MAN FACTORS IN TRANSPORTATION

ries of Volumes Edited by
y H. Kantowitz

ngs• *Aviation Automation: The Search for a Human-Centered Approach*

suraman/Mouloua • *Automation and Human Performance:*
neory and Application

eld/Dingus • *Human Factors in Intelligent Transportation Systems*
n press)

• *Ergonomics and Safety of Intelligent Driver Interfaces (In press)*

bridge • *Complex Cognition and the Implications for Design*
n preparation)

and/Wise/Hopkin • *Aviation Human Factors (In preparation)*

AVIATION AUTOMATION:
THE SEARCH FOR A HUMAN-CENTERED APPROACH

by

Charles E. Billings
The Ohio State University

LEA LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS
1997 Mahwah, New Jersey

CHAPTER 9

BENEFITS AND COSTS OF AVIATION AUTOMATION

INTRODUCTION

The NASA Aviation Safety/Automation research initiative (NASA, 1990), the work of Wiener and Curry that preceded it (Wiener & Curry, 1980; Curry, 1985; Wiener, 1985a, 1989, 1993); studies by Rouse and colleagues (1980, 1988; Rouse & Rouse, 1983; Rouse, Geddes, & Curry, 1987), research by Sarter and Woods (1991; 1992a, 1992b; 1994), Sheridan's (1984; 1987; 1988) studies of supervisory control, and contributions by Rasmussen (1988), Reason (1990), and many others form the theoretical and empirical foundations for these comments on humans and automation. There is now a substantial body of data concerning human cognitive function in complex, dynamic environments. I hope that this chapter demonstrates to designers and operators working in the aviation domain that there is considerable knowledge that can help them to do their respective jobs more effectively.

As noted before, it is necessary to examine the unwanted behavior both of automation and people in an automated system, because it is only through such study that we can minimize the costs, while increasing the already considerable benefits, of this technology. It is important that we not lose sight of the benefits (see next section), for aviation cannot advance without automation if we are to meet future challenges, which will tax our ingenuity to the utmost. We must not throw the baby out with the bath water.

But it is equally important that we not ignore the potential costs of yet more sophisticated automation, for if it is not designed and used properly it can make the future aviation system less flexible, less effective, and less able to meet those challenges. In recent years, it has become evident that our operators do not always understand or properly manage the automation they now have at their disposal. It is essential that we make every effort to understand why this is true, if we are to design future automation so that it will be more effective and error tolerant than what we now have.

BENEFITS OF AVIATION AUTOMATION

have referred in several places to the benefits derived from aviation automation o date. Let me summarize explicitly what these benefits are, to keep this discussion f problems in context. In a landmark paper, Wiener and Curry (1980) discussed ystem goals. Paraphrased, they are:

- Safety
- Reliability
- Economy
- Comfort

I briefly cite demonstrated benefits with respect to each of these system goals. his list is not inclusive, but it will provide some insights into the extent to which e rely upon automation to accomplish our objectives.

Safety has always been proclaimed by the aviation industry as its primary jective. An examination of air carrier accidents by Lautmann and colleagues autmann & Gallimore, 1987) suggests that newer, more highly automated aircraft ve had substantially fewer accidents than earlier aircraft (Fig. 9.1). In their first cade of operation, the widely used Boeing 757/767 models were involved in only e fatal accident (Thailand, 1991). (Two recent accidents have marred this record, wever: A B757 suffered a controlled flight into terrain accident near Cali, lombia on December 20, 1995; this accident involved both human error and human–machine interaction problems. Another 757 crashed into the sea after takeoff from Puerto Plata, Dominican Republic, on February 6, 1996. The mishap involved faulty airspeed indications presented to the pilot flying; these incorrect indications may also have affected autopilot operation and may have been due to a plugged pitot tube. The first officer's airspeed system was probably reading correctly. Both accidents are under investigation.) Other new types have been involved in more accidents, but the record is still generally good. (For a balanced discussion of this question, see Automated cockpits, 1995a, 1995b)

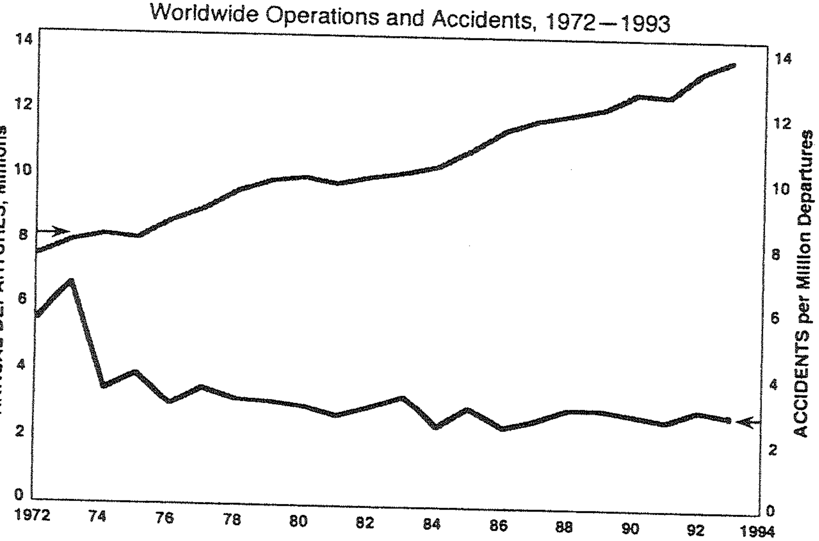

FIG. 9.1. Scheduled air transport operations and accident rates (Boeing, 1994).

Reliability has been improved; autoland-capable automation, head-up displays, and other innovations have increased the number of flights able to operate at destinations obscured by very low visibility. Newer systems (GNSS, enhanced vision) have the potential to improve approach and landing safety worldwide. Improvements in ATC also have the potential to increase reliability, as well as efficiency, in the future system.

Economy has been improved by flight management systems that can take costs into account in constructing flight plans (although the benefits possible from such computations have been diluted by the inability of the present ATC system to permit aircraft to operate routinely on most cost-efficient profiles). Despite this limitation, significant economies are being achieved in the United States by coordination of nonpreferred and direct routes between air carrier systems operations centers and the FAA's System Command Center.

Comfort has been improved by gust alleviation algorithms in the newest aircraft, as well as by the ability of newer aircraft to fly at higher altitudes, above most weather. Greater flexibility enabled by ATC automation will permit pilots to utilize a wider range of options to achieve more comfortable flight paths.

In what respects are we still deficient with respect to these system goals? Most of our accidents can be traced to the human operators of the system, and increasing numbers can be traced to the interactions of humans with automated systems. More can be done to make aircraft automation more human-centered, but perhaps even more important, advanced automation can be used to make the system as a whole more resistant to and tolerant of human errors, be they in the implementation or the operation of these systems.

COSTS OF AVIATION AUTOMATION

The 1989 ATA Human Factors Task Force report stated,

> During the 1970s and early 1980s, the concept of automating as much as possible was considered appropriate. The expected benefits were a reduction in pilot workload and increased safety.... Although many of these benefits have been

> realized, serious questions have arisen and incidents/accidents have occurred which question the underlying assumption that the maximum available automation is always appropriate, or that we understand how to design automated systems so that they are fully compatible with the capabilities and limitations of the humans in the system. (pp. 4–5)

Let us examine this statement, which was largely responsible for the inquiry described in this book and its predecessors (Billings, 1991, 1996).

At the time the ATA report was prepared, the outlines of the A320 and B-747-400 automation suites were just becoming visible to the knowledgeable observers on the Task Force. The MD-11 was at an early stage of development and its cockpit design was not yet firm. It is clear that in the A320 and MD-11, the concept of automating as much as possible, with the intent of reducing flight crew workload and minimizing human errors, was in fact considered appropriate, though the two design teams took different approaches. The 747-400 was more conservative in its automation philosophy and more evolutionary than revolutionary in its application.

It is clear, with the hindsight afforded by 5 years of operational experience, that at least some pilots have found certain of the automation features in this new generation of aircraft difficult to understand and to manage. The difficulties that have been experienced appear to me to have been due in large part to five factors. Four are design factors: complexity, brittleness, opacity, and literalism. A fifth related factor is training, which in turn is related to understanding. Each is considered in more detail here. A discussion of other relevant factors follows.

Complexity

As indicated in chapters 3 and 4, today's aircraft automation suites are very capable, increasingly flexible, and very complex. Tactical control automation (enabled through a mode control panel, as in Fig. 5.11) is tightly coupled to strategic flight management automation (the FMS, with its CDU interface) in ways that are not always obvious. The FMS itself is capable of autonomous operation through several phases of flight. Both parts of the system are "mode-rich" (Sarter & Woods, 1994); default and reversion options vary among modes.

When these interactions cause unwanted behavior (from the pilot's viewpoint), the pilot may not have a mental model that allows him or her to correct the situation short of reverting to a lower level of management (see chapter 10) or turning the automation off, which is not always desirable and may not be possible in some circumstances. "Turning it off" (Curry, 1985), for instance, may disable certain protective features such as FMS knowledge of altitude restrictions during a descent into a terminal area, or the automation's intent to level the aircraft at a given altitude during a climb. Pilots of recent, very powerful aircraft have become concerned about the rate at which the airplane was approaching a level-off altitude and have reverted to autopilot vertical speed mode to slow the climb as they approached the new altitude, unaware that this reversion also canceled the altitude capture mode. The result has often been a deviation above assigned altitude.

Another aspect of automation complexity is the great flexibility found in the modern flight management and autoflight system. Modern systems may have several modes for each of several control elements (Fig. 9.2). These modes interact in ways not always obvious to pilots. Operators must learn about, remember, and be able to access information concerning each mode in order to use it effectively; this imposes a considerable cognitive burden, makes it less likely that the operator will have an appropriate mental model of the automation, and increases the likelihood that modes may be used improperly. In addition, the capability of the modern FMS means that the system may direct the airplane through successive modes of operation autonomously, in ways that may leave the pilots uncertain of exactly why the automation is behaving in a certain manner at a particular point in time.

Sarter and Woods (1995) and Sarter (1994) discussed mode errors and mode awareness. Figure 9.2 is adapted from their paper. It illustrates the mode flexibility (and complexity) in a modern transport aircraft. Compare this with the relatively small number of flight modes in the Lockheed L-1011 automation shown in Fig. 5.12.

AIRCRAFT FLIGHT MODES: A320

Autothrust Modes	Vertical Modes	Lateral Modes
TOGA	SRS	RWY
FLX 42	CLB	NAV
MCT	DES	HDG/TRK
CLB	OPEN CLB	LOC*
IDLE	OPEN DES	LOC/APP NAV
THR	EXPEDITE	LAND
SPD/MACH	ALT	ROLLOUT
ALPHA FLOOR	V/S-FPA	
TOGA LK	G/S-FINAL	
	FLARE	

FIG. 9.2. FMS and autoflight modes in the Airbus A320. From "FMS and Autoflight Modes in the Airbus A320," by N. B. Sarter and D. D. Woods, 1995, *Human Factors* 37(1), p. 13. Copyright 1995 Human Factors and Ergonomics Society. All rights reserved.

Each of the modes listed represents a different set of operating instructions for he automation. The mode in use (or armed, ready for use) is displayed in an lphanumeric legend on a flight-mode annunciator panel, usually located at the top f the primary flight display. In their conclusions, the authors of this very useful aper stated,

> As technology allows for the proliferation of more automated modes of operation...human supervisory control faces new challenges. The flexibility achieved through these mode-rich systems has a price: it increases the need for mode awareness—human supervisory controllers tracking what their machine counterparts are doing, what they will do next, and why they are doing it.... While we understand a great deal about mode problems, the research to examine specific classes of countermeasures in more detail and to determine what is required to use them effectively, singly or in combination, is just beginning. (Sarter & Woods, 1994)

Hollnagel (1993) suggested that increasing system complexity leads to increasing ask complexity. This leads to an increasing opportunity for malfunctions and errors, hich leads to an increasing number of unwanted consequences, which in turn eads to solutions that ultimately increase system complexity still further. He noteed hat this is sometimes humorously referred to as the "law of unintended conse-uences." The "law" states that the effort to fix things sometimes worsens the amage. Although we are perhaps not there yet in this domain, the quantum ncrease in complexity of aircraft automation has unquestionably created new pportunities for human errors, both slips or mistakes by the operator and those hat result from deficient or "buggy" knowledge of the system being utilized.

I believe that automation complexity has been at least part of the problem in everal incidents and accidents involving this new generation of aircraft (in Appen-ix 1, see Mulhouse-Habsheim, 1988; Bangalore, 1990; Strasbourg, 1992; Man-hester, 1994; Paris, 1994; Toulouse, 1994). This is not to say that the automation as not functioned as it was intended to function; it has usually done exactly what s designers and programmers told it to do. The problem has been rather that the uman operators have not understood its intended functioning and consequently ave used it either beyond its capabilities or without regard to its constraints or ules. In another recent example of this problem, an A300-600 crashed at Nagoya, apan (1994), after the pilot flying inadvertently engaged an autopilot mode TOGA), then provided opposing inputs to the airplane's autoflight systems, which ere counteracted by the autopilot when it was engaged to stabilize the flight path Mecham, 1994).

The likelihood that all of the subtleties of such complex systems will be fully omprehended by pilots, even after considerable line experience with the systems, is not high (Sarter, 1994; Sarter & Woods, 1992b; Wiener, 1989); the likelihood that they will be understood after a few weeks of training is very small indeed. Uchtdorf and Heldt (1989), studying pilot understanding of the A310, indicated that a year or so of line experience may be required before pilots feel fully comfortable with the automation features—and this does not guarantee that they understand the entire system, only that they feel comfortable enough with its modes to operate it effectively.

Brittleness

As software becomes more and more complex, it becomes more and more difficult to verify that it will always function as desired throughout the full operating range of the aircraft in which it will be placed. The reason for this is that there is an almost infinite variety of circumstances that can affect its operation, only a subset of which can be evaluated prior to certification even if they are known to the evaluators. Even then, there will be conditions not thought of by the designers, that will inevitably arise at some point in the course of the airplane's operation. Brittleness is an attribute of a system that works well under normal or usual conditions but that does not have desired behavior at or close to some margin of its operating envelope.

An example might be a pitch control system that was selected, then reverted or defaulted to "vertical speed" mode while an airplane was climbing. The autoflight system would attempt to maintain constant vertical speed by increasing pitch angle at the expense of airspeed, which would gradually decay to unsafe levels. One of several examples was an Aeromexico DC-10 whose autoflight system maintained a climb at constant vertical speed until the airplane stalled; the pilots were thought to have improperly programmed the autopilot for constant vertical speed instead of constant airspeed and subsequently failed to notice the decaying airspeed until too late to maintain control (Luxembourg, 1979). Another example would be a descent mode that involved idle power without safeguards to ensure that such a descent could not continue all the way to the ground (see Bangalore, 1990), or an autothrust system that permitted power to remain at idle after descending onto and capturing the glide slope followed by a decrease in descent rate and a consequent decrease in airspeed to unsafe levels.

An example of brittle automation was present in the TCAS software when it was first implemented in civil transports. Under certain circumstances, the TCAS logic was able to recognize a hazard but was unable to advise a safe maneuver to resolve the conflict. When this occurred, the system simply gave up and indicated to the pilot that there was a conflict but the system could not resolve it. FAA certification pilots raised serious objections to such a mode and the software was modified to exclude this problem, although at the expense of commanding much more drastic avoidance actions under such circumstances, which has caused greater altitude excursions. This problem has still not been fully resolved, although the TCAS automation is no longer able to "walk away" from a conflict that requires a resolution advisory.

Yet another example of brittleness at the margins of the operating envelope was seen, I believe, in the crash of an A320 at Mulhouse-Habsheim after an experienced pilot made a low pass over the airfield at minimum airspeed during an air show (1988). During this maneuver, he descended below 100 ft above ground level and was unable to obtain full engine power in time to avoid trees at the far end of the runway. The automation prevented the airplane from stalling, but when the pilot descended below 100 ft the automation disabled the angle-of-attack protection also built into the airplane's flight control system. This feature, which under any other circumstances would have applied full power and rotated the airplane into a climb, must be disabled to permit the machine to land.

Opacity

Three questions with which Wiener (1989) paraphrased the frequent responses of pilots to automation surprises—"What is it doing?", "Why is it doing that?", and "What's it going to do next?"—may be indicative of either or both of two problems. One is a deficient mental model of the automation—a lack of understanding of how and why it functions as it does. This can be due to automation complexity, or to inadequate training, or both.

Another problem, however, is not that the operators do not understand the behavior being observed, but rather that the automation does not help them by telling them what it's doing (and if necessary, why). Sarter and Woods (1994, p. 24) observed that "The interpretation of data on the automation as process is apparently a cognitively demanding one rather than a mentally economical one given the 'strong and silent' character of the machine agent."

This problem represents a failure in communication or coordination between the machine and human elements of the system. Regardless of the cause, the net effect is diminished awareness of the situation, a serious problem in a dynamic environment.

In earlier times, less capable automation simply controlled the airplane's attitude and path; pilots could usually understand exactly what it was doing by observing the same instruments they used when they were controlling the airplane manually. Today's automation may use a combination of several modes to accomplish the objectives it has been ordered to reach. The information about what it is doing is almost always available somewhere in some form, although not necessarily in terms that the pilot can easily decipher. Why it is behaving in that manner is often not available except in the requirements document that motivated it. What it is going to do next is often, although not always, unavailable on the instrument panel.

In short, as automation complexity increases, it becomes more difficult for the designer to provide obvious, unambiguous information about its processes to the monitoring pilot (even if the designer believes that the pilot needs this information and therefore tries to provide it). I call this *opacity*. Others have referred to it as a *lack of transparency*; the two terms are synonymous in this context. Norman (1989) argued that the problem is not automation complexity, but lack of feedback to its operators.

As noted in chapter 4, automation opacity may be deliberate: One sure way to keep the operator from intervening in a process is to deny him or her the information necessary to permit intervention in that process. Much more commonly, I think, it is the desire, and need, to avoid overburdening the operator with information that is not essential to the performance of his or her necessary functions (as those functions are understood by the designer). The capabilities of the computer and its screens have made it possible for designers to overwhelm pilots with information and data. Opacity at some level is required to avoid overwhelming the pilot with data. We know that the ability of pilots to assimilate information is context dependent, and that when we provide more data without adequate consideration of context we simply make it less certain that they will attend to that which they really need to know (Woods, 1993c).

The mode awareness problems cited by Sarter and Woods (1992a) are in part due to opacity, although most modes are announced on mode annunciator panels. In part, the problem is one of salience: Alphanumeric symbols must not only be attended to, but must be read, to convey information. Hutchins (1993) has attempted to ease this problem by using iconic representations, with some experimental success (see Automated cockpits, 1995b, for an illustration of this approach). But Woods (1996) wrote of "*apparent simplicity* (of the system as represented), *real complexity* (of the system's actual behavior)" as one of our more serious problems with advanced automation.

There have been some notable examples of the effects of opacity on advanced flight decks, although it must be noted that in most of the cases, the information could have been found had there been time to look for it. This tends to reinforce the notion that drowning the operator in information isn't a wise way to design a system. Perhaps the most notable recent example is an accident that occurred during an approach to Strasbourg (1992), when the flight crew inadvertently commanded the autopilot to descend at a 3300 ft/min vertical speed rather than at a 3.3° flight path angle.

The FCU display read "–33" instead of "–3.3," although smaller letters on the LCD display also read "HDG/VS" instead of "T/FPA" and the symbology on the primary flight display was different in the two modes. The fact remains that the pilots, already heavily loaded because of late ATC instructions and inexperience in the airplane, missed these discrepancies and descended into the ground several miles from their destination. Changes have been made in later cockpits of this type to display "–3300" versus "–3.3" in the hope of eliminating this possible source of confusion. Another example is the TOGA (takeoff/go-around) indication in the A300 at Nagoya (1994), which was initially missed by the pilot flying. (It is worth noting that in both these cases, the flight crew provided the autoflight system with an incorrect indication of intent; see chapter 3.)

Literalism

A fourth attribute of automation (and of computers in general) could be described as its literalism or "narrow-mindedness" (S. W. A. Dekker, personal communication, January 1994). Automation is able only to do exactly what it is programmed to do, as it did in the two cases just cited. Human problem solvers are creative in their reasoning and their search for solutions to a problem. They can and will draw knowledge or evidence from any available source (either in memory or external to themselves: reference books, manuals, contact with others by radio, etc.), as long as that knowledge is relevant to the problem to be solved. Automation, on the other hand, is constrained by its instructions and is insensitive to unanticipated changes in goals and world states that may fall well within its usual operating range but were unanticipated by the designers of its software. It is in this sense that computer literality contrasts with brittleness; the latter term refers to undesired automation behavior at the margins of the operating envelope.

As an example of this, some flight management systems with vertical navigation capability will calculate an optimal descent point, based on cost factors, that is closer to a destination airport than pilots may wish for a smooth, gradual descent. The pilots may be unaware of the logic that drives this decision and action, but they learn through experience that they can "trick" the automation by programming a higher tailwind than is actually present. This false information causes the automation to begin the airplane's descent at an earlier point in time, thus achieving the pilots' desired ends. Human operators have always shaped the tools at hand to assist in accomplishing their objectives, but this shaping also increases ask demand and cognitive workload, and increases the opportunity for errors.

Training

indicated earlier that a fifth relevant factor is training. Let me preface this iscussion by saying that if we cannot *show* the pilot what he or she needs to know in a given situation, then the pilot needs to *know* what he or she needs to know. The only way this knowledge can be acquired is through education and training.

In the early 1960s, Trans World Airlines ordered its first DC-9 aircraft, also its first jets with a two-person crew complement. The airline decided to undertake a major revision of its training philosophy for the new airplane; its new, and highly successful, training program emphasized the specific behavioral objectives (SBOs) required of pilots, rather than the older (and until then universal) approach of "teaching the pilot how to build the airplane." (Previous training programs emphasized detailed knowledge of how airplane systems were constructed, how the various parts contributed to the whole, and based, on this knowledge, how to operate them.) The new approach provided significant economies in training time, which is expensive, and appeared to be fully as successful in teaching pilots how to operate the new airplanes without burdening them with more systems knowledge than they "needed to know." United Airlines later adopted a similar philosophy, with similar success, and a training revolution was underway.

There has been continual pressure to minimize training time for the last 30 years. Pilots are paid virtually the same amount for training as for line flying, and when they are in training they are not flying trips that produce revenue for their company. There is no question that the SBO concept has been effective and efficient. Until recently, there has been no reason to question the concept.

The complexity of advanced automation, however, gives rise to questions about this approach to training. As indicated earlier, pilots must have an adequate mental model of the behavior of the equipment they are flying. I believe that our experience to date with advanced automated aircraft suggests that the training we now provide does not always give them a sufficient basis for forming such models. One example of this, in the MD-11, was that takeoff speeds could be incorrectly calculated by the FMS if engine anti-ice significantly warmed certain sensors. An error message was generated, but this message was inhibited by flap extension. If flaps were lowered at the beginning of taxi, before airflow over the sensors had time to cool them, the erroneous speeds were locked in and takeoff speeds were incorrectly displayed on the speed tape of the PFD.

There is no question about the growing complexity and opacity of automated systems in these aircraft. I believe that questions must be raised about whether present training in *how to operate* these more complex and less transparent systems, as opposed to *how they operate*, is sufficient to provide pilots with the information they need when the systems reach their limits or behave unpredictably. If a pilot does not have an adequate internal model about how the computer works when it is functioning properly, it will be far more difficult for the pilot to detect a subtle failure. We cannot always predict failure modes in these complex digital systems, so we must provide pilots with adequate understanding of how and why aircraft automation functions as it does.

Comments about automation (McClumpha, James, Green, & Belyavin, 1991; udisill, 1994, 1995) make it plain that many pilots do not understand the reasons vhy aircraft and avionics manufacturers have built their automation as they ave—and there are usually very good reasons, although they may not be known o the users of the automation. This, again, represents a failure of training to explain ow the system operates and why, rather than simply how to operate the system.

OTHER OBSERVED PROBLEMS WITH AVIATION AUTOMATION

everal other problems, some associated with or caused by those just enumerated, eserve mention here. Each has been associated with undesired outcomes in line perations; all can be mitigated to some extent by effective human–machine nterface design.

eliance on Automation

everal examples showed that pilots given highly reliable automated devices (and ıost are) will come, over time, to rely on the assistance they provide. They rely on ne correct function of configuration warning systems, altitude alerters, and other nformation automation to which they have become accustomed. When GPWS vas first introduced, the nuisance warnings to which it was prone caused pilots to istrust it; conformance with its warnings had to be mandated by company standard perating procedures. Later models have proved themselves more trustworthy, and ney are relied on. Pilots have long been served reliably by autopilots and are ometimes less alert in monitoring their behavior than they should be, as evidenced y the failure to detect a few uncommanded roll inputs in early 747s (e.g., Nakina, ntario, 1991). In some cases, pilots have continued to use automation even when ney had every reason to mistrust it. This misplaced reliance has led to at least one ccident, a runway overrun on landing at J. F. Kennedy Airport (New York, 1984). he NTSB discussed overreliance on automation at length in its report on this ishap (National Transportation Safety Board, 1984).

Air traffic controllers likewise rely on the data presented to them on their CRTs, ven though much automation is required to present the synthetic images with hich they work. They are surprised by occasional "tag swaps" and other misrepesentations of the data when they occur.

It does little good to remind human operators that automation is not always reliable or ustworthy when their own experience tells them it can be trusted to perform correctly ver long periods of time. Many pilots have never seen these automation elements il, just as many of them have never had to shut down a malfunctioning engine

except in a simulator, and in any case, humans are not good monitors of infrequent events. The solutions to the "human failings" of trust, and of inattentiveness, must be found elsewhere. If we are to continue to provide operators with automation aids, we must make the system in which they are embedded more error tolerant so that such "failings" will not compromise safety of flight. In this area, there is much more we can do, even though much has been accomplished in the past.

Clumsy Automation

Wiener (1989) coined this descriptor to denote automation that lightens crew workload when it is already low, but requires more attention and interaction at times when workload is already high (see also Tenney, Rogers, & Pew, 1995). He and others have cited today's flight management systems as having this characteristic, as I noted in chapter 5. In the aviation context, it is in locations where traffic density is highest that ATC will most often have to change clearances to adjust to unexpected problems. It is also in these areas that aircraft are often climbing or descending and preparing to land.

These are the phases of flight that involve the highest likelihood of conflicts with other aircraft and that therefore demand that as much attention as possible be devoted to scanning for such traffic. Programming a flight management computer requires that the nonflying pilot's attention be inside the cockpit and focused on the CDU for some period of time. This is an attentional requirement that directly competes with outside surveillance and monitoring the activities of the flying pilot. Although efforts have been made in the newest FMSs to lighten this burden, reprogramming, often required to meet ATC requirements during transition to terminal areas, can still be cumbersome. Flights into Los Angeles, which may well be the world's most heavily traveled airspace, are often cited by pilots as perhaps the most taxing example of this problem.

Digital Versus Analog Control

I mentioned earlier the criticism by pilots of automation that makes it necessary for them to enter new navigation radio frequencies through alphanumeric keystrokes on the CDU rather than by turning rotary selectors as they did on older radio control units. Whether digital frequency entry actually takes longer has not been studied, to my knowledge, but I must confess that I share the bias of these pilots. At this time, communications frequencies are still accessed through the older types of control devices, most of which also show and make available both the old and new frequencies. This is a help to pilots if they are unable to establish radio contact on a new channel, but communication frequencies also may be accessed in future through the FMS.

In the autoflight control wheel steering (CWS) mode, pilots manipulate their ontrol columns to instruct the automation what rates of change are desired for a naneuver. Once placed in a certain attitude, the autopilot will hold that attitude intil other control instructions are received. This "rate command" function is all ccomplished digitally in newer aircraft, but the pilot perceives a graded input that roduces a continuous response. In contrast, the command mode of the autopilot s controlled by providing it with digital numeric targets representing airspeed, lesired altitude and heading, and sometimes desired vertical rate. In today's aircraft, hese digital values can be specified either through rotary switches on the mode ontrol panel in a manner quite similar to the selection of new radio frequencies in lder aircraft, or by digital numeric input to the FMS.

The control wheel steering mode can be a trap, as was evidenced in a DC-10 ncident in which, after a close-in turn to final approach, the flying pilot, who was eavily loaded, forgot that he was in that mode, continued to command an ncreasing pitch *rate*, and incurred a tail strike during the subsequent landing NASA ASRS, 1976). It is the normal mode of autopilot control in older Boeing 37-200 series aircraft, however; it permits quick tactical changes to flight path, nd it therefore represents a potentially useful intermediate between fully manual nd fully automatic flight. It is shown as assisted control in my control and nanagement continuum (see chapter 10). In at least one new airplane, the MD-11, ll longitudinal control is carried out through the CWS function of the autoflight ystem, and full-time CWS for lateral (roll) control is also available as a customer-pecified option.

ully Autonomous Automation

ome automation elements have been essentially autonomous for a long time. No ilot would think of hand-flying a jet throughout cruise, as one instance. Many rlines require the use of the autobraking function for all landings, and autospoilers e also used routinely. Several other automatic functions that are used at all times ave been mentioned in chapters 5 and 6. Despite this, concern has been expressed various quarters about more complex functions that are now essentially autono-ous, several of which can be turned off only with difficulty or not at all.

Among these functions is the full-time envelope protection system in the A320, hich in effect prevents pilots from exceeding certain flight control parameters. his could more accurately be called an *envelope limiting* system. Several current d planned aircraft have systems that fulfill similar functions, although in a mewhat different manner. The MD-11's automatic systems control computers, noted in chapter 5, will reconfigure aircraft subsystems autonomously if they nse specific malfunctions in those systems. Systems such as these give rise to questions concerning pilot authority and responsibility (Billings, 1996; Tenney et al., 1995. These questions are discussed in more detail in chapter 10.

Skill Degradation

One potentially serious problem in human–machine systems with highly capable automation is a loss of certain skills by the human when the automation routinely performs tasks that require such skills. This effect has been observed in numerous contexts (e.g., Cooley, 1987; see also chapter 15, case 1). It may be due largely to lack of practice of the particular skill by the human operator, although in certain contexts, other factors may play a part.

Psychomotor skill decrements were observed by pilots transitioning from copilot positions in the DC-10, a fairly automated airplane, to command positions in less automated aircraft such as the 727. After some failures to complete this transition, air carrier training personnel suggested to pilots approaching transition that they should forego the use of the automation for a couple of months prior to transition, in order to obtain more practice in manual control. The pilots took this advice and were able thereafter to complete transition training without difficulty. Note, in this example, that the pilots coming to transition all had extensive flying experience in older, relatively unautomated, aircraft. Their problem was to reacquire skills that they had already possessed in adequate measure before their transition to the more automated DC-10.

The advent of the new generation of highly automated aircraft, and the replacement of the older machines by such airplanes, implies that at some point in the near future, pilots may begin their airline careers flying as first officers on advanced aircraft that incorporate envelope protection and a variety of other control automation. Such automation may include limits on rate of roll, bank angle, pitch rate as a function of speed, gust alleviation, and other functions.

Will pilots who have never had to acquire the finely tuned manual skills that older pilots take for granted be able to demonstrate such skills at an acceptable level if they must transition to another aircraft that lacks these advanced features? Similarly, will they have learned the cognitive skills necessary for unassisted navigation if the flight management software fails? Finally, and perhaps most important given the high reliability of today's aircraft, will they acquire the judgmental skills and experience that alone can enable them to make wise decisions in the face of uncertainty or serious mechanical or environmental problems? At this point, no one knows the answers to such questions, but we do know that it is these skills, collectively called airmanship, that provide the last line of defense against catastrophes in aviation operations.

Similar questions can be asked about some air carriers that effectively require their pilots of advanced aircraft to utilize the automation on a full-time basis. *Flight International*, in its Letters columns, carried a brisk debate on this topic early in 1993: "Excessive reliance on equipment to help pilots fly 'smarter and safer' has become institutionalized to the point of becoming dangerous" (Hopkins, 1993, p. 40). "...I remember being admonished by the chief pilot for daring to hand-fly a raw-data standard instrument departure, and, worse still, for practising enroute VOR tracking by hand flying for 10 min in the cruise" (Laming, 1993, p. 140).

Some operators suggest to their pilots that they should exercise as many options as possible, and that they should fly at each level of automation on a periodic basis, to remain familiar with the systems and to maintain proficiency. Delta Airlines has stated these goals formally in its statement of automation philosophy: "Pilots must be proficient in operating their airplanes at all levels of automation. They must be knowledgeable in the selection of the appropriate degree of automation and must have the skills needed to move from one level of automation to another" (Byrnes & Black, 1993, p. 443). Many airline pilots make it a point to fly at least part of each flight segment manually to maintain their skills, regardless of the policies and preferences of their carriers.

Recall that similar questions were raised with respect to the ability of air traffic controllers, trained only in a full radar environment, to transition to procedural control of air traffic in the event of a complete radar failure. The ability of the FAA System Command Center to offload controllers during such failures has lessened this concern to some extent, but it is still possible for controllers to be grossly overloaded by system contingencies such as occurred after ATC communications and data transfer were suddenly shut down by a massive failure of communications facilities in New York (Lee, 1992), or by a 1-hour total power failure during morning rush hour at Oakland Air Route Traffic Control Center on August 9, 1995.

Crew Coordination

Wiener (1993) discussed crew coordination and resource management in the context of automated aircraft. In his extensive cockpit observations in advanced aircraft (Wiener, 1989), he noted several crew coordination issues (pp. 177–178):

- "Compared to traditional models, it is physically difficult for one pilot to see what the other is doing [on the CDU].... Though some carriers have a procedure that requires the captain (or pilot flying) to approve any changes entered into the CDU before they are executed, this is seldom done; often he or she is working on the CDU on another page at the same time."

- "It is more difficult for the captain to monitor the work of the first officer and to understand what he is doing, and vice versa."
- "Automation tends to induce a breakdown of the traditional (and stated) roles and duties of the pilot-flying versus pilot-not-flying and a less clear demarcation of 'who does what' than in traditional cockpits. In aircraft in the past, the standardization of allocation of duties and functions has been one of the foundations of cockpit safety."
- "There is a tendency for the crew to 'help' each other with programming duties when workload increases. This may or may not be a good thing...but it clearly tends to dissolve the clear demarcation of duties."

Costley, Johnson, and Lawson (1989) found in flight observations in 737 and 757 aircraft that less communication occurred in more advanced cockpits. Wiener interpreted these findings in terms of extremely low workload during cruise in advanced automated aircraft, and expressed concern "because of the presumed vulnerability of crews to boredom and complacency" (Wiener et al., 1993; p. 26). Wiener's findings agree with others reported here: that our traditional models of the behavior of competent air transport pilots may be insufficient guides to behavior in automated aircraft, because the machines themselves are, in certain respects, qualitatively different from older aircraft. New cognitive models that emphasize the increased cognitive loading on pilots are needed to guide our designs and implementation in the future.

We may have been shielded to some extent from problems in this realm by the very high experience levels of many first officers, as well as captains, in today's system. Many former captains with extensive command experience are now flying as copilots after having been laid off by defunct or bankrupt carriers. This will lessen during coming years, however.

Monitoring Requirements

Pilots (and increasingly, air traffic controllers as well) must monitor flight progress closely, for others, human and machine, are monitoring as well, to an extent unprecedented in the history of the industry. One problem inherent in automation is that pilots cannot usually detect that it is not going to do what they expected it to do until after it has failed to do it. It is only after automation has "misbehaved" that operators can detect its "misbehavior" and correct it. This is another aspect of the opacity problem. Unfortunately, when this occurs in aviation, the airplane may already be in a position from which rapid reactions may be necessary to return it to nominal conditions.

During an idle power descent, an airplane may descend 50 ft during each second it takes the crew to recognize an anomaly, decide to take action, make a control input, and wait for an appropriate response. Aircraft are separated by only 1,000 ft vertically below 29,000 ft; deviations of 500 ft or more are not uncommon after an autopilot has failed to capture an altitude. Such a deviation can be easily observed by air traffic control personnel and, if there is a conflict, by ATC automated conflict alert software. If the deviation is reported, pilots may face disciplinary or enforcement action from FAA.

For these reasons as well as others, pilots must closely monitor the behavior of their automated systems, but if an anomaly occurs, they must sometimes take very prompt action. Present automation (except the ubiquitous altitude alerting system) provides no predictive or premonitory warning that a failure is likely to occur in the immediate future; such information would give pilots time to prevent, rather than correct, the problem. Fortunately or unfortunately, flight path automation is reliable enough so that pilots may be tempted to relax their guard on the (justified) assumption that it will almost always behave correctly. Moray, Lee and Hiskes (1994) even suggested that this is the logical and appropriate strategy for pilots to adopt, because it is rare for such malfunctions to occur; thus, pilots are better advised to spend more time monitoring aspects of their flight that involve more uncontrolled variability.

Without question, the most effective monitoring of pilots flying is by a nonflying pilot in the same cockpit. This redundancy is absolutely critical. The vast majority of errors in the cockpit are detected, announced, and corrected without adverse consequences, often before any sort of anomaly can occur. When this fails, air traffic controllers often detect and warn of small deviations, permitting the pilot to correct them at an early stage. All of this cross-monitoring assumes that the monitoring agents understand the intent of the monitored agents (see chapter 3). Newer automation can do more than it has thus far been called on to do to strengthen still further the redundancy and thus the error tolerance of the aviation system.

Automated System Navigation Problems

Although manufacturers of the latest flight management systems have gone to considerable effort to simplify the operation of these systems, they are still exceedingly complex, and all interaction with them must be through several displays brought up sequentially on a single small CDU screen containing a large amount of alphanumeric information. As more functions have been implemented, more and more screens have been designed, each requiring serial access by the operator (see Fig. 5.31). In today's system, a great deal of information must be accessed through a very small "keyhole." As a consequence, navigating among the many screens has become complicated. This requirement imposes yet another cognitive burden on operators, who must remember enough of the FMS architecture to recall how to get to specific information when it is needed.

One method that designers have utilized to lessen the memory burden is to increase the number of modes in the FMS itself. This simplifies the navigation problem within the FMS but increases the requirement to remember the various modes and what each is used for. As these remarkable devices become still more capable, this cognitive burden imposed by the need for mode awareness can be expected to increase, unless a different approach is taken to their design (Woods et al., 1994).

Data Overload

Automation and the glass cockpit have increased considerably the amount of information available to pilots. The information is of much higher quality than was available in the past, a blessing for it decreases ambiguity and uncertainty, but the quantity imposes much higher attentional demands than in the past. The flight navigation displays on today's panels integrate a great deal of data into a clear and intuitive representation of the aircraft's location, directional trend, and chosen course—but this screen may also contain data regarding severe weather, wind shears, waypoints, airfields, obstructions, and other traffic, almost none of which was explicit in earlier aircraft. Depending on the circumstances of the flight, any part or all of this information may be relevant. Much of it, fortunately, can be turned off when it is not needed. Nonetheless, pilots must now manage a potential glut of information, where in the past, they simply had to wonder about it.

Pilots have often demonstrated that they want access to *all* information that may be relevant to their decision processes in flight, and that they are willing to accept a higher workload to deal with it. Unfortunately, as Fadden noted, if they have too much data, it become less certain that they will be able to attend to and integrate the appropriate data in time to address the problem that is most important. Particularly when virtually all information is visual in form, this is a serious potential problem for designers. Some have suggested adaptive displays that can be automatically decluttered as the pilot becomes more heavily loaded, but this poses other problems relating to operator authority (see chapters 10 and 12).

COMMENT

I have tried here to summarize some attributes of contemporary aircraft automation that appear to have been associated with problems in pilot cognitive behavior. Few of these problems represent failures of the automation as such; most represent either conceptual failures at the design or operator level, or problems in the implementation of these concepts. As machines grow more complex and difficult to understand, operators are more likely to err in their operation, so the net effect of these problems is often seen as human error at the sharp end. As Reason (1990) and Woods et al. (1994) pointed out so clearly, to say this and stop is simply to insure that the latent organizational and other factors that lie behind human error will go unnoticed, and that attempts to insulate the system against such errors will not get at the systemic and conceptual problems that cause most of them.

It is for this reason that I have tried, in this chapter, to generalize from the particular problems cited in earlier chapters to the conceptual issues that appear to me to underlie many or most of those problems. These issues, I believe, are the "latent factors" that we must attack if we are to make aviation automation more human-centered.

I have said little here about problems associated with ATC automation, simply because at this time there is relatively little automation to help the controller perform the primary task of directing air traffic (although much of the data management in ATC is automated). Controllers still work in a largely linear system whose peculiarities and nonlinearities they understand. Perrow (1984) cited the ATC system, in which "interactive complexity and tight coupling have been reduced by better organization and 'technological fixes' " (p. 5).

> The goal of preventing mid-air collisions conflicts with the production demands placed upon the airways system ... The problem... for ATC has been to keep collision risks low while increasing the occasions for collisions. This they have done with remarkable success. The density increases steadily, but the number of mid-air collisions has been reduced to near-zero (especially those where both airplanes are controlled by ATC). (p. 158)

As I said in the preface, it is necessary that we look not only at the human or at the machines, but at the system, if we are to correct system faults or to design and implement more effective systems in the future. If we do not take this approach, our present systems, as tightly integrated as they are, will simply acquire more layers of "Band-Aids" as we attempt to solve specific problems one by one, without considering the effects of those solutions on the system as a whole, or on the competing demands upon both pilots and controllers. I am frankly worried that this may be what we are doing in our present attempts to improve TCAS, a very tightly coupled system, by adding more and more software to lessen nuisance warnings while trying to extend the basic usefulness of the device by placing new requirements on it.

CHAPTER 10

HUMAN AND MACHINE ROLES: RESPONSIBILITY AND AUTHORITY

INTRODUCTION

Much industrial automation has been implemented on the implicit assumption that machines could be substituted for humans in the workplace (chapter 4). The Fitts (1951) list of functions that are best performed by humans and those best performed by machines exemplifies this concept. Jordan (1963) proposed that humans and machines should be considered as complementary, rather than competitive. The design and operation of the modern transport airplane exemplifies the concept of complementarity, yet in certain respects its automation very much exemplifies the principle of the interchangeability of parts. There are good reasons for this, but we must question whether we should still be designing and operating machines in that manner and whether a different approach could solve some of the problems we now observe in the aviation system.

Today's aircraft automation controls an airplane more or less as the pilot does. It navigates as the pilot does, or would if pilots could carry out in real time the complex calculations now performed by the computer. It operates the systems as the pilots do, or would do if they do not forget or overlook any of the procedural steps. In the near future, it will communicate with ATC computers, accept and execute ATC clearances, and report its location when not under radar coverage, just as pilots do now. Some have noted that automation usually performs all of these functions correctly, that it does not become tired or distracted or bored or irritable, that it often speaks more clearly and succinctly than pilots do, that its data stream will be easily comprehended by ATC computers in any nation, and that it does all these things without complaints. They have concluded that automation is as capable as the human for these functions, and some air carriers have mandated that it be used whenever possible. Are these "parts" interchangeable? That is the subject of this chapter.

THE PILOT AS CONTROLLER AND MANAGER

It should be clear from chapters 5 and 6 that pilots may play any of a variety of roles in the control and management of a highly automated airplane. These roles range from direct manual control of flight path and aircraft systems to a largely autonomous operation in which the pilot's active role is minimal. This range of allocation of functions between human and machine can be expressed as a control–management continuum, as shown in Fig. 10.1.

None of today's aircraft can be operated entirely at either end of this spectrum of control and management. Indeed, a complex airplane operated even by *direct manual control* may incorporate several kinds of control automation such as yaw dampers, a Mach trim compensator, automated configuration warning systems, and so forth. Conversely, even remotely piloted vehicles are not fully autonomous; the locus of control of these aircraft has simply been moved to another location.

Most transport flying today is *assisted* to a greater or lesser extent, by hydraulic amplification of control inputs and often by computer-implemented flight control laws. Flight directors, stability augmentation systems, enhanced displays, and, in newer aircraft, various degrees of envelope protection assist the pilot in his or her manual control tasks. To some extent, pilots can specify the degree of assistance desired, but much of it operates full-time and some of it is not intended to be bypassed. The pilot remains in the control loop, although it is an intermediate rather than the inner loop.

VERY LOW ← LEVEL OF AUTOMATION → VERY HIGH

VERY HIGH ← LEVEL OF INVOLVEMENT → VERY LOW

MANAGEMENT MODE	AUTOMATION FUNCTIONS	HUMAN FUNCTIONS
AUTONOMOUS OPERATION	Fully autonomous operation; Pilot not usually informed; System may or may not be capable of being disabled.	Pilot generally has no role in operation Monitoring is limited to fault detection Goals are self-defined; pilot normally has no reason to intervene.
MANAGEMENT BY EXCEPTION	Essentially autonomous operation; Automatic reconfiguration. System informs pilot and monitors responses.	Pilot informed of system intent; Must consent to critical decisions; May intervene by reverting to lower level of management.
MANAGEMENT BY CONSENT	Full automatic control of aircraft and flight. Intent, diagnostic and prompting functions provided.	Pilot must consent to state changes, checklist execution, anomaly resolution; Manual execution of critical actions.
MANAGEMENT BY DELEGATION	Autopilot & autothrottle control of flight path. Automatic communications and nav following.	Pilot commands hdg, alt, speed; Manual or coupled navigation; Commands system operations, checklists, communications.
SHARED CONTROL	Enhanced control and guidance; Smart advisory systems; Potential flight path and other predictor displays.	Pilot in control through CWS or envelope-protected system; May utilize advisory systems; System management is manual.
ASSISTED MANUAL CONTROL	Flight director, FMS, nav modules; Data link with manual messages; Monitoring of flight path control and aircraft systems.	Direct authority over all systems; Manual control, aided by F/D and enhanced navigation displays; FMS is available; trend info on request.
DIRECT MANUAL CONTROL	Normal warnings and alerts; Voice communication with ATC; Routine ACARS communications performed automatically.	Direct authority over all systems; Manual control utilizing raw data; Unaided decision-making; Manual communications.

FIG. 10.1. A continuum of aircraft control and management for pilots.

Whether pilots of limited experience should be required to have and demonstrate direct manual control ability in today's airplanes, which incorporate highly redundant automated control assistance, is a reasonable question but beyond the scope of this document. Airbus has rendered this issue moot to some extent by providing *shared control* as the A320's basic control mode. Pilots' control inputs are considerably modified and shaped by the flight control computers; envelope limits prevent them from exceeding predetermined parameters. In this airplane, pilots are provided with considerable assistance even during control failure modes; true manual flight capability is limited to rudder control and horizontal stabilizer trim and is designed only to maintain controlled flight while the automated systems are restored to operation. Under normal circumstances, the aircraft automation is responsible for much of the inner-loop control, although control laws are tailored to respond in ways that seem natural to the pilot. In the MD-11, a combination of longitudinal stability augmentation and control wheel steering is in operation at all times.

When an autopilot is used to perform the flight path (or power) control tasks, the pilot becomes a manager rather than a controller (this is also true to some extent of the shared control option). The pilot may elect to have the autopilot perform only the most basic functions: pitch, roll, and yaw control (this most basic autoflight level is no longer available in all systems); he or she may command the automation to maintain or alter heading, altitude, or speed, or may direct the autopilot to capture and follow navigation paths, either horizontal or vertical. This is *management by delegation*, although at differing levels of management, from fairly immediate to fairly remote. In all cases, however, the aircraft is carrying out a set of tactical directions supplied by the pilot. It will not deviate from these directions unless it is incapable of executing them.

As always, there are exceptions to the generalizations. Several aircraft will not initiate a programmed descent from cruise altitude without an enabling action by the pilot. Other modern flight management systems require that the pilots provide certain inputs before they will accept certain conditional instructions. *Management by consent* describes a mode of operation in which automation, once provided with goals to be achieved, operates autonomously, but requires consent from its supervisor before instituting successive phases of flight, or certain critical procedures. The consent principle has important theoretical advantages, in that it keeps pilots involved and aware of system intent, and provides them the opportunity to intervene if they believe the intended action is inappropriate at that point in time.

This management mode may become more important as intelligent decision-aiding or decision-making systems come into use (see chapter 12). A protracted period of close monitoring of these systems will be necessary; requiring consent is one way to monitor and moderate the potential influence of these systems. Although management by consent is an attractive option worthy of further exploration, it must be *informed* consent. More fundamental human factors research is needed to identify how to implement it without the consent becoming perfunctory.

Management by exception refers to a management/control mode in which the automation possesses the capability to perform all actions required for mission completion and performs them unless the pilot takes exception. Today's very capable flight management systems will conduct an operation in accordance with preprogrammed instructions unless a change in goals is provided to the flight management system and is enabled by the pilots. Such revisions occur relatively frequently when air traffic control requires changes in the previously cleared flight path, most often during descent into a terminal area. Some FMS lateral and waypoint management tasks now operate by exception.

The desire to lighten the pilot's workload and decrease the required bandwidth of pilot actions led to much of the control automation now installed in transport aircraft. The more capable control and management automation now in service has certainly achieved this objective. It also has the capacity, however, to decrease markedly the pilot's involvement with the flying task and even with the mission. Today's aircraft can be operated for long periods of time with very little pilot activity. Flight path control, navigation, and in some aircraft subsystems management are almost entirely automatic. The capable, alert pilot will remain conversant with flight progress despite the low level of required activity, but even capable, motivated pilots get tired, lose their concentration and become diverted, or worry about personal problems unrelated to the flight. A critical task for designers is to find ways to maintain and enhance pilot involvement during operation at higher levels of automation.

This is less simple than it sounds, for pilots will both resent and find ways to bypass tasks that are imposed merely for the purpose of ascertaining that they are still "present in the cockpit." Tasks to maintain involvement must be flight-relevant and, equally important, must be *perceived by pilots* to be relevant. Designing pilot involvement into highly automated systems will not be easy but must be accomplished to minimize boredom and complacency, particularly in very-long-range aircraft that spend many hours in overwater cruise. The progress of avionics, satellite navigation and communications, and data link will very likely have an opposite effect unless this uniquely human factor receives more consideration than it has to date.

Fully autonomous operation denotes operation in accordance with instructions provided by system designers; no attention or management is required of the pilots. Until recently, relatively few complex systems operated fully autonomously. With the introduction of the A320 and MD-11, however, major systems operate in this way.

A fundamental question is how wide a range of control and management options should be provided. This may well vary across functions; indeed, pilots often prefer to operate using a mix of levels, for example, controlling thrust manually while managing the autopilot and using the flight director to monitor navigation. Pilot cognitive styles vary; their skill levels also vary somewhat as a function of the amount of recent flying they have done, how tired they are, and so on. These factors lead me to argue that a reasonable range of control/management options must be provided, but widening that range is expensive in terms of training time and time required to maintain familiarity with a broader spectrum of automation capabilities, as well as in terms of equipment costs.

One possible way to keep pilots involved in the operation of an aircraft is to limit their ability to withdraw from it by invoking very high levels of management. Another, perhaps preferable way is to structure those higher levels of management so that they still require planning, decision making, and procedural tasks. The use of a management by consent approach, rather than management by exception, could be structured to insure that pilots must enable each successive flight phase or aircraft change of status, as an instance. It has been suggested by one air carrier that long-haul pilots should be given the tools with which to become involved in flight planning for maximum economy on an ongoing basis; this is another approach to maintaining higher levels of involvement, but it is presently being implemented as a dispatcher/AOC function.

THE ROLE OF THE AIR TRAFFIC CONTROLLER

When a more highly automated ATC system is implemented, its computers will be able to search for traffic conflicts and to provide at least decision support in resolving them. This is the foundation of the FAA's automated en route air traffic control system (formerly referred to as AERA), and it is a key feature of the free flight proposal (chapter 8). Direct ATC computer-to-flight management computer data transfers, and direct negotiations between these computers, will likewise be a part of such a system, which opens the possibility of direct control of air traffic by ATC automation without involvement of either controllers or pilots.

I have discussed a control–management continuum in terms of pilot roles in an automated system. A similar construct can be proposed for air traffic controllers and their automation (Fig. 10.2), although it should be kept in mind that air traffic controllers actually *direct* and *coordinate* the movements of aircraft; only pilots control them. In this respect, the controller's task is fundamentally different from that of the pilot.

VERY LOW ← LEVEL OF AUTOMATION → VERY HIGH

MANAGEMENT MODE	AUTOMATION FUNCTIONS	HUMAN FUNCTIONS
AUTONOMOUS OPERATION	Fully autonomous operation; Controller not usually informed. System may or may not be capable of being bypassed.	Controller has no active role in operation Monitoring is limited to fault detection. Goals are self-defined; controller normally has no reason to intervene.
MANAGEMENT BY EXCEPTION	Essentially autonomous operation. Automatic decision selection. System informs controller and monitors responses.	Controller is informed of system intent; May intervene by reverting to lower level.
MANAGEMENT BY CONSENT	Decisions are made by automation. Controller must assent to decisions before implementation.	Controller must consent to decisions. Controller may select alternative decision options.
MANAGEMENT BY DELEGATION	Automation takes action only as directed by controller. Level of assistance is selectable.	Controller specifies strategy and may specify level of computer authority.
ASSISTED CONTROL	Control automation is not available. Processed radar imagery is available. Backup computer data is available.	Direct authority over all decisions; Voice control and coordination.
UNASSISTED CONTROL	Complete computer failure; No assistance is available.	Procedural control of all traffic. Unaided decision-making; Voice communications.

VERY LOW → LEVEL OF INVOLVEMENT ← VERY HIGH

FIG. 10.2. A continuum of system control and management for air traffic controllers.

As in the case of pilots, a broad range of roles is theoretically possible, ranging rom unassisted procedural control without visualization aids such as radar, all the vay to autonomous machine control of air traffic. Indeed, the former option will robably continue in some parts of the world, even while other areas adopt dvanced automation. The important point is that the role of the controller can ary greatly, from absolute direct authority over the entire operation to a relatively assive oversight function in which air traffic control tactics are purely the comuter's task.

Whether such a broad range of roles is desirable is another matter entirely. The rst principles of human-centered automation indicate that involvement is necessary f the human operator is to remain in command of the operation. I question the ontroller's ability to remain actively involved for very long if he or she has no active ole in the conduct of an almost entirely automated process. On the other hand, ome range of options should be permitted, to account for differences in cognitive yle, variations in workload, and a wide range of controller experience levels.

The potential for increased opacity of new air traffic management automation is reat, as indicated in chapters 8 and 13. Controllers cannot remain in command of ir traffic unless they are both informed and involved, not only when automation ils but when it is performing normally. This again argues against placing the human perator in a role at the high extreme of this control–management continuum.

HUMAN AND MACHINE ROLES

Present aircraft automation does not plan flights, although it is able to execute them and to assist in replanning (e.g., after an engine failure). It cannot configure an airplane for flight or start the engines. It knows with great precision where runways are, but not how to get to them from a gate, nor from a runway turnoff to a gate after landing. Automation does not, at this time, accomplish the checklists required before and during flight. Flight control automation is locked out during the takeoff sequence, although thrust is under automatic control from early in the process in some aircraft. Automation controls neither the landing gear nor the flaps during takeoff and approach. From shortly after takeoff until the airplane touches down at a destination, however, automation is fully capable of executing virtually all the required tasks in a flight.

There is no reason, of course, why automation could not perform taxi maneuvers, although implementing this function would be extremely costly. There is absolutely no reason why landing gear and flap actuation could not be automatic. The few aspects of subsystem management that are still manual in some of the newer aircraft (e.g., the MD-11) could certainly be automated as well. Why, then, have they not been? The answer does not lie in the inadequacies of technology, but in the intricate domains of sociology, psychology, and politics.

Pilots are perceived to be essential because passengers are not willing to fly in an autonomous, unmanned airplane—although millions entrust themselves every day to the Bay Area Rapid Transit, the Washington Metro, and other mass transit systems in which the locus of control has shifted from the on-board operator station to a central control room. The trains on these systems do carry a human operator, but under normal circumstances, the operator does not operate the vehicles and is proscribed from doing so. Airport "people-movers," some of which travel over several miles of dedicated track or roadway, do not have on-board operators; the voice announcements are recorded or synthetic. Note that these systems are not fully autonomous; humans control them, as they always did, but the control is supervisory and remote (Sheridan, 1984).

The flight environment, however, is far more complex and variable than that of a modern light-rail system, and many of the variables are not under the control of system managers. Pilots are essential because they are trained to compensate for unexpected variability. Automation does fail, and unlike surface vehicles, airplanes cannot simply come to a stop while the automation is fixed. Once in flight, they must be guided to a landing. In other words, pilots and air traffic controllers are essential because they are able to make good decisions and take appropriate actions in difficult situations. We have not yet devised a computer that can cope with the variability inherent in the flight and air traffic environment.

The human role, then, is to do what the automation cannot do: to plan, to oversee, to reflect and make intelligent decisions in the face of uncertainty, and to make passengers (and air carrier management, and the FAA) feel comfortable about air transportation.

RESPONSIBILITY AND AUTHORITY

If a controller fails to maintain separation because of a tag swap or a radar outage, is the computer "grounded"? No; the controller remains responsible for traffic separation regardless of the circumstances. There may be mitigating circumstances, but this responsibility cannot be delegated.

If an automated airplane gets lost and lands at the wrong airport, or runs out of fuel and crash lands, or violates regulations for whatever reason, is the flight management computer held to account? Not to my knowledge. The pilot, not the autopilot, is in command of the flight and is responsible for its safe conduct.

Does the pilot have the authority required to fulfill this responsibility? What responsibility, and how much authority, does the pilot have in today's system and today's airplanes? It is a maxim of military command that authority can be delegated by a commander. Responsibility for the outcome cannot be delegated to others. It remains with the commander.

These precepts are extremely important in aviation. Although aviation involves a widely distributed system in which no individual can get the job done alone, the roles of all the humans in the system come together in the process of flight. In that process, the pilot and dispatcher share responsibility for the plan that guides the flight. The pilot is solely responsible for its safe execution, and the air traffic controller is solely responsible for keeping the flight safely separated from other air traffic.

Part 91.3 of the Federal Aviation Regulations (Federal Aviation Administration) describes the responsibility and authority of the pilot in command. It is brief and succinct:

> (a) The pilot in command of an aircraft is directly responsible for, and is the final authority as to, the operation of that aircraft.
>
> (b) In an in-flight emergency requiring immediate action, the pilot in command may deviate from *any rule* [italics added] of this part to the extent required to meet that emergency.
>
> (c) Each pilot in command who deviates from a rule under paragraph (b) of this section shall, upon the request of the Administrator, send a written report of that deviation to the Administrator.

This regulation confers upon the pilot essentially unlimited authority to depart from the accepted rules for the conduct of flights if that pilot believes that an emergency exists. Under emergency authority, the pilot is permitted to request whatever assistance is necessary, to declare for his or her flight absolute priority for any maneuver, flight path, or action, and to take whatever steps are necessary, in the pilot's view, to protect the passengers. The pilot's decisions may be questioned afterward, but the authority remains and is recognized without question at the time.

It is a matter of record that pilots have sometimes not used their emergency authority when hindsight says they should have done so. Some situations, like the undeclared fuel emergency that led to the loss of Avianca flight 107 (Cove Neck, New York, 1990), seem obvious to anyone, although the NTSB raised the question of whether the pilot's very limited English competency may have permitted him to think that he had made such a declaration when the proper enabling words ("Mayday" or "Emergency") were not used. In other cases, pilots have been inhibited by fear of the paperwork and questions that inevitably follow such a declaration (although onerous questions after a safe landing are a great deal easier to walk away from than an aircraft accident).

Pilots, then, have as much authority as they need to permit them to fulfill their responsibility for flight safety—or do they? Does a pilot whose control authority is limited by software encoded in the flight control computer have full authority to do whatever is necessary to avoid an imminent collision, or ground contact? United States transport aviation involving jet aircraft was scarcely 4 months old in 1959 when a Boeing 707 entered a vertical dive over the North Atlantic Ocean. The pilots recovered from the dive and landed the airplane safely at Gander, Newfoundland. Postflight inspections revealed severe structural damage of the wing and horizontal stabilizer, but all the passengers survived and the airplane flew again after major repairs (NTSB, 1960). Would this have been possible if flight control software had limited the forces that could be applied to levels within the permitted flight envelope of the airplane?

LIMITATIONS ON PILOT AUTHORITY

In the A320/330/340 series aircraft, the flight control system incorporates envelope limitation. Certain parameters (bank angle, pitch, or angle of attack) cannot be exceeded by the pilot except by turning off portions of the flight control computer systems or flying outside their cutoff values, as was done during the low-altitude flyover prior to the Mulhouse-Habsheim accident (1988). Predetermined thrust parameters also cannot be exceeded.

Systems designed for autonomous operation pose serious philosophical questions with respect to pilot authority as well as pilot involvement. These questions arose first in the design of fighter aircraft and were discussed succinctly in an unsigned editorial in *Flight International* ("Hard limits, soft options," 1990). The American F-16 fighter's fly-by-wire control system incorporates "hard" limits that "preserve the aircraft's flying qualities right to the limit of its closely defined envelope" but do not permit the pilot to maneuver beyond those limits. The *Flight* editorial pointed out that

> There is, however, another approach available: to develop a "softer" fly-by-wire system which allows the aircraft to go to higher limits than before but with a progressive degradation of flying qualities as those higher limits are approached. It is this latter philosophy which was adopted by the Soviets with fighters like the MiG-29 and Sukhoi Su-27. It is not, as Mikoyan's chief test pilot ... admits, "necessarily a philosophy which an air force will prefer." [He] says, however: "Although this ... approach requires greater efforts ... it guarantees a significant increase in the overall quality of the aircraft-pilot combination. This method also allows a pilot to use his intellect and initiative to their fullest extent." (p. 3)

The softer approach was taken in the MD-11 (Hopkins, 1990) and Boeing 777, which permit pilots to override automatic protection mechanisms by application of additional control forces. The flying qualities are degraded under these circumstances, but the pilots retain control authority. The MD-11 incorporates angle of attach protection, as do the A320/330/340, but the MD-11's limits can be overridden by the pilot, as can the limits of its autothrust system. The 777 incorporates hard limits on engine power, for reasons that are not clear to me. In the MD-11, as noted earlier, many aircraft systems operate autonomously; subsystem reconfiguration after failures is also autonomous if the ASCs are enabled (the normal condition). Any of these systems can also be operated manually, but the protections provided by the ASC computers are not available during manual operation.

Although civil aircraft do not face the threat posed to a fighter under attack if its maneuverability is limited, their pilots do on occasion have to take violent evasive or corrective action, and they may on rare occasions need control or power authority up to (or even beyond) normal structural and engine limits (e.g., Pacific Ocean, 1985) to cope with very serious problems. The issue is whether the pilot, who is ultimately responsible for safe mission completion, should be permitted to operate to or even beyond airplane limits when he or she determines that a dire emergency requires such operation. The counterargument, raised elsewhere in this book, is of course that some pilot, sometime, for some reason, will exceed these limits unnecessarily, raising an equal safety hazard. This issue will not be simply resolved, and the rarity of such emergencies makes it difficult to obtain empirical support for one or the other philosophy. Nonetheless, the issue is a fundamental one.

COMMENT

The increasing capabilities of advanced automation pose a severe temptation to new aircraft design teams. They can decide that the safety of the airplane makes it important that they limit the authority of the pilots, and they can implement that limitation very easily in airplane software. They could match the software limits to the structural parameters of the airplane insofar as possible, although this is an approach that has not yet been implemented. Whether they have considered all of the circumstances that may confront a pilot in line operations is a question that may only be answered when totally unforeseen circumstances arise, perhaps years after the airplane has left the factory. It is in such circumstances that the "parts" of the human–machine system will be found to be not interchangeable but complementary.

Given that pilots bear the ultimate responsibility for the outcome, it would seem that their authority to do whatever is necessary to insure that the outcome is favorable should be foreclosed only with extreme reluctance. The concept of "soft limits" on control authority may represent one useful and constructive approach to this dilemma. What is important is to realize how easily the pilot's authority can be compromised, given the technologies that are now available. It may take only a line or two of software and may or may not be known or obvious to the pilot. (See Mårtensson, 1995, for a trenchant example of a critical automated function of which the pilots were unaware.)

The same dilemma will face us in the near future with respect to air traffic controllers, as the tools they use are automated in the AAS. This question has not received the attention it deserves, and the rarity of situations that define the issue makes it very difficult to provide good data in support of any extreme position. It is necessary that we realize, however, that issues involving such rare events must sometimes be handled on the basis of the best available a priori reasoning. The views of pilots and controllers on this issue are clear: If they have the responsibility, they want knowledge and the authority necessary to remain in command.

CHAPTER 11

INTEGRATION AND COUPLING IN THE FUTURE AVIATION SYSTEM

INTRODUCTION

The technical challenge of integrating advanced automation in aircraft pales in the face of the challenge posed by the need for a highly integrated air traffic management system. Simply developing a set of agreed-on standards for such a system has already taken 5 years, and the task is far from finished. FAA, ICAO and other organizations must produce standards and requirements for data link technologies, the aeronautical telecommunications network, automatic dependent surveillance, future ATC procedures, satellite surveillance, navigation and communications, ground communication links, integration of satellite and radar surveillance, the necessary airborne equipment, and assessment of the problems posed by a mix of airborne capabilities (Fitzsimons, 1993). "Harmonizing" all of the pieces needed for a truly integrated aviation system will be a staggering task.

Despite, or perhaps because of, the magnitude of this task, many system planners have proposed that the ATC and aircraft computers in the aviation system can function more effectively if they can exchange data directly and can negotiate clearance revisions automatically. This concept would tightly couple the various system elements.

Perrow (1984) discussed automation complexity and coupling at length. He characterized tightly-coupled systems as having more time-dependent processes, requiring invariant sequences, "unifinality" (little flexibility regarding ways to reach a goal state), little slack, and limited to those buffers and redundancy built in to the system. He pointed out the many nonlinearities in such systems, and the difficulty in modeling them. He also noted that more ATC automation "will lead to much tighter coupling—that is, less resources for recovery from incidents" (p. 161). In accordance with Perrow's cautions, I examine issues related to coupling and complexity, as well as integration, in this system.

The UK National Air Traffic Service (NATS) has supported studies to ensure that a variety of technologies can "play together" in a future environment. In October 1991 Eurocontrol and the UK CAA demonstrated the automatic delivery of clearance data, weather interrogation by pilots, and the transmission of ATC instructions and pilot acknowledgements using a BAC 1-11 airplane belonging to the Royal Aircraft Establishment.

> Downlinked autopilot settings were automatically checked against the controllers' original instructions, enhancing safety, while the downlinking of other avionics data (such as true airspeed, heading and vertical rates) reduced voice traffic and the controller's workload. The Volmet [weather] messages were printed in the cockpit, reducing the pilots' workload, and the downlinking of pilot acknowledgements gave the controllers assurance that the message had reached the correct recipient and was unlikely to be misinterpreted. (Fitzsimons, 1993, p. 23)

In 1991, I proposed that ATC clearances transmitted to aircraft by datalink be downlinked to ATC computers as they were executed, to provide confirmation of FMS and presumably pilot intentions and to provide positive confirmation that the aircraft would proceed in accordance with ATC intentions (Billings, 1991).

"Studies suggest that aircraft-derived data could provide additional inputs to ground-based trackers, reducing position uncertainty and enabling improved conflict alert algorithms to reduce the number of nuisance alerts while giving earlier warning of potential conflicts" (Fitzsimons, 1993, p. 23). Although limited, the UK experiments represent an encouraging start on the task of integrating the ground and airborne components of modern aviation systems. Since 1991, a number of other demonstrations have been conducted to examine various elements of an integrated system. In this chapter, I examine the implications of creating such a system for the humans who must operate within it.

ELEMENTS OF AN INTEGRATED AVIATION SYSTEM

A very large number of functional capabilities must be in place in a future aviation system if it is to accomplish the tasks assigned to it. Briefly, these functions are to facilitate the movement and tracking of large numbers of variably equipped aircraft on or over any part of the earth's surface, to assist them in landing and taking off from airports, and to provide all assistance necessary during contingency operations. These tasks must be accomplished in all extremes of weather, across national boundaries, with limited resources. The aviation system is information bound, and the complexity of the system results largely from the complexity of moving all necessary information in real time to all system participants who have a need for it.

Avionics data have been downlinked and processed automatically during the UK NATS mode S trials. At certain airports, predeparture clearance delivery is now routinely accomplished by ACARS. Two carriers successfully tested automatic dependent surveillance over the Pacific, transmitting data through satellites to air traffic control facilities on land. Other elements of the system have also been tested in simulation; some have had flight trials. Large-scale global positioning system testing has been performed, and A330 and A340 aircraft have been certificated for satellite navigation by the JAA in Europe.

There appear to be no insurmountable technological barriers to the implementation of the technologies required for a more highly integrated system. The barriers that remain are in the areas of standards, procedures, software, and harmonization across nations. The knotty issue of how ATC will cope with a broad mix of aircraft capabilities is more difficult in a constrained economic climate. ICAO's Required Navigation Performance concept may help to some extent, although retrofit of advanced equipment in a large number of regional and commuter aircraft may not be economically possible in the near term.

The software issue is critical; the elements of this system must be able to communicate, and the design and verification of software to make this happen throughout the system will be immensely difficult tasks. The AAS system will incorporate more than 2 million lines of code; a system for the ground support of free flight is likely to be substantially more complex. A long period of debugging will be required. Some verification work may not be able to be performed until the system is on line with live traffic, for the new system may be difficult to integrate with the present one. The overall system will be extremely complex, distributed across a great many nodes. Integration of such a system is far different from integration of the many control and display modules in even as complicated a process control system as a nuclear power plant.

COUPLING AND COMPLEXITY

In our present aviation system, the various automation elements are not necessarily coupled except by their information content. That is to say that the various elements operate *independently*. The coupling among them (more properly, the integration) is procedural: It is agreed among the various system participants that on receipt of a given instruction or request, a system component will take certain actions. The results of those actions may be visible in many parts of the system, but they are not predetermined. Although the various system components may be very complex in and of themselves, they are not physically or virtually linked at this time.

As noted, most officials in the air traffic system and an increasing number in the air carrier technical community envision direct communications between ATC

computers (and perhaps, in the future, AOC computers as well) and aircraft flight management computers, although it is generally accepted at this time that when clearance modifications are uplinked to an aircraft, they will be subject to consent by the pilots. Direct negotiation of such clearance modifications between computers is also envisioned by FAA, however, and forms a part of the free flight concept (IATA, 1994). Such a process could confront both controllers and pilots with decisions arrived at by processes that were opaque to them.

It is also planned to require acceptance of datalinked messages within a certain short time interval (40 sec has been mentioned), although presumably execution of an uplinked clearance could be delayed for some further period of time to permit more review by the pilots. Nonetheless, the clearance execution process can be time-critical under some conditions.

These proposals present potentially serious problems for human operators. It is not always easy to understand a complicated clearance, particularly if it involves waypoints or instructions that depart from expectations. The process may require, for instance, that the pilots consult navigation charts, their dispatchers, or the FMS map display, even though the FMC may contain sufficient data to comply with the clearance. A new clearance may not comport with the pilot's view of the environment; it may require the expenditure of extra fuel or may take the airplane too close to the limits of an operating envelope. These factors will sometimes require deliberation and decision making by the flight crew, which will take time.

Executing an uplinked flight plan is simple, requiring only a single keystroke on the FMS CDU. If procedures for voice or data link negotiations with ATC to secure a revision are difficult or time-consuming, a flight crew already busy with another problem may not have the time and may accept an undesirable clearance rather than argue about it. A controller may also need time to understand a complex recommended clearance revision and may not have the time at that moment due to the pressure of other tasks. These are problems that occur now; they can be dealt with by methods similar to those used now, but in a future more automated system these problems can be dealt with *only if provision is made for them in its design.*

Figure 11.1 is a schematic representation of the present air traffic control and management process (solid arrows), in which pilots control aircraft by direction from air traffic controllers, with strategic oversight by the ATC System Command Center. Airline dispatchers and airline operations centers may coordinate aircraft movements with the SCC; revised flight plans are worked out with ATC traffic movements units at ATC facilities.

In a future automated air traffic management system such as the AERA system, ATC computers would negotiate necessary flight plan revisions directly with aircraft (FMS) computers (the vertically shaded arrows at left on the diagram). In a free flight system, in contrast, the locus of control of the system would normally reside

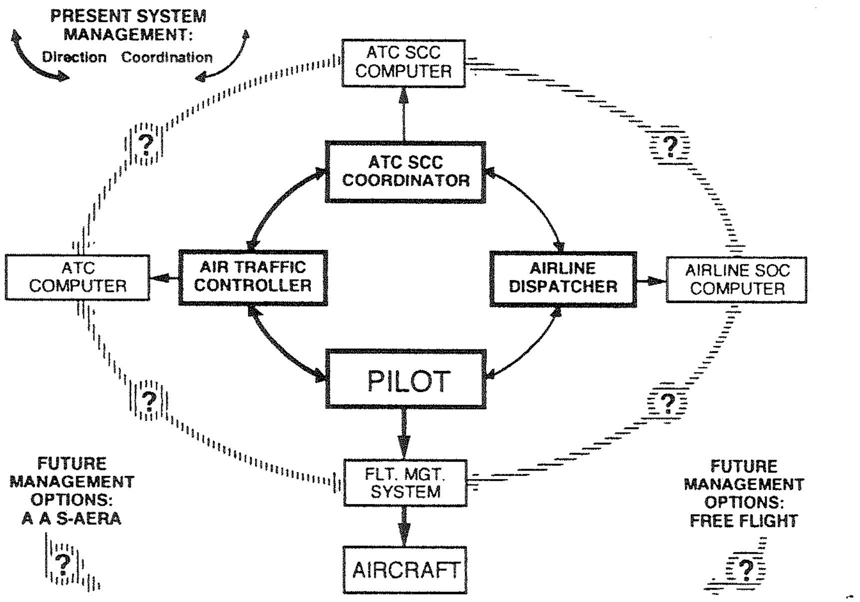

FIG. 11.1. Current process of air traffic control (inner ellipse), and future air traffic control options (AERA and Free Flight; outer ellipse).

in AOCs (horizontally shaded arrows at the right). The AOCs would communicate with aircraft FMSs and with the SCC or air traffic management facilities, which would monitor the flow of traffic for near-term conflicts. The communications technology for this system now exists. Together with data link, new flight management systems for the Boeing 737-700/800 will enable AOCs "to control data-link information from the ground. Operations could request information from the aircraft computer without assistance [from] or knowledge of the pilots" (Nordwall, 1995, p. 48).

At this time, it is envisioned that flight paths and routings will be executed only with pilot and controller consent. Future technology will be able to perform these functions autonomously, however, and automatic execution of such clearances might assist ATC by insuring prompt responses to ATC commands. In this case, the ATC and airborne components of the system would be *coupled* as well as integrated.

In an automated ground-centered system, airplane flight paths would be managed by exception rather than consent (pilots would presumably still be able to countermand the actions of the FMS, although they might not necessarily be given advance notice of its intent). The pilot's role in such a hypothetical situation begins to resemble the monitoring role of the air traffic manager under the free flight concept.

THE AUTOMATED AIR TRAFFIC MANAGEMENT SYSTEM CONCEPT

As noted in chapter 8, NASA and FAA are presently defining the elements of a new automated air transportation technology (AATT) research and development program devoted to advanced air traffic management. The objective of the program is to develop advanced "conflict-free, knowledge-based automated systems for real-time adaptive scheduling and sequencing, for global flow control of large numbers and varieties of aircraft, and for terminal area and ground operations that are compatible with 'free flight' enroute operations" (J. V. Lebacqz, personal communication, October 1994). This system will involve much tighter coupling, not merely integration, of the ground and airborne elements of the aviation system by virtue of the tighter linking of air and ground computers. These concepts run a very real (and very high) risk of infringing significantly on the authority of both air traffic controllers and pilots, despite their proponents' claims that the new automation will be human-centered.

Issues Raised by Tightly Coupled Systems Concepts

In a much more tightly coupled hypothetical system involving intercomputer negotiation and automatic execution of clearances, it would unquestionably be more difficult for pilots to understand how a clearance was arrived at and why it was given, because they would not have access to the ATC computer's reasoning. Similarly, it would be much more difficult for responsible controllers to understand the rules by which the clearance was derived, because they would not have access to the FMS data. This is the complexity-coupling problem discussed by Perrow (1984). It would certainly result in more surprises for the human operators, and would seriously diminish their ability to develop mental models of the ATC automation.

Although cruise flight is a comparatively low workload period for pilots of advanced aircraft, it is quite likely that the cognitive burdens, and workload, now placed on en route controllers will be transferred to pilots, not mitigated, if a free flight concept comes to fruition. This has happened before, when profile descents were imposed in busy terminal areas. Controllers found their workloads lightened by the new procedures, but pilots found their task demands to be considerably increased.

At this time, pilots do not have in their cockpits the information necessary to permit them to accomplish "air traffic control" other than short-range conflict avoidance using TCAS, which provides less than a fully adequate representation even of immediate potential threats (Fig. 6.7). Despite their limitations, which are considerable, TCAS displays are now being used on a test basis for in-trail climb separation over the Pacific Ocean. Other uses, to include lateral separation during closely spaced (1700 ft) parallel approaches to landing, are being actively considered (FAA, 1994), and displays for this function are in development. Note that none of this new functionality has been integrated into the cockpit task flow, nor have the displays and tasking been looked at in the larger context of cockpit and flight management, as so often happens when new functions are considered for retrofit on present flight decks. (See also discussion of FMS in chapter 5.)

COMMENT

Removing pilots or controllers from the command loop, under constrained conditions, would be a comparatively small step from a technical viewpoint. It would represent, however, a *qualitative* change in the rules by which the aviation system has been governed throughout its history. It would diminish the authority of the human operators appreciably, and it would change the dominant mode of system management as much as would the free flight concept. It would, however, be technologically feasible and implementable, and it could result in decreased workload for either pilots or controllers or possibly both—for which reasons, it will probably be seriously considered at some point in the future. This is the reason I have chosen to raise the specter here.

The differences between *integration of independent systems* and *coupling of interdependent systems* need to be clearly understood. The disadvantage of an uncoupled system is that its elements may, or may not, always behave predictably when particular instructions are issued. A pilot may turn too slowly after receiving a controller's request for an immediate maneuver, as an instance (and this is probably more likely when a data linked instruction is received than when a controller issues an urgent voice instruction). The most significant advantage of an integrated but uncoupled system is that operators are much more likely to understand it, and therefore less likely to be surprised by its behavior.

Given a system as complex as the future aviation system will be, however, attempts to couple its ground and airborne elements will inevitably make it more difficult for operators to predict its behavior, particularly under other than nominal conditions. I believe that this would be quite likely to result in less safe rather than more safe operations.

PART IV

ISSUES FOR FUTURE AVIATION AUTOMATION

The last part of this book deals with some issues facing system designers and operators, primarily in the aviation system but also in other domains. Chapter 12 contains a brief overview and discussion of newer computational concepts and techniques, including artificial intelligence (AI) and expert systems (ES), which have been proposed for use in future aviation system automation. In chapters 13 and 14, I attempt to summarize some "lessons learned" from the studies presented in chapters 1 through 11 and to suggest some requirements and guidance for future aviation system automation. Chapter 15 discusses other domains in which humans and complex computational machinery must solve difficult problems in real time, and considers whether the lessons learned in the aviation domain could be helpful in these domains. Chapter 16 contains some general comments and a brief conclusion.