generate lift through the same line, as needed.

Other travelers bound for a wedding board the airliners of tomorrow. They operate from launch and recovery areas built above the freeways coming into the cities. Ground transportation in all modes is available down the elevator from the air operation. These airliners fly along trajectories dynamically computed in response to the existing wind field to minimize the wind miles to destination while avoiding severe turbulence and optimizing fuel performance. Sensors on the aircraft continuously send current atmospheric data to the flight planning systems, and automated aircraft systems ensure separation from all other aircraft and restricted airspace.

Swarms of robotic vehicles protect migratory birds as they herd the birds away from approach and departure paths, keeping airliners safe from bird strikes. Other free flying, cooperative swarms of small robotic vehicles are called out to perform search and rescue missions, collectively covering large areas quickly with optical, radio frequency and other specialized sensor technologies to detect human signatures against varied terrain backgrounds.

Some surface traffic received an automatic detour notice this morning to avoid a bridge that had been closed for emergency repairs detected by a specialized crack sensing flyer patrolling under the bridge. These flying robots ensure safety while providing long life extensions to our infrastructure of roads, bridges, dams, buildings, power lines, and communications towers. While flying to work, some commuters noticed a formation of heavier lift robots below tending to crops. The once dangerous job of crop dusting is now performed safely and with extreme precision by these robotic flying tankers, dispersing precise amounts on only the desired areas. The national forests are similarly protected by heat seeking robotic flyers dispatched by satellite surveillance to investigate suspicious hot spots, identifying the sources as either harmless campfires or potential forest fires, and if necessary, to call in heavy lift tanker vehicles and direct them to drop water on the exact location of the threat before it has a chance to grow and spread. When fires do occur, autonomous swarms of specialized vehicles self-coordinate to fight the fire and evacuate people (Fig. 3).

Hybrid, lighter than air/direct lift vehicles carry heavy cargo loads to and from non-airport locations with very high efficiency, performing better than ground transport in both speed and gas mileage, especially where the roads between origin and destination are circuitous. The combination of buoyancy and direct lift through aerodynamic means enables operations from the rooftops and parking lots of warehouses, eliminating the need for a surface transportation mode between factories and distribution warehouses, and between warehouses and retail outlets.

IV. Future System Attributes

The *Share the Sky* vision assumes each flight is unique in its purpose and independent in control of its mission and associated flight trajectory. Real time safety functions are the responsibility of each operator to ensure the protection of persons and property in the air and on the ground, based upon common operating rules and procedures. Traditional airports continue to operate, but serve a minority of all flights while a multitude of vehicles, piloted and autonomous, operate from unprepared areas in close proximity to buildings and the people they serve. New sensors and wireless access to dynamic data bases permit avoidance of all operational hazards during flight, using common rule sets to coordinate actions among aircraft.

What type of airspace system is required to achieve such a vision? The system may need to have the following attributes. While many of these attributes cannot be achieved at all times and under all circumstances, they may serve as long-term goals that will maximize societal utilization of the National Airspace.




Fig. 1. Aerial windmill concept from Ref. 2



Fig. 3. Future forest fire fighting.

- To the maximum extent possible, air vehicles may operate without need for pre-authorization. The only limits to the use of the airspace are the physics of the vehicle itself, environmental impacts from the vehicle's use, and security concerns about where vehicles may be operated and by whom.
- Almost all traffic control is decentralized and vehicle-centric, thereby enabling the extremely high control bandwidths needed for small unmanned vehicles to fly in close proximity to each other or in close proximity to people and property.
- Vehicles communicate with each other and with ground-based systems as necessary to perform their missions with safety. Because extremely high-bandwidth control functions are local to each vehicle, the total communications bandwidth required is easily achievable and scalable to extremely high numbers of air vehicles.
- Fixed ground-based navigation and surveillance systems are minimal, thereby enabling a cost-effective scalable and demand-adaptive infrastructure.
- A centralized air navigation service provider exists to manage limited system resources equitably. These limited resources, such as landing slots at a major airport, are rare. The service provider manages by defining constraints rather than defining trajectories for aircraft to fly. All airspace users are free to do what they wish as long as they do not violate these constraints.
- Taking advantage of increased safety, environmental neutrality, and performance of future air vehicles, points of take-off and landing are numerous, ubiquitous, and in close proximity to people. Today's system bottlenecks, such as major hub-and-spoke airports that are capacity-limited during arrival and departure rushes, no longer exist. Because of their high number and close proximity to where people originate and terminate their trips, air travel is much more convenient than today, with the ratio of travel time to total trip time approaching unity.

- Increased machine intelligence enables increased access of people to personal air transportation. Obtaining a pilot’s license requires about the same level of training as is needed to obtain a driver’s license.
- System standards, defined to ensure interoperability and safety, are designed and established in a way that does not impede modernization. The standards favor accommodation of new technology over backward compatibility to the extent possible.
- The system is at least as safe as today’s system, even though it accommodates orders of magnitude more vehicles. Much of this safety is achieved through vastly increased machine intelligence, providing situation assessments, vehicle health monitoring, and strategies for mitigation and control.
- The system is secure, protecting privacy of system users to the extent possible.
- The system protects national security by identifying threats in time for countermeasures to be applied. While security functions are centralized, the system achieves high response bandwidth through the use of decentralized functional elements such as distributed sensor networks.

V. Enabling Operations Technologies

Share the Sky will rely on emerging technologies for information exchange, information processing, and real-time decision making and control. These “operations technologies” build on a foundation of enabling infrastructure technologies that include advanced communications, advanced miniature sensors, high computing speed and capacity, avionics miniaturization, flight path control, lightweight and morphing materials, high energy storage density, and high specific power. The vision also relies on cost reduction time profiles that are typical of advanced infrastructure technology. Some critical operations technologies are network-centric operations, decentralized control, multi-agent cooperation, and machine-based decision-making.

A. Network-Centric Operations

The concept of network-centric operations (NCO) relies on an undisputed prediction that tremendous capability increases are occurring for communications bandwidth (roughly, the amount of information that can be communicated between a sender and a receiver), data exchange reliability, and data currency, security, and customizability. Network-centric concepts make use of these capabilities to supply a very large number of distributed system elements (be they human or machine) with a very rich level of information. Benefits to be gained include the ability of each element to pull information from the network and to make its information available to the others. In theory, this leads to a common situational awareness that is also very deep. For many years, NCO has been popular among management theorists, especially within the business community, and it can be argued that the internet is a successful example of NCO, or “netcentricity.”

NCO may enable airspace operations to make use of *cloud*, *cluster*, and *network computing* concepts. These concepts are applied in large-scale computing networks to increase the efficiency of applications development, their upgrades, and to reduce duplication of development and data storage. They enable the location of use to be separated from the location of processing or storage, thereby allowing for effective and consistent access of the same capability to multiple locations or parties. Access to the information can be either synchronous or asynchronous. For widely-distributed applications or data, modifications and upgrades of the capabilities (such as a trajectory model of an aircraft) need to be done only once to a central application or data repository. This may reduce down time of aircraft, the air navigation service provider, and airline operations centers for upgrades of software, models, databases, and other functions. The separation of functions from location of use also allows minimalistic processing, storage, and software to be resident at the location of use. Because the capabilities must be accessible on-demand as needed, a reliable and secure networking capability is critical to the successful use of these network computing concepts for safety-critical applications.

The JPDO and the International Civil Aviation Organization (ICAO) recommend an initial NCO-like capability for multi-national airspace, termed System Wide Information Management (SWIM),¹ although available information indicates that SWIM’s purpose is different from that of netcentricity. SWIM advocates see it as a broker of information between producers and consumers of data. SWIM would also seek to reduce the number of interfaces between peer-to-peer systems. Each system will send its information to SWIM’s central repository, which then translates the data and makes it available to others. Many envision SWIM to enable “virtual facilities” – facilities

that do not change in function, but no longer need brick-and-mortar infrastructure, such as an airport tower that can be located anywhere because the information it needs can be available anywhere.

SWIM represents a limited instantiation of the full potential of NCO. This is understandable given that the full NCO vision raises substantial technical challenges in system safety, security, and reliability. In 2003, Rheingold introduced the concept of *smart mobs*: uncontrolled self-organizing intelligent groups, made possible by the ever-increasing network links and ever-lowering costs of communication.³ The air traffic management culture may also be concerned with NCO's potential for eroding centralized authority, possibly threatening safety, security, system stability, or simply making an air traffic controller's difficult job even more so. Controllers certainly do not appreciate having their decisions challenged by pilots today, and increased traffic awareness by flight crews in the future may be of concern to them. System gaming is also an issue long discussed in the air traffic management R&D community. Despite these issues, some of which are associated with merging transformational ideas with a legacy system, network centric operations will probably only enable *Share the Sky* if it is embraced in full, with only minor tailoring allowed to address unique issues posed by airspace operations.

B. Decentralized Operations and Distributed Agent Systems

In addition to a shared awareness, network centric communications may in fact enable a completely different approach to operations: decentralized decision making. Brynjolfsson and Hitt⁴ and many social network theorists have researched the potential for operations transformation enabled by the communications revolution. Table 1 lists the ways in which an operational organization can change if netcentricity is considered to be much more than an increase in situational awareness. Reference 5 gave these ideas the name *Power to the Edge* and revolutionized the way many think about netcentricity when it was published in 2003. Although the book's subject is military operations, the fundamental principles expressed apply to many types of operations, including air traffic management. The book's authors argue that true transformation is caused by distributing decision-making to those within an organization who have the timeliest and highest quality information, which are typically those who interact directly with the operating environment. Central managers control only that which should be kept centralized. For all other decisions, the managers provide general mission goals and allow the units on the edge of the organization to decide autonomously how to accomplish them. This is termed "self-synchronization" in the reference. As shown by the table, self-synchronization is achieved at a cost of having highly-skilled system participants. The use of machine-based decision-making may offset some of the costs of having such widespread skills, but the overall cost-benefit tradeoff should be a subject of future air traffic management research.

Table 1. Network centric organization as prescribed by Brynjolfsson and Hitt, using a military example (Ref. 4).

Mass Military	Netcentric Military
Centralized expertise and coordination	Dispersed expertise and self coordination
Vertical integration and channeled communication	Horizontal integration and extensive communication
Large formations	Small formations
Many layers	Few layers
Specialized functional units	General purpose units
Extreme division of labor	Extensive cross-training
Narrow skill requirements	Broad skill requirements
Low training requirements	Very high training requirements
Mass and firepower	Speed and precision
Ponderous	Flexible/Agile
Sequential action and maneuver	Continuous action and maneuver
Heavy reliance on resources held in reserve to deal with the unexpected	Capacity to redeploy quickly
Limited situation awareness	High degree of shared situation awareness
Formal relations with subordinates, supporting units, and suppliers	Long-term, trust-based relationships
Low emphasis on social learning and information sharing	High emphasis on social learning and information sharing
Slow to adapt	Quick to adapt

Ironically, self-synchronization is actually antithetical to a fundamental goal of netcentricity, which is to provide all with a shared level of deep awareness. If everyone has the same awareness, then decision making can remain centralized. In fact, military organizations began employing distributed authority concepts because they are more robust when communications break down.⁵ Shared situational awareness is still a critical element, however, because units on the edge do interact with each other and because shared awareness is critical for central functions to be effective. The book's authors understood this apparent contradiction in network centric operations theory and the danger that many will only want to implement the information exchange component. They address this as follows:

The path to NCO is forked. One road, often called "modernization," is the straightest and most clearly signed. Traveling this road is clearly within the comfort zone of the institution and most of its members. Unfortunately, this road will lead us only to incremental improvements and, ultimately, to a dead end. The improvements attained, however impressive, will fall short, not only of the potential of network-centricity, but more importantly, they will not enable us to meet the mission challenges of the 21st century. This is the road that many seem to have embarked upon, despite a high-level commitment to transformation. The other, less traveled road (actually it may appear more as a path) leads to a disruptive transformation of command and control (C2) that is central to all military organizations and processes... This transformation must focus on C2, where information is translated into actionable knowledge.

The potential benefits of the *Power to the Edge* paradigm are many, but its greatest benefit may prove to be a substantial increase in *system agility*. Agility is the capability to be flexible and adapt quickly to changing or unanticipated situations, and has obvious benefits for a military system. While airspace management does not involve defeating a deliberately unpredictable opponent, it does involve many forms of uncertainty management. *Power to the Edge* enables the system to deal with these uncertainties as they arise. A pilot or on-board machine intelligence can make decisions much faster and with much better local information than a centralized decision system can. While future high-speed computing and high-bandwidth communications may alleviate some system latencies, it is hard to envision a centralized system ever achieving the agility of a decentralized one. The future may see thousands of small aircraft flying in extremely close proximity to each other. Centralized system latencies will probably not be sufficient to enable the high-bandwidth control required. Local control may be especially important as long as human decision makers are in the loop. Experience-based expertise often cannot be translated into a communication message. Therefore, information processing based on decentralized human expertise cannot be performed by a central system, no matter how fast communications become.

For nearer-term applications in air traffic operations, centralization of control for the organization of flows to a constrained resource such as an airport runway makes sense for the group of aircraft that are directly affected. Centralized control makes less sense where the degrees of freedom are not so limited and the operations are neither as concentrated nor as intent-based, as in non-transportation uses of airspace.

A decentralized infrastructure based on the *Power to the Edge* paradigm may also facilitate *demand-adaptive* and *self-modernizing* system characteristics. Centralized control often requires a fixed infrastructure, such as ground-based radar surveillance. The network of ground-based radar installations is expensive to maintain and difficult to replace with more modern equipment due to issues associated with centralized funding in a democracy. It is also difficult to retire obsolete infrastructure if a politically significant segment of the user community relies on it. Today's ground-based radar provides full coverage of controlled airspace even when there are no aircraft within an installation's surveillance range. A surveillance system that adapts to demand has the potential to be much more affordable. If surveillance is distributed to the edge (in this case, to each aircraft), then surveillance capability becomes great when it is needed – at times when many aircraft are occupying a volume of airspace. It is nonexistent when it is not needed – when there are no aircraft in the airspace. Similarly, a distributed system has the potential to be largely self-modernizing. As aircraft with on-board surveillance are decommissioned, they will be replaced with new aircraft that have more modern surveillance capabilities.

C. Multi-Agent Cooperative System Control

A swarm is a massively parallel distributed system built from large numbers of autonomous elements. From the study of insect behavior, it appears to be one of nature's ways of building high-functioning systems from many simple parts. A swarm is characterized by the absence of a centralized control element; it has neither a single controller nor a hierarchy of managers directing its action. Yet in nature the swarm is capable of performing meaningful work. The key to the decentralized control is the flow of information between neighbors within the swarm. Complex behaviors emerge from the interactions of large numbers of agents applying relatively small sets of rules to local information.

The bee swarm is the classic example of a swarm, but the metaphor can be readily extended to other systems of similar architecture. Swarms also include distributed cooperating agents like computer networks and robots that join

to accomplish a task like climbing over an obstacle. In the most general sense, a swarm is any loosely structured collection of interacting units. In the future, multiple aircraft may be able to join formations to improve fuel efficiency for all. Fleets of small uninhabited aerial vehicles will operate as swarms for applications ranging from inspecting a bridge for structural integrity to exploring remote locations on Earth. A swarm of vehicles can provide enhanced area coverage with multiple look angles and a wide variety of sensors in a single flight. Large areas can be efficiently searched in parallel by swarm agents that can share information between themselves during the search process. A swarm can contain redundant elements to reduce the effects of failure of a single agent. And taking advantage of individual vehicle capabilities, a swarm could perform real-time division of labor and task allocation among its individual members to optimize the swarm application.

Currently, little is known about how to design a swarm to perform a specific task. Simple bird-like flocking has been demonstrated with computer animated agents, termed boids.⁶ Computer models which extend the flocking behavior to simulated aerial vehicles have also been implemented,⁷ but all of these simulations avoid the issues of directing a swarm en masse to accomplish something meaningful. The lack of a central controller is one of the strengths of a swarm, but the lack of a mission controller currently limits the swarm to only a few specific mission types. The behavior of a swarm results from interactions among individuals and between individuals and their environment as much as from the behaviors of the individuals themselves. Developing the understanding to direct a swarm will be a difficult, but necessary, task to accomplish.

D. Machine-Based Decision Making

The first flight called fully autonomous by the Department of Defense (DoD) took place in 1989, when the DARPA Condor took off and landed without human intervention. Despite the use of the word “fully,” this level of autonomy from human control actually requires a very rudimentary level of machine intelligence. The DoD has established autonomous capability levels to distinguish between current capabilities and those desired in the future, as shown by Fig. The Condor had some capability for redundancy management of systems, and it could choose between alternate runways. The Global Hawk has advanced autonomy to just below Autonomous Control Level 3 according to Ref. 8, with “real-time health and diagnostics and substantial improvements in adaptive behavior to flight conditions and in-flight failures.” The Air Force / Navy Joint Unmanned Combat Air Systems (J-UCAS) and the Army’s Unmanned Combat Armed Rotorcraft (UCAR) projects, when active, had goals to increase the autonomous capability levels significantly. J-UCAS planned to have multi-aircraft coordination for tactical replanning. It remains to be seen whether these goals will be reached in the time frame shown by the figure or how extensive and robust they will be. If the goals are fully met, intelligent machine-based decision making may be available sooner than most experts expect. For the nearer-term, the human and the machine automation functions will complement each other. Automation is excellent for well-understood tasks that require precision, high speed computation, high repetition, vigilance, monitoring, and mathematics-based optimization. Humans, on the other hand, excel in dealing with the unexpected.

Three approaches have been used in most machine-based intelligence progress to date: recursion, neural nets, and evolutionary algorithms. Recursion involves the modeling of a problem by calling upon the model itself, thereby creating trees of possible answers, which are then compared to achieve a solution. Chess programs that examine all future possibilities of moves work in this way. Neural nets emulate the brain by “learning” a pattern based on the value of individual neural connections in a very large parallel network of synapses. Evolutionary algorithms are a form of iterative search which emulates evolution by forcing random changes to each member of a “population” in successive “generations,” and allowing only the best members to survive based on a pre-defined metric.

These approaches may alone be sufficient to achieve *Share the Sky* as computing capabilities achieve sufficient size and speed. Neural nets will be much more effective than they are today when computers are constructed with the computer equivalent of neural synapses approaching the numbers of a human. Currently the largest computers have roughly a billion times fewer neural connection equivalents than the human brain, which has about 100 trillion connections.⁹ Extrapolating the rate of growth of computers, a personal computer may have capacity that is equivalent to the human brain in a few decades. In theory, computers could then be trained to develop their neural patterns that would give them decision-making capabilities equal to those of a human. However, there is much scientific debate regarding whether the capability will ever be achieved for a computer to make decisions as a human can. Humans learn by sensory interactions with the world around them, and some interactions may be much more difficult for a machine. Also, neuroscience and neurophysiology are still in their infancy; we do not yet know how thought occurs. Our current ability to look at a working brain is limited to a few single neurons or a macroscopic map of generalized activity in the brain. The most advanced techniques (e.g., functional MRI) measure blood flow in the brain. It is assumed that increased blood flow means increased activity – more neurons firing.




Fig. 4: The DoD autonomous control levels (Ref. 8).

What is not known is what these neurons are doing, and which neurons are firing. In summary, in the science of understanding human thought, we are still at the stage of observing macro-level patterns. If these issues are resolved, machines do have a significant advantage over humans in that their learned patterns, data, and programming can be mass-produced, whereas each new human must learn anew. Given the incentives that are already in place for improving machine-based intelligence, however, it is safe to predict that artificial intelligence research will continue to advance. Therefore, new approaches to artificial intelligence can be expected to be invented, augmenting the approaches that currently exist.

Because flight crew error is a major source of aviation accidents, some hope to replace many or all of the in-flight human decision-making tasks with machine intelligence. Increases in automation have reduced errors, but introduced new ones. It is probably not possible to eliminate all human error. As hardware capabilities increase, a critical question is whether or not it is possible for software capabilities to increase sufficiently, assuming the computer equivalent of a human brain does not arrive in time. There are many challenges ahead. Software code can only be proven to have bugs; it can never be proven bug-free. As software complexity increases, it becomes more difficult to verify all possible output states. Formal methods, which provide mathematical proof of correct operation for some software problems, will probably provide a part of the solution. Perhaps an even more difficult problem is that as we approach the capabilities of the human in making routine operational decisions that deal with a multitude of possible situations, the software will become increasingly complex; perhaps so complex that limits in verifiability and reliability will be reached.

VI. The First R&D Steps

The road to the *Share the Sky* vision may be a long one, and several steps should be initiated or accelerated as soon as possible. These steps involve gaining an increased understanding of operational needs, capabilities offered by new technology, and how to implement change. Thus, *the R&D community must take responsibility for these first steps*, leading change while maintaining the vision.

A. Decoupling tactical traffic control from strategic traffic management

Much air traffic management R&D worldwide rests upon an implicit assumption that tactical and strategic management of traffic must be performed simultaneously or iteratively. Two examples of this are 1) strategic conflict resolution, which builds resolution trajectories that also comply with flow management constraints such as

scheduled time of arrival, and 2) simultaneous conflict resolution for all trajectories – the human controller’s situational “picture” – to prevent cascading conflicts. This coupling assumption is not proven, and there is some evidence that indicates it is not universally true.¹⁰ If it is possible to decouple tactical control such as separation assurance from strategic issues such as throughput management, both problems may become simple enough to resolve. Research should be focused on determining if the coupling assumption is valid, and if it is shown to be invalid, R&D should provide clear definitions of a minimum set of constraints for each form of control that enables decoupled solutions. Reference 11 provides a concept of operations that addresses this issue, although the R&D is yet to be performed that will provide specific values for the constraints.

B. Understanding the relationships between system determinism, safety, and performance

Current NAS modernization approaches do not address the complex relationship between how deterministic the future NAS must be in order to make it safe and achieve desired performance levels. There appears to be an implicit assumption that a more deterministic system is safer and higher-performing than one with high levels of uncertainty. In this paradigm, order and safety are perceived to be tightly coupled, so order is more important than many other issues. The need for order drives a need for extensive pre-flight planning and approval, which results in rigidity of flight trajectories. This may not necessarily be the only solution for safety, and it is most likely not the best solution for system performance. A robust system with the flexibility and control power to overcome disturbances may be preferable to a deterministic system that reduces uncertainty using sophisticated data gathering, prediction, and planning. A lack of system-level understanding remains regarding when to reduce uncertainty and when to compensate for unplanned events. At a minimum, the relationship should be understood so that the safety, cost, and performance trade-offs in reducing uncertainty can be assessed.

C. Understanding cost/benefit trade-offs between distributed and centralized control

There has long been recognition that aircraft and flight crews have great potential for assisting in air traffic management, and that they may be neglected assets within our current system. Proponents of centralized control argue that its costs are much lower than decentralized concepts because 1) decentralized concepts incur additional costs to install sophisticated airborne equipment, maintain it, and train pilots to operate it; 2) decentralized concepts incur the costs of airworthiness certification of airborne equipment; and 3) airborne equipment is subject to a much harsher environment than ground-based equipment. Proponents of decentralized control argue that its costs may be lower because 1) the system is inherently more linearly scalable – by minimizing fixed infrastructure (equipment and controllers), the system infrastructure available is proportional to demand for airspace use; 2) the system may be in part self-modernizing because obsolete equipment is retired as aircraft are retired, thus reducing the required support of archaic infrastructure; 3) the automation required for each aircraft is much simpler than the automation required for a centralized system, thereby lowering software development and maintenance costs; and 4) economies of scale may enable a reduction in the costs of airborne equipment relative to ground-based equipment. These cost/benefit issues rarely go beyond the discussion phase because there are long-standing unresolved issues regarding feasibility and safety of decentralized control. Long-term cost/benefit issues associated with decentralized control should be addressed with focused R&D.

D. Establishing new methods for safety prediction and assessment for new types of operations

Today, the NAS is so complex that it is hard to prove that any changes will maintain safety. Flight hardware and software must be certified before it is legal to operate in the NAS. Any new operational procedure must obtain operational approval, whether or not it involves new flight hardware or software. As a result, the system has changed historically due either to major accidents, political events, or to lengthy, cumbersome, and sometimes overly restrictive authorization processes. The MIT International Center for Air Transportation has developed a model of NAS modernization¹² as shown in Fig. 5. The model is based on the analysis of thirteen cases of technology and policy changes, driven by needs for safety and capacity. The figure uses the introduction of Automatic Dependent Surveillance - Broadcast (ADS-B) as an example of a proposed change. As illustrated, the transition process involves an awareness-building process and a change process, in which the various stakeholders develop a negotiated common solution based on the needs and values of each. There is then a need for one or more commitment decisions by the FAA and the stakeholders. Some transitions have been thwarted in the past because one party did not follow through with commitment decisions, leading to significant mistrust and reluctance to commit among other stakeholders. The implementation process then involves operational approval and flight certification based on safety and environmental considerations. The implementation process often causes a revisiting of the change process, leading to significant delays.

Because of the high complexity of operations and technology, both certification and operational approval currently rely heavily on relating a proposed change to an existing operation or operational system. Incremental changes to the existing system are much simpler to characterize and analyze than those for an entire new system. Therefore, it is difficult to implement system improvements that are not evolutionary in nature. As system capabilities increase, science-based and model-based approaches to predicting safety of operations will become a critical function. Radical new types of operational concepts that have little or no historical precedent will mature, and demand for their operational approval will rise. The lack of new methods and processes for addressing operational safety and certification may be the single greatest hindrance to the implementation of new uses of the airspace.

JPDO is sponsoring an activity through its Safety Working Group to assess the safety of trajectory based operations. Concepts under consideration would place responsibility for separation in ground-based automation for conflict detection and resolution. This assessment activity will provide a test of our current ability to analyze the complex air traffic management system by comparing safety of a completely new automated approach with that of existing operations. Methods of analysis developed for this activity may prove useful in other similar assessments of proposed major changes to the National Airspace System.

E. Understanding limits of current and near-term automation technology, and establishing clear roles for human operators and automation

Moore's and Kurzweil's Laws suggest that computer processing capability will continue to increase exponentially. Some predict that within a generation, artificial intelligence that rivals or even surpasses the human brain is highly probable. If true, utilization of the NAS will certainly change dramatically. If these visions prove to be premature, it is critical to understand the appropriate roles for automation in the future NAS. Near-term designs have been proposed that remove human decision making and intervention from critical functions, even though no safety basis has been established in the designs for rare-normal events and failures. When given a choice, system operators often reject decision support automation because they do not trust its correctness, its reliability, or whether they will be able to recover from an unsafe system state resulting from an automation failure. Even without failures, there may be negative system impacts resulting from the inappropriate use of automation. Human experience is difficult to automate, and its value is not fully understood in air traffic management.

F. Determining how to implement change

Most air traffic management R&D relies heavily on evolutionary development and test approaches to modernization. It is possible that there is no way to achieve the *Share the Sky* vision using this approach.

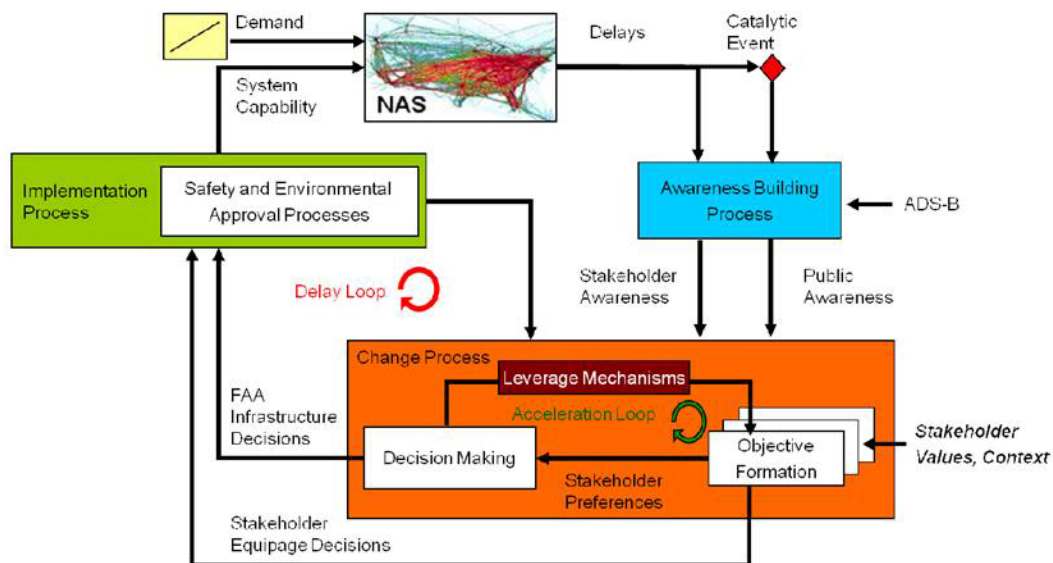


Fig. 5. NAS Transition Process Model, using the example of ADS-B implementation. (Ref. 12)

Approaches currently employed are a natural outcome of attempts to improve a system that is so complex it defies characterization. When faced with modernizing a system without understanding it, evolutionary “build-a-little, test-a-little” approaches appear to be the only option. One criticism of Ref. 5 is that the book is not realistic about changing an entrenched bureaucracy that wants to keep power in current hands. Even if this is true, it is possible that the opposite is in fact the dominant effect: *Power to the Edge* may be the only way to change an entrenched bureaucracy. The decentralization of a large system is actually a simplification of one large and complex problem into many smaller decoupled problems. It may be possible that management of the NAS is such a large and complex problem that the only possibility of solving it is to simplify it through a decoupling of issues, and thereby opening opportunities for early adopters. Traditional top-down engineering and design may not be the best approach to tackling airspace modernization, even though it is attractive to many because it facilitates system-level optimization. System optimization is often thought of as “an integrated solution,” or a “simultaneous solution” of the many issues, and the optimal service of multiple objectives. However, system-level optimization has no utility at all if such a system cannot be realized. As stated by Ref. 5:

There is a price to be paid for simultaneity. In our terms, the prompt suboptimization that created the desired effects was clearly preferable to the slow, ponderous processes that sought to optimize the use of weapons systems and platforms.

Even if top-down engineering and design is found not to be the best approach, compatibility and interoperability of all system elements will be necessary during all phases of modernization. This *synchronicity of change* requirement will certainly be more challenging in the absence of a top-down design. System operating standards will need to be developed that can accommodate change, providing a framework for technology insertion. Introducing an advanced operating procedure while minimally disrupting those who choose to operate under legacy procedures may be a critical to modernization. The advanced separation assurance concept of Ref. 11 relies on this approach. For some types of new operations technologies, it may be possible to build a fully parallel system that operates in a “shadow” mode until enough experience is gained to establish the necessary confidence in the new technology. This mode will utilize actual NAS flight plans, weather, and airspace and airport constraints in real-time, but with advanced technologies executing in a parallel virtual environment. Such a capability may facilitate faster development of functional and technical requirements and lower the risks of implementation.

Principles of complexity science may need to be more fully understood and applied to achieve the needed breakthroughs. Complex systems are characterized by numerous interacting agents, agent behavior that is affected by memory or feedback, adaptive strategies employed by agents, emergent system behavior, and self-evolving behavior.¹³ The NAS exhibits all of these characteristics. It is possible that the best approach to modernization is to seek ways of reducing rather than increasing system complexity. Most airspace system modernization approaches anticipate high-bandwidth rich-message-content communications and their associated standardization requirements, massive centralized data repositories, and sophisticated software to advise decision makers. Concepts such as *Power to the Edge* increase some forms of system complexity and decrease others. Complexity science, resilience engineering, and other system design principles should be explored and if applicable, applied to modernization strategies. A full understanding of the potential power and impacts of policy changes is also needed. Powerful new forms of agent-based simulation that can be applied to the analysis of policy, as described in Ref. 14, may be critical tools in gaining the understanding needed to achieve airspace modernization.

VII. Conclusion

Share the Sky is a long-term paradigm that promotes the future use of airspace for the public good, while accommodating freedom of airspace access by any segment of the population that has a need or desire to use it. It will facilitate a cost-effective, scalable, and equitable use of airspace, and it will ensure that airspace can safely accommodate both existing and all new flight operations as they are introduced. While much current R&D is focused on transformation of the NAS to enable three times today’s system capacity, the *Share the Sky* vision requires the accommodation of many times more vehicles in the long term.

How might we attain this vision? Because absolute system safety is paramount, a full understanding of every proposed change is critical, whether that change is large or small. The high costs involved with modernizing the system make it imperative that we know the new system will perform as expected *before* we start building it. Therefore, changing the way we operate in the National Airspace defines an almost unique class of problem. How do we improve a massive and complex system that has become extremely safe through evolutionary modernization processes, and is strictly regulated in order to maintain that safety? Unlike the way many innovations enter the mainstream, the need for absolute safety results in rules restricting operational use, thereby restricting trials of new

ideas. In other areas of modernization, these rules of use often lag the new technology, whereas in airspace operations, the rules lead the technology.

To resolve this modernization dilemma, we may need to approach the problem in a different way. Conventional evolutionary modernization cannot apply to many transformational concepts. Although there are valid reasons why change tends to occur slowly, perhaps some justifications for evolutionary modernization are based on invalid assumptions. For example, some assume that decentralized approaches are necessarily long term because aircraft of the next 30 years are already being delivered today. There are counter-examples of very fast modernization happening when disruptive new technologies are combined with economic incentives. The introduction of advanced in-flight entertainment systems may be such a counter-example. Perhaps approaching the modernization question in the opposite way is just as valid: What are the *advantages* of building system infrastructure that will have a defined retirement date? Elements of the infrastructure that reside on aircraft have that characteristic, while ground-based elements often do not. To break the modernization dilemma, we may also need to think about system complexity in a different way. Most NAS modernization approaches under consideration today tend to increase dependencies between system elements, require more information exchange, and require longer planning horizons. Many other systems advance by reducing these system characteristics rather than increasing them. An increased understanding of these complex issues will ultimately break the modernization dilemma and perhaps lead to a future where society truly shares the sky.

Acknowledgment

The authors thank Vincent Houston of NASA Langley Research Center for his contributing ideas and insights.

References

- ¹Joint Planning and Development Office, Joint Planning Environment website, URL: <http://jpe.jpdo.gov/ee/request/home>.
- ²Sky WindPower Corporation. URL: <http://skywindpower.com> [cited 9 August 2011]
- ³Rheingold, H., *Smart Mobs: The Next Social Revolution*, Basic Books, 2003.
- ⁴Brynjolfsson, E. and Hitt, L., "Beyond Computation: Information Technology, Organizational Transformation and Business Performance," *Journal of Economic Perspectives*, Vol. 14, No. 4, pp. 23-48, 2000.
- ⁵Alberts, D. S. and Hayes, R. E., *Power to the Edge: Command... Control... in the Information Age*, Department of Defense Command and Control Research Program, CCRP Publications, 2003. (http://www.dodccrp.org/html4/books_downloads.html.)
- ⁶Reynolds, C. W., "Flocks, Herds, and Schools: A Distributed Behavioral Model," *Computer Graphics*, 21(4), pp. 25-34., 1987.
- ⁷Crowther, W.J., "Flocking of Autonomous Unmanned Air Vehicles," *Aeronautical Journal*, Vol. 107, No. 1068, pp. 99-110, 2003.
- ⁸Office of the Secretary of Defense, *Unmanned Aircraft Systems Roadmap, 2005-2030*, August 4, 2005.
- ⁹Kurzweil, R., *The Age of Spiritual Machines*, New York, Viking Press, 1999.
- ¹⁰Hoekstra, J. M., "Designing for Safety: The Free Flight Air Traffic Management Concept," ISBN 90-806343-2-8, Doctoral Thesis, Delft University of Technology, November 2001.
- ¹¹Wing, D.J. and Cotton, W.B., "Autonomous Flight Rules: An ADS-B Application Description for Self Separation in US Domestic Airspace," NASA TP-2011-217174, 2011.
- ¹²Mozdzanowska, A. L.; Weibel, R. E.; Lester, E.; Hansman, R. J.; Weigel, A.; and Marais, K., "Dynamics of Air Transportation System Transition and Implications for ADS-B Equipage," AIAA -2007-7776, *AIAA 7th Aviation Technology, Integration, and Operations (ATIO) Forum*, Belfast, September 18-20, 2007.
- ¹³Johnson, N., *Two's Company, Three is Complexity*, Oneworld Publications. ISBN 13:978-1-85168-488-5, 2007.
- ¹⁴Epstein, J. M., (ed.), *Generative Social Science: Studies in Agent-Based Computational Modeling*, Princeton University Press, 2007.